

Statistically consistent inverse optimal control for discrete-time indefinite linear-quadratic systems [★]

Han Zhang ^a, Axel Ringh ^b

^a*Department of Automation, School of Electronic Information and Electrical Engineering, Shanghai Jiao Tong University, Shanghai, China*

^b*Department of Mathematical Sciences, Chalmers University of Technology and University of Gothenburg, 41296 Gothenburg, Sweden*

Abstract

The Inverse Optimal Control (IOC) problem is a structured system identification problem that aims to identify the underlying objective function based on observed optimal trajectories. This provides a data-driven way to model experts' behavior. In this paper, we consider the case of discrete-time finite-horizon linear-quadratic problems where: the quadratic cost term in the objective is not necessarily positive semi-definite; the planning horizon is a random variable; we have both process noise and observation noise; the dynamics can have a drift term; and where we can have a linear cost term in the objective. In this setting, we first formulate the necessary and sufficient conditions for when the forward optimal control problem is solvable. Next, we show that the corresponding IOC problem is identifiable. Using the conditions for existence of an optimum of the forward problem, we then formulate an estimator for the parameters in the objective function of the forward problem as the globally optimal solution to a convex optimization problem, and prove that the estimator is statistical consistent. Finally, the performance of the algorithm is demonstrated on two numerical examples.

Key words: Inverse optimal control, Indefinite linear quadratic regulator, System identification, Time-varying system matrices, Convex optimization, Semidefinite programming, Inverse reinforcement learning

1 Introduction

Optimal control is a powerful framework in which control decisions are performed in order to minimize some given objective function; see, e.g., one of the monographs [3, 8]. In fact, many processes in nature can be modelled as optimal control problems with respect to some criteria [2]. However, in applications of optimal control, a fundamental problem is to design an appropriate objective function. In order to induce an appropriate control response, the object function needs to be adapted to the contextual environment in which the system is operating. This is a difficult task, which relies heavily on the designers' experience and imagination.

Instead of designing the cost criteria, one way to overcome this difficulty would be to identify the cost function from the observations of an expert system that behaves “optimally” in the environment and thus “imitating” the expert behaviour. The latter is known as Inverse Optimal Control (IOC) [22], and has received considerable attention. In particular, IOC reconstructs the objective function of the expert system and hence predicts the closed-loop system's behaviour using observed data as well as the knowledge of underlying system dynamics. The problem can be categorized as gray-box system identification [27, p. 13], and is also closely related to inverse reinforcement learning [30].

As one of the most classical optimal controller designs, linear-quadratic optimal regulators has been widely used in engineering. Though most of the literature considers the case of Q (the penalty parameter for the states) being positive semi-definite, indefinite linear-quadratic optimal control [10, 13, 14, 34, 35] has found applications in, e.g., mathematical finance [50], crowd evacuation [41, 42], and controller design for automatic steering of ships [37]. To this end, we are motivated to de-

[★] This work was partially supported by National Natural Science Foundation (NNSF) of China under Grant 62103276, and partially by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

Email addresses: zhanghan_tc@sjtu.edu.cn (Han Zhang), axelri@chalmers.se (Axel Ringh).

velop an IOC framework for general indefinite linear-quadratic optimal control. Linear-quadratic IOC problem has been studied under many different settings, including the infinite-horizon case in both continuous time [3, 9] and discrete time [33], respectively, as well as the finite-horizon case in both continuous time [24, 25] and discrete time [23, 46–49], respectively. IOC is also closely connected to inverse reinforcement learning [30], and this perspective has been used in [26, 44, 45] to consider infinite horizon discrete-time and continuous-time linear-quadratic set-ups for regulation, tracking, and adversary scenarios, respectively. However, to the best of our knowledge, IOC frameworks for general indefinite linear-quadratic optimal control has not been considered yet. Furthermore, there are also other important aspects that have not been fully investigated in the aforementioned literature. More precisely, any real-world data would inevitably contain noise: it can either be process noise, observation noise, or both. Therefore, from robustness and accuracy perspectives, it is important to have a statistically consistent estimator, i.e., that converges to the true underlying parameter values as the number of observation grows. Moreover, many of the aforementioned IOC algorithms that are based on optimization either suffer from the fact that the estimation problems are nonconvex [23, 46, 49], and can therefore have issues with local minima, or suffer from the fact the estimation procedure needs to know the control gain a priori [3, 9, 24, 25, 33]. In the latter case, the estimation normally needs to be done in a two-stage procedure and the information is thus not used in the most efficient way. Furthermore, most of the literature on linear-quadratic IOC consider the regulation problem. However, in many experimental set-ups, an expert agent may have more complicated tasks than regulation, e.g., tracking a reference signal. Finally, real-world data can be of different time lengths, and this needs to be handled in a systematic way in order not to deteriorate the estimates.

In this work, we address these issues. More specifically, we extend our previous work [47] and the conference version [48] of this work, and consider the generalized linear-quadratic, indefinite, discrete-time IOC problem. In particular, in [47], we only consider the IOC problem of linear quadratic regulator with no process noise, and there is no linear state terms in the objective function. Moreover, in that paper the parameter Q is positive semi-definite, and the time-horizon length is fixed. On the other hand, compared to the conference version [48], we further consider the case that involves observation noise and indefinite matrix Q in the objective function.

The contribution of this work is three-fold:

- (1) We give necessary and sufficient conditions for the well-posedness of the generalized discrete-time finite-horizon indefinite linear-quadratic optimal control problem. Specifically, we do not assume that the running cost is positive semi-definite, and

- we include a linear term in the objective function, a forcing term in the dynamics, and process noise.
- (2) We prove the identifiability of the parameters in the objective function.
- (3) We construct an IOC algorithm that works for both the positive semi-definite and the indefinite case. The algorithm is based on convex optimization, and we show that its unique optimal solution is the “true” underlying parameters in the objective function of the forward problem. Moreover, the constructed optimization problem, which contains an expectation in the objective function, is approximated by the empirical mean in practice and we also show that the estimator based on this approximation is statistically consistent. This means that the estimator converge in probability to the true underlying parameter when the number of observations goes to infinity. In addition, the convex optimization formulation guarantees that the statistically consistent estimate that corresponds to the global optimum can actually be attained in practice.

