

# Near Optimality of Finite Memory Feedback Policies in Partially Observed Markov Decision Processes

**Ali Devran Kara**

*Department of Mathematics  
University of Michigan  
Ann Arbor, MI 48109-1043, USA*

ALIKARA@UMICH.EDU

**Serdar Yüksel**

*Department of Mathematics and Statistics  
Queen's University  
Kingston, ON, Canada,*

YUKSEL@QUEENSU.CA

**Editor:** Csaba Szepesvári

## Abstract

In the theory of Partially Observed Markov Decision Processes (POMDPs), existence of optimal policies have in general been established via converting the original partially observed stochastic control problem to a fully observed one on the belief space, leading to a belief-MDP. However, computing an optimal policy for this fully observed model, and so for the original POMDP, using classical dynamic or linear programming methods is challenging even if the original system has finite state and action spaces, since the state space of the fully observed belief-MDP model is always uncountable. Furthermore, there exist very few rigorous value function approximation and optimal policy approximation results, as regularity conditions needed often require a tedious study involving the spaces of probability measures leading to properties such as Feller continuity. In this paper, we study a planning problem for POMDPs where the system dynamics and measurement channel model are assumed to be known. We construct an approximate belief model by discretizing the belief space using only finite window information variables. We then find optimal policies for the approximate model and we rigorously establish near optimality of the constructed finite window control policies in POMDPs under mild non-linear filter stability conditions and the assumption that the measurement and action sets are finite (and the state space is real vector valued). We also establish a rate of convergence result which relates the finite window memory size and the approximation error bound, where the rate of convergence is exponential under explicit and testable exponential filter stability conditions. While there exist many experimental results and few rigorous asymptotic convergence results, an explicit rate of convergence result is new in the literature, to our knowledge.

**Keywords:** POMDPs, nonlinear filters, stochastic control

## 1. Introduction

For Partially Observed Stochastic Control, also known as Partially Observed Markov Decision Problems (POMDPs), existence of optimal policies have in general been established via converting the original partially observed stochastic control problem to a fully observed Markov Decision Problem (MDP) one on the belief space, leading to a belief-MDP. However, computing an optimal policy for this fully observed model, and so for the original POMDP, using classical methods (such as

dynamic programming, policy iteration, linear programming) is challenging even if the original system has finite state and action spaces, since the state space of the fully observed model is always uncountable.

In the MDP theory, various methods have been developed to compute approximately optimal policies by reducing the original problem into a simpler one. A partial list of these techniques is as follows: approximate dynamic programming, approximate value or policy iteration, simulation-based techniques, neuro-dynamic programming (or reinforcement learning), state aggregation, etc. (Dufour and Prieto-Rumeau, 2012; Bertsekas, 1975; Chow and Tsitsiklis, 1991). Saldi et al. (2017) investigated finite action and state approximations of fully observed stochastic control problems with general state and action spaces under the discounted cost and average cost optimality criteria where weak continuity conditions were shown to be sufficient for near optimality of finite state approximations in the sense that optimal policies obtained from these models asymptotically achieve the optimal cost for the original problem under the weak continuity assumption on the controlled transition kernel.

On POMDPs, however, the problem of approximation is significantly more challenging. Most of the studies in the literature are algorithmic and computational contributions (Porta et al., 2006; Zhou and Hansen, 2001). These studies develop computational algorithms, utilizing structural convexity/concavity properties of the value function under the discounted cost criterion. Vlassis and Spaan (2005) provide an insightful algorithm which may be regarded as a quantization of the belief space; however, no rigorous convergence results are provided. Smith and Simmons (2012); Pineau et al. (2006) also present quantization based algorithms for the belief state, where the state, measurement, and the action sets are finite. Zhang et al. (2014) also provides a computationally efficient approximation scheme by quantizing the belief space uniformly under  $L_1$  distance.

For partially observed setups, Saldi et al. (2020, 2017) introduce a rigorous approximation analysis after establishing weak continuity conditions on the transition kernel defining the (belief-MDP) via the non-linear filter (Feinberg et al., 2012; Kara et al., 2019), and show that finite model approximations obtained through quantization are asymptotically optimal and the control policies obtained from the finite model can be applied to the actual system with asymptotically vanishing error as the number of quantization bins increases. Another rigorous set of studies is by Zhou et al. (2008, 2010) where the authors provide an explicit quantization method for the set of probability measures containing the belief states, where the state space is parametrically representable under strong density regularity conditions. The quantization is done through the approximations as measured by Kullback-Leibler divergence (relative entropy) between probability density functions. Further recent studies include Mao et al. (2020); Subramanian and Mahajan (2019). Subramanian and Mahajan (2019) present a notion of approximate information variable and studies near optimality of policies that satisfies the approximate information state property. Mao et al. (2020) analyzes a similar problem under a decentralized setup. Our explicit approximation results in this paper will find applications in both of these studies.

We refer the reader to the survey papers by Lovejoy (1991); White (1991); Hansen (2013) and the recent book by Krishnamurthy (2016) for further structural results as well as algorithmic and computational methods for approximating POMDPs. Notably, for POMDPs Krishnamurthy (2016) presents structural results on optimal policies under monotonicity conditions of the value function in the belief variable.

For our work, we specifically focus on finite memory approximations. With regard to approximations based on finite memory, the following two papers are particularly relevant to our paper:

Yu and Bertsekas (2008) study near optimality of finite window policies for average cost problems where the state, action and observation spaces are finite; under the condition that the  $\liminf$  and  $\limsup$  of the average cost are equal and independent of the initial state, the paper establishes the near-optimality of (non-stationary) finite memory policies. Here, a concavity argument building on a work of Feinberg (1982) (which becomes consequential by the equality assumption) and the finiteness of the state space is crucial. The paper shows that for any given  $\epsilon > 0$ , there exists an  $\epsilon$ -optimal finite window policy. However, the authors do not provide a performance bound related to the length of the window, and in fact the proof method builds on convex analysis. Nonetheless, the constant property of the value functions over initial priors is related to unique ergodicity, and thus the stability problem of non-linear filters, which is a topic of current investigation particularly in the controlled setup.

In another related direction, White-III and Scherer (1994) study finite memory approximation techniques for POMDPs with finite state, action and measurements. The POMDP is reduced to a belief MDP and the worst and best case predictors prior to the  $N$  most recent information variables are considered to build an approximate belief MDP. The original value function is bounded using these approximate belief MDPs that use only finite memory, where the finiteness of the state space is critically used. Furthermore, a loss bound is provided for a suboptimally constructed policy that only uses finite history, where the bound depends on a more specific ergodicity coefficient (which requires restrictive, sample pathwise, contraction properties). In our paper, we consider more general signal spaces and more relaxed filter stability conditions, and establish explicit rates of convergence results. We also rigorously establish the relation of the loss bound to nonlinear filter stability and state space reduction techniques for MDPs.

A recent work by the authors Kara and Yüksel (2021) introduces a different finite history approximation technique where the approximation is done via an alternative belief MDP-reduction method rather than direct discretization of the space of probability measures. It is shown that the approximation error can be related to the controlled filter stability in terms of the total variation distance, whereas in this paper, the error bound is also shown to be related to more general, and in particular weak convergence inducing, metrics. Although, the approximation technique introduced in Kara and Yüksel (2021) provides an error upper bound in terms of the more stringent total variation distance, it proves to be numerically efficient as it is shown that a finite history Q learning algorithm converges to the optimality equation of an approximate model. Our analysis here requires less stringent conditions on filter stability, however the use of the bounded-Lipschitz metric on probability measures leads to a significantly more tedious analysis.

**Contributions.** In this paper, we rigorously establish near optimality of finite memory feedback control policies for the case where the actions and measurements are finite (with the state being real vector valued), provided that the controlled non-linear filter is stable in a sense to be presented in the paper. We also explicitly relate the approximation error with the window size. This is the first rigorous result, to our knowledge, where finite window policies are shown to be  $\epsilon$ -optimal with an explicit rate of convergence with respect to the window size.

## 1.1 Preliminaries and the Main Results

Let  $\mathbb{X} \subset \mathbb{R}^m$  denote a Borel set which is the state space of a partially observed controlled Markov process. Here and throughout the paper  $\mathbb{Z}_+$  denotes the set of non-negative integers and  $\mathbb{N}$  denotes the set of positive integers. Let  $\mathbb{Y}$  be a finite set denoting the observation space of the model, and let

the state be observed through an observation channel  $Q$ . The observation channel,  $Q$ , is defined as a stochastic kernel (regular conditional probability) from  $\mathbb{X}$  to  $\mathbb{Y}$ , such that  $Q(\cdot|x)$  is a probability measure on the power set  $P(\mathbb{Y})$  of  $\mathbb{Y}$  for every  $x \in \mathbb{X}$ , and  $Q(A|\cdot) : \mathbb{X} \rightarrow [0, 1]$  is a Borel measurable function for every  $A \in P(\mathbb{Y})$ . A decision maker (DM) is located at the output of the channel  $Q$ , and hence it only sees the observations  $\{Y_t, t \in \mathbb{Z}_+\}$  and chooses its actions from  $\mathbb{U}$ , the action space which is a finite subset of some Euclidean space. An *admissible policy*  $\gamma$  is a sequence of control functions  $\{\gamma_t, t \in \mathbb{Z}_+\}$  such that  $\gamma_t$  is measurable with respect to the  $\sigma$ -algebra generated by the information variables  $I_t = \{Y_{[0,t]}, U_{[0,t-1]}\}$ ,  $t \in \mathbb{N}$ ,  $I_0 = \{Y_0\}$ , where

$$U_t = \gamma_t(I_t), \quad t \in \mathbb{Z}_+, \quad (1)$$

are the  $\mathbb{U}$ -valued control actions and  $Y_{[0,t]} = \{Y_s, 0 \leq s \leq t\}$ ,  $U_{[0,t-1]} = \{U_s, 0 \leq s \leq t-1\}$ . We define  $\Gamma$  to be the set of all such admissible policies. The update rules of the system are determined by (1) and the following relationships:

$$\Pr((X_0, Y_0) \in B) = \int_B \mu(dx_0)Q(dy_0|x_0), \quad B \in \mathcal{B}(\mathbb{X} \times \mathbb{Y}),$$

where  $\mu$  is the (prior) distribution of the initial state  $X_0$ , and

$$\Pr\left((X_t, Y_t) \in B \mid (X, Y, U)_{[0,t-1]} = (x, y, u)_{[0,t-1]}\right) = \int_B \mathcal{T}(dx_t|x_{t-1}, u_{t-1})Q(dy_t|x_t),$$

$B \in \mathcal{B}(\mathbb{X} \times \mathbb{Y})$ ,  $t \in \mathbb{N}$ , where  $\mathcal{T}$  is the transition kernel of the model which is a stochastic kernel from  $\mathbb{X} \times \mathbb{U}$  to  $\mathbb{X}$ . Note that, although  $\mathbb{Y}$  is finite, we use integral sign instead of the summation sign for notation convenience by letting the measure to be sum of dirac-delta measures. We let the objective of the agent (decision maker) be the minimization of the infinite horizon discounted cost,

$$J_\beta(\mu, \mathcal{T}, \gamma) = E_\mu^{\mathcal{T}, \gamma} \left[ \sum_{t=0}^{\infty} \beta^t c(X_t, U_t) \right]$$

for some discount factor  $\beta \in (0, 1)$ , over the set of admissible policies  $\gamma \in \Gamma$ , where  $c : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}$  is a Borel-measurable stage-wise cost function and  $E_\mu^{\mathcal{T}, \gamma}$  denotes the expectation with initial state probability measure  $\mu$  and transition kernel  $\mathcal{T}$  under policy  $\gamma$ . Note that  $\mu \in \mathcal{P}(\mathbb{X})$ , where we let  $\mathcal{P}(\mathbb{X})$  denote the set of probability measures on  $\mathbb{X}$ . We define the optimal cost for the discounted infinite horizon setup as a function of the priors and the transition kernels as

$$J_\beta^*(\mu, \mathcal{T}) = \inf_{\gamma \in \Gamma} J_\beta(\mu, \mathcal{T}, \gamma).$$

For partially observed stochastic problems, the optimal policies use all the available information in general. The question we ask is the following one: suppose we define an  *$N$ -memory admissible policy*  $\gamma^N$  so that  $\gamma^N$  is a sequence of control functions  $\{\gamma_t, t \in \mathbb{Z}_+\}$  such that  $\gamma_t$  is measurable with respect to the  $\sigma$ -algebra generated by the information variables

$$\begin{aligned} I_t^N &= \{Y_{[t-N,t]}, U_{[t-N,t-1]}\}, \text{ if } t \geq N, \\ I_t^N &= \{Y_{[0,t]}, U_{[0,t-1]}\}, \text{ if } 0 < t < N, \\ I_0 &= \{Y_0\}, \end{aligned} \quad (2)$$

that is the controller can only have access to the information variables through a window whose length is  $N$ . We define  $\Gamma^N$  to be the set of all such  $N$ -memory admissible policies. Similarly, we define the optimal cost function under  $N$ -memory admissible policies as

$$J_\beta^N(\mu, \mathcal{T}) = \inf_{\gamma^N \in \Gamma^N} J_\beta(\mu, \mathcal{T}, \gamma^N).$$

Under this setup, we will study the following problem.

*Problem:* Under suitable conditions, can we find explicit bounds on  $J_\beta^N(\mu, \mathcal{T}) - J_\beta^*(\mu, \mathcal{T})$  in terms of  $N$  and a constructive approximate solution achieving this bound?

Our goal is to find the best possible control policy in  $\Gamma^N$  that is, in the set of policies that use only a finite history of information variables, in an offline setting by reducing the problem to a simpler approximate setup where we assume that the system dynamics are known to the designer. A general summary of the approach we will follow to answer this problem is as follows: We first define the belief MDP counterpart of the partially observed system. Then, we construct a finite subset of the belief state space using the probability distributions that can be achieved using the finite window information variables ( $I_t^N$ 's) from a fixed probability distribution. This finite subset leads to an approximate MDP model for which we find optimal policies. The calculation of the policies is greatly simplified compared to the calculation of optimal policies for the original POMDP model. Finally, we show that the loss occurring from applying this approximate policy to the original model can be upper bounded by the expected error of the discretization of the belief space. The loss is evaluated compared to the best possible *admissible* policy in the set  $\Gamma$ . The accumulating error then, can be represented in relation to the filter stability problem, that is, how fast the controlled process forgets its initial distribution as it observes the information variables from the system.

Note that although we take the infimum over all  $N$ -memory admissible policies, we will explicitly construct finite window policies which will be *time-invariant*, or a *finite-state probabilistic automaton* (as is also referred to by Yu and Bertsekas, 2008) that accepts as inputs the finite window of observations and actions, and produces as outputs the control actions in a time-invariant/stationary fashion. Accordingly, the infimum above for  $J_\beta^N(\mu, \mathcal{T})$  can be replaced with the minimum over such policies.

We answer the problem above affirmatively in Theorems 12, 16, and 17 under complementary conditions.

## 2. Regularity and Stability Properties of the Belief-MDP

In this section, we introduce the belief MDP reduction of POMDPs and provide regularity properties of the belief MDPs.

### 2.1 Convergence Notions for Probability Measures

For the analysis of the technical results, we will use different notions of convergence for sequences of probability measures.

Two important notions of convergences for sequences of probability measures are weak convergence, and convergence under total variation. For some  $N \in \mathbb{N}$  a sequence  $\{\mu_n, n \in \mathbb{N}\}$  in  $\mathcal{P}(\mathbb{R}^N)$  is said to converge to  $\mu \in \mathcal{P}(\mathbb{R}^N)$  *weakly* if  $\int_{\mathbb{R}^N} c(x) \mu_n(dx) \rightarrow \int_{\mathbb{R}^N} c(x) \mu(dx)$  for every continuous and bounded  $c : \mathbb{R}^N \rightarrow \mathbb{R}$ . One important property of weak convergence is that the space of

probability measures on a complete, separable, and metric (Polish) space endowed with the topology of weak convergence is itself complete, separable, and metric (Parthasarathy, 1967). One such metric is the bounded Lipschitz metric (Villani, 2008, p.109), which is defined for  $\mu, \nu \in \mathcal{P}(\mathbb{X})$  as

$$\rho_{BL}(\mu, \nu) := \sup_{\|f\|_{BL} \leq 1} \left| \int f d\mu - \int f d\nu \right| \quad (3)$$

where

$$\|f\|_{BL} := \|f\|_{\infty} + \sup_{x \neq y} \frac{|f(x) - f(y)|}{d(x, y)}$$

and  $\|f\|_{\infty} = \sup_{x \in \mathbb{X}} |f(x)|$ .

For probability measures  $\mu, \nu \in \mathcal{P}(\mathbb{R}^N)$ , the *total variation* metric is given by

$$\|\mu - \nu\|_{TV} = 2 \sup_{B \in \mathcal{B}(\mathbb{R}^N)} |\mu(B) - \nu(B)| = \sup_{f: \|f\|_{\infty} \leq 1} \left| \int f(x) \mu(dx) - \int f(x) \nu(dx) \right|,$$

where the supremum is taken over all measurable real  $f$  such that  $\|f\|_{\infty} = \sup_{x \in \mathbb{R}^N} |f(x)| \leq 1$ . A sequence  $\mu_n$  is said to converge in total variation to  $\mu \in \mathcal{P}(\mathbb{R}^N)$  if  $\|\mu_n - \mu\|_{TV} \rightarrow 0$ .

## 2.2 Ergodicity and Filter Stability Properties of Partially Observed MDPs

Given a prior  $\mu \in \mathcal{P}(\mathbb{X})$  and a policy  $\gamma \in \Gamma$ , we define the filter and predictor for a POMDP in the following.

**Definition 1** *The one step predictor process is defined as the sequence of conditional probability measures*

$$\pi_{n-}^{\mu, \gamma}(\cdot) = P^{\mu, \gamma}(X_n \in \cdot | Y_{[0, n-1]}, U_{[0, n-1]} = \gamma_n(Y_{[0, n-1]}, U_{[0, n-2]})) = P^{\mu, \gamma}(X_n \in \cdot | Y_{[0, n-1]}), \quad n \in \mathbb{N}$$

where  $P^{\mu, \gamma}$  is the probability measure induced by the prior  $\mu$  and the policy  $\gamma$ , when  $\mu$  is the probability measure on  $X_0$ .

**Definition 2** *The filter process is defined as the sequence of conditional probability measures*

$$\pi_n^{\mu, \gamma}(\cdot) = P^{\mu, \gamma}(X_n \in \cdot | Y_{[0, n]}, U_{[0, n-1]} = \gamma_n(Y_{[0, n-1]}, U_{[0, n-2]})) = P^{\mu, \gamma}(X_n \in \cdot | Y_{[0, n]}), \quad n \in \mathbb{N} \quad (4)$$

where  $P^{\mu, \gamma}$  is the probability measure induced by the prior  $\mu$  and the policy  $\gamma$ .

**Definition 3** (Dobrushin, 1956, Equation 1.16) *For a kernel operator  $K : S_1 \rightarrow \mathcal{P}(S_2)$  (that is a regular conditional probability from  $S_1$  to  $S_2$ ) for standard Borel spaces  $S_1, S_2$ , we define the Dobrushin coefficient as:*

$$\delta(K) = \inf \sum_{i=1}^n \min(K(x, A_i), K(y, A_i)) \quad (5)$$

where the infimum is over all  $x, y \in S_1$  and all partitions  $\{A_i\}_{i=1}^n$  of  $S_2$ .

We note that this definition holds for continuous or finite/countable spaces  $S_1$  and  $S_2$  and  $0 \leq \delta(K) \leq 1$  for any kernel operator.

**Example 1** Assume for a finite setup, we have the following stochastic transition matrix

$$K = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ 0 & \frac{1}{2} & \frac{1}{2} \\ \frac{3}{4} & 0 & \frac{1}{4} \end{pmatrix}$$

The Dobrushin coefficient is the minimum over any two rows where we sum the minimum elements among those rows. For this example, the first and the second rows give  $\frac{2}{3}$ , the first and the third rows give  $\frac{7}{12}$  and the second and the third rows give  $\frac{1}{4}$ . Then the Dobrushin coefficient is  $\frac{1}{4}$ .

Let

$$\tilde{\delta}(\mathcal{T}) := \inf_{u \in \mathbb{U}} \delta(\mathcal{T}(\cdot, u)).$$

**Definition 4** For  $\mathbb{X} \subset \mathbb{R}^m$  for some  $m \in \mathbb{N}$ , and for two probability measures  $\mu, \nu \in \mathcal{P}(\mathbb{X})$ ,  $\mu$  is said to be absolutely continuous with respect to  $\nu$  if  $\mu(A) = 0$  for every set  $A \in \mathcal{B}(\mathbb{X})$  for which  $\nu(A) = 0$ . We denote the absolute continuity of  $\mu$  with respect to  $\nu$  by  $\mu \ll \nu$ .

**Theorem 5** (McDonald and Yüksel, 2020, Theorem 3.3) Assume that for  $\mu, \nu \in \mathcal{P}(\mathbb{X})$ , we have  $\mu \ll \nu$ . Then we have (exponential filter stability)

$$E^{\mu, \gamma} [\|\pi_{n+1}^{\mu, \gamma} - \pi_{n+1}^{\nu, \gamma}\|_{TV}] \leq (1 - \tilde{\delta}(\mathcal{T}))(2 - \delta(Q))E^{\mu, \gamma} [\|\pi_n^{\mu, \gamma} - \pi_n^{\nu, \gamma}\|_{TV}].$$

In particular, defining  $\alpha := (1 - \tilde{\delta}(\mathcal{T}))(2 - \delta(Q))$ , we have

$$E^{\mu, \gamma} [\|\pi_n^{\mu, \gamma} - \pi_n^{\nu, \gamma}\|_{TV}] \leq 2\alpha^n.$$

This result will be a key ingredient for our main results. It provides conditions on when the belief-state processes for a given POMDP under different priors get closer when they are fed with same observation processes, in expectation under the true probability space. In a vague sense, if the state process is tracked only using a finite window of recent measurement and control variables (and forgets the past observations and actions), then the amount of mismatch from the true filter can be bounded with an error that is exponentially diminishing with the window size. The relationship is via the term  $\sup_{\pi \in \mathcal{P}(\mathbb{X})} \sup_{\gamma \in \Gamma} E_{\pi}^{\gamma} [\rho_{BL}(P^{\pi}(X_N \in \cdot | Y_{[0, N]}), P^{\hat{\pi}}(X_N \in \cdot | Y_{[0, N]}))]$  which appears crucially in (9). We note that if one does not wish to have an explicit rate of convergence result, one could have more relaxed conditions for filter stability which will still lead to rigorous approximation results on the performance of finite window policies via the controlled filter stability analysis in McDonald and Yüksel (2022).

