

ROBUSTNESS TO INCORRECT PRIORS AND CONTROLLED FILTER STABILITY IN PARTIALLY OBSERVED STOCHASTIC CONTROL*

CURTIS McDONALD[†] AND SERDAR YÜKSEL[‡]

Abstract. We study controlled filter stability and its effects on the robustness properties of optimal control policies designed for systems with incorrect priors applied to a true system. Filter stability refers to the correction of an incorrectly initialized filter for a partially observed stochastic dynamical system (controlled or control-free) with increasing measurements. This problem has been studied extensively in the control-free context, and except for the standard machinery for linear Gaussian systems involving the Kalman filter, few studies exist for the controlled setup. One of the main differences between control-free and controlled partially observed Markov chains is that the filter is always Markovian under the former, whereas under a controlled model the filter process may not be Markovian since the control policy may depend on past measurements in an arbitrary (measurable) fashion. This complicates the dependency structure and therefore results from the control-free literature do not directly apply to the controlled setup. In this paper, we study the filter stability problem for controlled stochastic dynamical systems and provide sufficient conditions for when a falsely initialized filter merges with the correctly initialized filter over time. These stability results are applied to robust stochastic control problems: under filter stability, we bound the difference in the expected cost incurred for implementing an incorrectly designed control policy compared to an optimal policy. A conclusion is that filter stability leads to stronger robustness results to incorrect priors (compared with results in [A. D. Kara and S. Yüksel, *SIAM J. Control Optim.*, 57 (2019), pp. 1929–1964] without controlled filter stability). Furthermore, if the optimum cost is the same for each prior, the cost of mismatch between the true prior and the assumed prior is zero.

Key words. partially observed Markov decision processes (POMDPs), robustness, nonlinear filtering

AMS subject classifications. 93E20, 93E11, 93B07

DOI. 10.1137/21M1417442

1. Introduction. In this paper we study partially observed Markov decision problems (POMDPs) where a controller/decision maker (DM) has an incorrect prior on the initial state of the system. (i) We first study (controlled) filter stability: the effect this incorrect prior has on the sequence of conditional probabilities on the state variable given the measurements, and when the increasingly available measurement data over time leads the filter to correct itself. (ii) Then we study robustness, when the controller selects a control policy for some cost minimization criterion based on this incorrect prior. We investigate the performance loss due to the mismatch error, that is, the difference in expected cost induced by this policy under the true prior compared to that induced by an optimal control policy.

As we note below, neither of these evidently much-related problems has been considered in the literature in the generality considered in this paper. Filter stability has

*Received by the editors May 4, 2021; accepted for publication (in revised form) December 21, 2021; published electronically April 4, 2022. A preliminary version of some of the results was reported in *Proceedings of the 58th Annual IEEE Conference on Decision and Control* [34].

<https://doi.org/10.1137/21M1417442>

Funding: This research was supported in part by the Natural Sciences and Engineering Research Council (NSERC) of Canada.

[†]Department of Statistics and Data Science, Yale University, New Haven, CT 06511 USA (curtis.mcdonald@yale.edu).

[‡]Department of Mathematics and Statistics, Queen's University, Kingston, ON, K7L 3N6 Canada, (yüksel@mast.queensu.ca).

TABLE 1.1

Some notation for random variables, measures, and sigma fields.

Notation	
$Y_{[0,n]}$	A finite collection of random variables $\{Y_0, Y_1, \dots, Y_n\}$.
$Y_{[0,\infty)}$	An infinite sequence of random variables Y_0, Y_1, \dots .
$\mathcal{P}(\mathcal{X})$	The space of probability measures on the space \mathcal{X} with Borel sigma field $\mathcal{B}(\mathcal{X})$.
$\mu \ll \nu$	The measure μ is absolutely continuous with respect to ν .
$\mathcal{F}_{a,b}^{\mathcal{X}}$	The sigma field generated by (X_a, \dots, X_b) .
$\mathcal{F}_n^{\mathcal{X}}$	The sigma field generated by X_n .
$\mathcal{F}_{0,\infty}^{\mathcal{X}} \vee \mathcal{F}_{0,\infty}^{\mathcal{Y}}$	The sigma field generated by all state and measurement sequences.
$P^{\mu,\gamma}$	The strategic measure on the full sigma field $\mathcal{F}_{0,\infty}^{\mathcal{X}} \vee \mathcal{F}_{0,\infty}^{\mathcal{Y}}$.
$P^{\mu,\gamma}(X_{[0,n]}) \equiv P^{\mu,\gamma} _{\mathcal{F}_{0,n}^{\mathcal{X}}}$	The measure $P^{\mu,\gamma}$ restricted to the sigma field $\mathcal{F}_{0,n}^{\mathcal{X}}$.
$P^{\mu,\gamma}(X_{[0,\infty)}, Y_{[0,\infty)} Y_{[0,n]}) \equiv P^{\mu,\gamma} _{\mathcal{F}_{0,n}^{\mathcal{Y}}}$	The conditional measure of $P^{\mu,\gamma}$ with respect to the sigma field $\mathcal{F}_{0,n}^{\mathcal{Y}}$.

primarily been studied for control-free systems, and robustness (to incorrect priors) has been studied without the consideration of filter stability.

We now describe the probabilistic setup in more detail; see Table 1.1 for specific notation used to describe random variables, measures, and sigma fields. Let $\mathcal{X}, \mathcal{Y}, \mathcal{U}$ be Polish spaces (that is, complete, separable, and metric spaces) equipped with their Borel sigma fields $\mathcal{B}(\mathcal{X}), \mathcal{B}(\mathcal{Y}), \mathcal{B}(\mathcal{U})$; \mathcal{X} will be called the state space, \mathcal{Y} the measurement space, and \mathcal{U} the control action space. We also assume that \mathcal{X} is σ -compact.

Define the transition kernel T and observation kernel Q as the mappings

$$\begin{aligned}
 T : \mathcal{X} \times \mathcal{U} &\rightarrow \mathcal{P}(\mathcal{X}), & Q : \mathcal{X} &\rightarrow \mathcal{P}(\mathcal{Y}), \\
 (x, u) &\mapsto T(dx'|x, u), & x &\mapsto Q(dy|x).
 \end{aligned}$$

The POMDP is initialized with a state X_0 distributed according to a prior $\mu \in \mathcal{P}(\mathcal{X})$. However, the DM does not have access to the state realizations, but instead sees Y_0, Y_1, \dots, Y_n at time n . An admissible (deterministic) control policy $\gamma = \{\gamma_n\}_{n=0}^\infty$ is a sequence of measurable mappings $\gamma_n : \mathcal{Y}^n \times \mathcal{U}^{n-1} \rightarrow \mathcal{U}$ that maps past measurements and control actions to a new control action. At each time stage, the DM uses its control policy to apply a control action $U_n = \gamma_n(Y_{[0,n]}, U_{[0,n-1]})$ which affects the transition kernel to the next state, $X_n \sim T(dx|X_{n-1}, U_{n-1})$. A new observation is made and the process repeats. Recursively, it follows that U_0 is a function of Y_0 and U_1 is a function of Y_0, Y_1 , and U_0 . Yet since U_0 is itself a function of Y_0 , we have that U_1 is essentially just a function of $Y_{[0,1]}$. In other words, we can restrict ourselves to control policies that are only functions of measurements $Y_{[0,n]}$, though explicit dependence on past actions will be useful for some of our results to follow. We will denote the collection of *admissible* control policies as Γ .

Consider the measurable space $\Omega = \mathcal{X}^{\mathbb{Z}^+} \times \mathcal{Y}^{\mathbb{Z}^+}$, endowed with the product topology (that is, $\omega \in \Omega$ is a sequence of states and measurements $\omega = \{(x_i, y_i)\}_{i=0}^\infty$).

DEFINITION 1.1. For a fixed initial measure $\mu \in \mathcal{P}(\mathcal{X})$ and a policy $\gamma \in \Gamma$, we define the strategic measure $P^{\mu,\gamma}$ as the probability measure on $(\Omega, \mathcal{B}(\Omega))$ such that we have the following:

(i) For all $A \in \mathcal{B}(\mathcal{X} \times \mathcal{Y})$ we have

$$(1.1) \quad P^{\mu, \gamma}((X_0, Y_0) \in A) = \int_A Q(dy|x)\mu(dx).$$

(ii) For every $n \geq 1$, for all $A \in \mathcal{B}(\mathcal{X} \times \mathcal{Y})$, let $u_{n-1} = \gamma_{n-1}(y_{[0, n-1]}, u_{[0, n-2]})$; then we have

$$(1.2) \quad \begin{aligned} P^{\mu, \gamma}((X_n, Y_n) \in A | (X, Y)_{[0, n-1]} = (x, y)_{[0, n-1]}) \\ = \int_A Q(dy|x)T(dx|x_{n-1}, u_{n-1}). \end{aligned}$$

Remark 1.1. Note that $(X, Y)_{[0, \infty)}$ is, in general, not a Markov chain under $P^{\mu, \gamma}$ as u_{n-1} depends on the past measurements in (1.2).

1.1. Filter stability. Given a prior $\mu \in \mathcal{P}(\mathcal{X})$ and a policy $\gamma \in \Gamma$, we can then define the filter and predictor for a POMDP using the strategic measure $P^{\mu, \gamma}$.

DEFINITION 1.2.

(i) We define the one step predictor process as the sequence of conditional probability measures

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

(ii) We define the filter process as the sequence of conditional probability measures

$$(1.3) \quad \pi_n^{\mu, \gamma}(\cdot) = P^{\mu, \gamma}(X_n \in \cdot | Y_{[0, n]}, U_{[0, n-1]}) = P^{\mu, \gamma}(X_n \in \cdot | Y_{[0, n]}), \quad n \in \mathbb{Z}_+.$$

Remark 1.2. Recall that the $U_{[0, n-1]}$ are all functions of the $Y_{[0, n-1]}$, so conditioning on the control actions is not necessary in the above definitions. Yet this conditional probability would be *policy dependent*. If we condition on the past actions, this conditioning would be *policy independent*. This distinction is important for some of our discussions to follow.

Remark 1.3. It will be useful to note that the filter is the strategic measure conditioned on the sigma field $\mathcal{F}_{0, n}^{\mathcal{Y}}$ and restricted to the sigma field $\mathcal{F}_n^{\mathcal{X}}$.

$$\pi_n^{\mu, \gamma}(\cdot) = P^{\mu, \gamma}(X_n \in \cdot | Y_{[0, n]}) = P^{\mu, \gamma}|_{\mathcal{F}_n^{\mathcal{X}} | \mathcal{F}_{0, n}^{\mathcal{Y}}}.$$

Say a prior $\mu \in \mathcal{P}(\mathcal{X})$ and a policy $\gamma \in \Gamma$ are chosen; an observer sees measurements $Y_{[0, \infty)}$ generated via the strategic measure $P^{\mu, \gamma}$. The observer is aware that the policy applied is γ , but incorrectly thinks the prior is $\nu \neq \mu$. The observer will then compute the incorrectly initialized filter $\pi_n^{\nu, \gamma}$ while the true filter is $\pi_n^{\mu, \gamma}$. The filter stability problem is concerned with the merging of $\pi_n^{\nu, \gamma}$ and $\pi_n^{\mu, \gamma}$ as n goes to infinity.

In the literature, there are a number of merging notions when one considers stability which we enumerate here. Let $C_b(\mathcal{X})$ represent the set of continuous and bounded functions from $\mathcal{X} \rightarrow \mathbb{R}$.

DEFINITION 1.3. Two sequences of probability measures P_n, Q_n merge weakly if $\forall f \in C_b(\mathcal{X})$ we have $\lim_{n \rightarrow \infty} |\int f dP_n - \int f dQ_n| = 0$.

DEFINITION 1.4. For two probability measures P and Q we define the total variation norm as $\|P - Q\|_{TV} = \sup_{\|f\|_{\infty} \leq 1} |\int f dP - \int f dQ|$, where f is assumed measurable. We say two sequences of probability measures P_n, Q_n merge in total variation if $\|P_n - Q_n\|_{TV} \rightarrow 0$ as $n \rightarrow \infty$.

DEFINITION 1.5.

- (i) For two probability measures P and Q we define the relative entropy as $D(P\|Q) = \int \log \frac{dP}{dQ} dP = \int \frac{dP}{dQ} \log \frac{dP}{dQ} dQ$, where we assume $P \ll Q$ and $\frac{dP}{dQ}$ denotes the Radon–Nikodym derivative of P with respect to Q .
- (ii) Let X and Y be two random variables, and let P and Q be two different joint measures for (X, Y) with $P \ll Q$. Then we define the (conditional) relative entropy between $P(X|Y)$ and $Q(X|Y)$ as

$$\begin{aligned}
 D(P(X|Y)\|Q(X|Y)) &= \int \log \left(\frac{dP_{X|Y}}{dQ_{X|Y}}(x, y) \right) dP(x, y) \\
 (1.4) \qquad \qquad \qquad &= \int \left(\int \log \left(\frac{dP_{X|Y}}{dQ_{X|Y}}(x, y) \right) dP(x|Y = y) \right) dP(y).
 \end{aligned}$$

Total variation merging implies weak merging, and relative entropy merging (i.e., $D(P_n\|Q_n) \rightarrow 0$) implies total variation merging via Pinsker’s inequality [19].

As (1.4) shows, the relative entropy of two conditional measures is the expectation of the relative entropy of the conditional measures at a fixed conditional value $D(P(X|Y = y)\|Q(X|Y = y))$, where the expectation is taken over the marginal of P on Y . In the case of the filters, $\pi_n^{\mu, \gamma}$ and $\pi_n^{\nu, \gamma}$ play the role of the inner conditional measures, and the strategic measure $P^{\mu, \gamma}$ is the outer marginal measure over $Y_{[0, n]}$. We write this as $E^{\mu, \gamma}[D(\pi_n^{\mu, \gamma}\|\pi_n^{\nu, \gamma})]$ where $D(\pi_n^{\mu, \gamma}\|\pi_n^{\nu, \gamma})$ plays the role of the inner integral in (1.4).