The article is organized as follows. In Section 2, we formulate the forward and inverse problem, and specify the assumptions we use. Next, in Section 3, we analyze the forward problem and prove the necessary and sufficient conditions for when it is well-posed. Section 4 investigates the identifiability issue of the formulated IOC problem, and in particular prove that the formulated IOC problem is identifiable. In Section 5, we construct the IOC algorithm, construct the estimator, and prove that the latter is statistical consistent. To illustrate the performance of the proposed algorithm, in Section 6 we include some numerical examples. Finally, we conclude the paper in Section 7. For improved readability, some proofs are deferred to the Appendix.

Notation: \mathbb{S}^n denotes the set of $n \times n$ symmetric matrices, while \mathbb{S}_+^n denotes the set of $n \times n$ positive semi-definite matrices. $\|\cdot\|$ denotes l_2 -norm and $\|\cdot\|_F$ denotes Frobenius norm. $\mathcal{B}_\varphi^n(x) := \{y \in \mathbb{R}^n \mid \|x-y\| < \varphi\}$ denotes the open ball of radius φ centered at x , and $\bar{\mathcal{B}}_\varphi^n(x)$ denotes its closure. Moreover, $\mathbb{S}_+(\varphi) := \{G \in \mathbb{S}_+^n \mid \|G\|_F \leq \varphi\}$. For $G_1, G_2 \in \mathbb{S}^n$, we denote $G_1 \succeq G_2$ as Loewner partial order of G_1 and G_2 , i.e., $G_1 - G_2 \in \mathbb{S}_+^n$. Further, G^\dagger denotes the Moore-Penrose pseudo-inverse of G . For a square matrix $G = \begin{bmatrix} G_{11} & G_{12} \\ G_{21} & G_{22} \end{bmatrix}$, where G_{11} and G_{22} are square, we define the Schur complements $G \setminus G_{22} = G_{11} - G_{21}G_{22}^\dagger G_{12}$ and $G \setminus G_{11} = G_{22} - G_{12}G_{11}^\dagger G_{21}$ (see, e.g., [19, Sec. 1.6]). For integers ν_1, ν_2 , we let $\nu_1 : \nu_2$ denote $\nu_1, \nu_1 + 1, \dots, \nu_2$, and define the expression as empty if $\nu_2 < \nu_1$. In addition, by Im and ker we denote the image space and kernel, respectively, by $^\perp$ we denote the orthogonal complement, and by \neg we denote the (mathematical) negation of a statement. Finally, we use *italic bold font* to denote stochastic elements, and use \xrightarrow{P} to denote convergence in probability, i.e., for a sequence of random elements $\{\mathbf{a}^i\}_{i=1}^\infty$ and a

random element \mathbf{a} , $\mathbf{a}^i \xrightarrow{P} \mathbf{a}$ means that for all $\varepsilon > 0$, $\lim_{i \rightarrow \infty} \mathbb{P}(\|\mathbf{a}^i - \mathbf{a}\| > \varepsilon) = 0$.

2 Problem formulation

In this section, we introduce the forward problem as well as the inverse problem. To solve the inverse problem, we use measured (noisy) optimal trajectories from an expert that performs the given task multiple times. This gives multiple demonstration trajectories that can be used to learn the cost. However, while the underlying “decision principle” is the same in all the demonstration trajectories, these trajectories may have different lengths. We are hence motivated to cover those demonstrations with one mathematical formulation.

We start by introducing the mathematical formulation of the forward optimal control problem. To this end, let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space that carries a random vector $\bar{\mathbf{x}} \in \mathbb{R}^n$, stochastic processes $\{\mathbf{w}_t \in \mathbb{R}^n\}_{t=1}^{\infty}$, $\{\mathbf{v}_t \in \mathbb{R}^n\}_{t=1}^{\infty}$ (the measurement noise to appear in (2)), and a random variable $N \in \{2, 3, \dots, \nu\} \subset \mathbb{Z}_+$. It is assumed that for each realization $(\bar{\mathbf{x}}, N)$ of the random element $(\bar{\mathbf{x}}, N)$ (corresponding to the initial position and planning horizon length), the agent’s control decision \mathbf{u}_t is determined by a stochastic generalized linear-quadratic control problem, namely,

$$\min_{\substack{\mathbf{x}_{1:\nu}, \\ \mathbf{u}_{1:\nu}}} J_N := \mathbb{E}_{\mathbf{w}_{\nu-N+1:\nu-1}} \left[\frac{1}{2} \mathbf{x}_{\nu}^T \bar{Q} \mathbf{x}_{\nu} + \bar{q}^T \mathbf{x}_{\nu} \right. \\ \left. + \sum_{t=\nu-N+1}^{\nu-1} \left[\frac{1}{2} \mathbf{x}_t^T \bar{Q} \mathbf{x}_t + \bar{q}^T \mathbf{x}_t + \frac{1}{2} \mathbf{u}_t^T \bar{R} \mathbf{u}_t \right] \right] \quad (1a)$$

$$\text{s.t. } \mathbf{x}_{t+1} = A \mathbf{x}_t + B \mathbf{u}_t + d + \mathbf{w}_t, \quad (1b)$$

$$t = \nu - N + 1 : \nu - 1, \quad (1b)$$

$$\mathbf{x}_{t+1} = \mathbf{x}_t, \quad t = 1 : \nu - N \quad (1c)$$

$$\mathbf{x}_1 = \bar{\mathbf{x}}, \quad (1d)$$

$$\mathbf{u}_1 = \dots = \mathbf{u}_{\nu-N} = 0, \quad (1e)$$

where $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $\bar{Q} \in \mathbb{S}^n$, $\bar{R} \in \mathbb{S}^m$, and $\bar{q}, d \in \mathbb{R}^n$. More specifically, the minimization in (1) is over admissible control strategies with complete state information, i.e., \mathbf{u}_t is a function that maps from \mathbb{R}^n to \mathbb{R}^m , $\mathbf{u}_t : \mathbf{x}_t \mapsto \mathbf{u}_t(\mathbf{x}_t)$ (see, e.g., [36, Chp. 8]).