### 2.3 Reduction to Fully Observed Models and Regularity Properties of Belief-MDPs

It is by now a standard result that, for optimality analysis, any POMDP can be reduced to a completely observable Markov decision process (Yushkevich, 1976; Rhenius, 1974), whose states are the posterior state distributions or *beliefs* of the observer or the filter process as defined in (4); that is, the state at time  $n$  is

$$\Pr\{X_n \in \cdot | Y_0, \dots, Y_n, U_0, \dots, U_{n-1}\} \in \mathcal{P}(\mathbb{X}).$$

We call this equivalent process the filter process. The filter process has state space  $\mathcal{Z} = \mathcal{P}(\mathbb{X})$  and action space  $\mathbb{U}$ . Here,  $\mathcal{Z}$  is equipped with the Borel  $\sigma$ -algebra generated by the topology of weak convergence (Billingsley, 1999). As noted earlier, under this topology,  $\mathcal{Z}$  is a standard Borel space (Parthasarathy, 1967). Then, the transition probability  $\eta$  of the filter process can be constructed as follows (see also Hernández-Lerma, 1989). If we define the measurable function

$$F(z, u, y) := F(\cdot | y, u, z) = \Pr\{X_{n+1} \in \cdot | Z_n = z, U_n = u, Y_{n+1} = y\}$$

from  $\mathcal{P}(\mathbb{X}) \times \mathbb{U} \times \mathbb{Y}$  to  $\mathcal{P}(\mathbb{X})$  and use the stochastic kernel  $P(\cdot | z, u) = \Pr\{Y_{n+1} \in \cdot | Z_n = z, U_n = u\}$  from  $\mathcal{P}(\mathbb{X}) \times \mathbb{U}$  to  $\mathbb{Y}$ , we can write  $\eta$  as

$$\eta(\cdot | z, u) = \int_{\mathbb{Y}} 1_{\{F(z, u, y) \in \cdot\}} P(dy | z, u). \quad (6)$$

The one-stage cost function  $\tilde{c} : \mathcal{P}(\mathbb{X}) \times \mathbb{U} \rightarrow [0, \infty)$  of the filter process is given by

$$\tilde{c}(z, u) := \int_{\mathbb{X}} c(x, u) z(dx), \quad (7)$$

which is a Borel measurable function. Hence, the filter process is a completely observable Markov process with the components  $(\mathcal{Z}, \mathbb{U}, \tilde{c}, \eta)$ .

For the filter process, the information variables is defined as

$$\tilde{I}_t = \{Z_{[0,t]}, U_{[0,t-1]}\}, \quad t \in \mathbb{N}, \quad \tilde{I}_0 = \{Z_0\}.$$

It is well known that an optimal control policy of the original POMDP can use the belief  $Z_t$  as a sufficient statistic for optimal policies (see Yushkevich, 1976; Rhenius, 1974), provided they exist. More precisely, the filter process is equivalent to the original POMP in the sense that for any optimal policy for the filter process, one can construct a policy for the original POMP which is optimal. On existence, we note the following.

With the recent results by Feinberg et al. (2016); Kara et al. (2019) the transition model of the belief-MDP can be shown to satisfy weak continuity conditions on the belief state and action variables, and accordingly we have that the measurable selection conditions (Hernandez-Lerma and Lasserre, 1996, Chapter 3) apply. Notably, we state the following.

**Assumption 1** (i) *The transition probability  $\mathcal{T}(\cdot | x, u)$  is weakly continuous in  $(x, u)$ , i.e., for any  $(x_n, u_n) \rightarrow (x, u)$ ,  $\mathcal{T}(\cdot | x_n, u_n) \rightarrow \mathcal{T}(\cdot | x, u)$  weakly.*

(ii) *The observation channel  $Q(\cdot | x, u)$  is continuous in total variation, i.e., for any  $(x_n, u_n) \rightarrow (x, u)$ ,  $Q(\cdot | x_n, u_n) \rightarrow Q(\cdot | x, u)$  in total variation.*

**Assumption 2** (i) *The transition probability  $\mathcal{T}(\cdot | x, u)$  is continuous in total variation in  $(x, u)$ , i.e., for any  $(x_n, u_n) \rightarrow (x, u)$ ,  $\mathcal{T}(\cdot | x_n, u_n) \rightarrow \mathcal{T}(\cdot | x, u)$  in total variation.*

(ii) *The observation channel  $Q(\cdot | x)$  is independent of the control variable.*

**Theorem 6** (i) *(Feinberg et al., 2016) Under Assumption 1, the transition probability  $\eta(\cdot | z, u)$  of the filter process is weakly continuous in  $(z, u)$ .*



(ii) (Kara et al., 2019) Under Assumption 2, the transition probability  $\eta(\cdot|z, u)$  of the filter process is weakly continuous in  $(z, u)$ .

Under the above weak continuity conditions, the measurable selection conditions (Hernandez-Lerma and Lasserre, 1996, Chapter 3) apply and a solution to the discounted cost optimality equation exists, and accordingly an optimal control policy exists. This policy is stationary (in the belief state). Thus there exists a function  $\Phi : \mathcal{P}(\mathbb{X}) \rightarrow \mathbb{U}$  such that for any policy  $\gamma$  for a prior  $\mu$

$$\gamma(y_{[0,n]}) = \Phi \left( P^{\mu, \gamma}(X_n \in \cdot | Y_{[0,n]} = y_{[0,n]}) \right) = \Phi(\pi_n^{\mu, \gamma})$$

In particular, we have that  $J_\beta^*(\mu, \mathcal{T}, Q) = J_\beta^*(\mu, \eta)$ . This will be the case in our paper, under the assumptions we will work with.

For the rest of the paper, we will use  $\gamma$  for the belief process policy  $\Phi$  for consistency of notation.

The following supporting result, to be used later in the paper, provides further regularity properties on the transition model for the belief model under mild conditions on the fully observed model. This result may be useful for POMDP theory beyond the application considered in this paper. We also note that the bounds in (i) and (iii) below are applicable when we only have filter stability, but not exponential filter stability (McDonald and Yüksel, 2022), whereas items (ii)-(iv) will be used under exponential filter stability in this paper.

**Theorem 7** *i. Assume that*

$$\|\mathcal{T}(\cdot|x, u) - \mathcal{T}(\cdot|x', u)\|_{TV} \leq \alpha_{\mathbb{X}}|x - x'|$$

for some  $\alpha_{\mathbb{X}} < \infty$  for all  $u \in \mathbb{U}$ . We have

$$\rho_{BL}(\eta(\cdot|z, u), \eta(\cdot|z', u)) \leq 3(1 + \alpha_{\mathbb{X}})\rho_{BL}(z, z').$$

*ii. Assume that*

$$\|\mathcal{T}(\cdot|x, u) - \mathcal{T}(\cdot|x', u)\|_{TV} \leq \alpha_{\mathbb{X}}|x - x'|$$

for some  $\alpha_{\mathbb{X}} < \infty$  for all  $u \in \mathbb{U}$ . Then, under the conditions of Theorem 5 we have

$$\rho_{BL}(\eta(\cdot|z, u), \eta(\cdot|z', u)) \leq (3 - 2\delta(Q))(1 + \alpha_{\mathbb{X}})\rho_{BL}(z, z').$$

*iii. Without any assumption*

$$\rho_{BL}(\eta(\cdot|z, u), \eta(\cdot|z', u)) \leq 3\|z - z'\|_{TV}.$$

*iv. Under the conditions of Theorem 5*

$$\rho_{BL}(\eta(\cdot|z, u), \eta(\cdot|z', u)) \leq (3 - 2\delta(Q))(1 - \tilde{\delta}(\mathcal{T}))\|z - z'\|_{TV}.$$

**Proof** See Section 6.1. ■

### 3. Approximate Model Construction: Finite Belief-MDP through Finite Memory

In this section, we will construct a finite state space by quantizing the belief state space so that the approximate finite model is obtained using only a finite memory.

Our construction builds on but significantly differs from the approach by Saldi et al. (2018, 2017). As we will explain, we cannot afford to use uniform quantization in our setup, which was a crucial tool used by Saldi et al. (2018, 2017).

As we discussed in the previous section, we can write the infinite horizon cost as

$$\begin{aligned} J_\beta(\mathcal{T}, Q, \gamma, \mu) &= \sum_{t=0}^{\infty} \beta^t E_\mu [\tilde{c}(\pi_t, \gamma(\pi_t))] \\ &= \sum_{t=0}^{N-1} \beta^t E_\mu [\tilde{c}(\pi_t, \gamma(\pi_t))] + \sum_{t=N}^{\infty} \beta^t E_\mu [\tilde{c}(\pi_t, \gamma(\pi_t))]. \end{aligned}$$

Now we focus on the second term:

$$\begin{aligned} &\sum_{t=N}^{\infty} \beta^t E_\mu [\tilde{c}(\pi_t, \gamma(\pi_t))] \\ &= \sum_{t=N}^{\infty} \beta^t E_\mu [\tilde{c}(P^{\mu, \gamma}(X_t \in \cdot | Y_{[0,t]}, U_{[0,t-1]}), \gamma(P^{\mu, \gamma}(X_t \in \cdot | Y_{[0,t]}, U_{[0,t-1]})))] . \end{aligned}$$

Notice that for any time step  $t \geq N$  and for a fixed observation realization sequence  $y_{[0,t]}$  and for a fixed control action sequence  $u_{[0,t-1]}$ , the state process can be viewed as

$$\begin{aligned} &P^\mu(X_t \in \cdot | Y_{[0,t]} = y_{[0,t]}, U_{[0,t-1]} = u_{[0,t-1]}) \\ &= P^{\pi_{t-N}}(X_t \in \cdot | Y_{[t-N,t]} = y_{[t-N,t]}, U_{[t-N,t-1]} = u_{[t-N,t-1]}) \end{aligned}$$

where

$$\pi_{t-N}(\cdot) = P^\mu(X_{t-N} \in \cdot | Y_{[0,t-N-1]} = y_{[0,t-N-1]}, U_{[0,t-N-1]} = u_{[0,t-N-1]}).$$

That is, we can view the state as the Bayesian update of  $\pi_{t-N}$ , the predictor at time  $t-N$ , using the observations  $Y_{t-N}, \dots, Y_t$ . Notice that with this representation only the most recent  $N$  observation realizations are used for the update and the past information of the observations is embedded in  $\pi_{t-N}$ .

Hence, we can view the state space for time stages  $t \geq N$  as

$$\mathcal{Z} = \{P^\pi(X_N \in \cdot | Y_{[0,N]}, U_{[0,N-1]}); \pi \in \mathcal{P}(\mathbb{X}), Y_{[0,N]} \in \mathbb{Y}^{N+1}, U_{[0,N-1]} \in \mathbb{U}^N\}$$

Consider the following finite set  $\mathcal{Z}_{\hat{\pi}}^N$  defined by, for a fixed probability measure  $\hat{\pi} \in \mathcal{P}(\mathbb{X})$ ,

$$\mathcal{Z}_{\hat{\pi}}^N := \{P^{\hat{\pi}}(X_N \in \cdot | Y_{[0,N]}, U_{[0,N-1]}); Y_{[0,N]} \in \mathbb{Y}^{N+1}, U_{[0,N-1]} \in \mathbb{U}^N\}.$$

Define the map  $F : \mathcal{Z} \rightarrow \mathcal{Z}_{\hat{\pi}}^N$ , which will serve as the quantizer, with

$$F(z) := \arg \min_{y \in \mathcal{Z}_{\hat{\pi}}^N} \rho_{BL}(z, y) \tag{8}$$

This map separates the set  $\mathcal{Z}$  into  $|\mathbb{Y}|^{N+1} \times |\mathbb{U}|^N$  sets, accordingly this quantizes the set of probability measures.

To complete our approximate controlled Markov model, we now define the one-stage cost function  $c^N : \mathcal{Z}_{\hat{\pi}}^N \times \mathbb{U} \rightarrow [0, \infty)$  and the transition probability  $\eta^N$  on  $\mathcal{Z}_{\hat{\pi}}^N$  given realizations in  $\mathcal{Z}_{\hat{\pi}}^N \times \mathbb{U}$ : For a given  $z_i = P^{\hat{\pi}}(X_N \in \cdot | y_{[0,N]}^i, u_{[0,N-1]}^i)$  and control action  $u$

$$c^N(z_i, u) = c^N(P^{\hat{\pi}}(X_N \in \cdot | y_{[0,N]}^i, u_{[0,N-1]}^i), u) := \tilde{c}(P^{\hat{\pi}}(X_N \in \cdot | y_{[0,N]}^i, u_{[0,N-1]}^i), u),$$

where  $\tilde{c}$  is defined in (7) and

$$\begin{aligned} \eta^N(\cdot | z_i, u) &= \eta_N(\cdot | P^{\hat{\pi}}(X_N \in \cdot | y_{[0,N]}^i, u_{[0,N-1]}^i), u) \\ &:= F * \eta(\cdot | P^{\hat{\pi}}(X_N \in \cdot | y_{[0,N]}^i, u_{[0,N-1]}^i), u), \end{aligned}$$

where

$$F * \eta(z' | z, u) = \eta(\{z \in \mathcal{P}(\mathbb{X}) : F(z) = z'\} | z, u),$$

for all  $z' \in \mathcal{Z}_{\hat{\pi}}^N$ .

We thus have defined a finite state MDP with the state space  $\mathcal{Z}_{\hat{\pi}}^N$ , action space  $\mathbb{U}$ , cost function  $c_N$  and the transition probability  $\eta_N$ .

An optimal policy  $\gamma_N^*$  for this finite state model is a function taking values from the finite state space and hence at some time step  $t \geq N$  it only uses  $N$  most recent observations and control action variables that is,  $\gamma_N^*$  is a measurable function of the information set  $I_t^N$  defined in (2) for all  $t \geq N$ .

We list the steps to construct the approximate model informally as follows:

- Fix a probability distribution  $\hat{\pi} \in \mathcal{P}(\mathbb{X})$  as an estimator for the predictor of  $N$  step back ( $\pi_{t-N}^-$ ).
- Calculate Bayesian updates for all possible realizations  $y_{[0,N]}, u_{[0,N-1]}$  ( $|\mathbb{Y}|^{N+1} \times |\mathbb{U}|^N$  many realizations) starting from the prior  $\hat{\pi}$  to form the finite subset  $\mathcal{Z}_{\hat{\pi}}^N$ .
- Calculate  $\eta^N$  (approximate transition model) and  $c^N$  (approximate cost function) using a nearest neighbor map. Note that from any  $z_i \in \mathcal{Z}_{\hat{\pi}}^N$  there are only  $|\mathbb{Y}|$  many possible transitions under the true dynamics. To construct the  $\eta^N$ , these possible transitions are mapped to the closest element in  $\mathcal{Z}_{\hat{\pi}}^N$ .
- Calculate the value functions and the optimal policies for the finite model with state space  $\mathcal{Z}_{\hat{\pi}}^N$ , transition kernel  $\eta^N$  and the cost function  $c^N$ .

As we have noted before, the complexity of POMDPs in general arises from the structure of the belief state space  $\mathcal{Z}$  which is a set of probability measures on  $\mathbb{X}$ . This set is always uncountable and needs to be associated with proper topologies to make the analysis feasible. Approximations for POMDPs are usually done by choosing a finite subset, say  $\hat{\mathcal{Z}}$ , of the belief state space  $\mathcal{Z}$  (Smith and Simmons, 2012; Pineau et al., 2006; Saldi et al., 2017; Zhou et al., 2008, 2010; Subramanian and Mahajan, 2019; Zhang et al., 2014), and finding an approximate MDP model for this finite set. To choose the finite set, the aforementioned works use a uniform quantization scheme, in various topologies on  $\mathcal{Z}$ . In other words, the quantization is made such that for any  $z \in \mathcal{Z}$ , there exists an element  $\hat{z} \in \hat{\mathcal{Z}}$  with  $\|z - \hat{z}\| \leq \epsilon$  for a fixed  $\epsilon > 0$ . The metric to measure distances of the belief

states varies for different works, although for finite  $\mathbb{X}$ ,  $L_1$  distance of distributions is what is used in general for the quantization of  $\mathcal{Z}$ , which coincides with total variation and weak convergence topology on  $\mathcal{Z}$  when  $\mathbb{X}$  is finite; for general  $\mathbb{X}$ , a more appropriate and natural topology is the weak convergence topology on  $\mathcal{Z}$  which is what we work with in the paper since  $\rho_{BL}$  metrizes the weak convergence.

In this paper, instead of quantizing  $\mathcal{Z}$  directly and uniformly, we use finite window information variables ( $I_t^N$ 's) to construct the finite subset of  $\mathcal{Z}$  since our goal is to analyze the effect of the window size on the approximation performance. That is, we use the finite set

$$\mathcal{Z}_{\hat{\pi}}^N := \left\{ P^{\hat{\pi}}(X_N \in \cdot | Y_{[0,N]}, U_{[0,N-1]}); Y_{[0,N]} \in \mathbb{Y}^{N+1}, U_{[0,N-1]} \in \mathbb{U}^N \right\}$$

constructed using  $Y_{[0,N]}, U_{[0,N-1]}$ . For this set, we cannot afford a uniform discretization scheme. A uniform quantization would mean that for a fixed  $\epsilon > 0$

$$\rho_{BL} \left( P^{\hat{\pi}}(X_N \in \cdot | y_{[0,N]}, u_{[0,N-1]}), P^{\pi}(X_N \in \cdot | y_{[0,N]}, u_{[0,N-1]}) \right) < \epsilon$$

uniformly for any  $\pi \in \mathcal{Z}$  and for any  $y_{[0,N]} \in \mathbb{Y}^{N+1}, u_{[0,N-1]} \in \mathbb{U}^N$ . However, this is in general inapplicable for filter stability problems as it requires for the processes with different starting points to converge uniformly for any realizations of information variables (even for highly unlikely ones). That is why, we follow a different approach and show that we do not have to force uniform quantization, and the error of approximate value functions can be related to the expected error of the form

$$E_{\pi}^{\gamma} \left[ \rho_{BL} \left( P^{\pi}(X_N \in \cdot | Y_{[0,N]}, U_{[0,N-1]}), P^{\hat{\pi}}(X_N \in \cdot | Y_{[0,N]}, U_{[0,N-1]}) \right) \right]$$

which in turn can be bounded using Theorem 5. Our technical analysis, accordingly, is slightly more tedious at the benefit of arriving at a practical and intuitive finite-memory method whose near optimality is rigorously established.

**Remark 8** *The finite subset  $\mathcal{Z}_{\hat{\pi}}^N$  is crucial for the approximation of the belief MDP we construct in this section. This set depends on the choice of  $\hat{\pi}$  as the starting probability distribution and the finite window information variables  $Y_{[0,N]}, U_{[0,N-1]}$ . Kara and Yüksel (2021) consider a similar approach for construction of an approximate POMDP model by using the same finite set  $\mathcal{Z}_{\hat{\pi}}^N$ . However, instead of using a nearest neighbor map for the correspondence between the original belief space and the finite set  $\mathcal{Z}_{\hat{\pi}}^N$ , there, the states with matching finite history are used for the correspondence by putting a different topology on the belief space and the set  $\mathcal{Z}_{\hat{\pi}}^N$ . The method used in this paper naturally results in a smaller approximation error due to the nature of the nearest neighbor map, and lets us work with the weak convergence topology and  $\rho_{BL}$  metric. On one hand, Kara and Yüksel (2021) provide a greater error term in terms of the total variation distance, which always upper bounds the  $\rho_{BL}$  metric, on the other hand, the method used there, is computationally more efficient and the approximate model can actually be learned with a reinforcement learning algorithm that uses finite window information variables.*

*On the choice of  $\hat{\pi}$ , we note that Theorem 12 provides a bound that is independent of the choice of  $\hat{\pi}$ . However, since this is only an upper bound, it is apparent that the true value of the error may depend on different choices of  $\hat{\pi}$ . For the approximate model, we use  $\hat{\pi}$  as an estimator for the true predictor  $\pi_{t-}^{\mu, \gamma}(\cdot) = P^{\mu, \gamma}(X_t \in \cdot | Y_{[0,t-1]})$  at any given time  $t$  under some policy  $\gamma$ . Since*

$\hat{\pi}$  is a fixed probability distribution and does not vary with time, we can argue that a reasonable choice would be the time averages of  $\pi_{t-}^{\mu, \gamma^*}(\cdot)$  under an optimal policy  $\gamma^*$  and if the hidden state process  $X_t$  admits an invariant measure under this optimal policy  $\gamma^*$ , the time averages of  $\pi_{t-}^{\mu, \gamma^*}(\cdot)$  will converge to the invariant measure of  $X_t$  under the optimal policy. However, one issue with this approach is that the designer does not have access to the optimal policy and thus, cannot compute the invariant measure under  $\gamma^*$ . The designer, instead, can use the approximate optimal policy  $\gamma^N$ , but the choice of  $\hat{\pi}$  also affects the approximate policy  $\gamma^N$ , and hence, this approach may not be feasible for practical purposes. In any case, our result provides a bound uniform over the prior selections.