1.2. Control cost and robustness. In addition to filter stability the second, and perhaps more operational, goal of this paper is to utilize stability to study robustness of optimal control to incorrect priors. We will consider two optimal stochastic control criteria:

- 1. The infinite horizon discounted cost

$$J_\beta(\mu, \gamma) = E^{\mu, \gamma} \left[\sum_{i=0}^{\infty} \beta^i c(X_i, U_i) \right], \quad \beta \in [0, 1).$$

- 2. The infinite horizon average cost

$$J_\infty(\mu, \gamma) = \limsup_{N \rightarrow \infty} \frac{1}{N} E^{\mu, \gamma} \left[\sum_{i=0}^{N-1} c(X_i, U_i) \right].$$

For a given prior μ , we will also consider the optimal costs:

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

An optimal control policy γ^μ for a given prior μ is the control policy which achieves the infimum of the expected cost over all admissible control policies,

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

Consider then whether a controller falsely thinks the prior of a system is ν , when in reality the prior is μ . Then the controller will implement the policy γ^ν which is optimal with respect to ν but incurs an expected cost of $J_\beta(\mu, \gamma^\nu)$. If the controller

had utilized the correctly designed policy, the cost could have been $J_\beta^*(\mu)$. In studying robustness, we are interested in studying this difference:

$$(1.5) \quad J_\beta(\mu, \gamma^\nu) - J_\beta^*(\mu),$$

$$(1.6) \quad J_\infty(\mu, \gamma^\nu) - J_\infty^*(\mu)$$

for the discounted cost problem and the average cost problem. We will show that bounds can be derived by utilizing filter stability.

1.3. Relations between the two problems. We now present the connections between filter stability and robustness (in the context of the discounted cost criterion, this also applies to the average cost setup). Consider a prior μ and a control policy γ^ν which is optimal with respect to a different prior ν . For any $n \in \mathbb{N}$, we have

$$(1.7) \quad \begin{aligned} J_\beta(\mu, \gamma^\nu) &= E^{\mu, \gamma^\nu} \left[\sum_{i=0}^{\infty} \beta^i c(X_i, U_i) \right] \\ &= E^{\mu, \gamma^\nu} \left[\sum_{i=0}^{n-1} \beta^i c(X_i, U_i) \right] + E^{\mu, \gamma^\nu} \left[E^{\mu, \gamma^\nu} \left[\sum_{i=n}^{\infty} \beta^i c(X_i, U_i) \mid Y_{[0, n-1]} \right] \right] \\ &= E^{\mu, \gamma^\nu} \left[\sum_{i=0}^{n-1} \beta^i c(X_i, U_i) \right] + (\beta^n) E^{\mu, \gamma^\nu} \left[E^{\mu, \gamma^\nu} \left[\sum_{i=0}^{\infty} \beta^i c(X_{n+i}, U_{n+i}) \mid Y_{[0, n-1]} \right] \right]. \end{aligned}$$

As will be discussed in more detail in the literature review, under mild regularity conditions for a POMDP an optimal control policy γ^μ is a stationary function of the filter realization at time n , π_n^{μ, γ^μ} . That is, consider some measurable mapping $\Phi : \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{U}$ that maps a probability measure to a control action. Recall an optimal control policy $\gamma^\mu = \{\gamma_n^\mu\}_{n=0}^\infty$ is a sequence of mappings $\gamma_n^\mu : \mathcal{Y}^{n+1} \rightarrow \mathcal{U}$. Then there exists some operator Φ such that each γ_n^μ mapping is equivalent to this operator acting on the filter realization at time n :

$$\gamma_n^\mu(y_{[0, n]}) = \Phi(P^{\mu, \gamma^\mu}(X_n \in \cdot \cdot \mid Y_{[0, n]} = y_{[0, n]})) = \Phi(\pi_n^{\mu, \gamma^\mu}).$$

Consider the 0th time stage in the problem. X_0 is distributed according to μ , but the DM thinks the initial prior is ν . An observation Y_0 is made, and the control action u_0 is a function of the filter realization believing the prior is ν ,

$$u_0 = \gamma_0^\nu(y_0) = \Phi(\pi_0^{\nu, \gamma^\nu}).$$

Now consider the n th time stage. Conditioned on $Y_{[0, n-1]}$, X_n is distributed according to π_n^{μ, γ^ν} (the true predictor), but the DM thinks the distribution is the false predictor π_n^{ν, γ^ν} . An observation Y_n is made, and the control action u_n is the same stationary function of the filter realization believing the prior is ν ,

$$u_n = \gamma_n^\nu(y_{[0, n]}) = \Phi(\pi_n^{\nu, \gamma^\nu}).$$

Therefore, the optimal policy under prior ν at time n is the same as the optimal policy under the prior $\nu' = \pi_n^{\nu, \gamma^\nu}$ at time 0 since it is a stationary function of the filter realization:

$$\gamma_n^\nu(y_{[0, n]}) = \Phi(\pi_n^{\nu, \gamma^\nu}) = \Phi(\pi_0^{\nu', \gamma^{\nu'}}) = \gamma_0^{\nu'}(y_n).$$

We see that the true predictor at time n , $X_n|Y_{[0,n-1]} \sim \pi_{n-}^{\mu, \gamma^\nu}$, acts as a new prior for a restarted control problem, and the control policy implemented is optimal with respect to the false predictor $\nu' = \pi_{n-}^{\nu, \gamma^\nu}$. We then have

$$E^{\mu, \gamma^\nu} \left[\sum_{i=0}^{\infty} \beta^i c(X_{i+n}, U_{i+n}) | Y_{[0,n-1]} \right] = J_\beta(\pi_{n-}^{\mu, \gamma^\nu}, \gamma^{\nu'}).$$

Therefore,

$$J_\beta(\mu, \gamma^\nu) = E^{\mu, \gamma^\nu} \left[\sum_{i=0}^{n-1} \beta^i c(X_i, U_i) \right] + (\beta^n) E^{\mu, \gamma^\nu} \left[J_\beta(\pi_{n-}^{\mu, \gamma^\nu}, \gamma^{\nu'}) \right].$$

If we instead apply the correctly designed policy γ^μ and let $\mu' = \pi_{n-}^{\mu, \gamma^\mu}$ be the correctly initialized predictor, we have

$$\begin{aligned} J_\beta^*(\mu) &= E^{\mu, \gamma^\mu} \left[\sum_{i=0}^{n-1} \beta^i c(X_i, U_i) \right] + (\beta^n) E^{\mu, \gamma^\mu} \left[J_\beta(\pi_{n-}^{\mu, \gamma^\mu}, \gamma^{\mu'}) \right] \\ &= E^{\mu, \gamma^\mu} \left[\sum_{i=0}^{n-1} \beta^i c(X_i, U_i) \right] + (\beta^n) E^{\mu, \gamma^\mu} \left[J_\beta^*(\pi_{n-}^{\mu, \gamma^\mu}) \right]. \end{aligned}$$

The difference satisfies

$$\begin{aligned} &J_\beta(\mu, \gamma^\nu) - J_\beta^*(\mu) \\ &= E^{\mu, \gamma^\nu} \left[\sum_{i=0}^{n-1} \beta^i c(X_i, U_i) \right] - E^{\mu, \gamma^\mu} \left[\sum_{i=0}^{n-1} \beta^i c(X_i, U_i) \right] \\ &\quad + \beta^n \left(E^{\mu, \gamma^\nu} \left[J_\beta(\pi_{n-}^{\mu, \gamma^\nu}, \gamma^{\nu'}) \right] - E^{\mu, \gamma^\mu} \left[J_\beta^*(\pi_{n-}^{\mu, \gamma^\mu}) \right] \right) \\ &= E^{\mu, \gamma^\nu} \left[\sum_{i=0}^{n-1} \beta^i c(X_i, U_i) \right] - E^{\mu, \gamma^\mu} \left[\sum_{i=0}^{n-1} \beta^i c(X_i, U_i) \right] \\ &\quad + \beta^n \left(E^{\mu, \gamma^\nu} \left[J_\beta(\pi_{n-}^{\mu, \gamma^\nu}, \gamma^{\nu'}) \right] + J_\beta^*(\pi_{n-}^{\mu, \gamma^\nu}) - J_\beta^*(\pi_{n-}^{\mu, \gamma^\nu}) \right] - E^{\mu, \gamma^\mu} \left[J_\beta^*(\pi_{n-}^{\mu, \gamma^\mu}) \right] \end{aligned} \tag{1.8}$$

$$\begin{aligned} &= E^{\mu, \gamma^\nu} \left[\sum_{i=0}^{n-1} \beta^i c(X_i, U_i) \right] - E^{\mu, \gamma^\mu} \left[\sum_{i=0}^{n-1} \beta^i c(X_i, U_i) \right] \\ &\quad + \beta^n \left(E^{\mu, \gamma^\nu} \left[J_\beta^*(\pi_{n-}^{\mu, \gamma^\nu}) \right] - E^{\mu, \gamma^\mu} \left[J_\beta^*(\pi_{n-}^{\mu, \gamma^\mu}) \right] \right) \end{aligned} \tag{1.9}$$

$$\begin{aligned} &+ \beta^n \left(E^{\mu, \gamma^\nu} \left[J_\beta(\pi_{n-}^{\mu, \gamma^\nu}, \gamma^{\nu'}) \right] - J_\beta^*(\pi_{n-}^{\mu, \gamma^\nu}) \right). \end{aligned} \tag{1.10}$$

Therefore, we see that there are, in general, three costs associated with applying an incorrectly designed control policy to a control system. The first cost (1.8) can be thought of as the “transient” cost. At time $[0, n - 1]$ the control policy γ^ν makes different control decisions than the optimal policy γ^μ . As such, the costs incurred from time $[0, n - 1]$ will be different for the optimal control policy and the incorrectly designed policy. The second cost (1.9) is the “strategic measure cost”: P^{μ, γ^ν} and

P^{μ, γ^μ} place different expected distributions on X_n , which acts as the new prior for the system for future stages. Therefore, even if the DM implements optimal control from time n onwards, the “initial condition” the DM finds itself in at time n may have a poor optimal cost over the future time stages. The third cost (1.10) is the “filter approximation cost,” which is *the term related to filter stability*. Under predictor stability the falsely initialized predictor $\pi_{n-}^{\nu, \gamma^\nu}$ and the true predictor $\pi_{n-}^{\mu, \gamma^\mu}$ merge as time goes on. These three costs form the fundamentals of studying the losses a DM incurs for using an incorrectly designed policy.

The rest of this paper is organized as follows. In section 2, we review existing literature related to filter stability, the existence and structure of optimal control policies, and robustness for stochastic control problems. In section 3, we present our main results for filter stability in control systems, as well as our results for robustness derived from filter stability. Proofs and further discussion of the filter stability results are in section 4, and proofs and discussion of the robustness results are in section 5. Section 6 discusses generalizations, some extensions, and directions for future research.

2. Literature review.

2.1. Controlled filter stability and nonlinear observability. Observability is one of the foundational concepts of modern linear systems theory [15, 1, 24, 29]. For linear systems, exact recovery of any initial condition with measurements available until some finite time is defined as observability and is characterized by an observability rank condition in both continuous- and discrete-time [16]. For linear systems, such an observability definition is global (as it applies for all initial states) and is universal in the control policies applied, as the control policy does not affect the estimation errors (known as the *no-dual effect* [3] property). For nonlinear systems, however, due to the challenges in the analysis which prevent globality as well as control-dependence, more modest and localized definitions are to be imposed: For deterministic continuous-time nonlinear systems [22, 39] present local indistinguishability conditions with subtle differences, and establish relations with Lie-theoretic characterizations which generalize observability rank conditions for nonlinear systems defined locally. For discrete-time deterministic models, observability has also been defined by invertibility or exact recovery of an initial state, locally, given measurements with finitely many observations. Nijmeier [36] developed discrete-time analogues of the observability notions presented in [22] (see also [39] for sampled continuous-time systems).

Such definitions, either with local or with exact invertibility, are often too restrictive for stochastic systems driven by noise. Liu and Bitmead [32, 31] introduced a nonlinear stochastic observability definition through entropy, where the conditional entropy of the hidden state given measurements not being the same as the unconditional entropy implies their observability notion. Ugrinovskii [42] also presents an information-theoretic formulation, and defines observability as an informativeness, but not invertibility, condition, which would lead to a mild notion of observability for nonlinear systems.

In the filtering literature for control systems, the classical setup involves the linear Gaussian system. The filter in this case is the celebrated Kalman filter, where the finite-dimensional Kalman filter is computed recursively using the Riccati equation. Under linear observability and controllability conditions, the Riccati equation admits a unique solution [29, 30, 15], which is the unique limit of the Riccati recursions regardless of the initialization. Thus, the Kalman filter is stable with respect to incorrect, though still Gaussian, priors under the aforementioned conditions (we note that partial convergence and robustness results on the asymptotic equivalence of con-

ditional expectations and linear estimates for non-Gaussian priors for linear systems are reported in [40]). This concept of being insensitive to incorrect initializations is called *filter stability*.

Aside from the above, much of the results on filter stability involves control-free systems. Thus, results have considered partially observed Markov processes (POMPs) as opposed to partially observed Markov *decision* processes (POMDPs). Since there is no control in such systems, there is no past dependency in the system and the pair $(X_n, Y_n)_{n=0}^\infty$ is always a Markov chain. For such control-free models, filter stability has been studied extensively, and we refer the reader to [18] for a comprehensive review and a collection of different approaches. As discussed in [18], filter stability may arise via two separate mechanisms: (i) The transition kernel is in some sense *sufficiently* ergodic, forgetting the initial measure and therefore passing this insensitivity (to incorrect initializations) onto the filter process. (ii) The measurement channel is sufficiently informative (or the system is sufficiently observable) about the underlying state. Much of the control-free literature has focused on the first of the two mechanisms noted above.

For the second type of mechanisms noted above, Chigansky and Liptser [17] and a series of papers by van Handel [43, 44, 46] have presented the first conditions, to the best of our knowledge, where various observability/informativeness conditions lead to filter stability for control-free systems (see also an earlier study, [41], for the linear measurements setup). A related fundamental result which pairs with observability, intrinsically connected with the analysis in the papers noted above, is that of Blackwell and Dubins [7]: If P and Q are two measures on a fully observed stochastic process $\{X_n\}_{n=0}^\infty$ with $P \ll Q$, then the conditional distributions on the future based on the past merge in total variation P almost surely (a.s.), that is, $\|P(X_{[n+1,\infty)} \in \cdot | X_{[0,n]}) - Q(X_{[n+1,\infty)} \in \cdot | X_{[0,n]})\|_{TV} \rightarrow 0$ P a.s.. This result immediately shows that

$$\|P^{\mu,\gamma}(Y_n | Y_{[0,n-1]}) - P^{\nu,\gamma}(Y_n | Y_{[0,n-1]})\|_{TV} \rightarrow 0 \quad P^{\mu,\gamma}$$

when $\mu \ll \nu$ and when combined with uniform observability [44] results in predictor stability. We refer the reader to [17, 44] for related results utilizing such a convergence argument though [17] arrives at the convergence result without [7]. Van Handel also provides a definition of observability for compact state spaces in [44]: A (control-free) system is observable if every prior results in a unique probability measure on the measurement sequences, $P^\mu(Y_{[0,\infty)} \in \cdot) = P^\nu(Y_{[0,\infty)} \in \cdot) \implies \mu = \nu$. For noncompact spaces, there are further characterizations [44], such as uniform observability.