The formulation (1) is motivated by the fact that an agent can have different time horizon lengths (and initial values) to complete different tasks, while the underlying decision principle (i.e., running cost) remains unchanged since the principle is connected to the agent’s characteristics. In particular, given a realization of the time-horizon length $N = N$ and initial value $\bar{\mathbf{x}} = \bar{\mathbf{x}}$, the agent starts to apply its control from the initial value $\bar{\mathbf{x}}$ at the time instant $t = \nu - N + 1$ and the agent maintains *the same running cost* for each control execution.

This formulation gives a systematic way to handle real-world data with different time lengths (see [48]). Moreover, note that since the dynamics (1b) and the running cost in (1a) are time-invariant, by Bellman’s principle of optimality, (1) can be reformulated to optimal control problems with planning horizon length N and that start to control from an initial state at time point $t = 1$. However, it turns out to be convenient to align the optimal demonstration trajectories with different lengths at the end time point, and view the demonstration trajectories as if the expert starts to control the system at different time instants, i.e., to formulate the problem as in (1).

Remark 2.1 *The reason why we only consider a time-invariant external forcing term d in (1b) is to simplify the presentation. The results of the paper also hold for time-varying forcing terms d_t , provided that the agent knows all the future forcing terms for $t = \nu - N + 1 : \nu - 1$. Similarly, this formulation can be extended to tracking-problems by letting the linear cost term be time-varying, $q_t = -Qx_t^r$ where x_t^r is the reference signal, and the results in the paper follows analogously (cf. [48]). Notably, with time-varying d_t or q_t , the arguments in the paragraph just before (using time-invariance of running cost and dynamics) does not hold. Nevertheless, the formulation (1) is still of interest in application such as the tracking in rehabilitation trainings, see [48].*

In the remainder of the paper, we make the following assumptions.

Assumption 2.2 (Controllability and full rank)

The system (A, B) is controllable, A is invertible, and B is of full-column rank.

Assumption 2.3 (Independent white random elements)

The discrete time stochastic processes $\{\mathbf{w}_t\}_{t=1}^{\infty}$ and $\{\mathbf{v}_t\}_{t=1}^{\infty}$ are independent zero-mean white-noise processes. More specifically, this means that $\mathbb{E}[\mathbf{w}_t] = 0$, $\mathbb{E}[\mathbf{v}_t] = 0$, $\forall t$, and $\text{cov}(\mathbf{w}_t, \mathbf{w}_s) = \Sigma_w \delta(t - s)$, $\text{cov}(\mathbf{v}_t, \mathbf{v}_t) = \Sigma_v \delta(t - s)$, where $\delta(t)$ is the Dirac-delta function, and where $\Sigma_w \succeq 0$ and $\Sigma_v \succeq 0$ are a priori known. Moreover, the random element $(\bar{\mathbf{x}}, N)$ is independent of the two stochastic processes.

The rationale behind the first assumption is that we are considering a controllable discrete-time linear system that is not over-actuated,¹ and that is sampled from a continuous-time system. Next, make the following mild assumptions regarding the stochastic planning horizon N .

Assumption 2.4 (Support of the planning horizon)

The constant $\nu \in \mathbb{Z}_+$ is known, and $\nu \geq n + 1$.

¹ It means that, given the system state’s evolution from time step t to $t + 1$, there is only one possible control input to realize that.

Moreover, the probability distribution for \mathbf{N} satisfies $\mathbb{P}(\mathbf{N} \in [2, \nu]) = 1$, and $\mathbb{P}(\mathbf{N} = \nu) > 0$.

The above assumption means that the longest possible planning horizon is known, that it is sufficiently long, and that the longest horizon ν can be realized, i.e., it has a nonzero probability. Next, note that since \bar{Q} might not be positive semi-definite, (1) might not admit an optimal solution (see, e.g. [13, 14]). We analyze the well-posedness of the forward problem (1) in depth in Section 3. However, before that, we have the following proposition which illustrates the reason why we emphasize the longest time-horizon length in Assumption 2.4.

Proposition 2.5 *Under Assumptions 2.2 and 2.3, if the optimal control problem (1) with the objective function given by $(\bar{Q}, \bar{q}, \bar{R})$ admits a solution for planning horizon $N = \nu$ for any $\bar{x} \in \mathbb{R}^n$, then it admits a solution for all $N = 2 : \nu$ for any $\bar{x} \in \mathbb{R}^n$.*

PROOF. See appendix. \square

With Assumption 2.4 and Proposition 2.5 in mind, we are thus interested in parameters that belong to the following set:

$\mathcal{F}(\bar{R}) = \{(Q, q) \in \mathbb{S}^n \times \mathbb{R}^n \mid \text{the optimal control problem (1) with the objective function given by } (Q, q, \bar{R}) \text{ admits solutions a.s. under the distribution of } \bar{x} \text{ and for all } N \in \{2, 3, \dots, \nu\}\}$.

More specifically, if $(\bar{Q}, \bar{q}) \in \mathcal{F}(\bar{R})$, then the forward problem is well-posed for any $N \in \{2, \dots, \nu\}$ under the distribution of \bar{x} almost surely. In addition, for the inverse problem we assume that the observation of the optimal states $\mathbf{x}_{1:\nu}$ are contaminated by observation noise:

$$\mathbf{y}_t = \mathbf{x}_t + \mathbf{v}_t, \quad t = 1 : \nu, \quad (2)$$

and that we observe M trials of the agent. More precisely, let \mathbf{y}_t^i have the same distribution as \mathbf{y}_t , for all $i = 1 : M$ and all $t = 1 : \nu$. Then the observed M trajectories of the agent's trials, $\{\mathbf{y}_t^i\}_{i=1}^M$, are just realizations of the I.I.D. random vectors $\{\mathbf{y}_t^i\}_{i=1}^M$.