We now discuss the role of the realizations of  $Y_{[0,N]}, U_{[0,N-1]}$ . Notice that, to construct  $\mathcal{Z}_{\hat{\pi}}^N$ , we consider all possible  $|\mathbb{Y}|^{N+1} \times |\mathbb{U}|^N$  realizations which reduces the uncountable state space to a finite one but the size of this finite subset grows fast with the size of window,  $N$ , and with the size of spaces  $\mathbb{Y}$  and  $\mathbb{U}$ . The bound we provide in Theorem 9, reveals that the approximation error is related to the expectation over the realizations of  $Y_{[0,N]}, U_{[0,N-1]}$  under the true dynamics, which suggests that if a realization  $y_{[0,N]}, u_{[0,N-1]}$ , is highly unlikely under the true dynamics, these realizations can be discarded when constructing  $\mathcal{Z}_{\hat{\pi}}^N$  for computational ease by reducing the size of  $\mathcal{Z}_{\hat{\pi}}^N$ .

In fact this question motivates the following direction. A close look at Definition 3 reveals that the Dobrushin coefficient of a channel and a channel obtained by quantizing the measurements of that same channel are so that the coefficient of the former is a lower bound for the latter. Therefore, one can always quantize the measurement channels further if the original channel satisfies the contraction condition presented for filter stability, with the quantized channel also satisfying the contraction property. This presents a recipe for dealing with both low probability channel outcomes or even continuous space measurement channels; however one would additionally need to show that the value function of the approximate model with quantized measurements would be close to the value function of the original model. This is a possible future direction (using the weak Feller property of the belief-MDP) (as shown e.g. by Saldi et al., 2017) and continuity properties of filter updates in measurement realizations McDonald and Yüksel (2022).

## 4. Approximation Error Analysis and Rates of Convergence

An optimal policy for the constructed finite model,  $\gamma_N^*$  can be extended to  $\mathcal{P}(\mathbb{X})$  and can be used for the original MDP.

### 4.1 Analysis via Expected Filter Approximation Error

The next result, which is related to the construction given by Saldi et al. (2018, Theorem 4.38) (see also Saldi et al., 2017), provides a mismatch error for using this policy. This result is going to be the key supporting tool for the main theorem of the paper, which will be presented immediately after. The proof requires multiple technical lemmas and are presented in Section 6.2 (with some supporting but tedious technical steps moved to the Appendix).

**Assumption 3** •  $\rho_{BL}(\eta(\cdot|z, u), \eta(\cdot|z', u)) \leq \alpha_Z \rho_{BL}(z, z')$  (see Theorem 7).

- $|\tilde{c}(z, u) - \tilde{c}(z', u)| \leq \alpha_{\tilde{c}} \rho_{BL}(z, z')$  for some  $\alpha_{\tilde{c}}$  for all  $u \in \mathbb{U}$ .

**Theorem 9** Under Assumption 3, for all  $z \in \mathcal{P}(\mathbb{X})$  of the form

$$z = P^\pi(X_N \in \cdot | y_{[0,N]}, u_{[0,N-1]}),$$

the following holds if  $\beta < \frac{1}{4\alpha_Z + 1}$ :

$$\begin{aligned} & \sup_{\gamma \in \Gamma} E \left[ J_\beta(\eta, \gamma_N^*, z) - J_\beta(\eta, \gamma^*, z) \middle| Y_{[0,N]}, \gamma(Y_{[0,N-1]}) \right] \\ & \leq K \sup_{\pi \in \mathcal{P}(\mathbb{X})} \sup_{\gamma \in \Gamma} E_\pi^\gamma \left[ \rho_{BL} \left( P^\pi(X_N \in \cdot | Y_{[0,N]}, U_{[0,N-1]}), P^{\hat{\pi}}(X_N \in \cdot | Y_{[0,N]}, U_{[0,N-1]}) \right) \right]. \end{aligned}$$

Where  $K$  is a constant that depends on  $\beta, \alpha_Z, \alpha_{\tilde{c}}$  and  $\|\tilde{c}\|_\infty$ . (The exact expression for the constant can be found in the proof).

**Remark 10** We note that the upper bound for  $\beta$  can be chosen as  $\frac{1}{(2+\|L\|_\infty)\alpha_Z+1}$ , instead of  $4\alpha_Z+1$  (see Remark 21 in Appendix), where

$$\|L\|_\infty = \sup_{\pi \in \mathcal{P}(\mathbb{X})} \sup_{\gamma \in \Gamma} \sup_{Y_{[0,N]}, U_{[0,N-1]}} \rho_{BL} \left( P^\pi(\cdot | Y_{[0,N]}, U_{[0,N-1]}), P^{\hat{\pi}}(\cdot | Y_{[0,N]}, U_{[0,N-1]}) \right).$$

Hence, under a uniform filter stability bound, the upper bound can be chosen near  $\frac{1}{2\alpha_Z+1}$ .

The proof of this theorem is rather long, and accordingly, it is presented in Section 6.2.

Before presenting the main result of the paper, we provide further supporting results that will let us work with a probability density function.

**Lemma 11** Assume that the transition kernel  $\mathcal{T}(dx_1|x_0, u_0)$  admits a density function  $f$  with respect to a reference measure  $\phi$  such that  $\mathcal{T}(dx_1|x_0, u_0) = f(x_1, x_0, u_0)\phi(dx_1)$ . If  $|f(x_1, x_0, u_0) - f(x_1, x_0', u_0)| \leq \alpha_{\mathbb{X}}|x_0 - x_0'|$  for all  $x_1, x_0, x_0' \in \mathbb{X}$  and  $u_0 \in \mathbb{U}$  then

$$\|\mathcal{T}(\cdot|x_0, u_0) - \mathcal{T}(\cdot|x_0', u_0)\|_{TV} \leq \alpha_{\mathbb{X}}|x_0 - x_0'|.$$

We note that using this result, assumptions of Theorem 7 can be expressed with the Lipschitz condition on the density function noted above. We now restate the assumptions that will be used for the main result.

**Assumption 4**

1. The transition kernel  $\mathcal{T}(dx_1|x_0, u_0)$  admits a density function  $f$  with respect to a reference measure  $\phi$  such that  $\mathcal{T}(dx_1|x_0, u_0) = f(x_1, x_0, u_0)\phi(dx_1)$ .
2. There exists some  $\alpha_{\mathbb{X}} < \infty$  such that  $|f(x_1, x_0, u_0) - f(x_1, x_0', u_0)| \leq \alpha_{\mathbb{X}}|x_0 - x_0'|$ .
3. There exists some  $\alpha_c < \infty$  such that for all  $u \in \mathbb{U}$ ,  $|c(x, u) - c(x', u)| \leq \alpha_c|x - x'|$ .
4.  $\alpha := (1 - \tilde{\delta}(\mathcal{T}))(2 - \delta(Q)) < 1$ ,
5. The transition kernel  $\mathcal{T}$  is dominated, i.e. there exists a dominating measure  $\hat{\pi} \in \mathcal{P}(\mathbb{X})$  such that for every  $x \in \mathbb{X}$  and  $u \in \mathbb{U}$ ,  $\mathcal{T}(\cdot|x, u) \ll \hat{\pi}(\cdot)$  that is  $\mathcal{T}(\cdot|x, u)$  is absolutely continuous with respect to  $\hat{\pi}$  for every  $x \in \mathbb{X}$  and  $u \in \mathbb{U}$ .

Before the result, we discuss the Lipschitz constants of interest and their relation to each other. First, note that by Lemma 11, Assumption 4 (second item) implies that  $\|\mathcal{T}(\cdot|x_0, u_0) - \mathcal{T}(\cdot|x_0', u_0)\|_{TV} \leq \alpha_{\mathbb{X}}|x - x'|$ . Hence, by Theorem 7, we can have various bounds on  $\alpha_{\mathcal{Z}}$ , where  $\rho_{BL}(\eta(\cdot|z, u), \eta(\cdot|z', u)) \leq \alpha_{\mathcal{Z}}\rho_{BL}(z, z')$ . In particular, we have  $\alpha_{\mathcal{Z}} \leq (3 - 2\delta(Q))(1 + \alpha_{\mathbb{X}})$ .

**Theorem 12** *Assume that we let the system start running for  $N$  time steps under a known policy  $\gamma$  (which may not be optimal), and the finite window policy starts acting after observing  $N$  information variables at time  $t = N$ . Under Assumption 4, if  $\beta < \frac{1}{4\alpha_{\mathcal{Z}}+1}$  we have*

$$E_{\mu}^{\gamma} [J_{\beta}(\pi_{N-}, \mathcal{T}, \gamma_N^*) - J_{\beta}(\pi_{N-}, \mathcal{T}, \gamma^*) | Y_{[0,N]}, U_{[0,N-1]}] \leq K(\beta, \alpha_{\mathbb{X}}, \alpha_c, \|c\|_{\infty})\alpha^N.$$

where  $\gamma_N^*$  is the optimal finite window policy and where  $K(\beta, \alpha_{\mathbb{X}}, \alpha_c, \|c\|_{\infty})$  is a constant that depends on  $\beta, \alpha_{\mathbb{X}}, \alpha_c$  and  $\|c\|_{\infty}$  (The exact formula for the constant can be found in the appendix in (28)).

In the statement,  $J_{\beta}(\pi_{N-}, \mathcal{T}, \gamma_N^*)$  (respectively  $J_{\beta}(\pi_{N-}, \mathcal{T}, \gamma^*)$ ) denotes the cost under the policy  $\gamma_N^*$  (respectively  $\gamma^*$ ) when the initial prior distribution is  $\pi_{N-}$ , where  $\pi_{N-} = \Pr(X_N \in \cdot | Y_{[0,N-1]}, U_{[0,N-1]})$  and the expectation is with respect to the realizations of  $Y_{[0,N-1]}, U_{[0,N-1]}$ .

**Remark 13** *We note that Theorem 12 applies to all finite state, measurement, action models as long as  $\alpha = (1 - \delta(\mathcal{T}))(2 - \delta(Q)) < 1$  and  $\beta$  satisfies the condition noted. Example 1 shows how to calculate the Dobrushin coefficient for transition matrices in finite setups. All other conditions apply since all probability measures on a finite/countable set are majorized by a probability measure which places a positive mass on every single point. Notice that the third condition only requires that the difference in the cost is bounded for every fixed control action. Condition 1 and 5 of Assumption 4 coincide in the finite case.*

**Remark 14** *We note that the error bound of the result is independent of the chosen  $\hat{\pi}$ . As we will see in the following proof, the first upper bound is a result of Theorem 9 which indeed depends on the  $\hat{\pi}$  chosen by the user. However, thanks to Theorem 5, we can get a further upper bound on the error which is uniform over any  $\hat{\pi}$  as long as  $\hat{\pi}$  is a dominating measure.*

One can interpret the absolute continuity assumption, that is  $\pi \ll \hat{\pi}$ , as follows: assume that the starting distribution of the process is  $\pi$  but we start the update from the fixed prior  $\hat{\pi}$ . The information,  $y_{[0,t]}, u_{[0,t-1]}$  can eventually fix the starting error uniformly for all such  $\hat{\pi}$ , as long as, the fixed starting distribution  $\hat{\pi}$ , puts on a positive measure to every event that the real starting distribution  $\pi$  puts on a positive measure. However, if it is not the case, that is, if the incorrect starting distribution  $\hat{\pi}$  puts 0 measure to some event, that  $\pi$  puts positive measure to, information variables are not sufficient to fix the starting error occurring from that 0 measure event. Of course this would not be feasible as the prior would not be compatible with the measured data. In any case, in our setup, the fixed prior  $\hat{\pi}$  serves as an approximation and this can be made to satisfy the absolute continuity condition by design.

**Proof** When we reduce a partially observed MDP to a fully observed process, the initial state of the belief process becomes the Bayesian update of the prior distribution of the state process of the POMDP. Hence, we can write that

$$E_{\mu}^{\gamma} [J_{\beta}(\pi_{N-}, \mathcal{T}, \gamma_N^*) - J_{\beta}(\pi_{N-}, \mathcal{T}, \gamma^*) | Y_{[0,N]}, U_{[0,N-1]}]$$

$$= E_\mu^\gamma [J_\beta(\pi_N, \eta, \gamma_N^*) - J_\beta(\pi_N, \eta, \gamma^*) | Y_{[0,N]}, U_{[0,N-1]}].$$

Notice that, using Theorem 7, Condition 2 of Assumption 4 implies that

$$\rho_{BL}(\eta(\cdot|z, u), \eta(\cdot|z', u)) \leq 3(1 + \alpha_{\mathbb{X}})\rho_{BL}(z, z')$$

and we also have that

$$|\tilde{c}(z, u) - \tilde{c}(z', u)| = \left| \int c(x, u)z(dx) - \int c(x, u)z'(dx) \right| \leq (\alpha_c + \|c\|_\infty)\rho_{BL}(z, z').$$

Thus, we can use Theorems 7 and 9 to write

$$\begin{aligned} & E_\mu^\gamma [J_\beta(\pi_N, \eta, \gamma_N^*) - J_\beta(\pi_N, \eta, \gamma^*) | Y_{[0,N]}, U_{[0,N-1]}] \\ & \leq K(\beta, \alpha_{\mathbb{X}}, \alpha_c, \|c\|_\infty) \sup_{\pi \in \mathcal{P}(\mathbb{X})} \sup_{\gamma \in \Gamma} E_\pi^\gamma \left[ \rho_{BL} \left( P^\pi(X_N \in \cdot | Y_{[0,N]}), P^{\hat{\pi}}(X_N \in \cdot | Y_{[0,N]}) \right) \right] \end{aligned} \quad (9)$$

We set  $\hat{\pi}$  in Assumption 4 as our representative probability measure for the quantization of the belief space. Notice that by our choice of  $\hat{\pi}$  is a dominating measure and therefore  $\pi_{t-} \ll \hat{\pi}$  for any time step  $t$ , where  $\pi_{t-}$  is the predictor at time  $t$ . Thus, Theorem 5 yields that under Assumption 4

$$E_\pi^\gamma \left[ \|P^\pi(\cdot | Y_{[0,N]}) - P^{\hat{\pi}}(\cdot | Y_{[0,N]})\|_{TV} \right] \leq \alpha^N,$$

for any predictor  $\pi$ . We also have by definition that

$$\rho_{BL} \left( P^{\pi, \gamma}(\cdot | Y_{[0,N]}), P^{\hat{\pi}, \gamma}(\cdot | Y_{[0,N]}) \right) \leq \|P^{\pi, \gamma}(\cdot | Y_{[0,N]}) - P^{\hat{\pi}, \gamma}(\cdot | Y_{[0,N]})\|_{TV}.$$

Hence the result follows. ■

**Remark 15** *By studying (9) and utilizing Theorem 7(i)-(iii), we can arrive at complementary results when we only have filter stability but not exponential filter stability.*

We note that the initial  $N$  time steps of the control problem should be treated separately as our approximation technique uses information variables with size  $N$ , but for the initial  $N$  steps, there is not enough observation or control action variables to make use of the approximate finite window policies. In the following result, for the first  $N$  time steps, we use a control policy that is found with a policy iteration type argument where the terminal cost estimated with  $\beta^N E[J_\beta(\gamma_N^*)]$  with  $\gamma_N^*$  being the approximate finite window policy.

The following result is stated for the best possible policy in the set  $\Gamma^N$  (see Equation 2), since  $\inf_{\gamma^N \in \Gamma^N} J_\beta(\mu, \mathcal{T}, \gamma^N) = J_\beta^N(\mu, \mathcal{T})$ , however, in the proof, instead of working with the policy achieving this infimum, we construct a policy (possibly sub-optimal) which achieves the stated upper bound, and this upper bound naturally becomes an upper bound for the best possible finite window policy in  $\Gamma^N$ .

**Theorem 16** *Under Assumption 4, if  $\beta < \frac{1}{4\alpha_z + 1}$  we have*

$$J_\beta^N(\mu, \mathcal{T}) - J_\beta^*(\mu, \mathcal{T}) \leq K(\beta, \alpha_{\mathbb{X}}, \alpha_c, \|c\|_\infty) \alpha^N \beta^N$$

where  $K(\beta, \alpha_{\mathbb{X}}, \alpha_c, \|c\|_\infty)$  is a constant that depends on  $\beta, \alpha_{\mathbb{X}}, \alpha_c$  and  $\|c\|_\infty$ . (The exact formula for the constant can be found in the appendix).



**Proof** Recall the optimal policy for the finite model constructed in Section 3 is denoted by  $\gamma_N^*$ . Notice that  $\gamma_N^*$  is a stationary policy and optimal for any initial point since it solves the discounted cost optimality equation.

Now, we construct the following policy. Use the policy  $\gamma_N^*$ , after time  $N$  unaltered, but modify the first  $N$  time-stage policies, as a batch update, which can be solved via a finite dynamic programming algorithm.

$$\tilde{\gamma}_0, \dots, \tilde{\gamma}_{N-1} = \operatorname{argmin}_{\gamma_1, \dots, \gamma_{N-1}} E \left[ \sum_{k=0}^{N-1} \beta^k c(x_k, u_k) \right] + \beta^N E[J_\beta(\gamma_N^*) | I_N]$$

where  $I_N$  is the history by time  $N$ . We denote this policy by  $\gamma^{++} := \{\tilde{\gamma}_0, \dots, \tilde{\gamma}_{N-1}, \gamma_N^*, \gamma_N^*, \dots\}$ .

Note that we denote the true optimal policy for the original model by  $\gamma^*$ . We now define the policy  $\gamma^+ := \{\gamma_{[0, N-1]}^*, \gamma_N^*, \gamma_N^*, \dots\}$ : Apply  $\gamma^*$  until time  $N$ , and then use our finite window policy  $\gamma^{*N}$ . Note that this policy is not practical as it assumes that the controller already knows the true optimal policy  $\gamma^*$ , however, we will make use this hypothetical policy for the analysis.

From the way  $\gamma^+$  is constructed, we clearly,  $J_\beta(\gamma^{++}) \leq J_\beta(\gamma^+)$ . And thus,

$$J_\beta(\gamma^{++}) - J_\beta(\gamma^*) \leq J_\beta(\gamma^+) - J_\beta(\gamma^*)$$

Furthermore, because of the way we constructed them, we have  $\gamma^{++} \in \Gamma^N$ . Hence, we can write

$$\begin{aligned} J_\beta^N(\mu, \mathcal{T}) - J_\beta^*(\mu, \mathcal{T}) &= \inf_{\gamma \in \Gamma^N} J_\beta(\mu, \mathcal{T}, \gamma) - J_\beta(\mu, \mathcal{T}, \gamma^*) \\ &\leq J_\beta(\mu, \mathcal{T}, \gamma^{++}) - J_\beta(\mu, \mathcal{T}, \gamma^*) \\ &\leq J_\beta(\mu, \mathcal{T}, \gamma^+) - J_\beta(\mu, \mathcal{T}, \gamma^*) \end{aligned}$$

Then, we start analyzing the error term: because  $\gamma^+$  and  $\gamma^*$  use the same policy before time  $N$ , we have

$$\begin{aligned} &J_\beta(\mu, \mathcal{T}, \gamma^+) - J_\beta(\mu, \mathcal{T}, \gamma^*) \\ &= \sum_{t=N}^{\infty} \beta^t E_\mu [c(X_t, \gamma_N^*(Y_{[t-N, t]}, U_{[t-N, t-1]}))] - \sum_{t=N}^{\infty} \beta^t E_\mu [c(X_t, \gamma^*(Y_{[0, t]}, U_{[0, t-1]}))] \\ &= \beta^N \sum_{t=N}^{\infty} \beta^{t-N} E_\mu^{\gamma^*} \left[ E_\mu [c(X_t, \gamma_N^*(Y_{[t-N, t]}, U_{[t-N, t-1]}))] \right. \\ &\quad \left. - E_\mu [c(X_t, \gamma^*(Y_{[0, t]}, U_{[0, t-1]}))] | Y_{[0, N]}, U_{[0, N-1]} \right] \\ &= \beta^N E_\mu^{\gamma^*} [J_\beta(\pi_{N-}, \mathcal{T}, \gamma_N^*) - J_\beta(\pi_{N-}, \mathcal{T}, \gamma^*) | Y_{[0, N]}, U_{[0, N-1]}] \end{aligned} \tag{10}$$

The last step follows from the observation that conditioning on the observations and control actions  $Y_{[0, N]}, U_{[0, N-1]}$ , the state process can be thought as if it starts at time  $t = N$  whose prior measure is  $\pi_{N-}$  and from the fact that the probability measure of  $Y_{[0, N]}, U_{[0, N-1]}$  is determined by the initial measure  $\mu$  and the policy  $\gamma^*$  since  $\gamma_t^N = \gamma_t^*$  for  $t \leq N$ .

The result then follows from Theorem 12. ■

## 4.2 Analysis via Uniform Bounds on Filter Approximation Error

The following result provides an alternative setup where our analysis is applicable when the filter stability error is presented in terms of a uniform bound on the filter approximation error under the total variation distance. We define this bound, arising from filter stability error, as follows:

$$\bar{L}_{TV} := \sup_{\pi \in \mathcal{P}(\mathbb{X})} \sup_{\gamma \in \Gamma} \sup_{y_{[0,N]}, u_{[0,N-1]}} \left\| P^\pi(\cdot | y_{[0,N]}, u_{[0,N-1]}) - P^{\hat{\pi}}(\cdot | y_{[0,N]}, u_{[0,N-1]}) \right\|_{TV}. \quad (11)$$

**Theorem 17** *If  $\beta < \frac{1}{\alpha_Z}$ , where  $\alpha_Z$  can be chosen as  $(3 - 2\delta(Q))(1 - \tilde{\delta}(\mathcal{T}))$  by Theorem 7, we have that*

(i)

$$\sup_z |J_\beta^N(z) - J_\beta^*(z)| \leq \frac{(\alpha_Z - 1)\beta + 1}{(1 - \beta)^2(1 - \alpha_Z\beta)} \|c\|_\infty \bar{L}_{TV}.$$

(ii)

$$\sup_z |J_\beta(z, \gamma_N) - J_\beta^*(z)| \leq \frac{2(1 + (\alpha_Z - 1)\beta)}{(1 - \beta)^3(1 - \alpha_Z\beta)} \|c\|_\infty \bar{L}_{TV}.$$

Before the proof, we note that the bound applies for all  $\beta < 1$ , if  $(3 - 2\delta(Q))(1 - \tilde{\delta}(\mathcal{T})) < 1$ . Furthermore, since  $(3 - 2\delta(Q))(1 - \tilde{\delta}(\mathcal{T})) = 2\alpha - (1 - \tilde{\delta}(\mathcal{T}))$  where  $\alpha = (2 - \delta(Q))(1 - \tilde{\delta}(\mathcal{T}))$  and, therefore, if  $\alpha < 1$  we have that for all  $\beta < 1/(1 + \tilde{\delta}(\mathcal{T}))$  the result applies independent of  $\delta(Q)$ , though this is a conservative bound. If  $\delta(Q) = 1$ , the bound applies for all  $\beta < 1$  as long as  $\alpha < 1$ .