Again, for control-free systems, [33] introduces an explicit notion of nonlinear observability that we continue to work with in this paper. A control-free model is called one step observable if for every $f \in C_b(\mathcal{X})$ and every $\epsilon > 0$ there exists a measurable and bounded function g such that

$$(2.1) \quad \left\| f(\cdot) - \int_{\mathcal{Y}} g(y)Q(dy|\cdot) \right\|_\infty < \epsilon.$$

We will show one step observability has an identical application in control systems. However, for control-free models [33] also has a notion of multiple step observability: A POMP is N -step observable if for every $f \in C_b(\mathcal{X})$ and every $\epsilon > 0$ there exists a measurable and bounded function g such that

$$(2.2) \quad \left\| f(\cdot) - \int_{\mathcal{Y}} g(y_{[1,N]})Q(dy_{[1,N]}|X_1 = \cdot) \right\|_\infty < \epsilon.$$

A further notion is observability: A POMDP is *observable* if for every $f \in C_b(\mathcal{X})$ and every $\epsilon > 0$ there exist $N \in \mathbb{N}$ and a measurable and bounded function g such that (2.2) applies. We will show that such multistep definitions of observability are more intricate in a controlled setup (POMDP); see section 6.1.

Building on a functional analytic duality result between controllability and observability, and along the same spirit as in [44] with respect to an analysis involving the null space of an appropriate linear map, a recent work on nonlinear filter stability is [28].

We outline the rigorous argument for controlled filter stability in section 4, but here we summarize the main ideas. Blackwell and Dubins [7] ensure that the conditional measures on the observed chains $\{Y_n\}_{n=0}^\infty$ merge as time goes to infinity. We would like to trace this back to the unobserved variable, which is where a notion of observability allows us to conclude the $X_n|Y_{[0,n]}$ conditionals merge as well. However, in a control system the past reliance of control policies means $\{X_n, Y_n\}_{n=0}^\infty$ is not a Markov chain as it is in the uncontrolled case. Nonetheless, the observation channel $Y_n|X_n \sim Q$ is unaffected by control and is time invariant; thus with our notion of “one step” observability we are able to recreate similar merging results in the controlled setup.

For both control-free and controlled setups, [35] studied filter stability through the use of Dobrushin’s coefficient involving both of the kernels T and Q . Aside from [35], we are not aware of studies which also consider the controlled setup; the following from [35] will be utilized in our robustness analysis in the paper (for the discounted cost criterion analysis).

DEFINITION 2.1 (see [20, eq. 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*

$$(2.3) \quad \delta(K) = \inf \sum_{i=1}^n \min(K(x, A_i), K(y, A_i))t,$$

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

THEOREM 2.2 (see [35, Theorem 4.1]). *Let $\tilde{\delta}(T) := \inf_{u \in \mathcal{U}} \delta(T(\cdot|\cdot, u))$. Assume that for $\mu, \nu \in \mathcal{P}(\mathcal{X})$, we have $\mu \ll \nu$. Then we have*

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

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

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

2.2. On regularity of POMDPs and optimal policies. It is known that any POMDP can be reduced to a (completely observable) belief-MDP [47, 37], whose states are the posterior state distributions or *beliefs* of the observer; that is, the state at time t is the filter realization given in (1.3). We call the resulting MDP the belief-MDP. If one reduces a POMDP to such a belief-MDP and if one can establish the measurable selection criteria [23, Chap. 3] (e.g., via showing the weak Feller property), existence of optimal solutions for discounted cost problems would follow. Indeed, recently [21, 25] have established complementary conditions on when the kernel for the belief-MDP is weakly continuous.

Thus, for the discounted cost criterion, existence of optimal solutions, which are also stationary (in the belief state), follows. Existence also holds for the average cost

problem under some further technical conditions facilitating the vanishing discount method [10, 13, 12, 14]. The convex analytic method of Borkar [11] can also be utilized under the weak continuity conditions presented above (also see [23]); however, for the average cost setup some restrictions may need to be imposed on the initial state [2].

Robustness in control refers to the property that a control policy operates well under misspecified system dynamics or models, yet this concept does not have a singular definition in the literature. Robustness problems often involve studies on when a controller has an incorrect system model, or has some uncertainty about the specifics of the actual system model. A common approach in the literature is to design controllers that work with guaranteed performance bounds over a class of uncertain systems under some structured constraints (see [5, 48]). However, these studies are different from what we consider here in that the controller is *aware* that she may have the wrong system specifications and has some limited idea of what the possible models may be. In our problem, we start with two disparate priors μ and ν and show that filter stability leads to robustness (and not just continuity as shown in [26]). Thus, our paper provides some unification between controlled filter stability and optimal robust stochastic control.

We also note that controlled filter stability (in total variation and weak convergence/merging senses) has been utilized in rigorously establishing near-optimality of finite memory (finite window) control policies recently in [27].

2.3. Contributions.

- (i) *Controlled filter stability.* We study the filter stability problem for nonlinear stochastic systems driven by control. Towards this goal, we present a definition of observability, inspired from a related definition from our prior work for control-free systems. Using this definition, we provide sufficient conditions for when a falsely initialized filter merges with the correctly initialized filter over time. As noted, a primary difference between control-free and controlled partially observed Markov chains is that the filter is always Markovian under a control-free model but not so in a controlled setup as the control policy may depend on past measurements in an arbitrary, measurable manner. We also establish relations between various notions of stability for controlled nonlinear filters.
- (ii) *Robustness to incorrect priors under filter stability.* Building on our results on filter stability, we study the robustness problem in terms of the mismatch loss in the expected cost incurred for implementing an incorrectly designed control policy (optimal for an incorrect initialization) compared to an optimal policy (for a correct initialization), and relate filter stability and robustness. In particular, we show that, unlike [26] where only continuity properties of mismatch in the incorrect priors were established, under filter stability even distant incorrect initial priors may lead to correct costs; and under mild conditions, filter stability leads to insensitivity of mismatch to incorrect priors, and not just continuity in them.

3. Statements of main results. We define here the different notions of stability for the filter.

DEFINITION 3.1.

- (i) *A filter process is said to be stable in the sense of weak merging with respect to a policy γ $P^{\mu,\gamma}$ a.s. if there exists a set of measurement sequences $A \subset \mathcal{Y}^{\mathbb{Z}^+}$ with $P^{\mu,\gamma}$ probability 1 such that for any sequence in A , for any $f \in C_b(\mathcal{X})$*

and any prior ν with $\mu \ll \nu$ (i.e., for all Borel B $\nu(B) = 0 \implies \mu(B) = 0$), we have $\lim_{n \rightarrow \infty} \left| \int f d\pi_n^{\mu, \gamma} - \int f d\pi_n^{\nu, \gamma} \right| = 0$.

- (ii) A filter process is said to be stable in the sense of total variation in expectation with respect to a policy γ if for any measure ν with $\mu \ll \nu$ we have $\lim_{n \rightarrow \infty} E^{\mu, \gamma} [\|\pi_n^{\mu, \gamma} - \pi_n^{\nu, \gamma}\|_{TV}] = 0$.
- (iii) A filter process is said to be stable in the sense of total variation with respect to a policy γ $P^{\mu, \gamma}$ a.s. if there exists a set of measurement sequences $A \subset \mathcal{Y}^{\mathbb{Z}^+}$ with $P^{\mu, \gamma}$ probability 1 such that for any sequence in A , for any measure ν with $\mu \ll \nu$ we have $\lim_{n \rightarrow \infty} \|\pi_n^{\mu, \gamma} - \pi_n^{\nu, \gamma}\|_{TV} = 0$ $P^{\mu, \gamma}$ a.s.
- (iv) A filter process is said to be stable in the sense of relative entropy with respect to a policy γ if for any measure ν with $\mu \ll \nu$ we have $\lim_{n \rightarrow \infty} E^{\mu, \gamma} [D(\pi_n^{\mu, \gamma} \| \pi_n^{\nu, \gamma})] = 0$.
- (v) The filter is said to be universally stable in one of the above notions if the notion holds with respect to every admissible policy $\gamma \in \Gamma$.

Predictor stability is defined in an analogous fashion for each of the criteria above. These notions of stability are asymptotic notions of stability; they do not imply a rate of convergence. In some of the robustness results to be presented, we need a stronger notion of stability (see Theorem 2.2).

DEFINITION 3.2. A POMP is said to be universally exponentially stable in total variation if there exists a coefficient $0 < \alpha < 1$ such that for any $\mu \ll \nu$ and any policy γ we have

$$E^{\mu, \gamma} [\|\pi_{n+1}^{\mu, \gamma} - \pi_{n+1}^{\nu, \gamma}\|_{TV}] \leq \alpha E^{\mu, \gamma} [\|\pi_n^{\mu, \gamma} - \pi_n^{\nu, \gamma}\|_{TV}], \quad n \in \{0, 1, \dots\}.$$

3.1. Controlled filter stability results. A key condition that will drive the stability results is the following definition of observability.

DEFINITION 3.3 (a definition of observability for controlled stochastic systems). A POMDP is called one step observable (universal in admissible control policies) if for every $f \in C_b(\mathcal{X})$ and every $\epsilon > 0$ there exists a measurable and bounded function g such that

$$(3.1) \quad \left\| f(\cdot) - \int_{\mathcal{Y}} g(y) Q(dy|\cdot) \right\|_{\infty} < \epsilon.$$

THEOREM 3.4. Assume that $\mu \ll \nu$ and that the POMDP is one step observable. Then the predictor is universally stable weakly a.s.

In section 6, we present some generalizations but also explain why the proof method used for control-free systems does not apply for controlled systems for more relaxed (in particular, an N -step generalization) definitions of observability. A number of examples for measurement channels satisfying Definition 3.3 have been reported in [33, section 3].

The observability notion defined above only results in stability of the predictor in the weak sense $P^{\mu, \gamma}$ a.s. We next extend this stability to total variation $P^{\mu, \gamma}$ a.s. Let the measurement channel Q be dominated in the sense that there exists a reference measure λ such that $\forall x \in \mathcal{X}, Q(Y \in \cdot | x_n = x) \ll \lambda(\cdot)$. Then, we define the Radon–Nikodym derivative

$$(3.2) \quad q(x, y) := \frac{dQ(Y_n \in \cdot | x_n = x)}{d\lambda}(y),$$

which serves as a likelihood function (a conditional probability density function (pdf's)). We will consider one of the following assumptions.

ASSUMPTION 3.1.

- (i) $T(\cdot|x, u)$ is absolutely continuous with respect to a dominating measure ϕ for every $x \in \mathcal{X}, u \in \mathcal{U}$, so that $t(x_1, x, u) = \frac{dT(\cdot|x, u)}{d\phi}(x_1)$, where t is continuous in x for every $x_1 \in \mathcal{X}$ and $u \in \mathcal{U}$.
- (ii) $q(x, y)$ is bounded and continuous in x for every fixed y . Furthermore, $q(x, y) > 0$ for all $x \in \mathcal{X}, y \in \mathcal{Y}$.

ASSUMPTION 3.2. $T(\cdot|x, u)$ is absolutely continuous with respect to a dominating measure ϕ for every $x \in \mathcal{X}, u \in \mathcal{U}$, so that $t(x_1, x, u) = \frac{dT(\cdot|x, u)}{d\phi}(x_1)$. The family of (conditional densities) $\{t(\cdot, x, u)\}_{x \in \mathcal{X}, u \in \mathcal{U}}$ is uniformly bounded and equicontinuous.

THEOREM 3.5. Let $\mu \ll \nu$. Let Assumption 3.1 or Assumption 3.2 hold. If the predictor is universally stable in the weak sense a.s., then it is also universally stable in total variation a.s.

One of the key steps in the proof of Theorem 3.4 is that $P^{\mu, \gamma}(Y_n \in \cdot | Y_{[0, n-1]})$ and $P^{\nu, \gamma}(Y_n \in \cdot | Y_{[0, n-1]})$ merge in total variation $P^{\mu, \gamma}$ a.s. as $n \rightarrow \infty$. To achieve this in a POMDP, we apply Blackwell and Dubins [7] to the measurement process $\{Y_n\}_{n=0}^\infty$. However, [7] is fundamentally about predictive measures of the future given the past, and hence only directly implies predictor stability results, not the filter. Filter stability is studied next.

ASSUMPTION 3.3. The measurement channel Q is continuous in total variation. That is, for any sequence $a_n \rightarrow a \in \mathcal{X}$ we have $\|Q(\cdot|a_n) - Q(\cdot|a)\|_{TV} \rightarrow 0$ or in other words $\|P(Y_0 \in \cdot | X_0 = a_n) - P(Y_0 \in \cdot | X_0 = a)\|_{TV} \rightarrow 0$.

Assumption 3.1(ii), together with the related domination condition (3.2), implies Assumption 3.3 (see [25, section 2.3]); see also [25, Theorem 3] for a partial converse result.

THEOREM 3.6.

- (i) Let Assumption 3.3 hold. If the predictor is universally stable in weak merging a.s., then the filter is universally stable in weak merging in expectation.
- (ii) The filter is universally stable in total variation in expectation if and only if the predictor is universally stable in total variation in expectation.
- (iii) The filter is universally stable in total variation in expectation if and only if it is universally stable in total variation a.s.
- (iv) Let $\mu \ll \nu$, and assume for any policy γ there exists some finite n such that $E^{\mu, \gamma}[D(\pi_n^{\mu, \gamma} || \pi_n^{\nu, \gamma})] < \infty$ and some m such that $E^{\mu, \gamma}[D(P^{\mu, \gamma}|_{\mathcal{F}_{0, m}^y} || (P^{\nu, \gamma}|_{\mathcal{F}_{0, m}^y}))] < \infty$. Then the filter is universally stable in relative entropy if and only if it is universally stable in total variation in expectation.

The proofs are given in section 4.

3.2. Results on robustness to incorrect priors. Robustness and continuity properties in incorrect priors have been studied in [26] for the single stage cost problem and the infinite horizon discounted cost problem. This paper provides conditions for when robustness difference in (1.6) goes to 0 as $\nu \rightarrow \mu$ in either weak convergence or total variation topologies. This paper provides the following useful result.

THEOREM 3.7 (see [26, Theorem 3.2]). Assume the cost function c is bounded, nonnegative, and measurable. Let γ^ν be the optimal control policy designed with respect

to a prior ν . Then we have

$$(3.3) \quad J_\beta(\mu, \gamma^\nu) - J_\beta^*(\mu) \leq 2 \frac{\|c\|_\infty}{1-\beta} \|\mu - \nu\|_{TV}.$$

While the average cost case is not studied in [26], we can adapt the technique to achieve a similar result for the average cost problem (see the appendix for a proof).