Before we formulate the IOC problem, we also make the following assumption.

Assumption 2.6 (Initial value distribution) *The random element (\bar{x}, \mathbf{N}) is such that $\mathbb{E}[\|\bar{x}\|^2] < \infty$. Moreover, for all $N \in \{2, \dots, \nu\}$ such that $\mathbb{P}(\mathbf{N} = N) > 0$, it holds that for all $\chi \in \mathbb{R}^n$, there exists a $\rho > 0$ such that $\mathbb{P}(\bar{x} \in \mathcal{B}_\epsilon^n(\rho\chi) \mid \mathbf{N} = N) > 0$ for all $\epsilon > 0$.*

Intuitively speaking, the above assumption states that, for each planning horizon length of interest, the initial value for the forward problem can be in any “direction” from the origin. This turns out to be important for both the forward and the inverse problem. The latter will be discussed in Section 4.

With the setup presented in this section, we summarize the IOC problem to be considered in this paper. For the sake of simplicity, we consider the case $\bar{R} = I$ when designing the IOC algorithm.

Problem 2.7 (General stochastic linear-quadratic IOC)

Suppose the unknown $(\bar{Q}, \bar{q}) \in \mathcal{F}(I)$. Given the optimal state trajectory observations $\{y_t^i\}_{t=1}^\nu$ of the agent's trials $i = 1 : M$ that are governed by (1), estimate the corresponding (\bar{Q}, \bar{q}) in the objective function (1a) that governs the agents' motion.

3 Forward problem analysis

Before we present the IOC algorithm set-up, we first need to analyze the forward problem. More precisely, we need to characterize the set $\mathcal{F}(\bar{R})$ and find the necessary and sufficient optimality conditions for the existence of such generalized indefinite linear-quadratic optimal control. This is not only because we want to ensure that the forward-problem is well-behaved, but also to construct the IOC algorithm based on the optimality conditions. Moreover, we also analyze the properties of the time-varying closed-loop system matrices that will be useful in developing the IOC algorithm. For the theoretical development in this section, we do not assume that $\bar{R} = I$.

3.1 Necessary and sufficient conditions for existence of optimal control

To this end, we first derive the necessary and sufficient conditions for existence of optimal control to (1). The results build on [13]; in particular, some of the proof ideas are inspired by the proof in [13, Thm. 2.1]. However, we not only extend the result to a more general setting of stochastic linear-quadratic problems, but also show that solvability of the forward problem, i.e., that $(\bar{Q}, \bar{q}) \in \mathcal{F}(\bar{R})$, can be characterized in different, but equivalent, ways. The main result of this section is the following.

Theorem 3.1 (Boundedness of forward problem)

Under Assumptions 2.3 and 2.6, the following statements are equivalent:

- (1) $(\bar{Q}, \bar{q}) \in \mathcal{F}(\bar{R})$.
- (2) Let $P_{1:\nu}$ and $\bar{\eta}_{1:\nu}$ be generated by the following Ric-

cati iterations

$$\bar{P}_\nu = \bar{Q}, \quad (3a)$$

$$\bar{P}_t = A^T \bar{P}_{t+1} A + \bar{Q} - A^T \bar{P}_{t+1} B (B^T \bar{P}_{t+1} B + \bar{R})^\dagger \\ \times B^T \bar{P}_{t+1} A, \quad t = 1 : \nu - 1; \quad (3b)$$

$$\bar{\eta}_\nu = \bar{q}, \quad (3c)$$

$$\bar{\eta}_t = (A - B(B^T \bar{P}_{t+1} B + \bar{R})^\dagger B^T \bar{P}_{t+1} A)^T \\ \times (\bar{\eta}_{t+1} + \bar{P}_{t+1} d) + \bar{q}, \quad t = 1 : \nu - 1. \quad (3d)$$

Denote

$$\bar{\mathfrak{S}}_t := B^T \bar{P}_{t+1} A, \quad (4a)$$

$$\bar{\mathfrak{R}}_t := B^T \bar{P}_{t+1} B + \bar{R}, \quad (4b)$$

$$\bar{g}_t := B^T \bar{\eta}_{t+1} + B^T \bar{P}_{t+1} d. \quad (4c)$$

It holds that

$$\bar{\mathfrak{R}}_t \succeq 0, \quad (4d)$$

$$\ker(\bar{\mathfrak{R}}_t) \subset \left[\ker(\bar{\mathfrak{S}}_t^T) \cap \ker(\bar{g}_t^T) \right], \quad (4e)$$

for all $t = 1 : \nu - 1$.

(3) There exists $\{\bar{P}_t \in \mathbb{S}^n\}_{t=1:\nu}$, $\{\bar{\eta}_t \in \mathbb{R}^n\}_{t=1:\nu}$, and $\{\bar{\xi}_t \in \mathbb{R}\}_{t=1:\nu}$ such that

$$\bar{P}_\nu = \bar{Q}, \quad \bar{\eta}_\nu = \bar{q}, \quad (5a)$$

$$\bar{H}_t := \begin{bmatrix} \bar{\mathfrak{R}}_t & \bar{\mathfrak{S}}_t & \bar{g}_t \\ \bar{\mathfrak{S}}_t^T & A^T \bar{P}_{t+1} A + \bar{Q} - \bar{P}_t & \bar{\beta}_t \\ \bar{g}_t^T & \bar{\beta}_t^T & \bar{\xi}_t \end{bmatrix} \succeq 0 \quad (5b)$$

$$\text{rank}(\bar{H}_t) = \text{rank}(\bar{\mathfrak{R}}_t) \quad (5c)$$

where $\bar{\beta}_t := \bar{q} + A^T \bar{P}_{t+1} d + A^T \bar{\eta}_{t+1} - \bar{\eta}_t$ and where $\bar{\mathfrak{S}}_t$, $\bar{\mathfrak{R}}_t$, and \bar{g}_t are as in (4a), (4b), and (4c), respectively, for all $t = 1 : \nu - 1$.