**Proof** (i) We start by writing the fixed point equations

$$\begin{aligned} J_\beta^*(z) &= \min_u \left( \tilde{c}(z, u) + \beta \int J_\beta^*(z_1) \eta(dz_1 | z, u) \right) \\ J_\beta^N(z) &= \min_u \left( \tilde{c}(F(z), u) + \beta \int J_\beta^N(z_1) \eta(dz_1 | F(z), u) \right). \end{aligned}$$

Hence we can write that

$$\begin{aligned} \sup_z |J_\beta^*(z) - J_\beta^N(z)| &\leq \sup_u |\tilde{c}(z, u) - \tilde{c}(F(z), u)| \\ &\quad + \beta \sup_u \left| \int J_\beta^*(z_1) \eta(dz_1 | z, u) - \int J_\beta^N(z_1) \eta(dz_1 | F(z), u) \right| \\ &\leq \|c\|_\infty \|z - F(z)\|_{TV} + \beta \sup_u \int |J_\beta^*(z_1) - J_\beta^N(z_1)| \eta(dz_1 | F(z), u) \\ &\quad + \beta \sup_u \left| \int J_\beta^*(z_1) \eta(dz_1 | z, u) - \int J_\beta^*(z_1) \eta(dz_1 | F(z), u) \right| \\ &\leq \|c\|_\infty \|z - F(z)\|_{TV} + \beta \sup_z |J_\beta^*(z) - J_\beta^N(z)| + \beta \|J_\beta^*\|_{BL} \alpha_Z \|z - F(z)\|_{TV}. \quad (12) \end{aligned}$$

Note that by Theorem 7,  $\alpha_{\mathcal{Z}} \leq (3 - 2\delta(Q))(1 - \tilde{\delta}(\mathcal{T}))$  when  $\mathcal{Z}$  is associated with total variation distance. We also have that  $\|J_{\beta}^*\|_{BL} \leq \frac{2-\beta}{(1-\beta)(1-\alpha_{\mathcal{Z}}\beta)}\|c\|_{\infty}$  when  $\mathcal{Z}$  metrized by total variation distance (see Lemma 24). Hence, by noting  $\|z - F(z)\|_{TV} \leq \bar{L}_{TV}$ , we can conclude that

$$\sup_z |J_{\beta}^*(z) - J_{\beta}^N(z)| \leq \frac{(\alpha_{\mathcal{Z}} - 1)\beta + 1}{(1 - \beta)^2(1 - \alpha_{\mathcal{Z}}\beta)} \|c\|_{\infty} \bar{L}_{TV}.$$

(ii) We start by writing

$$\sup_z |J_{\beta}(z, \gamma_N) - J_{\beta}^*(z)| \leq \sup_z |J_{\beta}(z, \gamma_N) - J_{\beta}^N(z)| + \sup_z |J_{\beta}^N(z) - J_{\beta}^*(z)|$$

where the second term is bounded by (i). We now focus on the first term:

$$\begin{aligned} |J_{\beta}(z, \gamma_N) - J_{\beta}^N(z)| &\leq \sup_u |\tilde{c}(z, u) - \tilde{c}(F(z), u)| \\ &\quad + \beta \left| \int J_{\beta}(z_1, \gamma_N) \eta(dz_1 | z, \gamma_N(z)) - \int J_{\beta}^N(z_1) \eta(dz_1 | F(z), \gamma_N(z)) \right| \\ &\leq \|c\|_{\infty} \|z - F(z)\|_{TV} + \beta \int |J_{\beta}(z_1, \gamma_N) - J_{\beta}^N(z_1)| \eta(dz_1 | z, \gamma_N(z)) \\ &\quad + \beta \int |J_{\beta}^N(z_1) - J_{\beta}^*(z_1)| \eta(dz_1 | z, \gamma_N(z)) \\ &\quad + \beta \left| \int J_{\beta}^*(z_1) \eta(dz_1 | z, \gamma_N(z)) - \int J_{\beta}^*(z_1) \eta(dz_1 | F(z), \gamma_N(z)) \right| \\ &\quad + \beta \int |J_{\beta}^*(z_1) - J_{\beta}^N(z_1)| \eta(dz_1 | F(z), \gamma_N(z)) \\ &\leq \|c\|_{\infty} \|z - F(z)\|_{TV} + \beta \sup_z |J_{\beta}(z, \gamma_N) - J_{\beta}^N(z)| \\ &\quad + 2\beta \sup_z |J_{\beta}^N(z) - J_{\beta}^*(z)| + \beta \|J_{\beta}^*\|_{BL} \alpha_{\mathcal{Z}} \|z - F(z)\|_{TV}. \end{aligned}$$

Thus, using (i), and  $\|J_{\beta}^*\|_{BL} \leq \frac{2-\beta}{(1-\beta)(1-\alpha_{\mathcal{Z}}\beta)}\|c\|_{\infty}$ , we can write

$$\sup_z |J_{\beta}(z, \gamma_N) - J_{\beta}^N(z)| \leq \frac{(1 + \beta)((\alpha_{\mathcal{Z}} - 1)\beta + 1)}{(1 - \beta)^3(1 - \alpha_{\mathcal{Z}}\beta)} \|c\|_{\infty} \bar{L}_{TV}.$$

We then conclude that

$$\begin{aligned} \sup_z |J_{\beta}(z, \gamma_N) - J_{\beta}^*(z)| &\leq \sup_z |J_{\beta}(z, \gamma_N) - J_{\beta}^N(z)| + \sup_z |J_{\beta}^N(z) - J_{\beta}^*(z)| \\ &\leq \frac{(\alpha_{\mathcal{Z}} - 1)\beta + 1}{(1 - \beta)^2(1 - \alpha_{\mathcal{Z}}\beta)} \|c\|_{\infty} \bar{L}_{TV} + \frac{(1 + \beta)((\alpha_{\mathcal{Z}} - 1)\beta + 1)}{(1 - \beta)^3(1 - \alpha_{\mathcal{Z}}\beta)} \|c\|_{\infty} \bar{L}_{TV} \\ &= \frac{2(1 + (\alpha_{\mathcal{Z}} - 1)\beta)}{(1 - \beta)^3(1 - \alpha_{\mathcal{Z}}\beta)} \|c\|_{\infty} \bar{L}_{TV}. \end{aligned}$$

■

### 4.3 A Discussion on the Controlled Filter Stability Problem

The above results suggest that the loss occurring from applying a finite window policy is mainly controlled by the term

$$\sup_{\pi \in \mathcal{P}(\mathbb{X})} \sup_{\gamma \in \Gamma} E_{\pi}^{\gamma} \left[ \rho_{BL} \left( P^{\pi}(\cdot | Y_{[0,N]}, U_{[0,N-1]}), P^{\hat{\pi}}(\cdot | Y_{[0,N]}, U_{[0,N-1]}) \right) \right], \quad (13)$$

or its variations where the  $BL$  metric is replaced with total variation, or the expectation is replaced with a uniform bound (given in (11)). All these terms describe how fast two different belief-state processes forget their initial priors when fed with the same observations/control actions under a true distribution. Thus, any bound for this term directly applies to the main results we presented for the loss caused by a finite window policy. This term is related to the filter stability problem and our approximation results point out the close relation between filter stability and the performance of finite window policies. In a way, the main result or message of this paper is perhaps to explicitly relate finite window approximations for POMDPs to the filter stability problem.

To bound the term (13), we use Theorem 5, to achieve an exponential convergence rate in the window size for a controlled setup. However, we should note that, Theorem 5 provides only a sufficient condition to bound the filter stability term geometrically fast in the total variation distance. This result can be seen as too strong, if one is only interested in making this  $\rho_{BL}$  distance (13) smaller with increasing window size. In fact, as we will see in Section 5, even when the assumptions of Theorem 5 are not satisfied, the filter stability term still converges to 0. In the literature, there are various set of assumptions to achieve filter stability. Two main approaches have been:

- The transition kernel is in some sense sufficiently ergodic, forgetting the initial measure and therefore passing this insensitivity (to incorrect initializations) on to the filter process. This condition is often tailored towards control-free models.
- The measurement channel provides sufficient information about the underlying state, allowing the filter to track the true state process. This approach is typically based on martingale methods and accordingly does not often lead to rates of convergence for the filter stability problem, but only asymptotic filter stability.

The result we use in this paper (Theorem 5) provides exponential filter stability, using a joint contraction property of the Bayesian filter update and measurement update steps through the Dobrushin coefficient. When these requirements are not satisfied, the filter stability can be checked via different assumptions from the literature. However, we also note that, for the controlled setup, filter stability results are limited compared to the control free setup. A comprehensive review on filter stability in the control-free case, is available by Chigansky et al. (2009). In the controlled case, recent additional studies include McDonald and Yüksel (2019); McDonald and Yüksel (2022) where martingale methods are used to arrive at controlled filter stability, which in turn can lead to weaker conditions (though without rates of convergence) for near optimality of finite window policies.

With regard to Theorem 5, the following example studies an additive Gaussian model (not necessarily linear) and provides different parameters for the condition  $(1 - \delta(\mathcal{T})) \times (2 - \delta(\hat{Q})) < 1$  to hold.

**Example 2** Consider a system where  $\mathbb{X} = \mathbb{Y} = \mathbb{R}$  and the transition and measurement kernels are given by

$$x_{n+1} = f(x_n, u_n) + N(0, \sigma_t^2), \quad y_n = g(x_n) + N(0, \sigma_q^2)$$

where the functions  $f$  and  $g$  are measurable and bounded such that  $f(x, u) \in [-t, t]$  and  $g(x) \in [-q, q]$ .

Note that, in our paper we assume and present our results from finite  $\mathbb{Y}$ . Therefore, to make the example compatible with our results, we discretize the observation space  $\mathbb{Y}$ . We provide two discretization schemes one with  $\hat{\mathbb{Y}}_1 = \{-q, q\}$  and the other with  $\hat{\mathbb{Y}}_2 = \{-q, 0, q\}$ . For discretization, we use a nearest neighbor mapping.

First, we study the observation space  $\hat{\mathbb{Y}}_1 = \{-q, q\}$ . Using the nearest neighbor mapping, we have that

$$\begin{aligned}\hat{y}_n &= -q & \text{if } y_n = g(x_n) + N(0, \sigma_q^2) \leq 0, \\ \hat{y}_n &= q & \text{if } y_n = g(x_n) + N(0, \sigma_q^2) > 0\end{aligned}$$

We then can write the following for the transition kernel  $\mathcal{T}(\cdot|x, u) \in \mathcal{P}(\mathbb{X})$

$$\mathcal{T}(dx_{t+n}|x_n, u_n) \sim N(f(x_n, u_n), \sigma_t^2)$$

and for the compound channel  $\hat{Q}(\cdot|x) \in \mathcal{P}(\hat{\mathbb{Y}}_1)$ :

$$\hat{Q}(q|x_n) = Pr(N(g(x_n), \sigma_q^2) > 0), \quad \hat{Q}(-q|x_n) = Pr(N(g(x_n), \sigma_q^2) \leq 0).$$

For these kernels, the Dobrushin coefficients can be calculated as

$$\delta(\mathcal{T}) \geq 2Pr(N(t, \sigma_t^2) \leq 0), \quad \delta(\hat{Q}) = 2Pr(N(q, \sigma_q^2) \leq 0).$$

Notice that these probabilities are fully determined by the ratio of the mean and standard deviation of the Gaussian in question,  $\sigma_t/t$  and  $\sigma_q/q$ . The higher the ratio, the higher the Dobrushin coefficient. Below, we see a list of the ratio of the transition kernel and lowest possible ratio of the measurement kernel such that  $(1 - \delta(\mathcal{T})) \times (2 - \delta(\hat{Q})) < 1$ . If the ratio of  $\sigma_q/q$  is higher than the stated value, we will get exponential stability for the given transition kernel. Note that if  $\sigma_t/t > 1.5$ , then  $\delta(\mathcal{T}) > 0.5$  which makes  $(1 - \delta(\mathcal{T})) \times (2 - \delta(\hat{Q})) < 1$  regardless of the channel for which we use 'any' in table to indicate that any channel would lead to exponential filter stability.

$\frac{\sigma_t}{t}$	1.5	1.4	1.3	1.2	1.1	1.0	0.9	0.8	0.7	0.6	0.5	0.4	0.3
$\frac{\sigma_q}{q}$	any	0.6	0.8	1.01	1.3	1.65	2.13	3.25	5.5	8.0	20.0	70.0	1000.0
$\delta(\mathcal{T})$	0.50	0.48	0.44	0.40	0.36	0.32	0.27	0.21	0.15	0.10	0.05	0.01	0.00
$\delta(\hat{Q})$	any	0.1	0.21	0.32	0.44	0.54	0.64	0.76	0.86	0.90	0.96	0.99	1.00

Table 1: Approximate minimum ratio of  $\frac{\sigma_q}{q}$  for exponential filter stability for  $\hat{\mathbb{Y}}_1 = \{-q, q\}$

We now analyze the problem for the observation space  $\hat{\mathbb{Y}}_2 = \{-q, 0, q\}$ . For this set, the nearest neighbor mapping yields that

$$\begin{aligned}\hat{y}_n &= -q & \text{if } y_n = g(x_n) + N(0, \sigma_q^2) \leq \frac{-q}{2}, \\ \hat{y}_n &= q & \text{if } y_n = g(x_n) + N(0, \sigma_q^2) > \frac{q}{2}, \\ \hat{y}_n &= 0 & \text{else.}\end{aligned}$$

For the compound channel, we then have

$$\begin{aligned}\hat{Q}(q|x_n) &= Pr(N(g(x_n), \sigma_q^2) > \frac{q}{2}), & \hat{Q}(-q|x_n) &= Pr(N(g(x_n), \sigma_q^2) \leq \frac{-q}{2}), \\ \hat{Q}(0|x_n) &= Pr(\frac{q}{2} \geq N(g(x_n), \sigma_q^2) > \frac{-q}{2}).\end{aligned}$$

For these kernels, the Dobrushin coefficients can be calculated as

$$\begin{aligned}\delta(\mathcal{T}) &= 2Pr(N(t, \sigma_t^2) \leq 0), \\ \delta(\hat{Q}) &= 2Pr(N(q, \sigma_q^2) \leq \frac{-q}{2}) + Pr(\frac{-q}{2} < N(q, \sigma_q^2) < \frac{q}{2}).\end{aligned}$$

Below, we now see a list of the ratio of the transition kernel and lowest possible ratio of the measurement kernel such that  $(1 - \delta(\mathcal{T})) \times (2 - \delta(\hat{Q})) < 1$  for the observation space  $\hat{\mathbb{Y}}_2 = \{-q, 0, q\}$ .

$\frac{\sigma_t}{t}$	1.5	1.4	1.3	1.2	1.1	1.0	0.9	0.8	0.7	0.6	0.5	0.4	0.3
$\frac{\sigma_q}{q}$	any	0.39	0.6	0.85	1.2	1.54	2.1	3.2	5.9	8.0	20.0	80.0	1000.0
$\delta(\mathcal{T})$	0.50	0.48	0.44	0.40	0.36	0.32	0.27	0.21	0.15	0.10	0.05	0.01	0.00
$\delta(\hat{Q})$	any	0.1	0.21	0.32	0.44	0.54	0.64	0.76	0.86	0.90	0.96	0.99	1.00

Table 2: Approximate minimum ratio of  $\frac{\sigma_q}{q}$  for exponential filter stability for  $\hat{\mathbb{Y}}_2 = \{-q, 0, q\}$

## 5. Numerical Study

In this section, we give an outline of the algorithm used to determine the approximate belief MDP and the finite window policy. We also present the performance of the finite window policies and the value function error for the approximate belief MDP in relation to the window size.

A summary of the algorithm is as follows:

- We determine a  $\hat{\pi} \in \mathcal{P}(\mathbb{X})$  such that it puts positive measure over the state space  $\mathbb{X}$ .
- Following Section 3, we construct the finite belief space  $\mathcal{Z}_{\hat{\pi}}^N$  by taking the Bayesian update of  $\hat{\pi}$  for all possible realizations of the form  $\{y_0, \dots, y_N, u_0, \dots, u_{N-1}\}$ . Hence we get a approximate finite belief space with size  $|\mathbb{Y}|^{N+1} \times |\mathbb{U}|^N$ .
- We calculate all transitions from each  $z \in \mathcal{Z}_{\hat{\pi}}^N$  by considering every possible observation  $y$  and control action  $u$  and we map them to  $\mathcal{Z}_{\hat{\pi}}^N$  using a nearest neighbor map and construct the transition kernels  $\eta^N$  for the finite model.
- For the finite models obtained, through value or policy iteration, we calculate the value functions and optimal policies.

The example we use is a machine repair problem. In this model, we have  $\mathbb{X}, \mathbb{Y}, \mathbb{U} = \{0, 1\}$  with

$$x_t = \begin{cases} 1 & \text{machine is working at time } t \\ 0 & \text{machine is not working at time } t. \end{cases} \quad u_t = \begin{cases} 1 & \text{machine is being repaired at time } t \\ 0 & \text{machine is not being repaired at time } t. \end{cases}$$

The probability that the repair was successful given initially the machine was not working is given by  $\kappa$ :

$$Pr(x_{t+1} = 1|x_t = 0, u_t = 1) = \kappa$$

The probability that the machine breaks down while in a working state is given by  $\theta$ :

$$Pr(x_t = 0|x_t = 1, u_t = 0) = \theta$$

The probability that the channel gives an incorrect measurement is given by  $\epsilon$ :

$$Pr(y_t = 1|x_t = 0) = Pr(y_t = 0|x_t = 1) = \epsilon$$

The one stage cost function is given by

$$c(x, u) = \begin{cases} R + E & x = 0, u = 1 \\ E & x = 0, u = 0 \\ 0 & x = 1, u = 0 \\ R & x = 1, u = 1 \end{cases}$$

where  $R$  is the cost of repair and  $E$  is the cost incurred by a broken machine.

We study the example with discount factor  $\beta = 0.8$ , and  $R = 5, E = 1$  and present three different results by changing the other parameters. For all different cases, we choose  $\hat{\pi}(\cdot) = 0.1\delta_0(\cdot) + 0.9\delta_1(\cdot)$ .

For the first case, we take  $\epsilon = 0.3, \kappa = 0.2, \theta = 0.1$ . We analyze the performance for  $N \in \{0, \dots, 5\}$ . To get a proper finite window policy for all  $N$ 's we use, we let the system run 5 time steps under the policy  $\gamma_0(I_t) = 0$  that is  $u_t = 0$  no matter what for the first 5 time steps. Then, we start applying the policy and start calculating the cost. In Figure 1, we plot the approximate value functions and the cost incurred by the finite window policies, that is we plot the terms  $E^{\gamma_0} [J_\beta^N(Z_5)|Y_{[0,5]}, U_{[0,4]}]$  and  $E^{\gamma_0} [J_\beta(Z_5, \gamma^N)|Y_{[0,5]}, U_{[0,4]}]$ .

Recall that our main result Theorem 12, emphasizes the relation between the filter stability term and 'value error' and 'robustness error', both of which are guaranteed to converge to zero with increasing window length, where

$$\text{Filter stability term: } \sup_{\pi} \sup_{\gamma} E_{\pi}^{\gamma} \left[ \rho_{BL} \left( P^{\pi}(\cdot|Y_{[0,N]}), P^{\hat{\pi}}(\cdot|Y_{[0,N]}) \right) \right]$$

$$\text{Value error: } E^{\gamma_0} [J_\beta^N(Z_5) - J_\beta^*(Z_5)|Y_{[0,5]}, U_{[0,4]}]$$

$$\text{Robustness error: } E^{\gamma_0} [J_\beta(Z_5, \gamma^N) - J_\beta^*(Z_5)|Y_{[0,5]}, U_{[0,4]}]$$

To get the errors, we simply subtract the cost values from their minimum (largest window) which serves as an approximation of the value function. Furthermore, we scale the errors according to the filter stability term to get a better understanding of the error rate in relation to the filter stability constant that is we make them start from the same value to see the decrease rates more clearly. Figure 2 shows the relation.




Figure 1: Approximate value function and the performance of finite window policy for different window sizes




Figure 2: Comparison of approximation error and filter stability term for different window sizes

It can be seen from Figure 2, that the error rate stays below the filter stability convergence rate and goes to 0 as the window size increases.

For the second case, we take  $\epsilon = 0.01$ ,  $\kappa = 0.3$ ,  $\theta = 0.1$ . For this case, we directly plot the scaled errors in relation to the filter stability term in Figure 3.

It can be seen that as the channel is more informative, the filter stability term gets smaller much faster and the recent information becomes more informative. As a result, we get a faster decrease in the error with the increasing window size.