THEOREM 3.8. *Assume the cost function c is bounded, nonnegative, and measurable. Let γ^ν be the optimal control policy designed with respect to a prior ν . Then we have*

$$J_\infty(\mu, \gamma^\nu) - J_\infty^*(\mu) \leq 2\|c\|_\infty \|\mu - \nu\|_{TV}.$$

However, in the following, we will refine these results by considering *the effects of filter stability on robustness to incorrect priors*. Unlike [26], where continuity in priors was studied, here we utilize filter stability to arrive at stronger insensitivity and robustness results.

For the average cost case, asymptotic filter stability (without a rate of convergence) may suffice for robustness, as we establish in the following.

THEOREM 3.9. *Assume the cost function c is bounded, nonnegative, and measurable and assume the predictor is universally stable in total variation in expectation. Consider the span seminorm $\|\cdot\|_{sp}$:*

$$\|J_\infty^*\|_{sp} := \sup_{\mu_1 \in \mathcal{P}(\mathcal{X})} J_\infty^*(\mu_1) - \inf_{\mu_2 \in \mathcal{P}(\mathcal{X})} J_\infty^*(\mu_2);$$

then we have

$$J_\infty(\mu, \gamma^\nu) - J_\infty^*(\mu) \leq \|J_\infty^*\|_{sp}.$$

In particular, if $\|J_\infty^*\|_{sp} = 0$, then the average cost optimization problem is completely robust to initialization errors (see [10, 13, 12, 14] for some setups where this holds).

We next study the discounted cost problem. This result, due to the discounted nature, is not as strong as the one above involving average cost control as it involves exponential filter stability to mitigate the effects of transient costs (which are absent in the average cost criterion setup).

THEOREM 3.10. *Assume the cost function c is bounded, nonnegative, and measurable and assume the filter is universally exponentially stable in total variation in expectation with coefficient α . Let*

$$(3.4) \quad \|J_\beta^*\|_{sp} = \sup_{\mu_1 \in \mathcal{P}(\mathcal{X})} J_\beta^*(\mu_1) - \inf_{\mu_2 \in \mathcal{P}(\mathcal{X})} J_\beta^*(\mu_2),$$

$$\rho = \left(\frac{\|c\|_\infty}{1-\beta} - \|J_\beta^*\|_{sp} \right) \left(\frac{\|c\|_\infty}{1-\beta} \right)^{-1},$$

$$(3.5) \quad f(n) = \beta^n (\rho - 4\alpha^n), \quad n^* = \frac{\ln \left(\left(\frac{\rho}{4} \right) \left(\frac{\ln(\beta)}{\ln(\alpha) + \ln(\beta)} \right) \right)}{\ln \alpha}.$$

For any priors $\mu \ll \nu$, we have

$$(3.6) \quad J_\beta(\mu, \gamma^\nu) - J_\beta^*(\mu) \leq \frac{\|c\|_\infty}{1-\beta} (1 - \max(f(\lfloor n^* \rfloor), f(\lceil n^* \rceil))).$$

This result may be considered a “prior independent” bound since the bound on the robustness distance does not depend on the actual priors μ and ν being considered. This is useful when we have no real knowledge of how “close” (in total variation distance) μ and ν are; thus, if they are far apart the bound will not change. Comparing (3.3) and (3.6), we see that when $\|\mu - \nu\|_{TV}$ is small, i.e., our priors start close together to begin with, the bound (3.3) is actually better than our prior independent bound. However, when they are far apart it may be that

$$2\|\mu - \nu\|_{TV} > 1 - \max(f(\lfloor n^* \rfloor), f(\lceil n^* \rceil))$$

and the prior independent bound is stronger.

4. Proofs on predictor and filter stability results.

4.1. Predictor stability results. Our approach to predictor stability is as follows. By the result of Blackwell and Dubins [7] we have that the conditional distributions on the measures merge in total variation almost surely for any control policy. Given any measurable and bounded function $g : \mathcal{Y}^{N+1} \rightarrow \mathbb{R}$ of multiple measurements we have

$$(4.1) \quad \int g(y_{[n,n+N]})P^{\mu,\gamma}(dy_{[n,n+N]}|Y_{[0,n-1]}) - \int g(y_{[n,n+N]})P^{\nu,\gamma}(dy_{[n,n+N]}|Y_{[0,n-1]})$$

goes to zero $P^{\mu,\gamma}$ almost surely as $n \rightarrow \infty$ for any control policy γ . Consider some continuous and bounded function f on the state space \mathcal{X} . In considering the weak merging of the predictor we take the limit of the terms

$$(4.2) \quad \left| \int f(x_n)P^{\mu,\gamma}(dx_n|y_{[0,n-1]}) - \int f(x_n)P^{\nu,\gamma}(dx_n|y_{[0,n-1]}) \right|.$$

Now, if we could replace f with a function $g : \mathcal{Y}^{N+1} \rightarrow \mathbb{R}$ integrated over a conditional measure of $P^{\mu,\gamma}(Y_{[0,N]} \in \cdot | X_0 = x)$, then we would have

$$\left| \int \left(\int g(\tilde{y}_{[0,N]})P^{\mu,\gamma}(d\tilde{y}_{[0,N]}|X_0 = x_n) \right) P^{\mu,\gamma}(dx_n|y_{[0,n-1]}) - \int \left(\int g(\tilde{y}_{[0,N]})P^{\nu,\gamma}(d\tilde{y}_{[0,N]}|X_0 = x_n) \right) P^{\nu,\gamma}(dx_n|y_{[0,n-1]}) \right|.$$

If we can then show that this is the same as expression (4.1), then the weak stability of the predictor follows. However, this will require some special properties of the conditional measure of $Y_{[n,n+N]}|X_n$. Let us consider a level of abstraction higher to better view these properties in more succinct notation. Consider three stochastic processes $A = \{A_n\}_{n=0}^\infty, B = \{B_n\}_{n=0}^\infty, C = \{C_n\}_{n=0}^\infty$ defined on the same measurable space (Ω, \mathcal{F}) mapping to spaces $\mathcal{A}, \mathcal{B}, \mathcal{C}$ with their respective Borel sigma fields. We can think of A and C as the “observed” processes and B as some “hidden” process whose realizations are not known by an observer. For possible measures P and P^* on (Ω, \mathcal{F}) we have the following definitions.

DEFINITION 4.1. *For a measure P we say the process C only depends on the process A through B if for every $n \in \mathbb{N}$ we have P a.s.*

$$P(C_n \in \cdot | A_n = a, B_n = b) = P(C_n \in \cdot | B_n = b).$$

DEFINITION 4.2. For a measure P we say the channel $C|B$ is time homogeneous if for every $n \in \mathbb{N}$, P a.s.

$$P(C_n \in \cdot | B_n = b) = P(C_0 \in \cdot | B_0 = b).$$

DEFINITION 4.3. For two measures P and P^* we say the channel $C|B$ is measure equivalent if for all $n \in \mathbb{N}$ we have, P, P^* a.s.

$$P(C_n \in \cdot | B_n = b) = P^*(C_n \in \cdot | B_n = b).$$

DEFINITION 4.4. For a measure P , the channel $C|B$ is observable if for every continuous and bounded function $f : \mathcal{B} \rightarrow \mathbb{R}$ and every $\epsilon > 0$ we can find a measurable and bounded function $g : \mathcal{C} \rightarrow \mathbb{R}$ such that

$$\sup_{b \in \mathcal{B}} \left| f(b) - \int_{\mathcal{C}} g(c_0) P(dc_0 | B_0 = b) \right| < \epsilon.$$

LEMMA 4.5. Let A, B, C be stochastic processes as above and assume measures P, P^* satisfy Definitions 4.1–4.4. Assume that

$$\lim_{n \rightarrow \infty} \|P(C_n | A_n) - P^*(C_n | A_n)\|_{TV} = 0 \quad P \text{ a.s.};$$

then we have that $P(B_n | A_n)$ and $P^*(B_n | A_n)$ merge weakly P a.s.

Proof. Consider any continuous and bounded function $f : \mathcal{B} \rightarrow \mathbb{R}$. Pick any $\epsilon > 0$; by observability (Definition 4.4) we can find a measurable and bounded function $g : \mathcal{C} \rightarrow \mathbb{R}$ such that

$$\tilde{f}(b) = \int_{\mathcal{C}} g(c_0) P(dc_0 | B_0 = b), \quad \|f - \tilde{f}\|_{\infty} < \frac{\epsilon}{3}.$$

Now consider

$$\begin{aligned} & \left| \int f(b_n) P(db_n | a_n) - \int f(b_n) P^*(db_n | a_n) \right| \\ & \leq \left| \int \tilde{f}(b_n) P(db_n | a_n) - \int \tilde{f}(b_n) P^*(db_n | a_n) \right| + \left| \int (f - \tilde{f})(b_n) P(db_n | a_n) \right| \\ & \quad + \left| \int (f - \tilde{f})(b_n) P^*(db_n | a_n) \right| \\ & \leq \left| \int \tilde{f}(b_n) P(db_n | a_n) - \int \tilde{f}(b_n) P^*(db_n | a_n) \right| + 2\|f - \tilde{f}\|_{\infty} \\ (4.3) \quad & \leq \left| \int \tilde{f}(b_n) P(db_n | a_n) - \int \tilde{f}(b_n) P^*(db_n | a_n) \right| + \frac{2}{3}\epsilon. \end{aligned}$$

We then have

$$\begin{aligned} & \left| \int_{\mathcal{B}} \tilde{f}(b_n) P(db_n | a_n) - \int_{\mathcal{B}} \tilde{f}(b_n) P^*(db_n | a_n) \right| \\ & = \left| \int_{\mathcal{B}} \int_{\mathcal{C}} g(c_0) P(dc_0 | B_0 = b_n) P(db_n | a_n) - \int_{\mathcal{B}} \int_{\mathcal{C}} g(c_0) P(dc_0 | B_0 = b_n) P^*(db_n | a_n) \right|. \end{aligned}$$

Furthermore, by measure equivalence (Definition 4.3) we can replace $P(dc_0|B_0 = b_n)$ with $P^*(dc_0|B_0 = b_n)$ in the second term. By time homogeneity (Definition 4.2), we can replace $P(dc_0|B_0 = b_n)$ with $P(dc_n|B_n = b_n)$ and the same for P^* . We then have

$$\left| \int_{\mathcal{B}} \int_{\mathcal{C}} g(c_n) P(dc_n|B_n = b_n) P(db_n|a_n) - \int_{\mathcal{B}} \int_{\mathcal{C}} g(c_n) P^*(dc_n|B_n = b_n) P^*(db_n|a_n) \right|.$$

By assumption, C only depends on A through B (Definition 4.1), so we can write $P(dc_n|B_n = b_n) = P(dc_n|B_n = b_n, A_n = a_n)$ for any a_n . We finally apply the chain rule for conditional probability, and we have

(4.4)

$$\left| \int_{\mathcal{C}} g(c_n) P(dc_n|a_n) - \int_{\mathcal{C}} g(c_n) P^*(dc_n|a_n) \right| \leq \|g\|_{\infty} \|P(C_n|A_n) - P^*(C_n|A_n)\|_{TV}.$$

By assumption, this goes to zero P a.s. as $n \rightarrow \infty$ (where the set of convergence of measure 1 applies for all g by total variation convergence) so we can find an N where $\forall n > N$ we have (4.4) is less than $\frac{\epsilon}{5}$, and therefore (4.3) is less than ϵ . \square

Proof of Theorem 3.4. We apply Lemma 4.5 to our one step observable POMDP. Denote the processes $A_n = Y_{[0,n-1]}$, $B_n = X_n$, and $C_n = Y_n$. Fix any control policy γ . We must then check that for the measures $P^{\mu,\gamma}$ and $P^{\nu,\gamma}$ these processes satisfy Definitions 4.1–4.4. The measurement channel $Y_n|X_n$ is unaffected by the presence of control and independent of the policy chosen. The distribution of Y_n is fully determined by Q and the realization of X_n :

$$P^{\mu,\gamma}(Y_n|Y_{[0,n-1]}, X_n) = Q(Y_n|X_n) = P^{\nu,\gamma}(Y_n|X_n, Y_{[0,n-1]}).$$

Therefore, the process Y_n only depends on $Y_{[0,n-1]}$ through X_n (Definition 4.1). Furthermore, the channel between $Y_n|X_n$ is Q regardless of the time index or the initial measure. Therefore, the channel is time homogeneous (Definition 4.2) and measure equivalent (Definition 4.3). The process is observable (Definition 4.4) by assumption of being one step observable. The assumption $\mu \ll \nu$ implies $P^{\mu,\gamma}|_{\mathcal{F}_{0,\infty}^{\mathcal{Y}}} \ll P^{\nu,\gamma}|_{\mathcal{F}_{0,\infty}^{\mathcal{Y}}}$. $\{Y_n\}_{n=0}^{\infty}$ is a fully observed stochastic process. Therefore, by Blackwell and Dubins [7] we have that

(4.5) $\|P^{\mu,\gamma}(Y_{[n,n+N-1]} \in \cdot | Y_{[0,n-1]}) - P^{\nu,\gamma}(Y_{[n,n+N-1]} \in \cdot | Y_{[0,n-1]})\|_{TV} \rightarrow 0.$

Therefore, we satisfy all the conditions of Lemma 4.5 and the weak merging of the predictor follows. This holds for any γ , so we have universal stability. \square

We first state the following supporting results.

LEMMA 4.6. *The (measurement-update) map*

$$(\pi_{n-}^{\mu,\gamma}, y_n) \mapsto \pi_n \quad : \quad \pi_n^{\mu,\gamma}(dx_n) := \frac{q(y_n|x_n)\pi_{n-}^{\mu,\gamma}(dx_n)}{\int_{\mathcal{X}} q(y_n|x_n)\pi_{n-}^{\mu,\gamma}(dx_n)},$$

which maps from $\mathcal{P}(\mathcal{X}) \times \mathcal{Y}$ to $\mathcal{P}(\mathcal{X})$, is weakly continuous in π_{n-} for almost every y , provided that $q(x, y)$ is positive, bounded, and continuous in x for every fixed y .