(4) The Hamilton-Jacobi-Bellman equation (HJBE)

$$V_\nu(\chi_\nu) := \frac{1}{2} \chi_\nu^T \bar{Q} \chi_\nu + \bar{q}^T \chi_\nu, \quad (6a)$$

$$V_t(\chi_t) = \min_{\mu_t} \left\{ \frac{1}{2} \chi_t^T \bar{Q} \chi_t + \bar{q}^T \chi_t + \frac{1}{2} \mu_t^T \bar{R} \mu_t \right. \quad (6b)$$

$$\left. + \mathbb{E}_{\mathbf{w}_t} [V_{t+1}(A\chi_t + B\mu_t + d + \mathbf{w}_t)] \right\}, \quad t = 1 : \nu - 1$$

has a solution. More precisely, this means that $V_t(\chi_t)$ is bounded from below for any $\chi_t \in \mathbb{R}^n$, for $t = 1 : \nu - 1$. Moreover, the solution has the form

$$V_t(\chi_t) = \frac{1}{2} \chi_t^T \bar{P}_t \chi_t + \bar{\eta}_t^T \chi_t + \bar{\gamma}_t, \quad t = 1 : \nu, \quad (7)$$

where $\bar{P}_{1:\nu}$ and $\bar{\eta}_{1:\nu}$ are generated by (3) and

$$\bar{\gamma}_\nu = 0 \quad (8a)$$

$$\bar{\gamma}_t = \bar{\gamma}_{t+1} - \frac{1}{2} \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t + \frac{1}{2} d^T \bar{P}_{t+1} d + \bar{\eta}_{t+1}^T d \\ + \frac{1}{2} \text{tr}(\bar{P}_{t+1} \Sigma_w), \quad t = 1 : \nu - 1, \quad (8b)$$

with $\bar{\mathfrak{R}}_{1:\nu-1}$ and $\bar{g}_{1:\nu-1}$ as in (4b) and (4c), respectively.

In addition, if any of the four above conditions hold, then the optimal control signal \mathbf{u}_t for (1) is parametrized by any arbitrary vector $\lambda_t \in \mathbb{R}^n$, and is given by

$$\mathbf{u}_t = -\bar{\mathfrak{R}}_t^\dagger (\bar{\mathfrak{S}}_t \mathbf{x}_t + \bar{g}_t) + \mathcal{P}_t^{\ker(\bar{\mathfrak{R}})} \lambda_t, \quad (9a)$$

$$t = \nu - N + 1 : \nu - 1, \quad (9b)$$

$$\mathbf{u}_t = 0, \quad t = 1 : \nu - N,$$

where $\mathcal{P}_t^{\ker(\bar{\mathfrak{R}})} = I - \bar{\mathfrak{R}}_t^\dagger \bar{\mathfrak{R}}_t$ is the projection operator onto the kernel space of $\bar{\mathfrak{R}}_t$.²

PROOF. First, assume that 4) holds. Note that for an agent with a planning horizon realization $N = \nu$, by the principle of optimality in dynamic programming (see, e.g., [8, p. 18]), we can easily show that 4) \implies 1) (cf. [8, Prop. 1.3.1]). For the case $N < \nu$, note that the HJBE is still valid for $t = \nu - N + 1 : \nu$, and thus the agents behavior is still optimal in $t = \nu - N + 1 : \nu$. Since the systems behavior for $t = 1 : \nu - N$ is completely determined by (1c), and (1e), by Bellman's principle of optimality, the solution to the HJBE still gives the optimal control. This implies 1), and hence shows that 4) \implies 1).

Next, we prove that 1) \implies 2) by proving the contraposition of the statement, i.e., by proving that $\neg 2) \implies \neg 1)$. Before we proceed, first note the fact that

$$\sum_{t=\nu-N+1}^{\nu-1} \left\{ \frac{1}{2} \mathbf{x}_{t+1}^T \bar{P}_{t+1} \mathbf{x}_{t+1} + \bar{\eta}_{t+1}^T \mathbf{x}_{t+1} \right. \\ \left. - \frac{1}{2} \mathbf{x}_t^T \bar{P}_t \mathbf{x}_t - \bar{\eta}_t^T \mathbf{x}_t \right\} + \frac{1}{2} \mathbf{x}_{\nu-N+1}^T \bar{P}_{\nu-N+1} \mathbf{x}_{\nu-N+1} \\ + \bar{\eta}_{\nu-N+1}^T \mathbf{x}_{\nu-N+1} - \frac{1}{2} \mathbf{x}_\nu^T \bar{P}_\nu \mathbf{x}_\nu - \bar{\eta}_\nu^T \mathbf{x}_\nu = 0, \quad (10)$$

for any $N \in \{2, \dots, \nu\}$. Taking the expectation with respect to $\mathbf{w}_{\nu-N+1:\nu-1}$ on both sides of (10), adding the latter expression to (1a), and using (1b), (3a), (3c), and

² For the fact that $\mathcal{P}_t^{\ker(\bar{\mathfrak{R}})} = I - \bar{\mathfrak{R}}_t^\dagger \bar{\mathfrak{R}}_t$ is indeed the projection operator onto the kernel space of $\bar{\mathfrak{R}}_t$, see, e.g., [12, Prop. 2.3].