One can notice that for the first two cases, the Dobrushin coefficient  $\alpha$ , we defined in Assumption 4 is greater than 1, however, the error terms still converge to 0 with increasing window. This is because the filter stability term gets smaller even though the Dobrushin constant  $\alpha > 1$ . For the last case, we select the parameters in a way to make  $\alpha < 1$ .

We take  $\epsilon = 0.3$ ,  $\kappa = 0.4$ ,  $\theta = 0.3$  so that  $\alpha = 0.7$ . In Figure 4, it can be observed that the error terms and the filter stability term converges to 0 at similar rates.  $\alpha^N$ , however, gets smaller at a slower rate and upper bounds the convergence rate of error terms.






Figure 3: Comparison of approximation error and filter stability term for different window sizes for an informative channel




Figure 4: Comparison of approximation error, filter stability term and Dobrushin term for different window sizes

## 6. Proofs of Main Technical Results

In this section, we provide the technical proofs.

### 6.1 Proof of Theorem 7

We will build on the proof by Kara et al. (2019). Let  $\mathbb{X}$  be a separable metric space. Another metric that metrizes the weak topology on  $\mathcal{P}(\mathbb{X})$  is the following:

$$\rho(\mu, \nu) = \sum_{m=1}^{\infty} 2^{-(m+1)} \left| \int_{\mathbb{X}} f_m(x) \mu(dx) - \int_{\mathbb{X}} f_m(x) \nu(dx) \right|, \quad (14)$$

where  $\{f_m\}_{m \geq 1}$  is an appropriate sequence of continuous and bounded functions such that  $\|f_m\|_{\infty} \leq 1$  for all  $m \geq 1$  (see Parthasarathy, 1967, Theorem 6.6, p. 47).

We equip  $\mathcal{P}(\mathbb{X})$  with the metric  $\rho$  to define bounded-Lipschitz norm  $\|f\|_{BL}$  of any Borel measurable function  $f : \mathcal{P}(\mathbb{X}) \rightarrow \mathbb{R}$ . With this metric, we can start the proof:

$$\begin{aligned}
 & \sup_{\|f\|_{BL} \leq 1} \left| \int_{\mathcal{P}(\mathbb{X})} f(z_1) \eta(dz_1 | z'_0, u) - \int_{\mathcal{P}(\mathbb{X})} f(z_1) \eta(dz_1 | z_0, u) \right| \\
 &= \sup_{\|f\|_{BL} \leq 1} \left| \int_{\mathbb{Y}} f(z_1(z'_0, u, y_1)) P(dy_1 | z'_0, u) - \int_{\mathbb{Y}} f(z_1(z_0, u, y_1)) P(dy_1 | z_0, u) \right| \\
 &\leq \sup_{\|f\|_{BL} \leq 1} \left| \int_{\mathbb{Y}} f(z_1(z'_0, u, y_1)) P(dy_1 | z'_0, u) - \int_{\mathbb{Y}} f(z_1(z'_0, u, y_1)) P(dy_1 | z_0, u) \right| \\
 &\quad + \sup_{\|f\|_{BL} \leq 1} \int_{\mathbb{Y}} |f(z_1(z'_0, u, y_1)) - f(z_1(z_0, u, y_1))| P(dy_1 | z_0, u) \\
 &\leq \|P(\cdot | z'_0, u) - P(\cdot | z_0, u)\|_{TV} \tag{15}
 \end{aligned}$$

$$\begin{aligned}
 & \sup_{\|f\|_{BL} \leq 1} \int_{\mathbb{Y}} |f(z_1(z'_0, u, y_1)) - f(z_1(z_0, u, y_1))| P(dy_1 | z_0, u) \tag{16}
 \end{aligned}$$

where, in the last inequality, we have used  $\|f\|_{\infty} \leq \|f\|_{BL} \leq 1$ .

We first analyze the first term:

$$\begin{aligned}
 \|P(\cdot | z'_0, u) - P(\cdot | z_0, u)\|_{TV} &= \sup_{\|f\|_{\infty} \leq 1} \left| \int f(y_1) P(dy_1 | z'_0, u) - \int f(y_1) P(dy_1 | z_0, u) \right| \\
 &= \sup_{\|f\|_{\infty} \leq 1} \left| \int_{\mathbb{X}} \int_{\mathbb{X}} \int_{\mathbb{Y}} f(y_1) Q(dy_1 | x_1) \mathcal{T}(dx_1 | x_0, u) z'_0(dx_0) \right. \\
 &\quad \left. - \int_{\mathbb{X}} \int_{\mathbb{X}} \int_{\mathbb{Y}} f(y_1) Q(dy_1 | x_1) \mathcal{T}(dx_1 | x_0, u) z_0(dx_0) \right|.
 \end{aligned}$$

For all  $x'_0, x_0$  and for any  $\|f\|_{\infty} \leq 1$ , we have

$$\begin{aligned}
 & \left| \int_{\mathbb{X}} \int_{\mathbb{Y}} f(y_1) Q(dy_1 | x_1) \mathcal{T}(dx_1 | x'_0, u) - \int_{\mathbb{X}} \int_{\mathbb{Y}} f(y_1) Q(dy_1 | x_1) \mathcal{T}(dx_1 | x_0, u) \right| \\
 & \leq \|\mathcal{T}(\cdot | x'_0, u) - \mathcal{T}(\cdot | x_0, u)\|_{TV} \leq \alpha_{\mathbb{X}} |x'_0 - x_0|.
 \end{aligned}$$

Then, we can write

$$\|P(\cdot | z'_0, u) - P(\cdot | z_0, u)\|_{TV} \leq (1 + \alpha_{\mathbb{X}}) \rho_{BL}(z'_0, z_0).$$

Recall the metric introduced in (14); using this metric, we now focus on the term (16):

$$\begin{aligned}
 & \sup_{\|f\|_{BL} \leq 1} \int_{\mathbb{Y}} |f(z_1(z'_0, u, y_1)) - f(z_1(z_0, u, y_1))| P(dy_1 | z_0, u) \\
 & \leq \int_{\mathbb{Y}} \rho(z_1(z'_0, u, y_1), z_1(z_0, u, y_1)) P(dy_1 | z_0, u) \\
 & = \int_{\mathbb{Y}} \sum_{m=1}^{\infty} 2^{-m+1} \left| \int_{\mathbb{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) \right|
 \end{aligned}$$

$$\begin{aligned}
 & - \int_{\mathbf{X}} f_m(x_1) z_1(z_0, u, y_1)(dx_1) \Big| P(dy_1 | z_0, u) \\
 = & \sum_{m=1}^{\infty} 2^{-m+1} \int_{\mathbf{Y}} \Big| \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) \\
 & - \int_{\mathbf{X}} f_m(x_1) z_1(z_0, u, y_1)(dx_1) \Big| P(dy_1 | z_0, u),
 \end{aligned}$$

where we have used Fubini's theorem with the fact that  $\sup_m \|f_m\|_{\infty} \leq 1$ . For each  $m$ , let us define

$$\begin{aligned}
 I_+ & := \left\{ y_1 \in \mathbf{Y} : \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) > \int_{\mathbf{X}} f_m(x_1) z_1(z_0, u, y_1)(dx_1) \right\} \\
 I_- & := \left\{ y_1 \in \mathbf{Y} : \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) \leq \int_{\mathbf{X}} f_m(x_1) z_1(z_0, u, y_1)(dx_1) \right\}. \quad (17)
 \end{aligned}$$

Then, we can write

$$\begin{aligned}
 & \int_{\mathbf{Y}} \Big| \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) - \int_{\mathbf{X}} f_m(x_1) z_1(z_0, u, y_1)(dx_1) \Big| P(dy_1 | z_0, u) \\
 = & \int_{I_+} \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) P(dy_1 | z_0, u) - \int_{I_+} \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) P(dy_1 | z'_0, u) \\
 & + \int_{I_+} \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) P(dy_1 | z'_0, u) - \int_{I_+} \int_{\mathbf{X}} f_m(x_1) z_1(z_0, u, y_1)(dx_1) P(dy_1 | z_0, u) \\
 & + \int_{I_-} \int_{\mathbf{X}} f_m(x_1) z_1(z_0, u, y_1)(dx_1) P(dy_1 | z_0, u) - \int_{I_-} \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) P(dy_1 | z'_0, u) \\
 & + \int_{I_-} \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) P(dy_1 | z'_0, u) - \int_{I_-} \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) P(dy_1 | z_0, u)
 \end{aligned}$$

For the first and the last term, we can write

$$\begin{aligned}
 & \int_{I_+} \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) P(dy_1 | z_0, u) - \int_{I_+} \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) P(dy_1 | z'_0, u) \\
 & + \int_{I_-} \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) P(dy_1 | z'_0, u) - \int_{I_-} \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) P(dy_1 | z_0, u) \\
 = & \int_{\mathbf{Y}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) (\mathbf{1}_{I_-} - \mathbf{1}_{I_+}) P(dy_1 | z'_0, u) \\
 & - \int_{\mathbf{Y}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) (\mathbf{1}_{I_-} - \mathbf{1}_{I_+}) P(dy_1 | z_0, u) \\
 \leq & \|P(\cdot | z'_0, u) - P(\cdot | z_0, u)\|_{TV} \leq (1 + \alpha_{\mathbf{X}}) \rho_{BL}(z'_0, z_0). \quad (18)
 \end{aligned}$$

For the second and the third term:

$$\begin{aligned}
 & \int_{I_+} \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) P(dy_1 | z'_0, u) - \int_{I_+} \int_{\mathbf{X}} f_m(x_1) z_1(z_0, u, y_1)(dx_1) P(dy_1 | z_0, u) \\
 & + \int_{I_-} \int_{\mathbf{X}} f_m(x_1) z_1(z_0, u, y_1)(dx_1) P(dy_1 | z_0, u) - \int_{I_-} \int_{\mathbf{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) P(dy_1 | z'_0, u)
 \end{aligned}$$

$$\begin{aligned}
 &= \int_{I_+} \int_{\mathbb{X}} f_m(x_1) Q(dy_1|x_1) \mathcal{T}(dx_1|z'_0, u) - \int_{I_+} \int_{\mathbb{X}} f_m(x_1) Q(dy_1|x_1) \mathcal{T}(dx_1|z_0, u) \\
 &\quad + \int_{I_-} \int_{\mathbb{X}} f_m(x_1) Q(dy_1|x_1) \mathcal{T}(dx_1|z_0, u) - \int_{I_-} \int_{\mathbb{X}} f_m(x_1) Q(dy_1|x_1) \mathcal{T}(dx_1|z'_0, u) \\
 &= \int_{\mathbb{X}} f_m(x_1) (Q(I_+|x_1) - Q(I_-|x_1)) \mathcal{T}(dx_1|z'_0, u) \\
 &\quad - \int_{\mathbb{X}} f_m(x_1) (Q(I_+|x_1) - Q(I_-|x_1)) \mathcal{T}(dx_1|z_0, u) \\
 &\leq (1 + \alpha_{\mathbb{X}}) \rho_{BL}(z'_0, z_0).
 \end{aligned} \tag{19}$$

Hence, we can write that

$$\begin{aligned}
 &\int_{\mathbb{Y}} \left| \int_{\mathbb{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) - \int_{\mathbb{X}} f_m(x_1) z_1(z_0, u, y_1)(dx_1) \right| P(dy_1|z_0, u) \\
 &\leq 2(1 + \alpha_{\mathbb{X}}) \rho_{BL}(z'_0, z_0)
 \end{aligned}$$

Now we go back to the term (16):

$$\begin{aligned}
 &\sup_{\|f\|_{BL} \leq 1} \int_{\mathbb{Y}} |f(z_1(z'_0, u, y_1)) - f(z_1(z_0, u, y_1))| P(dy_1|z_0, u) \\
 &\leq \int_{\mathbb{Y}} \rho(z_1(z'_0, u, y_1), z_1(z_0, u, y_1)) P(dy_1|z_0, u) \\
 &= \int_{\mathbb{Y}} \sum_{m=1}^{\infty} 2^{-m+1} \left| \int_{\mathbb{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) - \int_{\mathbb{X}} f_m(x_1) z_1(z_0, u, y_1)(dx_1) \right| P(dy_1|z_0, u) \\
 &= \sum_{m=1}^{\infty} 2^{-m+1} \int_{\mathbb{Y}} \left| \int_{\mathbb{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) - \int_{\mathbb{X}} f_m(x_1) z_1(z_0, u, y_1)(dx_1) \right| P(dy_1|z_0, u) \\
 &\leq 2(1 + \alpha_{\mathbb{X}}) \rho_{BL}(z'_0, z_0).
 \end{aligned}$$

Thus, combining all the results:

$$\begin{aligned}
 &\sup_{\|f\|_{BL} \leq 1} \left| \int_{\mathcal{P}(\mathbb{X})} f(z_1) \eta(dz_1|z'_0, u) - \int_{\mathcal{P}(\mathbb{X})} f(z_1) \eta(dz_1|z_0, u) \right| \\
 &\leq \|P(\cdot|z'_0, u) - P(\cdot|z_0, u)\|_{TV} \\
 &\quad + \sup_{\|f\|_{BL} \leq 1} \int_{\mathbb{Y}} |f(z_1(z'_0, u, y_1)) - f(z_1(z_0, u, y_1))| P(dy_1|z_0, u) \\
 &\leq 3(1 + \alpha_{\mathbb{X}}) \rho_{BL}(z'_0, z_0).
 \end{aligned}$$

We now prove part (ii); we note first that if  $\|P(\cdot|z'_0, u) - P(\cdot|z_0, u)\|_{TV} \leq 2(1 - \delta(Q)) \rho_{BL}(z'_0, z_0)$ , the result follows. Hence, we write (by Dobrushin, 1956)

$$\begin{aligned}
 &\|P(\cdot|z'_0, u) - P(\cdot|z_0, u)\|_{TV} = \sup_{\|f\|_{\infty} \leq 1} \left| \int f(y_1) P(dy_1|z'_0, u) - \int f(y_1) P(dy_1|z_0, u) \right| \\
 &= \sup_{\|f\|_{\infty} \leq 1} \left| \int f(y_1) Q(dy_1|x_1) \mathcal{T}(dx_1|z'_0, u) - \int f(y_1) Q(dy_1|x_1) \mathcal{T}(dx_1|z_0, u) \right|
 \end{aligned}$$

$$\leq (1 - \delta(Q)) \|\mathcal{T}(dx_1|z'_0, u) - \mathcal{T}(dx_1|z_0, u)\|_{TV} \leq (1 - \delta(Q))(1 + \alpha_X) \rho_{BL}(z'_0, z_0).$$

To prove part (iii) of the theorem, we start from the terms (15) and (16). For (15), we have

$$\begin{aligned} & \|P(\cdot|z'_0, u) - P(\cdot|z_0, u)\|_{TV} \\ &= \sup_{\|f\|_\infty \leq 1} \left| \int f(y_1) Q(dy_1|x_1) \mathcal{T}(dx_1|x_0, u) z'_0(dx_0) - \int f(y_1) Q(dy_1|x_1) \mathcal{T}(dx_1|x_0, u) z_0(dx_0) \right| \\ &\leq \sup_{\|f\|_\infty \leq 1} \left| \int f(x_0) z_0(dx_0) - \int f(x_0) z'_0(dx_0) \right| = \|z_0 - z'_0\|_{TV}. \end{aligned}$$

For (16), as before, we start by writing

$$\begin{aligned} & \sup_{\|f\|_{BL} \leq 1} \int_{\mathbb{Y}} |f(z_1(z'_0, u, y_1)) - f(z_1(z_0, u, y_1))| P(dy_1|z_0, u) \\ &\leq \sum_{m=1}^{\infty} 2^{-m+1} \int_{\mathbb{Y}} \left| \int_{\mathbb{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) \right. \\ &\quad \left. - \int_{\mathbb{X}} f_m(x_1) z_1(z_0, u, y_1)(dx_1) \right| P(dy_1|z_0, u). \end{aligned}$$

For each  $m$ , using the bounding terms (18) and (19), we have

$$\begin{aligned} & \int_{\mathbb{Y}} \left| \int_{\mathbb{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) - \int_{\mathbb{X}} f_m(x_1) z_1(z_0, u, y_1)(dx_1) \right| P(dy_1|z_0, u) \\ &\leq \|P(\cdot|z'_0, u) - P(\cdot|z_0, u)\|_{TV} + \|\mathcal{T}(\cdot|z'_0, u) - \mathcal{T}(\cdot|z_0, u)\|_{TV} \leq 2\|z'_0 - z_0\|_{TV} \end{aligned}$$

which completes the proof of part (iii).

For part (iv), for (15) we have (by Dobrushin, 1956)

$$\begin{aligned} & \|P(\cdot|z'_0, u) - P(\cdot|z_0, u)\|_{TV} \\ &= \sup_{\|f\|_\infty \leq 1} \left| \int f(y_1) Q(dy_1|x_1) \mathcal{T}(dx_1|z'_0, u) - \int f(y_1) Q(dy_1|x_1) \mathcal{T}(dx_1|z_0, u) \right| \\ &\leq (1 - \delta(Q)) \|\mathcal{T}(dx_1|z'_0, u) - \mathcal{T}(dx_1|z_0, u)\|_{TV} \\ &= (1 - \delta(Q)) \sup_{\|f\|_\infty \leq 1} \left| \int f(x_1) \mathcal{T}(dx_1|x_0, u) z'_0(dx_0) - \int f(x_1) \mathcal{T}(dx_1|x_0, u) z_0(dx_0) \right| \\ &\leq (1 - \delta(Q))(1 - \tilde{\delta}(\mathcal{T})) \|z'_0 - z_0\|_{TV}. \end{aligned}$$

For (16), (using Lemma 3.2 McDonald and Yüksel, 2020) and the analysis above,

$$\begin{aligned} & \int_{\mathbb{Y}} \left| \int_{\mathbb{X}} f_m(x_1) z_1(z'_0, u, y_1)(dx_1) - \int_{\mathbb{X}} f_m(x_1) z_1(z_0, u, y_1)(dx_1) \right| P(dy_1|z_0, u) \\ &\leq (2 - \delta(Q)) \|\mathcal{T}(\cdot|z'_0, u) - \mathcal{T}(\cdot|z_0, u)\|_{TV} \leq (2 - \delta(Q))(1 - \tilde{\delta}(\mathcal{T})) \|z'_0 - z_0\|_{TV} \end{aligned}$$

which completes the proof.

## 6.2 Proof of Theorem 9

Before presenting our proof program and the series of technical results needed, we introduce some notation.

We denote the loss function due to the quantization by  $L(z)$ , i.e.

$$L(z) = \rho_{BL}(z, F(z))$$

with  $F$  defined in Equation 8.

In the following, to specify the probability measures according to which the expectations are taken, we use the following notation; for the kernel  $\eta^N$  and a policy  $\gamma$ , we use  $E_N^\gamma$  and for the kernel  $\eta$  and a policy  $\gamma$ , we will use  $E^\gamma$ .

The last notation we introduce is the following: Recall that we denote the optimal cost for the finite model by  $J_\beta^N$  and the optimal policy by  $\gamma_N^*$ . These are defined on a finite set  $\mathcal{Z}_\pi^N$ , however, we can always extend them over the original state space  $\mathcal{Z}$  so that they are constant over the quantization bins. We denote the extended versions by  $\tilde{J}_\beta^N$  and  $\tilde{\gamma}_N^*$ .

Now, we introduce our value iteration approach for the original model and the finite model. We write for any  $k < \infty$

$$\begin{aligned} v_{k+1}(z) &= \min_u \left( \tilde{c}(z, u) + \beta \int v_k(z_1) \eta(dz_1|z, u) \right) \quad \forall z \in \mathcal{Z}, \\ v_{k+1}^N(z) &= \min_u \left( c^N(z, u) + \beta \sum_{z_1} v_k^N(z_1) \eta^N(z_1|z, u) \right) \quad \forall z \in \mathcal{Z}_\pi^N. \end{aligned}$$

where  $v_0, v_0^N \equiv 0$ . It is well known that the above operations define contractions under either model and hence the value functions converge to the optimal expected discounted cost. In particular, we have that

$$|J_\beta^N(z) - v_k^N(z)| \leq \|\tilde{c}\|_\infty \frac{\beta^k}{1-\beta}, \quad |J_\beta^*(z) - v_k(z)| \leq \|\tilde{c}\|_\infty \frac{\beta^k}{1-\beta}.$$

Notice that if we extend  $v_{k+1}^N(z)$  and  $c^N(z, u)$  over all of the state space then  $\tilde{v}_{k+1}^N(z)$  and  $\tilde{c}^N(z, u)$  becomes constant over the quantization bins. Then, the extended value functions for the finite model can be also obtained with

$$\tilde{v}_{k+1}^N(z) = \min_u \left( \tilde{c}(F(z), u) + \beta \int \tilde{v}_k^N(z_1) \eta(dz_1|F(z), u) \right) \quad \forall z \in \mathcal{Z}. \quad (20)$$

To see this, observe the following:

$$\begin{aligned} \tilde{v}_{k+1}^N(z) &= \min_u \left( c(F(z), u) + \beta \int \tilde{v}_k^N(z_1) \eta(dz_1|F(z), u) \right) \\ &= \min_u \left( \tilde{c}(F(z), u) + \beta \sum_{i=1}^{|\mathcal{Z}_\pi^N|} \int_{\mathcal{Z}_i^N} \tilde{v}_k^N(z_1) \eta(dz_1|F(z), u) \right) \\ &= \min_u \left( \tilde{c}(F(z), u) + \beta \sum_{i=1}^{|\mathcal{Z}_\pi^N|} \tilde{v}_k^N(z_1^i) \int_{\mathcal{Z}_i^N} \eta(dy|F(z), u) \right) \end{aligned}$$

$$= \min_u \left( \tilde{c}(F(z), u) + \beta \sum_{i=1}^{|\mathcal{Z}_{\hat{\pi}}^N|} \tilde{v}_k^N(z_1^i) \eta(\mathcal{Z}_i^N | F(z), u) \right)$$

where  $\mathcal{Z}_i^N$  denotes the  $i$ th quantization bin and  $\tilde{v}_k^N(z_1^i)$  denotes the value of  $\tilde{v}_k^N$  at that quantization bin and it is constant over the bin.