Proof. Consider a continuous and bounded f , and let $\pi_{n-}^m \rightarrow \pi_{n-}$ weakly. Then, $P^{\mu,\gamma}$ a.s.,

$$\int f(x_n) \frac{q(x_n, y_n)\pi_{n-}^m(dx_n)}{\int_{\mathcal{X}} q(x_n, y_n)\pi_{n-}^m(dx_n)} = \frac{\int f(x_n)q(x_n, y_n)\pi_{n-}^m(dx_n)}{\int_{\mathcal{X}} q(x_n, y_n)\pi_{n-}^m(dx_n)}.$$

Since q is bounded continuous, both the numerator and the denominator converge. \square

LEMMA 4.7. Let $T(dx_1|x, u) = t(x_1, x, u)\phi(dx_1)$, where t is continuous in x for every x_1 and u . Then, for any policy γ the (time-update) map

$$(\pi_n^{\mu, \gamma}, u_n) \mapsto \pi_{n+1-}^{\mu, \gamma} \quad : \quad \pi_{n+1-}^{\mu, \gamma}(dx_{n+1}) := \int_{\mathcal{X}} T(x_{n+1}|x_n, u_n)\pi_n^{\mu, \gamma}(dx_n),$$

which maps from $\mathcal{P}(\mathcal{X}) \times \mathcal{U}$ to $\mathcal{P}(\mathcal{X})$, is such that if $\pi_n^{\nu, \gamma} \rightarrow \pi_n^{\mu, \gamma}$ weakly, then $\pi_{n+1-}^{\nu, \gamma} \rightarrow \pi_{n+1-}^{\mu, \gamma}$ in total variation.

Proof. We will build on Scheffé's lemma [6]. For every given history and action, we have

$$\pi_{n+1-}^{\nu, \gamma}(dx_{n+1}) = \int T(dx_{n+1}|x_n, u_n)\pi_n^{\nu, \gamma}(dx_n).$$

Now, $\int T(dx_{n+1}|x_n, u_n)$ is such that for every fixed u_n , as $\pi_n^m \rightarrow \pi_n$ weakly

$$\begin{aligned} & \int_{x_{n+1} \in \cdot} \int_{x_n \in \mathcal{X}} t(x_{n+1}, x_n, u_n)\phi(dx_{n+1})\pi_n^m(dx_n) \\ & \rightarrow \int_{x_{n+1} \in \cdot} \int_{x_n \in \mathcal{X}} t(x_{n+1}, x_n, u_n)\phi(dx_{n+1})\pi_n(dx_n) \end{aligned}$$

in total variation since for every fixed z , the Radon–Nikodym derivative with respect to ϕ ,

$$\frac{d \int t(x_{n+1}, x_n, u_n)\phi(\cdot)\pi_n^m(dx_n)}{d\phi}(z) = \int t(z, x_n, u_n)\pi_n^m(dx_n),$$

satisfies pointwise convergence $\int t(z, x_n, u_n)\pi_n^m(dx_n) \rightarrow \int t(z, x_n, u_n)\pi_n(dx_n)$ (as $\pi_n^m \rightarrow \pi_n$ weakly), and Scheffé's lemma implies that convergence is in total variation. We apply this result to the sequence $\pi_n^{\nu, \gamma}$ converging to $\pi_n^{\mu, \gamma}$. \square

Proof of Theorem 3.5.

(i) Under Assumption 3.1, the proof follows from Lemmas 4.6 and 4.7. While in Lemmas 4.6 and 4.7 we consider convergence (and not merging), we note that the proof of Lemma 4.6 also implies weak merging of the posteriors as the priors weakly merge, and by considering the signed measure $\pi_n^{\nu, \gamma} - \pi_n^{\mu, \gamma}$ in the proof of Lemma 4.7, total variation merging is a result of a generalized Scheffé's lemma [8, Theorem 2.8.9].

(ii) Under Assumption 3.2, the proof technique follows the approach taken in [33, Theorem 2.10], which we will briefly recap and refer the reader to for a full consideration. Under Assumption 3.2, both predictors π_{n-}^{μ} and π_{n-}^{ν} admit pdf's f_{n-}^{μ} and f_{n-}^{ν} with respect to the dominating measure. Define the difference of the pdf's $f_{n-} = f_{n-}^{\mu} - f_{n-}^{\nu}$, which is then $P^{\mu, \gamma}$ a.s. a uniformly bounded and equicontinuous family. We can then apply the Arzelà–Ascoli theorem [38] to show over a compact set every subsequence of $\{f_{n-}\}_{n=0}^{\infty}$ admits a further subsequence that converges uniformly to 0. Due to the assumed σ -compactness, the analysis carries over to the entire state space pointwise. Since the measures merge weakly, by [9, Theorem 8.6.2] total variation merging $P^{\mu, \gamma}$ a.s. follows. \square

4.2. Filter stability results. We will present the proof of each of the results in Theorem 3.6 separately.

Proof of Theorem 3.6 (i). The proof is the same as [33, Theorem 2.9] with slight notational changes. \square

We now present some supporting results on the Radon–Nikodym derivatives. These results are proven in [33] for the control-free case (building in part on [45]), but

carry over with some additional work to the controlled environment. A sketch of the proof of the following lemma is given in the appendix.

LEMMA 4.8 (see [33, Lemma 4.8]). *Assume $\mu \ll \nu$, and fix any control policy γ . For any two sigma fields $\mathcal{G}_1, \mathcal{G}_2 \subset \mathcal{F}_{0,\infty}^X \vee \mathcal{F}_{0,\infty}^Y$ we have*

$$\begin{aligned} & \|P^{\mu,\gamma}|_{\mathcal{G}_1} \mathcal{G}_2 - P^{\nu,\gamma}|_{\mathcal{G}_1} \mathcal{G}_2\|_{TV} \\ &= \frac{E^{\nu,\gamma} \left[\left| E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) | \mathcal{G}_1 \vee \mathcal{G}_2 \right] - E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) | \mathcal{G}_2 \right] \right| \middle| \mathcal{G}_2 \right]}{E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) | \mathcal{G}_2 \right]} \quad P^{\mu,\gamma} \text{ a.s.} \end{aligned}$$

THEOREM 4.9. *Assume $\mu \ll \nu$ and fix any control policy γ . The predictor is stable in the sense of total variation in expectation with respect to the policy γ if and only if*

$$(4.6) \quad E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) \middle| \bigcap_{n \geq 1} \mathcal{F}_{0,\infty}^Y \vee \mathcal{F}_{n,\infty}^X \right] = E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) \middle| \mathcal{F}_{0,\infty}^Y \right] \quad P^{\nu,\gamma} \text{ a.s.}$$

Proof. Begin with the form laid out in Lemma 4.8. We recognize that for the predictor, $\mathcal{G}_1 = \sigma(X_n)$ and $\mathcal{G}_2 = \sigma(Y_{[0,n-1]})$, we have

$$\|\pi_{n-}^{\mu,\gamma} - \pi_{n-}^{\nu,\gamma}\|_{TV} = \frac{E^{\nu,\gamma} \left[\left| E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) | Y_{[0,n-1]}, X_n \right] - E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) | Y_{[0,n-1]} \right] \right| \middle| Y_{[0,n-1]} \right]}{E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) \middle| Y_{[0,n-1]} \right]}$$

Applying Lemma 8.1 leads to

$$\begin{aligned} E^{\mu,\gamma} [\|\pi_{n-}^{\mu,\gamma} - \pi_{n-}^{\nu,\gamma}\|_{TV}] &= E^{\nu,\gamma} \left[\frac{dP^{\mu,\gamma}|_{\mathcal{F}_{0,n-1}^Y}}{dP^{\nu,\gamma}|_{\mathcal{F}_{0,n-1}^Y}} \|\pi_{n-}^{\mu,\gamma} - \pi_{n-}^{\nu,\gamma}\|_{TV} \right] \\ &= E^{\nu,\gamma} \left[E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) \middle| Y_{[0,n-1]} \right] \|\pi_{n-}^{\mu,\gamma} - \pi_{n-}^{\nu,\gamma}\|_{TV} \right] \\ &= E^{\nu,\gamma} \left[E^{\nu,\gamma} \left[\left| E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) | Y_{[0,n-1]}, X_n \right] - E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) | Y_{[0,n-1]} \right] \right| \middle| Y_{[0,n-1]} \right] \right] \\ &= E^{\nu,\gamma} \left[\left| E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) | Y_{[0,n-1]}, X_n \right] - E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) | Y_{[0,n-1]} \right] \right| \right]. \end{aligned}$$

We now want to show that

$$(4.7) \quad E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) | Y_{[0,n-1]}, X_n \right] = E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) | Y_{[0,\infty]}, X_{[n,\infty)} \right].$$

Consider the diagram in Figure 4.1 which outlines the dependency structure in a POMDP. Say we are conditioning on X_2, Y_0, Y_1 . Since the control policy is known to us, we also know the realizations of U_0 and U_1 . If we focus on these nodes on the diagram, we see that they cordon off X_0 from $X_{[3,\infty)}, Y_{[2,\infty)}$. That, if we follow back any line from a node $X_n, n > 2$, or $Y_n, n > 1$, to X_0 , we must go through one of these nodes, and hence these future nodes do not add anything useful to the conditioning. Therefore, (4.7) holds and we have

$$(4.8) \quad \begin{aligned} & E^{\mu,\gamma} [\|\pi_{n-}^{\mu,\gamma} - \pi_{n-}^{\nu,\gamma}\|_{TV}] \\ &= E^{\nu,\gamma} \left[\left| E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) | Y_{[0,\infty)}, X_{[n,\infty)} \right] - E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) | Y_{[0,n-1]} \right] \right| \right]. \end{aligned}$$




FIG. 4.1. Diagram of dependence in a POMDP. When X_0 is conditioned on $X_2, Y_{[0,1]}$, then $X_{[3,\infty)}, Y_{[2,\infty)}$ do not add new information to the conditioning.

We then see that $A_n = E^\nu[\frac{d\mu}{d\nu}(X_0)|Y_{[0,n-1]}]$ is a nonnegative uniformly integrable martingale (with respect to the measure P^ν) adapted to the increasing filtration $\mathcal{F}_{0,n-1}^{\mathcal{Y}}$. Hence, the limit as $n \rightarrow \infty$ in $L^1(P^\nu)$ is $E^\nu[\frac{d\mu}{d\nu}(X_0)|\mathcal{F}_{0,\infty}^{\mathcal{Y}}]$. Similarly, we can view $B_n = E^\nu[\frac{d\mu}{d\nu}(X_0)|Y_{[0,\infty)}, X_{[n,\infty)}]$ as a backwards nonnegative uniformly integrable martingale (with respect to the measure P^ν) adapted to the decreasing sequence of filtrations $\mathcal{F}_{0,\infty}^{\mathcal{Y}} \vee \mathcal{F}_{n,\infty}^{\mathcal{X}}$. Then by the backwards martingale convergence theorem, the limit as $n \rightarrow \infty$ in $L^1(P^\nu)$ is $E^\nu[\frac{d\mu}{d\nu}(X_0)|\bigcap_{n=1}^\infty \mathcal{F}_{0,\infty}^{\mathcal{Y}} \vee \mathcal{F}_{n,\infty}^{\mathcal{X}}]$. \square

THEOREM 4.10. Assume $\mu \ll \nu$ and fix any control policy γ . The filter is stable in the sense of total variation in expectation with respect to the policy γ if and only if

$$(4.9) \quad E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) \middle| \bigcap_{n \geq 0} \mathcal{F}_{0,\infty}^{\mathcal{Y}} \vee \mathcal{F}_{n,\infty}^{\mathcal{X}} \right] = E^{\nu,\gamma} \left[\frac{d\mu}{d\nu}(X_0) \middle| \mathcal{F}_{0,\infty}^{\mathcal{Y}} \right] \quad P^{\nu,\gamma} \text{ a.s.}$$

Proof of Theorem 3.6 (ii). The sigma fields $\mathcal{F}_{n,\infty}^{\mathcal{X}} \vee \mathcal{F}_{0,\infty}^{\mathcal{Y}}$ are a decreasing sequence, that is, $\mathcal{F}_{n+1,\infty}^{\mathcal{X}} \vee \mathcal{F}_{0,\infty}^{\mathcal{Y}} \subset \mathcal{F}_{n,\infty}^{\mathcal{X}} \vee \mathcal{F}_{0,\infty}^{\mathcal{Y}}$. Therefore, when we take their intersection, removing the first or largest sigma field $\mathcal{F}_{0,\infty}^{\mathcal{X}} \vee \mathcal{F}_{0,\infty}^{\mathcal{Y}}$ from the intersection of a decreasing set of sigma fields does not change the overall intersection. Fix any control policy γ . The conditions in Theorems 4.9 and 4.10 are the same for each policy γ . Therefore, if either the filter or the predictor is universally stable in total variation in expectation, then the other is as well. \square

Proof of Theorem 3.6 (iii). Via Fatou’s lemma for the limit supremum, it is clear that almost sure stability implies stability in expectation. We focus on the other direction.