Assumption 2.3, we can write the objective function as

$$\begin{aligned}
J_N &= \mathbb{E}_{\mathbf{w}_{\nu-N+1:\nu-1}} \left[\frac{1}{2} \mathbf{x}_\nu^T \underbrace{(\bar{Q} - \bar{P}_\nu)}_{=0} \mathbf{x}_\nu + \underbrace{(\bar{q} - \bar{\eta}_\nu)^T}_{=0} \mathbf{x}_\nu \right. \\
&+ \sum_{t=\nu-N+1}^{\nu-1} \left\{ \frac{1}{2} (\mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t + d + \mathbf{w}_t)^T \bar{P}_{t+1} \right. \\
&\times (\mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t + d + \mathbf{w}_t) + \bar{\eta}_{t+1}^T (\mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t + d + \mathbf{w}_t) \\
&- \frac{1}{2} \mathbf{x}_t^T \bar{P}_t \mathbf{x}_t - \bar{\eta}_t^T \mathbf{x}_t + \frac{1}{2} \mathbf{x}_t^T \bar{Q} \mathbf{x}_t + \bar{q}^T \mathbf{x}_t + \frac{1}{2} \mathbf{u}_t^T \bar{R} \mathbf{u}_t \left. \right\} \\
&+ \frac{1}{2} \mathbf{x}_{\nu-N+1}^T \bar{P}_{\nu-N+1} \mathbf{x}_{\nu-N+1} + \bar{\eta}_{\nu-N+1}^T \mathbf{x}_{\nu-N+1} \left. \right] \\
&= \mathbb{E}_{\mathbf{w}_{\nu-N+1:\nu-1}} \left[\sum_{t=\nu-N+1}^{\nu-1} \left\{ \frac{1}{2} \begin{bmatrix} \mathbf{u}_t^T & \mathbf{x}_t^T & 1 \end{bmatrix} \bar{H}_t \begin{bmatrix} \mathbf{u}_t \\ \mathbf{x}_t \\ 1 \end{bmatrix} \right\} \right. \\
&+ \frac{1}{2} \mathbf{x}_{\nu-N+1}^T \bar{P}_{\nu-N+1} \mathbf{x}_{\nu-N+1} + \bar{\eta}_{\nu-N+1}^T \mathbf{x}_{\nu-N+1} \left. \right] \\
&+ \underbrace{\sum_{t=\nu-N+1}^{\nu-1} \frac{1}{2} d^T \bar{P}_{t+1} d + \frac{1}{2} \text{tr}(\bar{P}_{t+1} \Sigma_w) + \bar{\eta}_{t+1}^T d - \frac{1}{2} \bar{\xi}_t}_{=: \Upsilon_t}
\end{aligned} \tag{11}$$

with \bar{H}_t in the form of (5b) and $\bar{\xi}_t = \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t$ in \bar{H}_t . Note that the term Υ_t in the above equation is constant with respect to the state and the control and hence can be discarded in the optimization problem. On the other hand, since $\{\bar{P}_t\}_{t=1}^\nu$ and $\{\bar{\eta}_t\}_{t=1}^\nu$ are generated by (3), \bar{H}_t can also be written as

$$\bar{H}_t = \begin{bmatrix} \bar{\mathfrak{R}}_t & \bar{\mathfrak{S}}_t & \bar{g}_t \\ \bar{\mathfrak{S}}_t^T & \bar{\mathfrak{S}}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{\mathfrak{S}}_t & \bar{\mathfrak{S}}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t \\ \bar{g}_t^T & \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{\mathfrak{S}}_t & \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t \end{bmatrix}. \tag{12}$$

Now, we use the above trick to prove that $\neg 2) \implies \neg 1)$. Suppose (4d) and (4e) cease to hold at the N th backward iteration (3), where $N \in \{2, \dots, \nu\}$. Namely, $\bar{\mathfrak{R}}_{\nu-N+t} \succeq 0$, $\ker(\bar{\mathfrak{R}}_{\nu-N+t}) \subset [\ker(\bar{\mathfrak{S}}_{\nu-N+t}^T) \cap \ker(\bar{g}_{\nu-N+t}^T)]$ still holds for $t = 2 : N$ but not for $t = 1$. We proceed by showing that this implies that for this planning horizon length realization N , (1) is not bounded from below. In particular, by (12) and $\ker(\bar{\mathfrak{R}}_{\nu-N+t}) \subset [\ker(\bar{\mathfrak{S}}_{\nu-N+t}^T) \cap \ker(\bar{g}_{\nu-N+t}^T)]$, $\forall t = 2 : N$, it follows that

$$\begin{aligned}
\bar{H}_{\nu-N+t} &= \begin{bmatrix} \bar{\mathfrak{R}}_{\nu-N+t} \\ \bar{\mathfrak{S}}_{\nu-N+t}^T \\ \bar{g}_{\nu-N+t}^T \end{bmatrix} \bar{\mathfrak{R}}_{\nu-N+t}^\dagger \\
&\times \begin{bmatrix} \bar{\mathfrak{R}}_{\nu-N+t} & \bar{\mathfrak{S}}_{\nu-N+t} & \bar{g}_{\nu-N+t} \end{bmatrix} \succeq 0
\end{aligned} \tag{13}$$

for $t = 2 : N$. Hence, in view of (1d), (1c) and the fact that $\{\mathbf{w}_t\}_{t=1}^\infty$ is independent of other random elements from Assumption 2.3, for the given ‘‘initial state’’ realization $\mathbf{x}_{\nu-N+1} = \bar{x} \in \mathbb{R}^n$ from which the agent starts tracking, the objective function can be written as

$$\begin{aligned}
J_N &= \mathbb{E}_{\mathbf{w}_{\nu-N+1:\nu-1}} \left[\frac{1}{2} \begin{bmatrix} \mathbf{u}_{\nu-N+1}^T & \mathbf{x}_{\nu-N+1}^T & 1 \end{bmatrix} \bar{H}_{\nu-N+1} \right. \\
&\times \begin{bmatrix} \mathbf{u}_{\nu-N+1} \\ \mathbf{x}_{\nu-N+1} \\ 1 \end{bmatrix} + \sum_{t=2}^{N-1} \left\{ \frac{1}{2} \left\| \left(\bar{\mathfrak{R}}_{\nu-N+t}^\dagger \right)^{\frac{1}{2}} \left(\bar{\mathfrak{R}}_{\nu-N+t} \mathbf{u}_{\nu-N+t} \right. \right. \right. \\
&+ \left. \left. \bar{\mathfrak{S}}_{\nu-N+t} \mathbf{x}_{\nu-N+t} + \bar{g}_{\nu-N+t} \right) \right\|^2 \left. \right\} \\
&+ \frac{1}{2} \underbrace{\mathbf{x}_{\nu-N+1}^T}_{\bar{x}^T} \bar{P}_{\nu-N+1} \underbrace{\mathbf{x}_{\nu-N+1}}_{\bar{x}} + \bar{\eta}_{\nu-N+1}^T \underbrace{\mathbf{x}_{\nu-N+1}}_{\bar{x}} \left. \right]. \tag{14}
\end{aligned}$$