For the proof the goal is to bound the term  $|J_\beta(\tilde{\gamma}_N^*, z) - J_\beta(\gamma^*, z)|$ . We will see in the following that bounding this term is related to studying the term  $|\tilde{J}_\beta^N(\tilde{\gamma}_N^*, z) - J_\beta(\gamma^*, z)|$  (the value function approximation error). Therefore, in what follows, we first analyze  $|\tilde{J}_\beta^N(\tilde{\gamma}_N^*, z) - J_\beta(\gamma^*, z)|$  and observe that it can be bounded using the expected loss terms  $E[L(Z_t)]$ . Finally, we will observe that upper bounds on these expressions can be written through the filter stability term

$$\sup_{\pi \in \mathcal{P}(\mathbb{X})} \sup_{\gamma \in \Gamma} E_\pi^\gamma \left[ \rho_{BL} \left( P^\pi(X_N \in \cdot | Y_{[0,N]}, \gamma(Y_{[0,N-1]})), P^{\hat{\pi}}(X_N \in \cdot | Y_{[0,N]}, \gamma(Y_{[0,N-1]})) \right) \right].$$

Next, we present some supporting technical results.

**Lemma 18** *Under Assumption 3, we have*

$$\begin{aligned} & |\tilde{J}_\beta^N(\tilde{\gamma}_N^*, z) - J_\beta(\gamma^*, z)| \\ & \leq \limsup_{k \rightarrow \infty} \sup_{\gamma \in \Gamma} \left( \alpha_{\tilde{c}} \sum_{t=0}^{k-1} \beta^t \left( E_z^\gamma [L(Z_t)] + 3\alpha_{\mathcal{Z}} \sum_{m=0}^{t-1} (4\alpha_{\mathcal{Z}} + 1)^m E_z^\gamma [L(Z_{t-m-1})] \right) \right. \\ & \quad \left. + \alpha_{\mathcal{Z}} \sum_{t=0}^{k-1} \beta^{t+1} \|v_{k-t-1}\|_{BL} \left( E_z^\gamma [L(Z_t)] + 3\alpha_{\mathcal{Z}} \sum_{m=0}^{t-1} (4\alpha_{\mathcal{Z}} + 1)^m E_z^\gamma [L(Z_{t-m-1})] \right) \right). \end{aligned}$$

where  $v_k$  denotes the value function at a time step  $k$  that is produced by the value iteration with  $v_0 \equiv 0$ .

**Proof**

We follow a value iteration approach to write for any  $k < \infty$

$$\begin{aligned} v_{k+1}(z) &= \min_u \left( \tilde{c}(z, u) + \beta \int v_k(y) \eta(dy | z, u) \right) \quad \forall z \in \mathcal{Z}, \\ \tilde{v}_{k+1}^N(z) &= \min_u \left( \tilde{c}(F(z), u) + \beta \int \tilde{v}_k^N(y) \eta(dy | F(z), u) \right) \quad \forall z \in \mathcal{Z}. \end{aligned}$$

where  $v_0, \tilde{v}_0^N \equiv 0$ .

We then write

$$|\tilde{J}_\beta^N(z) - J_\beta^*(z)| \leq |\tilde{J}_\beta^N(z) - \tilde{v}_k^N(z)| + |\tilde{v}_k^N(z) - v_k(z)| + |v_k(z) - J_\beta^*(z)| \quad (21)$$

For the first and the last term, using the fact that the dynamic programming operator is a contraction, we have that

$$|\tilde{J}_\beta^N(z) - \tilde{v}_k^N(z)| \leq \|\tilde{c}\|_\infty \frac{\beta^k}{1-\beta}, \quad |J_\beta^*(z) - v_k(z)| \leq \|\tilde{c}\|_\infty \frac{\beta^k}{1-\beta}.$$

Now, we focus on the second term  $|\tilde{v}_k^N(z) - v_k(z)|$ .

**Step 1:** We show in the Appendix, Section A.1, that for all  $k \geq 1$ ,

$$\begin{aligned} & |\tilde{v}_k^N(z) - v_k(z)| \\ & \leq \sup_{\gamma \in \Gamma} \left( \alpha_{\tilde{z}} \sum_{t=0}^{k-1} \beta^t E_{z,N}^\gamma [L(Z_t)] + \sum_{t=0}^{k-1} \beta^{t+1} \|v_{k-t-1}\|_{BL} E_{z,N}^\gamma [L(Z_t)] \alpha_{\mathcal{Z}} \right) \end{aligned}$$

where  $L(z)$  is the loss function due to the quantization, i.e.  $L(z) = \rho_{BL}(z, \hat{z})$  where  $\hat{z}$  is the representative state  $z$  belongs to.

**Step 2:** We show in the Appendix, Section A.2, that

$$E_{z,N}^\gamma [L(Z_t)] \leq E_z^\gamma [L(Z_t)] + 3\alpha_{\mathcal{Z}} \sum_{m=0}^{t-1} (4\alpha_{\mathcal{Z}} + 1)^m E_z^\gamma [L(Z_{t-m-1})].$$

**Step 3:** Combining our results, we have that

$$\begin{aligned} & |\tilde{v}_k^N(z) - v_k(z)| \\ & \leq \sup_{\gamma} \left( \alpha_{\tilde{z}} \sum_{t=0}^{k-1} \beta^t E_{z,N}^\gamma [L(Z_t)] + \sum_{t=0}^{k-1} \beta^{t+1} \|v_{k-t-1}\|_{BL} E_{z,N}^\gamma [L(Z_t)] \alpha_{\mathcal{Z}} \right) \\ & \leq \sup_{\gamma} \left( \alpha_{\tilde{z}} \sum_{t=0}^{k-1} \beta^t \left( E_z^\gamma [L(Z_t)] + 3\alpha_{\mathcal{Z}} \sum_{m=0}^{t-1} (4\alpha_{\mathcal{Z}} + 1)^m E_z^\gamma [L(Z_{t-m-1})] \right) \right. \\ & \quad \left. + \alpha_{\mathcal{Z}} \sum_{t=0}^{k-1} \beta^{t+1} \|v_{k-t-1}\|_{BL} \left( E_z^\gamma [L(Z_t)] + 3\alpha_{\mathcal{Z}} \sum_{m=0}^{t-1} (4\alpha_{\mathcal{Z}} + 1)^m E_z^\gamma [L(Z_{t-m-1})] \right) \right). \end{aligned} \tag{22}$$

Hence, using (21) and (22), we write

$$\begin{aligned} & |\tilde{J}_\beta^N(z) - J_\beta^*(z)| \leq \lim_{k \rightarrow \infty} |\tilde{v}_k^N(z) - v_k(z)| \\ & \leq \lim_{k \rightarrow \infty} \sup_{\gamma} \left( \alpha_{\tilde{z}} \sum_{t=0}^{k-1} \beta^t \left( E_z^\gamma [L(Z_t)] + 3\alpha_{\mathcal{Z}} \sum_{m=0}^{t-1} (4\alpha_{\mathcal{Z}} + 1)^m E_z^\gamma [L(Z_{t-m-1})] \right) \right. \\ & \quad \left. + \alpha_{\mathcal{Z}} \sum_{t=0}^{k-1} \beta^{t+1} \|v_{k-t-1}\|_{BL} \left( E_z^\gamma [L(Z_t)] + 3\alpha_{\mathcal{Z}} \sum_{m=0}^{t-1} (4\alpha_{\mathcal{Z}} + 1)^m E_z^\gamma [L(Z_{t-m-1})] \right) \right). \end{aligned}$$

■

The next result, gives a bound on the loss function occurring from the quantization (on  $L(z)$ ) and relates the bound to the filter stability problem.

**Lemma 19** For  $Z_0 = P^{\pi_0}(X_N \in \cdot | Y_{[0,N]}, \gamma(Y_{[0,N-1]}))$ , where  $\pi_0 \in \mathcal{P}(\mathbb{X})$  is the prior distribution of  $X_0$ , we have that for any  $N < t < \infty$

$$\begin{aligned} & \sup_{\gamma \in \Gamma} E_{\pi_0}^\gamma \left[ E_{Z_0}^\gamma [L(Z_t)] | Y_{[0,N]}, \gamma(Y_{[0,N-1]}) \right] \\ & \leq \sup_{\pi \in \mathcal{P}(\mathbb{X})} \sup_{\gamma \in \Gamma} E_\pi^\gamma \left[ \rho_{BL} \left( P^\pi(X_N \in \cdot | Y_{[0,N]}, \gamma(Y_{[0,N-1]})), P^{\hat{\pi}}(X_N \in \cdot | Y_{[0,N]}, \gamma(Y_{[0,N-1]})) \right) \right]. \end{aligned}$$



**Proof** We first note that, since the quantizer is a nearest-neighbor quantization, the quantization error is almost surely upper bounded by  $\rho_{BL} \left( P^{\pi_{t-}}(X_{N+t} \in \cdot | Y_{[t,t+N]}, U_{[t,t+N-1]}), P^{\hat{\pi}}(X_{N+t} \in \cdot | Y_{[t,t+N]}, U_{[t,t+N-1]}) \right)$ . Using this and the law of total expectation:

$$\begin{aligned} & E_{\pi_0}^\gamma \left[ E_{z_0}^\gamma [L(Z_t) | Y_{[0,N]}, \gamma(Y_{[0,N-1]})] \right. \\ & \leq \sup_{\gamma \in \Gamma} E_{\pi_0}^\gamma \left[ E_{\pi_{t-}}^\gamma \left[ \rho_{BL} \left( P^{\pi_{t-}}(X_{N+t} \in \cdot | Y_{[t,t+N]}, U_{[t,t+N-1]}) \right. \right. \right. \\ & \qquad \qquad \qquad \left. \left. \left. , P^{\hat{\pi}}(X_{N+t} \in \cdot | Y_{[t,t+N]}, U_{[t,t+N-1]}) \right) \right] \middle| (U, Y)_{[0,t-1]} \right] \\ & \leq \sup_{\pi \in \mathcal{P}(\mathbb{X})} \sup_{\gamma \in \Gamma} E_\pi^\gamma \left[ \rho_{BL} \left( P^\pi(X_N \in \cdot | Y_{[0,N]}, U_{[0,N-1]}), P^{\hat{\pi}}(X_N \in \cdot | Y_{[0,N]}, U_{[0,N-1]}) \right) \right] \end{aligned}$$

where  $(U, Y)_{[0,t-1]} = U_{[0,t-1]}, Y_{[0,t-1]}$ . ■

**Lemma 20** *Under Assumption 3 we have that*

$$\|v_k\|_{BL} \leq \alpha_{\tilde{c}} \sum_{t=0}^{k-1} (\beta \alpha_{\mathcal{Z}})^t + \|\tilde{c}\|_\infty \sum_{t=0}^{k-1} (\beta \alpha_{\mathcal{Z}})^t \frac{1 - \beta^{k-t}}{1 - \beta}.$$

*In particular, we have that for all  $k$*

$$\|v_k\|_{BL} \leq \frac{1}{1 - \beta \alpha_{\mathcal{Z}}} \left( \frac{\|\tilde{c}\|_\infty}{1 - \beta} + \alpha_{\tilde{c}} \right).$$

**Proof** It is easy to see that  $\|v_k\|_\infty \leq \|\tilde{c}\|_\infty \sum_{t=0}^{k-1} \beta^t = \|\tilde{c}\|_\infty \frac{1 - \beta^k}{1 - \beta}$ . Then, we use an inductive approach and assume the claim holds for  $k$  and analyze the term  $k+1$

$$\|v_{k+1}\|_{BL} = \|v_{k+1}\|_\infty + \sup_{x \neq y} \frac{|v_{k+1}(x) - v_{k+1}(y)|}{\rho_{BL}(x, y)}$$

For the second term, we have

$$\begin{aligned} |v_{k+1}(x) - v_{k+1}(y)| & \leq \sup_u \left( |\tilde{c}(x, u) - \tilde{c}(y, u)| + \beta \left| \int v_k(z) \eta(dz|x, u) - \int v_k(z) \eta(dz|y, u) \right| \right) \\ & \leq \alpha_{\tilde{c}} \rho_{BL}(x, y) + \beta \|v_k\|_{BL} \alpha_{\mathcal{Z}} \rho_{BL}(x, y). \end{aligned}$$

Hence, using the induction hypothesis, we have that

$$\begin{aligned} \|v_{k+1}\|_{BL} & \leq \|v_{k+1}\|_\infty + \alpha_{\tilde{c}} + \beta \alpha_{\mathcal{Z}} \|v_k\|_{BL} \\ & \leq \|\tilde{c}\|_\infty \frac{1 - \beta^{k+1}}{1 - \beta} + \alpha_{\tilde{c}} + \beta \alpha_{\mathcal{Z}} \left( \alpha_{\tilde{c}} \sum_{t=0}^{k-1} (\beta \alpha_{\mathcal{Z}})^t + \|\tilde{c}\|_\infty \sum_{t=0}^{k-1} (\beta \alpha_{\mathcal{Z}})^t \frac{1 - \beta^{k-t}}{1 - \beta} \right) \\ & = \alpha_{\tilde{c}} \sum_{t'=0}^k (\beta \alpha_{\mathcal{Z}})^{t'} + \|\tilde{c}\|_\infty \sum_{t'=0}^k (\beta \alpha_{\mathcal{Z}})^{t'} \frac{1 - \beta^{k-t'}}{1 - \beta}, \quad t' = t + 1, \end{aligned}$$

which concludes the induction argument. Therefore, we can write

$$\begin{aligned} \|v_k\|_{BL} &\leq \alpha_{\tilde{c}} \sum_{t=0}^{k-1} (\beta \alpha_{\mathcal{Z}})^t + \|\tilde{c}\|_{\infty} \sum_{t=0}^{k-1} (\beta \alpha_{\mathcal{Z}})^t \frac{1 - \beta^{k-t}}{1 - \beta} \\ &= \left( \alpha_{\tilde{c}} + \frac{\|\tilde{c}\|_{\infty}}{1 - \beta} \right) \frac{1 - (\beta \alpha_{\mathcal{Z}})^k}{1 - \beta \alpha_{\mathcal{Z}}} - \frac{\|\tilde{c}\|_{\infty} \beta^k (1 - \alpha_{\mathcal{Z}}^k)}{(1 - \beta)(1 - \alpha_{\mathcal{Z}})}. \end{aligned}$$

■

Now, we can prove Theorem 9.

*Proof of Theorem 9*

Consider the following dynamic programming operators,  $\hat{T}_n, T_n : \mathcal{B}_b(\mathcal{Z}) \rightarrow \mathcal{B}_b(\mathcal{Z})$  ( $\mathcal{B}_b(\mathcal{Z})$  denotes the set of measurable and bounded functions on  $\mathcal{Z}$ ) such that for any  $f \in \mathcal{B}_b(\mathcal{Z})$  and for any  $z \in \mathcal{Z}$

$$\begin{aligned} (T_n(f))(z) &:= c(z, \tilde{\gamma}_N^*(z)) + \beta \int f(y) \eta(dy|z, \tilde{\gamma}_N^*(z)) \\ (\hat{T}_n(f))(z) &:= c(F(z), \tilde{\gamma}_N^*(z)) + \beta \int f(y) \eta(dy|F(z), \tilde{\gamma}_N^*(z)) \end{aligned}$$

where  $\tilde{\gamma}_N^*$  denotes the extension of the optimal policy for the finite state space model to the belief space  $\mathcal{Z}$ , i.e. it is constant over the quantization bins, furthermore  $F(z)$  is the nearest neighbor quantization map such that it maps  $z$  to the closest point from the finite state space  $\mathcal{Z}_{\tilde{\pi}}^N$  (recall that  $\rho_{BL}$  is used to metrize the belief space).

The optimal cost for the finite model,  $J_{\beta}^N$  is only defined on a finite set, we denote by  $\tilde{J}_{\beta}^N$  the extension of the optimal cost for the finite model to the belief space  $\mathcal{P}(\mathbb{X})$  by assigning the same values over the quantization bins, i.e. it is a piece-wise constant function over the quantization bins. Note that, we have

$$\begin{aligned} J_{\beta}(\tilde{\gamma}_N^*, z) &= T_n(J_{\beta}(\tilde{\gamma}_N^*, z)), \\ \tilde{J}_{\beta}^N(\tilde{\gamma}_N^*, z) &= \hat{T}_n(\tilde{J}_{\beta}^N(\tilde{\gamma}_N^*, z)). \end{aligned}$$

Using these equalities, we write

$$\begin{aligned} J_{\beta}(\tilde{\gamma}_N^*, z) - J_{\beta}(\gamma^*, z) &\leq |T_n(J_{\beta}(\tilde{\gamma}_N^*, z)) - T_n(J_{\beta}(\gamma^*, z))| \\ &\quad + |T_n(J_{\beta}(\gamma^*, z)) - \hat{T}_n(J_{\beta}(\gamma^*, z))| \\ &\quad + |\hat{T}_n(J_{\beta}(\gamma^*, z)) - \hat{T}_n(\tilde{J}_{\beta}^N(\tilde{\gamma}_N^*, z))| \\ &\quad + |\tilde{J}_{\beta}^N(\tilde{\gamma}_N^*, z) - J_{\beta}(\gamma^*, z)| \\ &\leq \beta \int |J_{\beta}(\tilde{\gamma}_N^*, z_1) - J_{\beta}(\gamma^*, z_1)| \eta(dz_1|z, \tilde{\gamma}_N^*(z)) \\ &\quad + |c(z, \tilde{\gamma}_N^*(z)) - c(F(z), \tilde{\gamma}_N^*(z))| \\ &\quad + \beta \left| \int J_{\beta}(\gamma^*, z_1) \eta(dz_1|z, \tilde{\gamma}_N^*(z)) - \int J_{\beta}(\gamma^*, z_1) \eta(dz_1|F(z), \tilde{\gamma}_N^*(z)) \right| \\ &\quad + \beta \int |\tilde{J}_{\beta}^N(\tilde{\gamma}_N^*, z_1) - J_{\beta}(\gamma^*, z_1)| \eta(dz_1|F(z), \tilde{\gamma}_N^*(z)) \end{aligned}$$

$$\begin{aligned}
 & + |\tilde{J}_\beta^N(\tilde{\gamma}_N^*, z) - J_\beta(\gamma^*, z)| \\
 & \leq \beta \sup_\gamma \left( E_z^\gamma [|J_\beta(\tilde{\gamma}_N^*, Z_1) - J_\beta(\gamma^*, Z_1)|] \right) + \alpha_c L(z) + \beta \alpha_Z L(z) \|J_\beta^*\|_{BL} \\
 & \quad + \beta E_{z,N}^\gamma [| \tilde{J}_\beta^N(\tilde{\gamma}_N^*, Z_1) - J_\beta(\gamma^*, Z_1) |] + |\tilde{J}_\beta^N(\tilde{\gamma}_N^*, z) - J_\beta(\gamma^*, z)| \quad (23)
 \end{aligned}$$

where  $E_{z,\gamma}^N$  denotes the expectation with respect to the kernel  $\eta(dy|F(z), \gamma(z))$  when the initial point is  $z$ .