Assume the filter is universally stable in total variation in expectation. Consider the measure $\rho = \frac{\mu + \nu}{2}$. We have that $\mu \ll \rho$ and $\nu \ll \rho$ and $\|\frac{d\mu}{d\rho}\|_\infty \leq 2, \|\frac{d\nu}{d\rho}\|_\infty \leq 2$. By Lemma 4.8 and the analysis in the proof of Theorem 4.9 but applied to the filter,

$$\|\pi_n^{\mu,\gamma} - \pi_n^{\rho,\gamma}\|_{TV} = \frac{E^{\rho,\gamma} \left[\left| E^{\rho,\gamma} \left[\frac{d\mu}{d\rho}(X_0) | Y_{[0,\infty)}, X_{[n,\infty)} \right] - E^{\rho,\gamma} \left[\frac{d\mu}{d\rho}(X_0) | Y_{[0,n]} \right] \right| \middle| Y_{[0,n]} \right]}{E^{\rho,\gamma} \left[\frac{d\mu}{d\rho}(X_0) \middle| Y_{[0,n]} \right]},$$

$$\|\pi_n^{\nu,\gamma} - \pi_n^{\rho,\gamma}\|_{TV} = \frac{E^{\rho,\gamma} \left[\left| E^{\rho,\gamma} \left[\frac{d\nu}{d\rho}(X_0) | Y_{[0,\infty)}, X_{[n,\infty)} \right] - E^{\rho,\gamma} \left[\frac{d\nu}{d\rho}(X_0) | Y_{[0,n]} \right] \right| \middle| Y_{[0,n]} \right]}{E^{\rho,\gamma} \left[\frac{d\nu}{d\rho}(X_0) \middle| Y_{[0,n]} \right]}.$$

Since the Radon–Nikodym derivatives are finite, the expressions in the denominators and numerators above are uniformly integrable martingales with respect to the

measure $P^{\rho,\gamma}$ and hence admit limits $P^{\rho,\gamma}$ a.s. This implies that they admit limits $P^{\mu,\gamma}$ a.s. Therefore, $\|\pi_n^{\mu,\gamma} - \pi_n^{\rho,\gamma}\|_{TV}$ and $\|\pi_n^{\nu,\gamma} - \pi_n^{\rho,\gamma}\|_{TV}$ admit limits $P^{\mu,\gamma}$ a.s. By the dominated convergence theorem,

$$E^\mu \left[\lim_{n \rightarrow \infty} \|\pi_n^{\mu,\gamma} - \pi_n^{\rho,\gamma}\|_{TV} \right] = \lim_{n \rightarrow \infty} E^\mu [\|\pi_n^{\mu,\gamma} - \pi_n^{\rho,\gamma}\|_{TV}] = 0$$

since we assume the filter is universally stable in expectation by assumption. Total variation is always nonnegative. Therefore, $\lim_{n \rightarrow \infty} \|\pi_n^{\mu,\gamma} - \pi_n^{\rho,\gamma}\|_{TV}$ must be zero a.s. Similarly,

$$\begin{aligned} E^\mu \left[\lim_{n \rightarrow \infty} \|\pi_n^{\nu,\gamma} - \pi_n^{\rho,\gamma}\|_{TV} \right] &= \lim_{n \rightarrow \infty} E^\mu [\|\pi_n^{\nu,\gamma} - \pi_n^{\rho,\gamma}\|_{TV}] \\ &\leq \lim_{n \rightarrow \infty} E^\mu [\|\pi_n^{\nu,\gamma} - \pi_n^{\mu,\gamma}\|_{TV}] + \lim_{n \rightarrow \infty} E^\mu [\|\pi_n^{\mu,\gamma} - \pi_n^{\rho,\gamma}\|_{TV}] = 0 \end{aligned}$$

so $\lim_{n \rightarrow \infty} \|\pi_n^{\nu,\gamma} - \pi_n^{\rho,\gamma}\|_{TV}$ must be 0 a.s. Then we have

$$\begin{aligned} \limsup_{n \rightarrow \infty} \|\pi_n^{\mu,\gamma} - \pi_n^{\nu,\gamma}\|_{TV} &\leq \limsup_{n \rightarrow \infty} \left(\|\pi_n^{\mu,\gamma} - \pi_n^{\rho,\gamma}\|_{TV} + \|\pi_n^{\rho,\gamma} - \pi_n^{\nu,\gamma}\|_{TV} \right) \\ &\leq \limsup_{n \rightarrow \infty} \|\pi_n^{\mu,\gamma} - \pi_n^{\rho,\gamma}\|_{TV} + \limsup_{n \rightarrow \infty} \|\pi_n^{\rho,\gamma} - \pi_n^{\nu,\gamma}\|_{TV} = 0 \end{aligned}$$

since the limit supremum of the sum of two nonnegative sequences is less than the sum of the limit supremum of the sequences. Therefore, the limit supremum of $\|\pi_n^\mu - \pi_n^\nu\|_{TV} = 0$ $P^{\mu,\gamma}$ a.s. Thus the limit exists and is 0. \square

Proof of Theorem 3.6 (iv). First, assume the filter is universally stable in relative entropy and fix some policy γ . Since the square root function is continuous and convex, we have

$$0 = \lim_{n \rightarrow \infty} \sqrt{\frac{2}{\log(e)} E^{\mu,\gamma} [D(\pi_n^{\mu,\gamma} \| \pi_n^{\nu,\gamma})]} \geq \lim_{n \rightarrow \infty} E^{\mu,\gamma} \left[\sqrt{\frac{2}{\log(e)} D(\pi_n^{\mu,\gamma} \| \pi_n^{\nu,\gamma})} \right],$$

where we have applied Jensen's inequality. We then apply Pinsker's inequality and have $\lim_{n \rightarrow \infty} E^{\mu,\gamma} [\|\pi_n^{\mu,\gamma} - \pi_n^{\nu,\gamma}\|_{TV}] = 0$.

For the converse direction, assume the filter is universally stable in total variation in expectation and fix a policy γ . If we apply the chain rule for relative entropy, we have

$$\begin{aligned} E^{\mu,\gamma} [D(\pi_n^{\mu,\gamma} \| \pi_n^{\nu,\gamma})] &= D(P^{\mu,\gamma}(X_n | Y_{[0,n]}) \| P^{\nu,\gamma}(X_n | Y_{[0,n]})) \\ &= D(P^{\mu,\gamma}(X_n, Y_{[0,n]}) \| P^{\nu,\gamma}(X_n, Y_{[0,n]})) \\ &\quad - D(P^{\mu,\gamma}(Y_{[0,n]}) \| P^{\nu,\gamma}(Y_{[0,n]})). \end{aligned}$$

We also have by the chain rule that

$$\begin{aligned} D(P^{\mu,\gamma}(X_n, Y_{[0,n]}) \| P^{\nu,\gamma}(X_n, Y_{[0,n]})) &= D(P^{\mu,\gamma}(X_{[n,\infty)}, Y_{[0,\infty]}) \| P^{\nu,\gamma}(X_{[n,\infty)}, Y_{[0,\infty]})) \\ &\quad - D(P^{\mu,\gamma}((X, Y)_{[n+1,\infty)} | X_n, Y_{[0,n]}) \| P^{\nu,\gamma}((X, Y)_{[n+1,\infty)} | X_n, Y_{[0,n]})). \end{aligned}$$

Consider now Figure 4.1. If we fix $X_2, Y_{[0,2]}$, then we also know the control actions $U_{[0,2]}$. Then the distribution of X_3 is fully determined by X_2, U_2 and the transition kernel; it is independent of the prior measure. Then the distribution of Y_3 is determined by X_3 and Q , and U_3 and X_4 and so on. Therefore, the conditional measures $P^{\mu,\gamma}((X, Y)_{[n+1,\infty)} | X_n, Y_{[0,n]})$ and $P^{\nu,\gamma}((X, Y)_{[n+1,\infty)} | X_n, Y_{[0,n]})$ are the same even though the priors are different, and therefore the relative entropy is zero.

We then have

$$\begin{aligned} E^{\mu,\gamma}[D(\pi_n^{\mu,\gamma} \parallel \pi_n^{\nu,\gamma})] &= D(P^{\mu,\gamma}(X_{[n,\infty)}, Y_{[0,\infty)}) \parallel P^{\nu,\gamma}(X_{[n,\infty)}, Y_{[0,\infty)}) \\ &\quad - D(P^{\mu,\gamma}(Y_{[0,n]}) \parallel P^{\nu,\gamma}(Y_{[0,n]})) \\ &= D(P^{\mu,\gamma}|_{\mathcal{F}_{n,\infty}^X \vee \mathcal{F}_{0,\infty}^Y} \parallel P^{\nu,\gamma}|_{\mathcal{F}_{n,\infty}^X \vee \mathcal{F}_{0,\infty}^Y}) \\ &\quad - D(P^{\mu,\gamma}|_{\mathcal{F}_{[0,n]}^Y} \parallel P^{\nu,\gamma}|_{\mathcal{F}_{[0,n]}^Y}). \end{aligned}$$

For each of these, building on [4, Theorem 3], we have that the convergence to zero follows the arguments presented in [33, Theorem 2.12]. \square

5. Proofs on robustness results.

Proof of Theorem 3.9. Fixing some finite n , we have

$$\begin{aligned} J_\infty(\mu, \gamma^\nu) &= \limsup_{T \rightarrow \infty} \frac{1}{T} \left(\sum_{i=0}^{n-1} E^{\mu,\gamma^\nu}[c(X_i, U_i)] + \sum_{i=n}^{T-1} E^{\mu,\gamma^\nu}[c(X_i, U_i)] \right) \\ &\leq \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{i=0}^{n-1} E^{\mu,\gamma^\nu}[c(X_i, U_i)] + \limsup_{T \rightarrow \infty} \frac{1}{T} E^{\mu,\gamma^\nu} \left[\sum_{i=n}^{T-1} c(X_i, U_i) \right] \\ &\leq \limsup_{T \rightarrow \infty} \frac{n \|c\|_\infty}{T} + \limsup_{T \rightarrow \infty} \frac{1}{T} E^{\mu,\gamma^\nu} \left[\sum_{i=0}^{T-n-1} c(X_{n+i}, U_{n+i}) \right] \\ &= \limsup_{T \rightarrow \infty} \frac{1}{T} E^{\mu,\gamma^\nu} \left[\sum_{i=0}^{T-n-1} c(X_{n+i}, U_{n+i}) \right]. \end{aligned}$$

Therefore, we see that no matter what decision the DM makes in the first n time stage, since n is finite and c bounded this cost will eventually be dominated by the denominator as $T \rightarrow \infty$ and there will be no transient cost associated with this robustness problem. We then claim that

$$\limsup_{T \rightarrow \infty} \frac{1}{T} E^{\mu,\gamma^\nu} \left[\sum_{i=0}^{T-n-1} c(X_{n+i}, U_{n+i}) \right] = \limsup_{T \rightarrow \infty} \frac{1}{T-n} E^{\mu,\gamma^\nu} \left[\sum_{i=0}^{T-n-1} c(X_{n+i}, U_{n+i}) \right].$$

All terms in the two limsup expressions are positive and bounded since c is a nonnegative bounded function. Therefore, we have in the following that (i) the difference of the limsups is less than or equal to the limsup of the difference, and (ii) the limsup of a product is less than or equal to the product of the limsups. Using these results,

$$\begin{aligned} &\limsup_{T \rightarrow \infty} \frac{1}{T-n} E^{\mu,\gamma^\nu} \left[\sum_{i=0}^{T-n-1} c(X_{n+i}, U_{n+i}) \right] - \limsup_{T \rightarrow \infty} \frac{1}{T} E^{\mu,\gamma^\nu} \left[\sum_{i=0}^{T-n-1} c(X_{n+i}, U_{n+i}) \right] \\ &\leq \limsup_{T \rightarrow \infty} \left(\frac{1}{T-n} - \frac{1}{T} \right) E^{\mu,\gamma^\nu} \left[\sum_{i=0}^{T-n-1} c(X_{n+i}, U_{n+i}) \right] \\ &\leq \left(\limsup_{T \rightarrow \infty} \frac{n}{T} \right) \left(\limsup_{T \rightarrow \infty} \frac{1}{T-n} E^{\mu,\gamma^\nu} \left[\sum_{i=0}^{T-n-1} c(X_{n+i}, U_{n+i}) \right] \right) = 0. \end{aligned}$$

We then apply iterated expectations and Fatou’s lemma, and we have

$$\begin{aligned} & \limsup_{T \rightarrow \infty} \frac{1}{T-n} E^{\mu, \gamma^\nu} \left[\sum_{i=0}^{T-n-1} c(X_{n+i}, U_{n+i}) \right] \\ &= \limsup_{T \rightarrow \infty} \frac{1}{T-n} E^{\mu, \gamma^\nu} \left[E^{\mu, \gamma^\nu} \left[\sum_{i=0}^{T-n-1} c(X_{n+i}, U_{n+i}) | Y_{[0, n-1]} \right] \right] \\ &\leq E^{\mu, \gamma^\nu} \left[\limsup_{T \rightarrow \infty} \frac{1}{T-n} E^{\mu, \gamma^\nu} \left[\sum_{i=0}^{T-n-1} c(X_{n+i}, U_{n+i}) | Y_{[0, n-1]} \right] \right]. \end{aligned}$$

The optimal control policy is a time invariant function of the filter realization so the predictor at time n acts as a new prior for the control problem. We then have

$$\begin{aligned} & E^{\mu, \gamma^\nu} \left[\limsup_{T \rightarrow \infty} \frac{1}{T-n} E^{\mu, \gamma^\nu} \left[\sum_{i=0}^{T-n-1} c(X_{n+i}, U_{n+i}) | Y_{[0, n-1]} \right] \right] \\ &= E^{\mu, \gamma^\nu} \left[J_\infty(\pi_{n-}^\mu, \gamma^{\pi_{n-}^\nu}) \right]. \end{aligned}$$

With this established, we can now move on to our robustness problem:

$$\begin{aligned} (5.1) \quad & J_\infty(\mu, \gamma^\nu) - J_\infty^*(\mu) \leq E^{\mu, \gamma^\nu} [J_\infty(\pi_{n-}^\mu, \gamma^{\pi_{n-}^\nu})] - \inf_{\tilde{\mu} \in \mathcal{P}(\mathcal{X})} J_\infty^*(\tilde{\mu}) \\ &= E^{\mu, \gamma^\nu} [J_\infty(\pi_{n-}^\mu, \gamma^{\pi_{n-}^\nu}) + J_\infty^*(\pi_{n-}^\mu) - J_\infty^*(\pi_{n-}^\mu)] - \inf_{\tilde{\mu} \in \mathcal{P}(\mathcal{X})} J_\infty^*(\tilde{\mu}) \\ &= E^{\mu, \gamma^\nu} [J_\infty(\pi_{n-}^\mu, \gamma^{\pi_{n-}^\nu}) - J_\infty^*(\pi_{n-}^\mu)] + E^{\mu, \gamma^\nu} [J_\infty^*(\pi_{n-}^\mu)] - \inf_{\tilde{\mu} \in \mathcal{P}(\mathcal{X})} J_\infty^*(\tilde{\mu}) \\ &\leq E^{\mu, \gamma^\nu} [J_\infty(\pi_{n-}^\mu, \gamma^{\pi_{n-}^\nu}) - J_\infty^*(\pi_{n-}^\mu)] + \sup_{\tilde{\mu} \in \mathcal{P}(\mathcal{X})} J_\infty^*(\tilde{\mu}) - \inf_{\tilde{\mu} \in \mathcal{P}(\mathcal{X})} J_\infty^*(\tilde{\mu}) \\ (5.2) \quad &= E^{\mu, \gamma^\nu} [J_\infty(\pi_{n-}^\mu, \gamma^{\pi_{n-}^\nu}) - J_\infty^*(\pi_{n-}^\mu)] + \|J_\infty^*\|_{sp}. \end{aligned}$$

By Theorem 3.8 we have

$$(5.3) \quad J_\infty(\mu, \gamma^\nu) - J_\infty^*(\mu) \leq 2\|c\|_\infty E^{\mu, \gamma^\nu} [\|\pi_{n-}^\mu - \pi_{n-}^\nu\|_{TV}] + \|J_\infty^*\|_{sp},$$

and this result holds for any n since our choice of n was arbitrary. By assumption, the predictor is universally stable in total variation in expectation. Therefore, for any $\epsilon > 0$ there exists an N such that for all $n > N$ we have $E^{\mu, \gamma^\nu} [\|\pi_{n-}^\mu - \pi_{n-}^\nu\|_{TV}] \leq \frac{\epsilon}{2\|c\|_\infty}$. Therefore, since (5.3) holds for every n , $J_\infty(\mu, \gamma^\nu) - J_\infty^*(\mu) \leq \|J_\infty^*\|_{sp} + \epsilon$ for any $\epsilon > 0$, yielding our result. \square