Notably, it holds that

$$\begin{aligned}
\frac{1}{2} \bar{x}^T \bar{P}_{\nu-N+1} \bar{x} + \bar{\eta}_{\nu-N+1}^T \bar{x} &\leq \frac{1}{2} \sigma_{\max}(\bar{P}_{\nu-N+1}) \|\bar{x}\|^2 \\
+ \|\bar{\eta}_{\nu-N+1}\| \cdot \|\bar{x}\| &:= \tau(\bar{x}),
\end{aligned}$$

where $\sigma_{\max}(\cdot)$ is the largest eigenvalue of a matrix. Also note that, for any $\mathbf{u}_{\nu-N+1}$ and $\mathbf{x}_{\nu-N+1} = \bar{x}$, by the dynamics (1b) we get an $\mathbf{x}_{\nu-N+2}$. Selecting $\mathbf{u}_{\nu-N+2} = -\bar{\mathfrak{R}}_{\nu-N+2}^\dagger (\bar{\mathfrak{S}}_{\nu-N+2} \mathbf{x}_{\nu-N+2} + \bar{g}_{\nu-N+2}) + \mathcal{P}_{\nu-N+2}^{\ker(\bar{\mathfrak{R}})} \lambda_{\nu-N+2}$, for some arbitrary $\lambda_{\nu-N+2} \in \mathbb{R}^n$, this would give the next state $\mathbf{x}_{\nu-N+3}$. Recursively selecting the other control signals for $t = \nu - N + 3 : \nu - 1$ accordingly, i.e., as $\mathbf{u}_t = -\bar{\mathfrak{R}}_t^\dagger (\bar{\mathfrak{S}}_t \mathbf{x}_t + \bar{g}_t) + \mathcal{P}_t^{\ker(\bar{\mathfrak{R}})} \lambda_t$, for some arbitrary $\lambda_t \in \mathbb{R}^n$, then all terms in the summation in (14) would become zero. Therefore, for the objective function J_N it holds that

$$\begin{aligned}
J_N &\leq \mathbb{E}_{\mathbf{w}_{\nu-N+1:\nu-1}} \left[\frac{1}{2} \mathbf{u}_{\nu-N+1}^T \bar{\mathfrak{R}}_{\nu-N+1} \mathbf{u}_{\nu-N+1} \right. \\
&+ \bar{x}^T \bar{\mathfrak{S}}_{\nu-N+1}^T \mathbf{u}_{\nu-N+1} + \bar{g}_{\nu-N+1}^T \mathbf{u}_{\nu-N+1} \left. \right] \\
&+ \bar{x}^T \bar{\mathfrak{S}}_{\nu-N+1}^T \bar{\mathfrak{R}}_{\nu-N+1}^\dagger \bar{\mathfrak{S}}_{\nu-N+1} \bar{x} \\
&+ \bar{g}_{\nu-N+1}^T \bar{\mathfrak{R}}_{\nu-N+1}^\dagger \bar{\mathfrak{S}}_{\nu-N+1} \bar{x} + \tau(\bar{x}).
\end{aligned} \tag{15}$$

If $\bar{\mathfrak{R}}_{\nu-N+1} \not\prec 0$, it is clear that we can choose $\mathbf{u}_{\nu-N+1} = \alpha v_{-}(\bar{\mathfrak{R}}_{\nu-N+1})$, where $v_{-}(\bar{\mathfrak{R}}_{\nu-N+1})$ is an eigenvector that corresponds to a negative eigenvalue of $\bar{\mathfrak{R}}_{\nu-N+1}$. In this case, since such choice of $\mathbf{u}_{\nu-N+1}$ does not depend on the random vectors $\mathbf{w}_{\nu-N+1:\nu-1}$, the expectation would be marginalized out. Letting $\alpha \rightarrow +\infty$ would make J_N tend to minus infinity. Thus, unless (4d) holds, irrespective of the initial condition \bar{x} there is a planning horizon length realization $N \in \{2, \dots, \nu\}$ such that the cost function is not bounded from below.

On the other hand, if (4d) holds but $\ker(\bar{\mathfrak{R}}_{\nu-N+1}) \not\subset [\ker(\bar{\mathfrak{S}}_{\nu-N+1}^T) \cap \ker(\bar{g}_{\nu-N+1}^T)]$, then there exists a vec-

tor $v \in \mathbb{R}^m$ with norm one, such that $\bar{\mathfrak{R}}_{\nu-N+1}v = 0$, and such that $\bar{\mathfrak{S}}_{\nu-N+1}^T v \neq 0$ or $\bar{g}_{\nu-N+1}^T v \neq 0$. Without loss of generality, assume that $\bar{g}_{\nu-N+1}^T v \geq 0$; otherwise we instead consider $-v$. By Assumption 2.6, it follows that we can find $\rho > 0$ such that $\mathbb{P}(\bar{\mathbf{x}} \in \mathcal{B}_\epsilon^n(\rho \bar{\mathfrak{S}}_{\nu-N+1}^T v)) > 0, \forall \epsilon > 0$. Now, consider $\mathbf{u}_{\nu-N+1} = \alpha v$ and $\mathbf{x}_{\nu-N+1} = \bar{\mathbf{x}} = \rho \bar{\mathfrak{S}}_{\nu-N+1}^T v + \tilde{v}$, where $\tilde{v} \in \mathcal{B}_{\epsilon_1}^n(0)$ and where $\epsilon_1 > 0$ will be determined shortly. Note that with such choice, the expectation in (15) would be again marginalized out. Then it holds for the objective function that