Now, using Lemma 18, we can write that:

$$\begin{aligned}
 & |\tilde{J}_\beta^N(\tilde{\gamma}_N^*, z) - J_\beta(\gamma^*, z)| \\
 & \leq \lim_{k \rightarrow \infty} \sup_\gamma \left( \alpha_c \sum_{t=0}^{k-1} \beta^t \left( E_z^\gamma [L(Z_t)] + 3\alpha_Z \sum_{m=0}^{t-1} (4\alpha_Z + 1)^m E_z^\gamma [L(Z_{t-m-1})] \right) \right. \\
 & \quad \left. + \alpha_Z \sum_{t=0}^{k-1} \beta^{t+1} \|v_{k-t-1}\|_{BL} \left( E_z^\gamma [L(Z_t)] + 3\alpha_Z \sum_{m=0}^{t-1} (4\alpha_Z + 1)^m E_z^\gamma [L(Z_{t-m-1})] \right) \right) \\
 & := g(z). \quad (24)
 \end{aligned}$$

We first introduce the following notation along with the above notation (24)

$$f(z) := |J_\beta(\tilde{\gamma}_N^*, z) - J_\beta(\gamma^*, z)|.$$

Notice that with the new notation, we can rewrite the bound on (23) as:

$$\begin{aligned}
 f(z) & \leq \beta \sup_\gamma E_z^\gamma [f(Z_1)] + \alpha_c L(z) + \beta \alpha_Z L(z) \|J_\beta^*\|_{BL} + \beta E_{z,N}^\gamma [g(Z_1)] + g(z) \\
 & \leq \beta \sup_\gamma E_z^\gamma [f(Z_1)] + \alpha_c L(z) + \beta \alpha_Z L(z) \|J_\beta^*\|_{BL} + \beta E_z^\gamma [g(Z_1)] + \beta \|g\|_{BL} \alpha_Z L(z) + g(z)
 \end{aligned}$$

where we used the fact that

$$\begin{aligned}
 E_{z,N}^\gamma [g(Z_1)] - E_z^\gamma [g(Z_1)] & = \int g(z_1) \eta(dz_1|F(z), \gamma(z)) - \int g(z_1) \eta(dz_1|z, \gamma(z)) \\
 & \leq \|g\|_{BL} \alpha_Z L(z).
 \end{aligned}$$

Using the same bound on  $f(Z_1)$  one can also write that for any initial point  $z_0$ :

$$\begin{aligned}
 f(z_0) & \leq \sup_{\gamma \in \Gamma} \beta^2 E_{z_0}^\gamma [f(Z_2)] + \alpha_c \sum_{t=0}^1 \beta^t E_{z_0}^\gamma [L(Z_t)] + \alpha_Z \|J_\beta^*\|_{BL} \sum_{t=0}^1 \beta^{t+1} E_{z_0}^\gamma [L(Z_t)] \\
 & \quad + \sum_{t=0}^1 \beta^{t+1} E_{z_0}^\gamma [g(Z_t)] + \|g\|_{BL} \alpha_Z \sum_{t=0}^1 \beta^{t+1} E_{z_0}^\gamma [L(Z_t)] + \sum_{t=0}^1 \beta^t E_{z_0}^\gamma [g(Z_t)].
 \end{aligned}$$

In general, for any  $k < \infty$ , we can write

$$f(z_0) \leq \sup_{\gamma \in \Gamma} \left( \beta^k E_{z_0}^\gamma [f(Z_k)] + \alpha_c \sum_{t=0}^{k-1} \beta^t E_{z_0}^\gamma [L(Z_t)] + \alpha_Z \|J_\beta^*\|_{BL} \sum_{t=0}^{k-1} \beta^{t+1} E_{z_0}^\gamma [L(Z_t)] \right)$$

$$+ \sum_{t=0}^{k-1} \beta^{t+1} E_{z_0}^\gamma [g(Z_t)] + \|g\|_{BL} \alpha_Z \sum_{t=0}^{k-1} \beta^{t+1} E_{z_0}^\gamma [L(Z_t)] + \sum_{t=0}^{k-1} \beta^t E_{z_0}^\gamma [g(Z_t)] \Big). \quad (25)$$

Recall that our main goal is to bound  $E_{\pi_0}[f(Z_0)|Y_{[0,N]}, \gamma(Y_{[0,N-1]})]$  where  $Z_0 = P_{\pi_0}(X_N \in \cdot | Y_{[0,N]}, \gamma(Y_{[0,N-1]}))$ . To this end, first notice that using Lemma 19, we have for any  $t$

$$\begin{aligned} & E_{\pi_0} \left[ E_{Z_0}^\gamma [L(Z_t)] | Y_{[0,N]}, \gamma(Y_{[0,N-1]}) \right] \\ & \leq \sup_{\pi \in \mathcal{P}(\mathbb{X})} \sup_{\gamma \in \Gamma} E^{\pi, \gamma} \left[ \rho_{BL} \left( P^\pi(X_N \in \cdot | Y_{[0,N]}), P^{\hat{\pi}}(X_N \in \cdot | Y_{[0,N]}) \right) \right] := B. \end{aligned} \quad (26)$$

Using this bound and Lemma 20 for  $\|v_{k-t-1}\|_{BL}$ , we can write that

$$\begin{aligned} & E \left[ E_{Z_0}^\gamma [g(Z_t)] | Y_{[0,N]}, \gamma(Y_{[0,N-1]}) \right] \\ & \leq \lim_{k \rightarrow \infty} \left( \alpha_c \sum_{t=0}^{k-1} \beta^t \left( B + 3\alpha_Z \sum_{m=0}^{t-1} B(4\alpha_Z + 1)^m \right) \right. \\ & \quad \left. + \alpha_Z \sum_{t=0}^{k-1} \beta^{t+1} \|v_{k-t-1}\|_{BL} \left( B + 3\alpha_Z \sum_{m=0}^{t-1} B(4\alpha_Z + 1)^m \right) \right) \\ & \leq B(\alpha_c + \beta\alpha_Z \|J_\beta^*\|_{BL}) \lim_{k \rightarrow \infty} \sum_{t=0}^{k-1} \beta^t \left( 1 + 3\alpha_Z \sum_{m=0}^{t-1} (4\alpha_Z + 1)^m \right) \\ & \leq B(\alpha_c + \beta\alpha_Z \|J_\beta^*\|_{BL}) \lim_{k \rightarrow \infty} \sum_{t=0}^{k-1} \beta^t \left( 1 + 3\alpha_Z \frac{(4\alpha_Z + 1)^t - 1}{4\alpha_Z} \right) \\ & \leq B(\alpha_c + \beta\alpha_Z \|J_\beta^*\|_{BL}) \lim_{k \rightarrow \infty} \sum_{t=0}^{k-1} \beta^t (4\alpha_Z + 1)^t \\ & = BK_0(\beta, \alpha_Z, \alpha_{\tilde{c}}, \|\tilde{c}\|_\infty) \end{aligned} \quad (27)$$

where

$$K_0(\beta, \alpha_Z, \alpha_{\tilde{c}}, \|\tilde{c}\|_\infty) = (\alpha_c + \beta\alpha_Z \|J_\beta^*\|_{BL}) \frac{1}{1 - \beta(4\alpha_Z + 1)}.$$

By Lemma 23 we also have that

$$\begin{aligned} \|g\|_{BL} & \leq (\alpha_c + \beta\alpha_Z \|J_\beta^*\|_{BL}) \left( \frac{2}{1 - \beta(4\alpha_Z + 1)} + \frac{3\alpha_Z}{1 - \beta\alpha_Z} + \frac{9\alpha_Z^2}{(1 - \beta(4\alpha_Z + 1))^2} \right) \\ & := \hat{K}_0(\beta, \alpha_Z, \alpha_{\tilde{c}}, \|\tilde{c}\|_\infty) \end{aligned}$$

Going back to (25), using (26) and (27) and taking the limit  $k \rightarrow \infty$ :

$$\begin{aligned} E[f(Z_0)|Y_{[0,N]}, \gamma(Y_{[0,N-1]})] & \leq \frac{B\alpha_c}{1 - \beta} + \frac{\beta\alpha_Z \|J_\beta^*\|_{BL} B}{1 - \beta} + \frac{\beta K_0(\beta, \alpha_Z, \alpha_{\tilde{c}}, \|\tilde{c}\|_\infty) B}{1 - \beta} \\ & \quad + \frac{\hat{K}_0(\beta, \alpha_Z, \alpha_{\tilde{c}}, \|\tilde{c}\|_\infty) \alpha_Z \beta B}{1 - \beta} + \frac{K_0(\beta, \alpha_Z, \alpha_{\tilde{c}}, \|\tilde{c}\|_\infty) B}{1 - \beta}. \end{aligned}$$

Thus, we can write that

$$E[f(Z_0)|Y_{[0,N]}, \gamma(Y_{[0,N-1]})] \leq K \sup_{\pi \in \mathcal{P}(\mathbb{X})} \sup_{\gamma \in \Gamma} E^{\pi, \gamma} \left[ \rho_{BL} \left( P^\pi(X_N \in \cdot | Y_{[0,N]}), P^{\hat{\pi}}(X_N \in \cdot | Y_{[0,N]}) \right) \right],$$

where

$$K = \frac{\alpha_c + \beta \alpha_Z \|J_\beta^*\|_{BL} + (\beta + 1)K_0(\beta, \alpha_Z, \alpha_{\tilde{c}}, \|\tilde{c}\|_\infty) + \hat{K}_0(\beta, \alpha_Z, \alpha_{\tilde{c}}, \|\tilde{c}\|_\infty)\beta \alpha_Z}{1 - \beta} \quad (28)$$

and using Lemma 24

$$\|J_\beta^*\|_{BL} \leq \frac{1}{1 - \beta \alpha_Z} \left( \frac{\|\tilde{c}\|_\infty}{1 - \beta} + \alpha_{\tilde{c}} \right).$$

## 7. Conclusion

In this paper, we studied performance bounds for policies that use only a finite-window of recent observation and action variables rather than the entire history in partially observed stochastic control problems. We have rigorously established approximation bounds that are easy to compute, and have shown that this bound critically depends on the ergodicity and stability properties of the belief-state process. We have provided the results for continuous-space valued state spaces and finite observation and action spaces, however our studies suggest that these results can also be generalized to real valued observation and actions also. Application to decentralized POMDPs is another direction that will benefit from the analysis presented.

## Appendix A. Technical Proofs of Supporting Results

In this section, we prove the proofs of some technical results.

### A.1 Proof of Step 1 in the Proof of Lemma 18

We claim that

$$|\tilde{v}_k^N(z) - v_k(z)| \leq \sup_{\gamma \in \Gamma} \left( \alpha_{\tilde{c}} \sum_{t=0}^{k-1} \beta^t E_{z, \gamma}^N [L(Z_t)] + \sum_{t=0}^{k-1} \beta^{t+1} \|v_{k-t-1}\|_{BL} E_{z, \gamma}^N [L(Z_t)] \alpha_Z \right)$$

where  $L(z)$  is the loss function due to the quantization, i.e.  $L(z) = \rho_{BL}(z, F(z))$  with  $F$  defined in (8).

We prove the claim using an inductive approach: for  $k = 1$  we have (noting  $v_0, \tilde{v}_0^N \equiv 0$ )

$$\begin{aligned} \tilde{v}_1^N(z) &= \min_u \left( \tilde{c}(F(z), u) + \beta \int \tilde{v}_0^N(y) \eta(dy|F(z), u) \right) = \min_u \tilde{c}(F(z), u) \\ v_1(z) &= \min_u \left( \tilde{c}(z, u) + \beta \int v_0(y) \eta(dy|z, u) \right) = \min_u \tilde{c}(z, u). \end{aligned}$$

Note that under the stated assumptions the measurable selection conditions hold, and the minimum can be achieved using a policy  $\gamma$  for the original model and a policy  $\gamma^N$  for the finite model which

defined only on a finite set. By extending the finite model policy  $\gamma^N$  over all state space  $\mathcal{P}(\mathbb{X})$ , we can write that

$$\begin{aligned} |\tilde{v}_1^N(z) - v_1(z)| &\leq \max \left( \tilde{c}(F(z), \gamma(z)) - \tilde{c}(z, \gamma(z)), \tilde{c}(z, \gamma^N(z)) - \tilde{c}(F(z), \gamma^N(z)) \right) \\ &\leq \sup_{\gamma} |\tilde{c}(F(z), \gamma(z)) - \tilde{c}(z, \gamma(z))| \leq \alpha_{\tilde{c}} L(z) \end{aligned}$$

which completes the proof for  $k = 1$ . Now, we assume that the claim holds for  $k$  and analyze the step  $k + 1$ . Similar to  $k = 1$  case, we can write

$$\begin{aligned} &|\tilde{v}_{k+1}^N(z) - v_{k+1}(z)| \\ &\leq \sup_{\gamma} \left| \tilde{c}(F(z), \gamma(z)) - \tilde{c}(z, \gamma(z)) + \beta \int \tilde{v}_k^N(y) \eta(dy|F(z), \gamma(z)) - \beta \int v_k(y) \eta(dy|z, \gamma(z)) \right| \\ &\leq \sup_{\gamma} \left( |\tilde{c}(F(z), \gamma(z)) - \tilde{c}(z, \gamma(z))| + \beta \int |\tilde{v}_k^N(y) - v_k(y)| \eta(dy|F(z), \gamma(z)) \right. \\ &\quad \left. + \beta \left| \int v_k(y) \eta(dy|F(z), \gamma(z)) - \int v_k(y) \eta(dy|z, \gamma(z)) \right| \right) \\ &\leq \alpha_{\tilde{c}} L(z) + \beta \|v_k\|_{BL} \alpha_{\mathcal{Z}} L(z) \\ &+ \beta \sup_{\gamma} \left( \int \left| \alpha_{\tilde{c}} \sum_{t=0}^{k-1} \beta^t E_{y,\gamma}^N [L(Z_t)] + \sum_{t=0}^{k-1} \beta^{t+1} \|v_{k-t-1}\|_{BL} E_{y,\gamma}^N [L(Z_t)] \alpha_{\mathcal{Z}} \right| \eta(dy|F(z), \gamma(z)) \right) \\ &\leq \sup_{\gamma} \left( \alpha_{\tilde{c}} L(z) + \alpha_{\tilde{c}} \sum_{t=0}^{k-1} \beta^{t+1} E_{z,\gamma}^N [L(Z_{t+1})] + \beta \|v_k\|_{BL} \alpha_{\mathcal{Z}} L(z) \right. \\ &\quad \left. + \sum_{t=0}^{k-1} \beta^{t+2} \|v_{k-t-1}\|_{BL} E_{z,\gamma}^N [L(Z_{t+1})] \right) \\ &= \sup_{\gamma} \left( \alpha_{\tilde{c}} \sum_{t'=0}^k \beta^{t'} E_{z,\gamma}^N [L(Z_{t'})] + \sum_{t'=0}^k \beta^{t'+1} \|v_{k-t'}\|_{BL} E_{z,\gamma}^N [L(Z_{t'})] \right), \quad (t' = t + 1). \end{aligned}$$

For the last two steps, note that  $E_{y,\gamma}^N [L(Z_t)]$  denotes the expected loss at time  $t$  when the initial state  $Z_0 = y$ , thus using the iterative expectation we have

$$\int E_{y,\gamma}^N [L(Z_t)] \eta(dy|F(z), \gamma(z)) = E_{z,\gamma}^N [L(Z_{t+1})].$$

Hence, we have proved that for all  $k \geq 1$

$$|\tilde{v}_k^N(z) - v_k(z)| \leq \sup_{\gamma} \left( \alpha_{\tilde{c}} \sum_{t=0}^{k-1} \beta^t E_{z,\gamma}^N [L(Z_t)] + \sum_{t=0}^{k-1} \beta^{t+1} \|v_{k-t-1}\|_{BL} E_{z,\gamma}^N [L(Z_t)] \alpha_{\mathcal{Z}} \right).$$

## A.2 Proof of Step 2 in the Proof of Lemma 18

The claim is that

$$E_{z,\gamma}^N [L(Z_t)] \leq E_{z,\gamma} [L(Z_t)] + 3\alpha_{\mathcal{Z}} \sum_{m=0}^{t-1} (4\alpha_{\mathcal{Z}} + 1)^m E_{z,\gamma} [L(Z_{t-m-1})].$$

We first write

$$E_{z,\gamma}^N [L(Z_t)] \leq E_{z,\gamma} [L(Z_t)] + \left( \|L\|_{BL} \right) \rho_{BL} \left( P_{z,t}^\gamma, P_{z,t}^{\gamma,N} \right) \quad (29)$$

where

$$P_{z,t}^\gamma, \quad P_{z,t}^{\gamma,N}$$

are the marginal distributions of the state  $z_t$  for the true model and approximate model respectively, with  $Z_0 = z$ .

We focus on the term  $\|L\|_{BL} \rho_{BL} \left( P_{z,t}^\gamma, P_{z,t}^{\gamma,N} \right)$ . We first claim that  $\|L\|_{BL} \leq 3$ , where  $L(z) = \rho(z, F(z))$ . Recall that  $F$  determines quantization given in (8).

$$\|L\|_{BL} \leq \|L\|_\infty + \sup_{z,z'} \frac{|\rho(z, F(z)) - \rho(z', F(z'))|}{\rho(z, z')} \leq 2 + \sup_{z,z'} \frac{|\rho(z, F(z)) - \rho(z', F(z'))|}{\rho(z, z')}$$

where we used the fact that

$$\|L\|_\infty = \sup_z \rho_{BL}(z, F(z)) \leq 2$$

as for any two probability measures  $\mu, \nu$ , we have that  $\rho_{BL}(\mu, \nu) \leq 2$  (see (3)).

Note that  $F$  is a nearest neighbor quantizer as defined in (8), thus we can write that

$$\begin{aligned} |\rho(z, F(z)) - \rho(z', F(z'))| &\leq \max(\rho(z, F(z')) - \rho(z', F(z')), \rho(z', F(z)) - \rho(z, F(z))) \\ &\leq \sup_{\hat{z}} |\rho(z, \hat{z}) - \rho(z', \hat{z})| \leq \rho(z, z') \end{aligned}$$

where for the last step we used the triangle inequality. Hence we can conclude that

$$\|L\|_{BL} \leq 2 + \sup_{z,z'} \frac{|\rho(z, F(z)) - \rho(z', F(z'))|}{\rho(z, z')} \leq 3. \quad (30)$$

Now, we show that

$$\rho_{BL} \left( P_{z,t}^\gamma, P_{z,t}^{\gamma,N} \right) \leq \alpha_Z \sum_{m=0}^{t-1} (4\alpha_Z + 1)^m E_{z,\gamma} [L(Z_{t-m-1})]. \quad (31)$$

We prove the claim by induction: for  $t = 1$

$$\begin{aligned} \rho_{BL} \left( P_{z,1}^\gamma, P_{z,1}^{\gamma,N} \right) &= \sup_{\|f\|_{BL} \leq 1} \left| \int f(z_1) \eta(dz_1 | F(z), u_1) - \int f(z_1) \eta(z_1) \eta(dz_1 | z, u_1) \right| \\ &\leq \alpha_Z L(z). \end{aligned}$$

Now assume the claim holds for  $t - 1$

$$\begin{aligned} \rho_{BL} \left( P_{z,t}^\gamma, P_{z,t}^{\gamma,N} \right) &= \sup_{\|f\|_{BL} \leq 1} \left| \int f(z_t) \eta(dz_t | z_{t-1}, \gamma(z_{t-1})) P_{z,t-1}^\gamma(dz_{t-1}) \right. \\ &\quad \left. - \int f(z_t) \eta(dz_t | F(z_{t-1}), \gamma(z_{t-1})) P_{z,t-1}^{\gamma,N}(dz_{t-1}) \right| \end{aligned}$$

$$\begin{aligned}
 &\leq \sup_{\|f\|_{BL} \leq 1} \left| \int f(z_t) \eta(dz_t | z_{t-1}, \gamma(z_{t-1})) P_{z,t-1}^\gamma(dz_{t-1}) \right. \\
 &\quad \left. - \int f(z_t) \eta(dz_t | z_{t-1}, \gamma(z_{t-1})) P_{z,t-1}^{\gamma,N}(dz_{t-1}) \right| \\
 &\quad + \sup_{\|f\|_{BL} \leq 1} \left| \int f(z_t) \eta(dz_t | z_{t-1}, \gamma(z_{t-1})) P_{z,t-1}^{\gamma,N}(dz_{t-1}) \right. \\
 &\quad \left. - \int f(z_t) \eta(dz_t | F(z_{t-1}), \gamma(z_{t-1})) P_{z,t-1}^{\gamma,N}(dz_{t-1}) \right| \\
 &\leq (\|f\|_\infty + \alpha_{\mathcal{Z}}) \rho_{BL} \left( P_{z,t-1}^\gamma, P_{z,t-1}^{\gamma,N} \right) + \alpha_{\mathcal{Z}} E_z^N [L(Z_{t-1})] \\
 &\leq (1 + \alpha_{\mathcal{Z}}) \rho_{BL} \left( P_{z,t-1}^\gamma, P_{z,t-1}^{\gamma,N} \right) + \alpha_{\mathcal{Z}} E_z^N [L(Z_{t-1})].
 \end{aligned}$$

Using (29), we can write

$$\begin{aligned}
 &\rho_{BL} \left( P_{z,t}^\gamma, P_{z,t}^{\gamma,N} \right) \\
 &\leq (1 + \alpha_{\mathcal{Z}}) \rho_{BL} \left( P_{z,t-1}^\gamma, P_{z,t-1}^{\gamma,N} \right) + \alpha_{\mathcal{Z}} \left( E_z [L(Z_{t-1})] + 3\rho_{BL} \left( P_{z,t-1}^\gamma, P_{z,t-1}^{\gamma,N} \right) \right) \\
 &= (4\alpha_{\mathcal{Z}} + 1) \rho_{BL} \left( P_{z,t-1}^\gamma, P_{z,t-1}^{\gamma,N} \right) + \alpha_{\mathcal{Z}} E_z [L(Z_{t-1})] \\
 &\leq (4\alpha_{\mathcal{Z}} + 1) \left( \alpha_{\mathcal{Z}} \sum_{m=0}^{t-2} (4\alpha_{\mathcal{Z}} + 1)^m E_z [L(Z_{t-m-2})] \right) + \alpha_{\mathcal{Z}} E_z [L(Z_{t-1})] \\
 &= \alpha_{\mathcal{Z}} \sum_{m=0}^{t-1} (4\alpha_{\mathcal{Z}} + 1)^m E [L_z(Z_{t-m-1})]
 \end{aligned}$$

which completes the proof of (31).

Now, we go back to the main claim and start from (29) to write

$$\begin{aligned}
 E_{z,\gamma}^N [L(Z_t)] &\leq E_{z,\gamma} [L(Z_t)] + \|L\|_{BL} \rho_{BL} \left( P_{z,t}^\gamma, P_{z,t}^{\gamma,N} \right) \leq E_{z,\gamma} [L(Z_t)] + 3\rho_{BL} \left( P_{z,t}^\gamma, P_{z,t}^{\gamma,N} \right) \\
 &\leq E_{z,\gamma} [L(Z_t)] + 3\alpha_{\mathcal{Z}} \sum_{m=0}^{t-1} (4\alpha_{\mathcal{Z}} + 1)^m E_{z,\gamma} [L(Z_{t-m-1})].
 \end{aligned}$$

**Remark 21** Note that for the previous proof,  $(4\alpha_{\mathcal{Z}} + 1)$  can be replaced by  $1 + \alpha_{\mathcal{Z}}(\|L\|_\infty + 2)$  where

$$\|L\|_\infty = \sup_{\pi \in \mathcal{P}(\mathbb{X})} \sup_{\gamma \in \Gamma} \sup_{y_{[0,N]}, u_{[0,N-1]}} \rho_{BL} \left( P^\pi(\cdot | y_{[0,N]}, u_{[0,N-1]}), P^{\hat{\pi}}(\cdot | y_{[0,N]}, u_{[0,N-1]}) \right)$$

which is upper bounded by 2, however, usually smaller than 2, provided there is a uniform filter stability.