Proof of Theorem 3.10. Pick any $n \in \mathbb{N}$. Starting from expressions (1.8), (1.9), and (1.10) we will consider the three costs. The transient cost is upper bound by

$$E^{\mu, \gamma^\nu} \left[\sum_{i=0}^{n-1} \beta^i c(x_i, u_i) \right] - E^{\mu, \gamma^\mu} \left[\sum_{i=0}^{n-1} \beta^i c(x_i, u_i) \right] \leq \|c\|_\infty \sum_{i=0}^{n-1} \beta^i = \|c\|_\infty \left(\frac{1 - \beta^n}{1 - \beta} \right),$$

the strategic measure cost is upper bound by

$$\beta^n \left(E^{\mu, \gamma^\nu} \left[J_\beta^*(\pi_{n-}^\mu) \right] - E^{\mu, \gamma^\mu} \left[J_\beta^*(\pi_{n-}^\mu) \right] \right) \leq \beta^n \|J_\beta^*\|_{sp},$$

and the approximation cost satisfies

$$\beta^n \left(E^{\mu, \gamma^\nu} \left[J_\beta(\pi_{n-}^{\mu, \gamma^\nu}, \gamma^{\pi_{n-}^{\nu, \gamma^\nu}}) - J_\beta^*(\pi_{n-}^{\mu, \gamma^\nu}) \right] \right) \leq 4 \frac{\|c\|_\infty}{1-\beta} (\alpha\beta)^n.$$

Putting these together,

$$\begin{aligned} J_\beta(\mu, \gamma^\nu) - J_\beta^*(\mu) &\leq \|c\|_\infty \left(\frac{1-\beta^n}{1-\beta} \right) + \beta^n \|J_\beta^*\|_{sp} + 4 \frac{\|c\|_\infty}{1-\beta} (\alpha\beta)^n \\ &= \frac{\|c\|_\infty}{1-\beta} + \beta^n \left(\|J_\beta^*\|_{sp} + 4 \frac{\|c\|_\infty}{1-\beta} \alpha^n - \frac{\|c\|_\infty}{1-\beta} \right) \\ &= \frac{\|c\|_\infty}{1-\beta} (1 + \beta^n (4\alpha^n - \rho)) = \frac{\|c\|_\infty}{1-\beta} (1 - f(n)). \end{aligned}$$

This holds for any n . Taking the derivative in n , at n^* , it follows that this will be maximized at some $n^* \in \mathbb{R}_+$,

$$\alpha^{n^*} = \left(\frac{\rho}{4} \right) \left(\frac{\ln(\beta)}{\ln(\alpha) + \ln(\beta)} \right), \quad n^* = \frac{\ln \left(\left(\frac{\rho}{4} \right) \left(\frac{\ln(\beta)}{\ln(\alpha) + \ln(\beta)} \right) \right)}{\ln \alpha},$$

and it can be shown that the maximum among the natural numbers $n \in \mathbb{N}$ will occur at the ceiling or floor of n^* . \square

6. Generalizations and discussion. In this section, we present a number of generalizations and discussions which will be summarized briefly with the purpose of making the paper more concise and accessible.

6.1. N -step observability and its limitations due to policy dependence.

Given the definition of one step observability in Definition 3.3, one could conceive of a definition of multiple step observability: for a policy γ , the POMDP could be called N step observable if for every $f \in C_b(\mathcal{X})$ and every $\epsilon > 0$ there exists a measurable and bounded function g of N measurements such that

$$\left\| f(\cdot) - \int_{\mathcal{Y}^N} g(y_{[1,N]}) P^{\mu, \gamma}(dy_{[1,N]} | X_1 = \cdot) \right\|_\infty < \epsilon,$$

adopting (2.2) (note that N can also depend on f, ϵ if we replace N -step observability with observability), which was shown to be an appropriate and consequential observability definition for control-free systems in [33].

However, this approach cannot be used to prove stability due to the past dependency of the control policy. In the following, we explain why the proof method we present is not applicable for such systems unless one restricts the control policies considered. To apply the general process Lemma 4.5 we need the following:

1. $Y_{[n, n+N-1]}$ can only depend on $Y_{[0, n-1]}$ through X_n :

$$(6.1) \quad P^{\mu, \gamma}(Y_{[n, n+N-1]} \in \cdot | X_n, Y_{[0, n-1]}) = P^{\mu, \gamma}(Y_{[n, n+N-1]} \in \cdot | X_n).$$

2. The channel $Y_{[n, n+N-1]} | X_n$ is measure equivalent for $P^{\mu, \gamma}$ and $P^{\nu, \gamma}$:

$$(6.2) \quad P^{\mu, \gamma}(Y_{[n, n+N-1]} \in \cdot | X_n) = P^{\nu, \gamma}(Y_{[n, n+N-1]} \in \cdot | X_n).$$

3. The channel $Y_{[n, n+N-1]} | X_n$ is time homogeneous:

$$(6.3) \quad P^{\mu, \gamma}(Y_{[n, n+N-1]} \in \cdot | X_n) = P^{\mu, \gamma}(Y_{[0, N-1]} \in \cdot | X_0) \quad \forall n \in \mathbb{N}.$$

If we take the left-hand side of (6.1), we have

$$\begin{aligned} & P^{\mu,\gamma}(Y_{[n,n+N-1]} \in \cdot | X_n, Y_{[0,n-1]}) \\ &= \int_{\mathcal{U}^{N-1}} P^{\mu,\gamma}(Y_{[n,n+N-1]} \in \cdot | X_n, Y_{[0,n-1]}, U_{[n,n+N-2]}) P^{\mu,\gamma}(du_{[n,n+N-2]} | X_n, Y_{[0,n-1]}) \end{aligned}$$

by chain rule of conditional probability. Now it is true that $Y_{[n,n+N-1]} | X_n, U_{[n,n+N-2]}$ is independent of $Y_{[0,n-1]}$ so we can stop conditioning on the past measurements in the inner argument. However, in the outer conditional measure the $U_{[n,n+N-2]}$ may still depend on the past and we have

$$(6.4) \quad \int_{\mathcal{U}^{N-1}} P^{\mu,\gamma}(Y_{[n,n+N-1]} \in \cdot | X_n, U_{[n,n+N-2]}) P^{\mu,\gamma}(du_{[n,n+N-2]} | X_n, Y_{[0,n-1]}).$$

If we take the right-hand side of (6.1), we have

$$\begin{aligned} P^{\mu,\gamma}(Y_{[n,n+N-1]} \in \cdot | X_n) &= \int_{\mathcal{U}^{N-1}} P^{\mu,\gamma}(Y_{[n,n+N-1]} \\ &\quad \in \cdot | X_n, U_{[n,n+N-2]}) P^{\mu,\gamma}(du_{[n,n+N-2]} | X_n). \end{aligned}$$

These two equations are not equal for a general control policy. Therefore, the process fails Definition 4.1 and $Y_{[n,n+N-1]}$ does not depend on $Y_{[0,n-1]}$ through X_n . As a result, a definition of $N > 1$ step observability is incompatible with the proof technique outlined in Lemma 4.5 and we cannot utilize $N > 1$ step observability in a controlled environment to prove stability unless control policies are restricted, e.g., to open-loop control policies where past dependence is avoided, such as sampled control policies where the control is open-loop in between sampling periods, effectively making the measurement a multistep one. This discussion reveals a further, but not surprising [3], layer of complexity for the theory of nonlinear controlled stochastic systems.

6.2. Localized definition of observability. As in the control-free case considered in [33], our observability definition can be generalized further to make its applicability for noncompact domains more general. While the definition of observability that we introduced is valid for both compact and noncompact state spaces, it is, in general, difficult to satisfy the definition in a noncompact state space; the concern is that when we have a noncompact space and consider functions f which are not bounded Lipschitz, it may, in general, be difficult to find a bounded function g which will approximate f through (3.1) sufficiently well over the entire state space. See [33, section 3] for some examples of observable and locally observable channels.

DEFINITION 6.1. *A POMDP is called locally predictable (universal in control policies) if there exists a sequence of $\mathcal{F}_{[0,n-1]}^{\mathcal{Y}}$ measurable mappings $a_n : \mathcal{Y}^n \rightarrow \mathcal{X}$ such that, for any $\gamma \in \Gamma$, the family of measures*

$$\tilde{\pi}_n^{\nu,\gamma}(\cdot) := \pi_n^{\nu,\gamma}(\cdot + a_n)$$

for every $\mu \ll \nu$ is a uniformly tight family of measures.

DEFINITION 6.2. *A POMDP is called locally observable (universal in control policies) if for every continuous and bounded function f , every compact set K , every sequence of numbers a_n , and every $\epsilon > 0$, there exists a sequence of uniformly bounded*

measurable functions g_n such that

$$\begin{aligned} \sup_{x \in K+a_n} \left| f(x) - \int_{\mathcal{Y}} g_n(y) Q(dy|x) \right| &\leq \epsilon \quad \forall n \in \mathbb{N}, \\ \sup_{x \notin K+a_n} \left| \int_{\mathcal{Y}} g_n(y) Q(dy|x) \right| &\leq 2\|f\|_{\infty} \quad \forall n \in \mathbb{N}. \end{aligned}$$

Then, we can state the following generalization of Theorem 3.4 (whose proof follows similarly).

THEOREM 6.3. *Assume $\mu \ll \nu$ and that the POMDP is locally predictable and locally observable (universal in control policies). Then the predictor is universally stable weakly a.s.*

6.3. Robustness under weak merging of priors. Our robustness result is also applicable under weak convergence of priors, provided that the channel is continuous under total variation, and the filter is stable subject to a modest rate of convergence condition. These would build on [26, Theorem 3.3] and the analysis presented in section 3.2.

7. Conclusion. Filter stability has been studied extensively in control-free contexts. In this paper, we studied the filter stability problem, developed new methods and results for controlled stochastic dynamical systems, and studied the implications of filter stability on robustness of optimal solutions for partially observed stochastic control problems.

8. Appendix.

8.1. Proof of Theorem 3.8. Consider the robustness difference

$$(8.1) \quad \begin{aligned} J_{\infty}(\mu, \gamma^{\nu}) - J_{\infty}^*(\mu) &= J_{\infty}^*(\nu) - J_{\infty}^*(\mu) + J_{\infty}(\mu, \gamma^{\nu}) - J_{\infty}^*(\nu) \\ &\leq |J_{\infty}^*(\nu) - J_{\infty}^*(\mu)| + J_{\infty}(\mu, \gamma^{\nu}) - J_{\infty}(\nu, \gamma^{\nu}). \end{aligned}$$

Consider now the difference $|J_{\infty}^*(\mu) - J_{\infty}^*(\nu)|$. If $J_{\infty}^*(\mu) < J_{\infty}^*(\nu)$, then the absolute value of their difference is the larger value $J_{\infty}^*(\nu)$ minus the smaller value $J_{\infty}^*(\mu)$. Then since $J_{\infty}^*(\nu) = J_{\infty}(\nu, \gamma^{\nu})$, if we replace γ^{ν} with γ^{μ} , we have $J_{\infty}(\nu, \gamma^{\mu}) > J_{\infty}^*(\nu)$ and the difference is even greater. Therefore, when $J_{\infty}^*(\mu) < J_{\infty}^*(\nu)$ we have

$$\begin{aligned} |J_{\infty}^*(\mu) - J_{\infty}^*(\nu)| &= J_{\infty}^*(\nu) - J_{\infty}^*(\mu) \\ &= J_{\infty}(\nu, \gamma^{\nu}) - J_{\infty}(\mu, \gamma^{\mu}) \\ &\leq J_{\infty}(\nu, \gamma^{\mu}) - J_{\infty}(\mu, \gamma^{\mu}). \end{aligned}$$

By a symmetric argument for the other direction in the inequality, we arrive at

$$|J_{\infty}^*(\mu) - J_{\infty}^*(\nu)| \leq \max(J_{\infty}(\mu, \gamma^{\nu}) - J_{\infty}(\nu, \gamma^{\nu}), J_{\infty}(\nu, \gamma^{\mu}) - J_{\infty}(\mu, \gamma^{\mu})),$$

and we can ultimately determine the robustness bound by studying the expected cost operator $J_{\infty}(\cdot, \gamma^{\nu})$ under different priors but the same control policy. By the same control policy, we mean γ^{ν} is the optimal control policy designed with respect to the prior ν . This means the DM sees observations $y_{[0,n]}$ and computes the filter believing the prior is ν (the mapping from measurements to control actions is the same for each POMDP with different priors).

Note that for two nonnegative bounded sequences $0 < a_n < m < \infty$ and $0 < b_n < k < \infty$, we have that the difference of their limsups is less than the limsup of the

difference $\limsup_{n \rightarrow \infty} a_n - \limsup_{n \rightarrow \infty} b_n \leq \limsup_{n \rightarrow \infty} (a_n - b_n)$. Therefore, since c is nonnegative and bounded, we have

$$\begin{aligned} & J_\infty(\mu, \gamma^\nu) - J_\infty(\nu, \gamma^\nu) \\ &= \limsup_{T \rightarrow \infty} \frac{1}{T} \left(\sum_{i=0}^{T-1} E^{\mu, \gamma^\nu} [c(X_i, U_i)] \right) - \limsup_{T \rightarrow \infty} \frac{1}{T} \left(\sum_{i=0}^{T-1} E^{\nu, \gamma^\nu} [c(X_i, U_i)] \right) \\ &\leq \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{i=0}^{T-1} \left(E^{\mu, \gamma^\nu} [c(X_i, U_i)] - E^{\nu, \gamma^\nu} [c(X_i, U_i)] \right) \\ &\leq \limsup_{T \rightarrow \infty} \frac{\|c\|_\infty}{T} \left(\sum_{i=0}^{T-1} \|P^{\mu, \gamma^\nu}((X_i, U_i) \in \cdot) - P^{\nu, \gamma^\nu}((X_i, U_i) \in \cdot)\|_{TV} \right); \end{aligned}$$

then we see that

$$\begin{aligned} & \|P^{\mu, \gamma^\nu}((X_i, U_i) \in \cdot) - P^{\nu, \gamma^\nu}((X_i, U_i) \in \cdot)\|_{TV} \\ &= \sup_{\|f\|_\infty \leq 1} \left| \int_{\mathcal{X} \times \mathcal{U}} f(x, u) P^{\mu, \gamma^\nu}(dx_i, du_i) - \int_{\mathcal{X} \times \mathcal{U}} f(x, u) P^{\nu, \gamma^\nu}(dx_i, du_i) \right| \\ &= \sup_{\|f\|_\infty \leq 1} \left| \int_{\mathcal{X}} \int_{\mathcal{X} \times \mathcal{U}} f(x, u) P^{\mu, \gamma^\nu}(dx_i, du_i | X_0) \mu(dx_0) \right. \\ (8.2) \quad & \left. - \int_{\mathcal{X}} \int_{\mathcal{X} \times \mathcal{U}} f(x, u) P^{\nu, \gamma^\nu}(dx_i, du_i | X_0) \nu(dx_0) \right|. \end{aligned}$$

As was discussed, both POMDP use the same control policy, which maps measurements to control actions in the same way. Once we fix the realization of $X_0 = x$ in either case, the distribution on Y_0 is the same, hence the distribution on U_0 is the same, and hence X_1, Y_1, U_1 and so on. Therefore, in the above the two inner integrals are the same function of x and we can upper bound by $\|\mu - \nu\|_{TV}$.