**Lemma 22** We introduce the following notation

$$\hat{E}_z [L(Z_t)] = \sup_{u_0} \int \dots \sup_{u_{t-2}} \int \sup_{u_{t-1}} \int L(z_t) \eta(dz_t | z_{t-1}, u_{t-1}) \eta(dz_{t-1} | z_{t-2}, u_{t-2}) \dots \eta(dz_1 | z, u_0).$$



Under Assumption 3, we have that

$$\sup_{\gamma \in \Gamma} E_{z,\gamma}[L(Z_t)] = \hat{E}_z[L(Z_t)].$$

**Proof** Recall that for a fixed policy  $\gamma = \{\gamma_t\}_t$ ,

$$E_{z,\gamma}[L(Z_t)] = \int L(z_t) \eta(dz_t|z_{t-1}, \gamma_{t-1}(z_{t-1})) \eta(dz_{t-1}|z_{t-1}, \gamma_{t-2}(z_{t-2})) \dots \eta(dz_1|z, \gamma_0(z_0))$$

It is easy to see that

$$\begin{aligned} & \int L(z_t) \eta(dz_t|z_{t-1}, \gamma_{t-1}(z_{t-1})) \eta(dz_{t-1}|z_{t-1}, \gamma_{t-2}(z_{t-2})) \dots \eta(dz_1|z, \gamma_0(z_0)) \\ & \leq \int \sup_{u_{t-1}} \int L(z_t) \eta(dz_t|z_{t-1}, u_{t-1}) \eta(dz_{t-1}|z_{t-1}, \gamma_{t-2}(z_{t-2})) \dots \eta(dz_1|z, \gamma_0(z_0)) \end{aligned}$$

By repeating this step we can have

$$\begin{aligned} & \int L(z_t) \eta(dz_t|z_{t-1}, \gamma_{t-1}(z_{t-1})) \eta(dz_{t-1}|z_{t-1}, \gamma_{t-2}(z_{t-2})) \dots \eta(dz_1|z, \gamma_0(z_0)) \\ & \leq \sup_{u_0} \int \dots \sup_{u_{t-2}} \int \sup_{u_{t-1}} \int L(z_t) \eta(dz_t|z_{t-1}, u_{t-1}) \eta(dz_{t-1}|z_{t-2}, u_{t-2}) \dots \eta(dz_1|z, u_0) \end{aligned}$$

Hence by taking supremum over all policies we can write

$$\sup_{\gamma \in \Gamma} E_{z,\gamma}[L(Z_t)] \leq \hat{E}_z[L(Z_t)].$$

For the other direction, we first focus on the term  $\sup_{u_{t-1}} \int L(z_t) \eta(dz_t|z_{t-1}, u_{t-1})$ , using Assumption 3, one can show that there exists a measurable map  $\gamma_{t-1}$  such that

$$\sup_{u_{t-1}} \int L(z_t) \eta(dz_t|z_{t-1}, u_{t-1}) = \int L(z_t) \eta(dz_t|z_{t-1}, \gamma_{t-1}(z_{t-1})).$$

Using the same argument, we can see that there exists a sequence of measurable functions  $\{\gamma_0, \gamma_1, \dots, \gamma_{t-1}\}$  such that

$$\begin{aligned} & \sup_{u_0} \int \dots \sup_{u_{t-2}} \int \sup_{u_{t-1}} \int L(z_t) \eta(dz_t|z_{t-1}, u_{t-1}) \eta(dz_{t-1}|z_{t-2}, u_{t-2}) \dots \eta(dz_1|z, u_0) \\ & = \int L(z_t) \eta(dz_t|z_{t-1}, \gamma_{t-1}(z_{t-1})) \eta(dz_{t-1}|z_{t-1}, \gamma_{t-2}(z_{t-2})) \dots \eta(dz_1|z, \gamma_0(z_0)). \end{aligned}$$

Hence we can write that

$$\hat{E}_z[L(Z_t)] \leq \sup_{\gamma \in \Gamma} E_{z,\gamma}[L(Z_t)]$$

which proves the main claim. ■

**Lemma 23** *Under Assumption 3*

$$\|g\|_{BL} \leq (\alpha_c + \beta\alpha_Z \|J_\beta\|_{BL}) \left( \frac{2}{1 - \beta(4\alpha_Z + 1)} + \frac{3\alpha_Z}{1 - \beta\alpha_Z} + \frac{9\alpha_Z^2}{(1 - \beta(4\alpha_Z + 1))^2} \right).$$

**Proof** Using Lemma 22, we can write that

$$\begin{aligned} g(z) &= \lim_{k \rightarrow \infty} \sup_{\gamma} \left( \alpha_c \sum_{t=0}^{k-1} \beta^t \left( E_{z,\gamma} [L(Z_t)] + 3\alpha_Z \sum_{m=0}^{t-1} (4\alpha_Z + 1)^m E_{z,\gamma} [L(Z_{t-m-1})] \right) \right. \\ &\quad \left. + \alpha_Z \sum_{t=0}^{k-1} \beta^{t+1} \|v_{k-t-1}\|_{BL} \left( E_{z,\gamma} [L(Z_t)] + 3\alpha_Z \sum_{m=0}^{t-1} (4\alpha_Z + 1)^m E_{z,\gamma} [L(Z_{t-m-1})] \right) \right) \\ &= \lim_{k \rightarrow \infty} \left( \alpha_c \sum_{t=0}^{k-1} \beta^t \left( \hat{E}_z [L(Z_t)] + 3\alpha_Z \sum_{m=0}^{t-1} (4\alpha_Z + 1)^m \hat{E}_z [L(Z_{t-m-1})] \right) \right. \\ &\quad \left. + \alpha_Z \sum_{t=0}^{k-1} \beta^{t+1} \|v_{k-t-1}\|_{BL} \left( \hat{E}_z [L(Z_t)] + 3\alpha_Z \sum_{m=0}^{t-1} (4\alpha_Z + 1)^m \hat{E}_z [L(Z_{t-m-1})] \right) \right). \end{aligned}$$

Using the fact that  $\|L\|_\infty \leq 2$  we can write

$$\begin{aligned} \|g\|_\infty &\leq \lim_{k \rightarrow \infty} \left( \alpha_c \sum_{t=0}^{k-1} \beta^t \left( 2 + 3\alpha_Z \sum_{m=0}^{t-1} 2(4\alpha_Z + 1)^m \right) \right. \\ &\quad \left. + \alpha_Z \sum_{t=0}^{k-1} \beta^{t+1} \|v_{k-t-1}\|_{BL} \left( 2 + 3\alpha_Z \sum_{m=0}^{t-1} 2(4\alpha_Z + 1)^m \right) \right) \\ &\leq 2(\alpha_c + \beta\alpha_Z \|J_\beta^*\|_{BL}) \frac{1}{1 - \beta(4\alpha_Z + 1)}. \end{aligned}$$

Next, we show that when defined as a function of  $z$ ,  $\|\hat{E}_z[L(Z_t)]\|_{BL} \leq 3\frac{\alpha_Z^{t+1}-1}{\alpha_Z-1}$ . We follow an inductive approach. For  $t = 1$ :

$$\begin{aligned} |\hat{E}_z[L(Z_1)] - \hat{E}_{\hat{z}}[L(Z_1)]| &\leq \sup_u \left| \int L(z_1)\eta(dz_1|z, u) - \int L(z_1)\eta(dz_1|\hat{z}, u) \right| \\ &\leq \|L\|_{BL} \alpha_Z \rho_{BL}(z, \hat{z}) \leq 3\alpha_Z \rho_{BL}(z, \hat{z}) \end{aligned}$$

where  $\|L\|_{BL} \leq 3$  follows from (30). Hence we have that  $\|\hat{E}_z[L(Z_1)]\|_{BL} \leq 3 + 3\alpha_Z$ . Now we assume the claim holds for  $t - 1$  and focus on the case for  $t$ :

$$\begin{aligned} &|\hat{E}_z[L(Z_t)] - \hat{E}_{\hat{z}}[L(Z_t)]| \\ &= \sup_{u_0} \int \dots \sup_{u_{t-2}} \int \sup_{u_{t-1}} \int L(z_t)\eta(dz_t|z_{t-1}, u_{t-1})\eta(dz_{t-1}|z_{t-2}, u_{t-2}) \dots \eta(dz_1|z, u_0) \\ &\quad - \sup_{u_0} \int \dots \sup_{u_{t-2}} \int \sup_{u_{t-1}} \int L(z_t)\eta(dz_t|z_{t-1}, u_{t-1})\eta(dz_{t-1}|z_{t-2}, u_{t-2}) \dots \eta(dz_1|\hat{z}, u_0) \\ &\leq \sup_{u_0} \left( \int \hat{E}_{z_1}[L(Z_t)]\eta(dz_1|z, u_0) - \int \hat{E}_{z_1}[L(Z_t)]\eta(dz_1|\hat{z}, u_0) \right) \end{aligned}$$

$$\leq \alpha_{\mathcal{Z}} \rho_{BL}(z, \hat{z}) \|\hat{E}_z[L(Z_{t-1})]\|_{BL} \leq \alpha_{\mathcal{Z}} \rho_{BL}(z, \hat{z}) 3 \frac{\alpha_{\mathcal{Z}}^t - 1}{\alpha_{\mathcal{Z}} - 1}.$$

We then have that

$$\|\hat{E}_z[L(Z_t)]\|_{BL} \leq \|L\|_{\infty} + 3\alpha_{\mathcal{Z}} \frac{\alpha_{\mathcal{Z}}^t - 1}{\alpha_{\mathcal{Z}} - 1} \leq 3 + 3\alpha_{\mathcal{Z}} \frac{\alpha_{\mathcal{Z}}^t - 1}{\alpha_{\mathcal{Z}} - 1} \leq 3 \frac{\alpha_{\mathcal{Z}}^{t+1} - 1}{\alpha_{\mathcal{Z}} - 1}.$$

Thus, we can write that

$$|\hat{E}_z[L(Z_t)] - \hat{E}_{\hat{z}}[L(Z_t)]| \leq 3\alpha_{\mathcal{Z}} \frac{\alpha_{\mathcal{Z}}^t - 1}{\alpha_{\mathcal{Z}} - 1} \rho_{BL}(z, \hat{z}).$$

Hence for any  $z, \hat{z}$

$$\begin{aligned} |g(z) - g(\hat{z})| &\leq \lim_{k \rightarrow \infty} \left( \alpha_c \sum_{t=0}^{k-1} \beta^t (|\hat{E}_z[L(Z_t)] - \hat{E}_{\hat{z}}[L(Z_t)]| \right. \\ &\quad \left. + 3\alpha_{\mathcal{Z}} \sum_{m=0}^{t-1} (4\alpha_{\mathcal{Z}} + 1)^m |\hat{E}_z[L(Z_{t-m-1})] - \hat{E}_{\hat{z}}[L(Z_{t-m-1})]| \right) \\ &\quad + \alpha_{\mathcal{Z}} \sum_{t=0}^{k-1} \beta^{t+1} \|v_{k-t-1}\|_{BL} \left( |\hat{E}_z[L(Z_t)] - \hat{E}_{\hat{z}}[L(Z_t)]| \right. \\ &\quad \left. + 3\alpha_{\mathcal{Z}} \sum_{m=0}^{t-1} (4\alpha_{\mathcal{Z}} + 1)^m |\hat{E}_z[L(Z_{t-m-1})] - \hat{E}_{\hat{z}}[L(Z_{t-m-1})]| \right) \\ &\leq \lim_{k \rightarrow \infty} \left( \alpha_c \sum_{t=0}^{k-1} \beta^t \left( 3\alpha_{\mathcal{Z}} \frac{\alpha_{\mathcal{Z}}^t - 1}{\alpha_{\mathcal{Z}} - 1} \rho_{BL}(z, \hat{z}) + 3\alpha_{\mathcal{Z}} \sum_{m=0}^{t-1} (4\alpha_{\mathcal{Z}} + 1)^m 3\alpha_{\mathcal{Z}} \frac{\alpha_{\mathcal{Z}}^{t-m-1} - 1}{\alpha_{\mathcal{Z}} - 1} \rho_{BL}(z, \hat{z}) \right) \right) \\ &\quad + \alpha_{\mathcal{Z}} \sum_{t=0}^{k-1} \beta^{t+1} \|v_{k-t-1}\|_{BL} \left( 3\alpha_{\mathcal{Z}} \frac{\alpha_{\mathcal{Z}}^t - 1}{\alpha_{\mathcal{Z}} - 1} \rho_{BL}(z, \hat{z}) + 3\alpha_{\mathcal{Z}} \sum_{m=0}^{t-1} (4\alpha_{\mathcal{Z}} + 1)^m 3\alpha_{\mathcal{Z}} \frac{\alpha_{\mathcal{Z}}^{t-m-1} - 1}{\alpha_{\mathcal{Z}} - 1} \rho_{BL}(z, \hat{z}) \right) \\ &\leq \frac{3\alpha_{\mathcal{Z}}}{\alpha_{\mathcal{Z}} - 1} (\alpha_c + \beta\alpha_{\mathcal{Z}} \|J_{\beta}\|_{BL}) \left( \frac{1}{1 - \beta\alpha_{\mathcal{Z}}} + \frac{3\alpha_{\mathcal{Z}}}{(1 - \beta(4\alpha_{\mathcal{Z}} + 1))^2} \right) \rho_{BL}(z, \hat{z}). \end{aligned}$$

Hence, we have that

$$\|g\|_{BL} \leq (\alpha_c + \beta\alpha_{\mathcal{Z}} \|J_{\beta}\|_{BL}) \left( \frac{2}{1 - \beta(4\alpha_{\mathcal{Z}} + 1)} + \frac{3\alpha_{\mathcal{Z}}}{1 - \beta\alpha_{\mathcal{Z}}} + \frac{9\alpha_{\mathcal{Z}}^2}{(1 - \beta(4\alpha_{\mathcal{Z}} + 1))^2} \right).$$

■

**Lemma 24** *i. Under Assumption 3, if  $\mathcal{Z}$  is metrized with  $\rho_{BL}$ , we have that*

$$\|J_{\beta}^*\|_{BL} \leq \frac{1}{1 - \beta\alpha_{\mathcal{Z}}} \left( \frac{\|\tilde{c}\|_{\infty}}{1 - \beta} + \alpha_{\tilde{c}} \right).$$

*ii. Without any assumption, if  $\mathcal{Z}$  is metrized with total variation distance, we have that*

$$\|J_{\beta}^*\|_{BL} \leq \frac{2 - \beta}{(1 - \beta)(1 - \beta\alpha_{\mathcal{Z}})} \|c\|_{\infty}$$

**Proof** We start with the first part, for any  $x, y \in Z$

$$\begin{aligned} |J_\beta^*(x) - J_\beta^*(y)| &\leq \sup_u \left( |\tilde{c}(x, u) - \tilde{c}(y, u)| + \beta \left| \int J_\beta^*(z) \eta(dz|x, u) - \int J_\beta^*(z) \eta(dz|y, u) \right| \right) \\ &\leq \alpha_{\tilde{c}} \rho_{BL}(x, y) + \beta \|J_\beta^*\|_{BL} \alpha_Z \rho_{BL}(x, y). \end{aligned}$$

Thus, we can write that

$$\|J_\beta^*\|_{BL} \leq \|J_\beta^*\|_\infty + \alpha_{\tilde{c}} + \beta \alpha_Z \|J_\beta^*\|_{BL} \leq \frac{\|\tilde{c}\|_\infty}{1 - \beta} + \alpha_{\tilde{c}} + \beta \alpha_Z \|J_\beta^*\|_{BL}.$$

Hence, rearranging the terms, we can write (provided  $\beta \alpha_Z < 1$ )

$$\|J_\beta^*\|_{BL} \leq \frac{1}{1 - \beta \alpha_Z} \left( \frac{\|\tilde{c}\|_\infty}{1 - \beta} + \alpha_{\tilde{c}} \right).$$

For the second part, similar arguments lead to the following bound

$$\|J_\beta^*\|_{BL} \leq \|J_\beta^*\|_\infty + \|c\|_\infty + \beta \alpha_Z \|J_\beta^*\|_{BL},$$

hence the result follows after rearranging the terms and noting that  $\|J_\beta^*\|_\infty \leq \frac{\|c\|_\infty}{1 - \beta}$ . ■

## References

- D.P. Bertsekas. Convergence of discretization procedures in dynamic programming. *IEEE Trans. Autom. Control*, 20(3):415–419, Jun. 1975.
- P. Billingsley. *Convergence of probability measures*. New York: Wiley, 2nd edition, 1999.
- P. Chigansky, R. Liptser, and R. van Handel. Intrinsic methods in filter stability. *Handbook of Nonlinear Filtering*, 2009.
- C.S. Chow and J. N. Tsitsiklis. An optimal one-way multigrid algorithm for discrete-time stochastic control. *IEEE transactions on automatic control*, 36(8):898–914, 1991.
- R.L. Dobrushin. Central limit theorem for nonstationary Markov chains. i. *Theory of Probability & Its Applications*, 1(1):65–80, 1956.
- F. Dufour and T. Prieto-Rumeau. Approximation of Markov decision processes with general state space. *J. Math. Anal. Appl.*, 388:1254–1267, 2012.
- E. A. Feinberg. Controlled Markov processes with arbitrary numerical criteria. *Th. Probability and its Appl.*, pages 486–503, 1982.
- E.A. Feinberg, P.O. Kasyanov, and N.V. Zadioanchuk. Average cost Markov decision processes with weakly continuous transition probabilities. *Math. Oper. Res.*, 37(4):591–607, Nov. 2012.
- E.A. Feinberg, P.O. Kasyanov, and M.Z. Zgurovsky. Partially observable total-cost Markov decision process with weakly continuous transition probabilities. *Mathematics of Operations Research*, 41(2):656–681, 2016.

- E. A. Hansen. Solving POMDPs by searching in policy space. *arXiv preprint arXiv:1301.7380*, 2013.
- O. Hernández-Lerma. *Adaptive Markov Control Processes*. Springer-Verlag, 1989.
- O. Hernandez-Lerma and J. B. Lasserre. *Discrete-Time Markov Control Processes: Basic Optimality Criteria*. Springer, 1996.
- A. D. Kara and S. Yüksel. Convergence of finite memory q-learning for pomdps and near optimality of learned policies under filter stability. *arXiv preprint arXiv:2103.12158*, 2021.
- A. D. Kara, N. Saldi, and S. Yüksel. Weak feller property of non-linear filters. *Systems & Control Letters*, 134:104–512, 2019.
- V. Krishnamurthy. *Partially observed Markov decision processes: from filtering to controlled sensing*. Cambridge University Press, 2016.
- W.S. Lovejoy. A survey of algorithmic methods for partially observed Markov decision processes. *Annals of Operations Research*, 28:47–66, 1991.
- W. Mao, K. Zhang, E. Miehling, and T. Başar. Information state embedding in partially observable cooperative multi-agent reinforcement learning. *arXiv preprint arXiv:2004.01098*, 2020.
- C. McDonald and S. Yüksel. Observability and filter stability for partially observed markov processes. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 1623–1628. IEEE, 2019.
- C. McDonald and S. Yüksel. Exponential filter stability via Dobrushin’s coefficient. *Electronic Communications in Probability*, 25, 2020.
- C. McDonald and S. Yüksel. Robustness to incorrect priors and controlled filter stability in partially observed stochastic control. *SIAM Journal on Control and Optimization*, to appear (also *arXiv:2110.03720*), 2022.
- K.R. Parthasarathy. *Probability Measures on Metric Spaces*. AMS Bookstore, 1967.
- J. Pineau, G. Gordon, and S. Thrun. Anytime point-based approximations for large pomdps. *Journal of Artificial Intelligence Research*, 27:335–380, 2006.
- J. M. Porta, N. Vlassis, M. T. J. Spaan, and P. Poupart. Point-based value iteration for continuous pomdps. *Journal of Machine Learning Research*, 7(Nov):2329–2367, 2006.
- D. Rhenius. Incomplete information in Markovian decision models. *Ann. Statist.*, 2:1327–1334, 1974.
- N. Saldi, S. Yüksel, and T. Linder. On the asymptotic optimality of finite approximations to markov decision processes with borel spaces. *Mathematics of Operations Research*, 42(4):945–978, 2017.
- N. Saldi, T. Linder, and S. Yuksel. *Finite Approximations in discrete-time stochastic control*. Birkhäuser, 2018.

- N. Saldi, S. Yüksel, and T. Linder. Finite model approximations for partially observed markov decision processes with discounted cost. *IEEE Transactions on Automatic Control*, 65, 2020.
- T. Smith and R. Simmons. Point-based pomdp algorithms: Improved analysis and implementation. *arXiv preprint arXiv:1207.1412*, 2012.
- J. Subramanian and A. Mahajan. Approximate information state for partially observed systems. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 1629–1636, 2019.
- C. Villani. *Optimal transport: old and new*. Springer, 2008.
- N. Vlassis and M. T. J. Spaan. Perseus: Randomized point-based value iteration for pomdps. *Journal of artificial intelligence research*, 24:195–220, 2005.
- C.C. White. A survey of solution techniques for the partially observed Markov decision process. *Annals of Operations Research*, 32:215–230, 1991.
- C. C. White-III and W. T. Scherer. Finite-memory suboptimal design for partially observed markov decision processes. *Operations Research*, 42(3):439–455, 1994.
- H. Yu and D.P. Bertsekas. On near optimality of the set of finite-state controllers for average cost pomdp. *Mathematics of Operations Research*, 33(1):1–11, 2008.
- A.A. Yushkevich. Reduction of a controlled Markov model with incomplete data to a problem with complete information in the case of Borel state and control spaces. *Theory Prob. Appl.*, 21: 153–158, 1976.
- Z. Zhang, D. Hsu, and W. S. Lee. Covering number for efficient heuristic-based pomdp planning. In *International conference on machine learning*, pages 28–36. PMLR, 2014.
- E. Zhou, M. C. Fu, and S. I. Marcus. A density projection approach to dimension reduction for continuous-state POMDPs. In *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, pages 5576–5581, 2008.
- E. Zhou, M. C. Fu, and S. I. Marcus. Solving continuous-state POMDPs via density projection. *IEEE Transactions on Automatic Control*, 55(5):1101 – 1116, 2010.
- R. Zhou and E.A. Hansen. An improved grid-based approximation algorithm for POMDPs. In *Int. J. Conf. Artificial Intelligence*, pages 707–714, Aug. 2001.