Thus, $J_\infty(\mu, \gamma^\nu) - J_\infty(\nu, \gamma^\nu) \leq \|c\|_\infty \|\mu - \nu\|_{TV}$. Both terms in (8.1) have this bound and the overall bound is multiplied by a factor of 2. \square

8.2. Proof of Lemma 4.8.

LEMMA 8.1 (see [33, Lemma 4.6]). *Assume $\mu \ll \nu$ and fix any control policy γ . For any sigma field $\mathcal{G} \subseteq \mathcal{F}_{0, \infty}^X \vee \mathcal{F}_{0, \infty}^Y$ we have*

$$\frac{dP^{\mu, \gamma}|_{\mathcal{G}}}{dP^{\nu, \gamma}|_{\mathcal{G}}} = E^{\nu, \gamma} \left[\frac{d\mu}{d\nu}(X_0) \middle| \mathcal{G} \right] \quad P^{\mu, \gamma} \text{ a.s.}$$

Proof. Note that conditioned on the value of X_0 , knowledge of the prior is irrelevant so for any set $A \in \mathcal{G}$ we have

$$E^\mu[1_A] = E^\nu \left[1_A \frac{d\mu}{d\nu}(X_0) \right];$$

then we can apply law iterated expectations, condition internally on \mathcal{G} , and move the indicator out in the inner expectation since it is \mathcal{G} measurable. This shows

$$E^\mu[1_A] = E^\nu \left[1_A E^\nu \left[\frac{d\mu}{d\nu}(X_0) | \mathcal{G} \right] \right]. \quad \square$$

LEMMA 8.2 (see [33, Lemma 4.7]). Assume $\mu \ll \nu$ and fix any control policy γ . For any two sigma fields $\mathcal{G}_1, \mathcal{G}_2 \subset \mathcal{F}_{0,\infty}^X \vee \mathcal{F}_{0,\infty}^Y$, let $P^{\mu,\gamma}|_{\mathcal{G}_1}|\mathcal{G}_2$ represent the probability measure $P^{\mu,\gamma}$ restricted to \mathcal{G}_1 , conditioned on field \mathcal{G}_2 . We then have

$$\frac{dP^{\mu,\gamma}|_{\mathcal{G}_1}|\mathcal{G}_2}{dP^{\nu,\gamma}|_{\mathcal{G}_1}|\mathcal{G}_2} = \frac{E^{\nu,\gamma}\left[\frac{d\mu}{d\nu}(X_0)|\mathcal{G}_1 \vee \mathcal{G}_2\right]}{E^{\nu,\gamma}\left[\frac{d\mu}{d\nu}(X_0)|\mathcal{G}_2\right]} \quad P^{\mu,\gamma} \text{ a.s.}$$

Proof. Note for any set $A \in \mathcal{G}_1$ we can write $E^\mu[1_A]$ in two ways:

$$\begin{aligned} E^\mu[1_A] &= E^\nu\left[E^\mu\left[1_A \frac{dP^\mu|_{\mathcal{G}_2}}{dP^\nu|_{\mathcal{G}_2}}|\mathcal{G}_2\right]\right] \\ &\leq E^\mu[1_A] = E^\nu\left[1_A E^\nu\left[\frac{dP^\mu}{dP^\nu}|\mathcal{G}_1 \vee \mathcal{G}_2\right]\right]. \end{aligned}$$

Applying Lemma 8.1,

$$\begin{aligned} E^\mu[1_A] &= E^\nu\left[E^\mu\left[1_A E^\nu\left[\frac{d\mu}{d\nu}(X_0)|\mathcal{G}_2\right]|\mathcal{G}_2\right]\right], \\ E^\mu[1_A] &= E^\nu\left[E^\nu\left[1_A \frac{d\mu}{d\nu}(X_0)|\mathcal{G}_1 \vee \mathcal{G}_2\right]\right]. \end{aligned}$$

By definition, the Radon–Nikodym derivative $\frac{dP^\mu|_{\mathcal{G}_1}|\mathcal{G}_2}{dP^\nu|_{\mathcal{G}_1}|\mathcal{G}_2} = f$ is the function that satisfies

$$E^\nu\left[E^\nu\left[1_A f E^\nu\left[\frac{d\mu}{d\nu}(X_0)|\mathcal{G}_2\right]|\mathcal{G}_2\right]\right] = E^\mu[1_A] = E^\nu\left[E^\nu\left[1_A \frac{d\mu}{d\nu}(X_0)|\mathcal{G}_1 \vee \mathcal{G}_2\right]\right].$$

One can show that substituting f for the function in the theorem has this property. \square

Proof of Lemma 4.8. For two measures μ_1 and μ_2 with $\mu_1 \ll \mu_2$, note that total variation can be written as

$$\|\mu_1 - \mu_2\|_{TV} = E^{\mu_2}\left[\left|\frac{d\mu_1}{d\mu_2} - 1\right|\right].$$

Applying Lemma 8.2 and cross multiplying to get a single fraction in the absolute value leads to the result. \square

REFERENCES

- [1] B. D. O. ANDERSON AND J. B. MOORE, *Optimal Filtering*, Prentice-Hall, Englewood Cliffs, NJ, 1979.
- [2] A. ARAPOSTATHIS, V. S. BORKAR, E. FERNANDEZ-GAUCHERAND, M. K. GHOSH, AND S. I. MARCUS, *Discrete-time controlled Markov processes with average cost criterion: A survey*, SIAM J. Control Optim., 31 (1993), pp. 282–344, <https://doi.org/10.1137/0331018>.
- [3] Y. BAR-SHALOM AND E. TSE, *Dual effect, certainty equivalence, and separation in stochastic control*, IEEE Trans. Automat. Control, 19 (1974), pp. 494–500.
- [4] A. R. BARRON, *Limits of information, Markov chains, and projections*, in Proceedings of the IEEE International Symposium on Information Theory, Sorrento, Italy, 2000, 25.
- [5] T. BAŞAR AND P. BERNHARD, *H-infinity Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*, Birkhäuser, Boston, 1995.
- [6] P. BILLINGSLEY, *Probability and Measure*, 2nd ed., John Wiley & Sons, New York, 1986.
- [7] D. BLACKWELL AND L. DUBINS, *Merging of opinions with increasing information*, Ann. Math. Statist., 33 (1962), pp. 882–886.
- [8] V. I. BOGACHEV, *Measure Theory*, Vol. I, Springer-Verlag, Berlin, 2007.

- [9] V. I. BOGACHEV, *Measure Theory*, Vol. II, Springer-Verlag, Berlin, 2007.
- [10] V. S. BORKAR, *Average cost dynamic programming equations for controlled Markov chains with partial observations*, SIAM J. Control Optim., 39 (2000), pp. 673–681, <https://doi.org/10.1137/S0363012998345172>.
- [11] V. S. BORKAR, *Convex analytic methods in Markov decision processes*, in Handbook of Markov Decision Processes, Internat. Ser. Oper. Res. Management Sci. 40, Kluwer Academic Publishers, Boston, 2002, pp. 347–375.
- [12] V. S. BORKAR, *Dynamic programming for ergodic control with partial observations*, Stochastic Process. Appl., 103 (2003), pp. 293–310.
- [13] V. S. BORKAR, *Dynamic programming for ergodic control of Markov chains under partial observations: A correction*, SIAM J. Control Optim., 45 (2007), pp. 2299–2304, <https://doi.org/10.1137/05064610X>.
- [14] V. S. BORKAR AND A. BUDHIRAJA, *A further remark on dynamic programming for partially observed Markov processes*, Stochastic Process. Appl., 112 (2004), pp. 79–93.
- [15] P. E. CAINES, *Linear Stochastic Systems*, John Wiley & Sons, New York, 1988.
- [16] C. T. CHEN, *Linear Systems Theory and Design*, Oxford University Press, Oxford, 1999.
- [17] P. CHIGANSKY AND R. LIPTSER, *On a role of predictor in the filtering stability*, Electron. Comm. Probab., 11 (2006), pp. 129–140.
- [18] P. CHIGANSKY, R. LIPTSER, AND R. VAN HANDEL, *Intrinsic methods in filter stability*, in Handbook of Nonlinear Filtering, Oxford University Press, Oxford, 2009, pp. 319–351.
- [19] I. CSISZÁR, *Information-type measures of difference of probability distributions and indirect observation*, Studia Sci. Math. Hungar., 2 (1967), pp. 229–318.
- [20] R. L. DOBRUSHIN, *Central limit theorem for nonstationary Markov chains. I*, Theory Probab. Appl., 1 (1956), pp. 65–80, <https://doi.org/10.1137/1101006>.
- [21] E. A. FEINBERG, P. O. KASYANOV, AND M. Z. ZGUROVSKY, *Partially observable total-cost Markov decision process with weakly continuous transition probabilities*, Math. Oper. Res., 41 (2016), pp. 656–681.
- [22] R. HERMANN AND A. KRENER, *Nonlinear controllability and observability*, IEEE Trans. Automat. Control, 22 (1977), pp. 728–740.
- [23] O. HERNÁNDEZ-LERMA AND J. B. LASSERRE, *Discrete-Time Markov Control Processes: Basic Optimality Criteria*, Springer, New York, 1996.
- [24] R. E. KALMAN, *A new approach to linear filtering and prediction problems*, Trans. ASME Ser. D. J. Basic Engrg., 82 (1960), pp. 35–45.
- [25] A. D. KARA, N. SALDI, AND S. YÜKSEL, *Weak Feller property of non-linear filters*, Systems Control Lett., 134 (2019), 104512.
- [26] A. D. KARA AND S. YÜKSEL, *Robustness to incorrect priors in partially observed stochastic control*, SIAM J. Control Optim., 57 (2019), pp. 1929–1964, <https://doi.org/10.1137/17M1157660>.
- [27] A. D. KARA AND S. YÜKSEL, *Near optimality of finite memory feedback policies in partially observed Markov decision processes*, J. Mach. Learn. Res., 23 (2022), pp. 1–46.
- [28] J.-W. KIM, P. G. MEHTA, AND S. P. MEYN, *What is the Lagrangian for nonlinear filtering?*, in Proceedings of the IEEE 58th Conference on Decision and Control (CDC), IEEE, 2019, pp. 1607–1614.
- [29] H. J. KUSHNER, *Introduction to Stochastic Control Theory*, Holt, Rinehart and Winston, New York, 1972.
- [30] H. J. KUSHNER, *A partial history of the early development of continuous-time nonlinear stochastic systems theory*, Automatica J. IFAC, 50 (2014), pp. 303–334.
- [31] A. R. LIU AND R. R. BITMEAD, *Observability and reconstructibility of hidden Markov models: Implications for control and network congestion control*, in Proceedings of the 49th IEEE Conference on Decision and Control (CDC), IEEE, 2010, pp. 918–923.
- [32] A. R. LIU AND R. R. BITMEAD, *Stochastic observability in network state estimation and control*, Automatica J. IFAC, 47 (2011), pp. 65–78.
- [33] C. McDONALD AND S. YÜKSEL, *Converse Results on Filter Stability Criteria and Stochastic Non-Linear Observability*, preprint, <https://arxiv.org/abs/1812.01772>, 2018.
- [34] C. McDONALD AND S. YÜKSEL, *Observability and filter stability for partially observed Markov processes*, in Proceedings of the 58th Annual IEEE Conference on Decision and Control (CDC), IEEE, 2019, pp. 1623–1628.
- [35] C. McDONALD AND S. YÜKSEL, *Exponential filter stability via Dobrushin’s coefficient*, Electron. Commun. Probab., 25 (2020), 53.
- [36] H. NIJMEIJER, *Observability of autonomous discrete time non-linear systems: A geometric approach*, Internat. J. Control, 36 (1982), pp. 867–874.
- [37] D. RHENIUS, *Incomplete information in Markovian decision models*, Ann. Statist., 2 (1974),

- pp. 1327–1334.
- [38] W. RUDIN, *Real and Complex Analysis*, Tata McGraw-Hill Education, 2006.
 - [39] E. D. SONTAG, *A concept of local observability*, Systems Control Lett., 5 (1984), pp. 41–47.
 - [40] R. B. SOWERS AND A. M. MAKOWSKI, *Discrete-time filtering for linear systems with non-Gaussian initial conditions: Asymptotic behavior of the difference between the MMSE and LMSE estimates*, IEEE Trans. Automat. Control, 37 (1992), pp. 114–120.
 - [41] W. STANNAT, *Stability of the filter equation for a time-dependent signal on \mathbb{R}^d* , Appl. Math. Optim., 52 (2005), pp. 39–71.
 - [42] V. A. UGRINOVSKII, *Observability of linear stochastic uncertain systems*, IEEE Trans. Automat. Control, 48 (2003), pp. 2264–2269.
 - [43] R. VAN HANDEL, *Discrete time nonlinear filters with informative observations are stable*, Electron. Commun. Probab., 13 (2008), pp. 562–575.
 - [44] R. VAN HANDEL, *Observability and nonlinear filtering*, Probab. Theory Related Fields, 145 (2009), pp. 35–74.
 - [45] R. VAN HANDEL, *The stability of conditional Markov processes and Markov chains in random environments*, Ann. Probab., 37 (2009), pp. 1876–1925.
 - [46] R. VAN HANDEL, *Uniform observability of hidden Markov models and filter stability for unstable signals*, Ann. Appl. Probab., 19 (2009), pp. 1172–1199.
 - [47] 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 Probab. Appl., 21 (1976), pp. 153–158.
 - [48] K. ZHOU, J. C. DOYLE, AND K. GLOVER, *Robust and Optimal Control*, Prentice-Hall, 1996.