$$\begin{aligned} J_N &\leq \alpha \underbrace{[\rho v^T \bar{\mathfrak{S}}_{\nu-N+1} \bar{\mathfrak{S}}_{\nu-N+1}^T v + \bar{g}_{\nu-N+1}^T v + \tilde{v}^T \bar{\mathfrak{S}}_{\nu-N+1}^T v]}_{=\rho \|\bar{\mathfrak{S}}_{\nu-N+1}^T v\|^2 + \bar{g}_{\nu-N+1}^T v > 0} \\ &+ (\rho v^T \bar{\mathfrak{S}}_{\nu-N+1} + \tilde{v}^T) \bar{\mathfrak{S}}_{\nu-N+1}^T \bar{\mathfrak{R}}_{\nu-N+1}^\dagger \bar{\mathfrak{S}}_{\nu-N+1} \\ &\times (\rho \bar{\mathfrak{S}}_{\nu-N+1}^T v + \tilde{v}) + \bar{g}_{\nu-N+1}^T \bar{\mathfrak{R}}_{\nu-N+1}^\dagger \bar{\mathfrak{S}}_{\nu-N+1} \\ &\times (\rho \bar{\mathfrak{S}}_{\nu-N+1}^T v + \tilde{v}) + \tau (\rho \bar{\mathfrak{S}}_{\nu-N+1}^T v + \tilde{v}). \end{aligned}$$

In particular, we can always make $\epsilon_1 > 0$ small enough so that $\rho \|\bar{\mathfrak{S}}_{\nu-N+1}^T v\|^2 + \bar{g}_{\nu-N+1}^T v + \tilde{v}^T \bar{\mathfrak{S}}_{\nu-N+1}^T v > 0, \forall \tilde{v} \in \mathcal{B}_{\epsilon_1}^n(0)$. By taking $\alpha \rightarrow -\infty$, the upper bound goes to minus infinity and hence pushes J_N to minus infinity. Recalling that $\mathbb{P}(\bar{\mathbf{x}} \in \mathcal{B}_\epsilon^n(\rho \bar{\mathfrak{S}}_{\nu-N+1}^T v)) > 0, \forall \epsilon \in (0, \epsilon_1)$, this shows that there exists a set of initial value realizations $\bar{\mathbf{x}}$ with non-zero probability such that the forward problem is ill-posed. This proves that $\neg 2) \implies \neg 1)$, and thus that $1) \implies 2)$.

Next, we prove that $2) \implies 4)$. We make an ansatz that the solution to the HJBE (6) is $V_t(\chi_t) = \frac{1}{2} \chi_t^T \bar{P}_t \chi_t + \bar{\eta}_t^T \chi_t + \bar{\gamma}_t$, where $\bar{P}_{1:\nu}$, $\bar{\eta}_{1:\nu}$ and $\bar{\gamma}_{1:\nu}$ are generated by (3) and (8), respectively. The ansatz fulfils (6a) and (8). Plugging the ansatz into (6b), we have that

$$\begin{aligned} \frac{1}{2} \chi_t^T \bar{P}_t \chi_t + \bar{\eta}_t^T \chi_t + \bar{\gamma}_t &= \min_{\mu_t} \left\{ \frac{1}{2} \chi_t^T \bar{Q} \chi_t + \bar{q}^T \chi_t \right. \\ &+ \frac{1}{2} \mu_t^T \bar{R} \mu_t + \mathbb{E}_{\mathbf{w}_t} \left[\frac{1}{2} (A \chi_t + B \mu_t + d + \mathbf{w}_t)^T \bar{P}_{t+1} \right. \\ &\times (A \chi_t + B \mu_t + d + \mathbf{w}_t) + \bar{\eta}_{t+1}^T (A \chi_t + B \mu_t + d + \mathbf{w}_t), \\ &\left. \left. + \bar{\gamma}_{t+1} \right] \right\}, \quad t = 1 : \nu - 1. \end{aligned}$$

By Assumption 2.3, we can expand the expectation regarding \mathbf{w}_t on the right hand side of the above equation, and removing constant terms that are irrelevant to the optimization we get

$$\begin{aligned} \min_{\mu_t} \left\{ \frac{1}{2} \mu_t^T \bar{\mathfrak{R}}_t \mu_t + (\bar{\mathfrak{S}}_t \chi_t + \bar{g}_t)^T \mu_t + (A \chi_t + d)^T \bar{P}_{t+1} \right. \\ \left. \times (A \chi_t + d) + \bar{\eta}_{t+1}^T (A \chi_t + d) + \frac{1}{2} \chi_t^T \bar{Q} \chi_t + \bar{q}^T \chi_t \right\}. \quad (16) \end{aligned}$$

(16) is an unconstrained quadratic optimization problem with respect to μ_t . By (4d) we know that it is a convex problem, and hence it has an optimal solution if and only if the gradient is zero in some point (see, e.g., [7,

Prop. 1.1.2]) To verify that it has a solution, and to work out the optimal control, we take the derivative of (16) with respect to μ_t and equate it to zero, which gives

$$\bar{\mathfrak{R}}_t \mu_t = -\bar{\mathfrak{S}}_t \chi_t - \bar{g}_t, \quad t = 1 : \nu - 1.$$

Since (4e) holds, we have that $\bar{\mathfrak{S}}_t \chi_t + \bar{g}_t \in \text{Im}(\bar{\mathfrak{S}}_t) \oplus \text{Im}(\bar{g}_t) = \left[\ker(\bar{\mathfrak{S}}_t^T) \cap \ker(\bar{g}_t^T) \right]^\perp \subset \ker(\bar{\mathfrak{R}}_t)^\perp = \text{Im}(\bar{\mathfrak{R}}_t)$ (see, e.g., [12, Prop. 2.3] for the last equality), where \oplus denotes the direct sum of the subspaces. Hence, the above equation has a solution, and the control signal takes the form of

$$\mu_t = -\bar{\mathfrak{R}}_t^\dagger (\bar{\mathfrak{S}}_t \chi_t + \bar{g}_t) + \mathcal{P}_t^{\ker(\bar{\mathfrak{R}})} \lambda_t, \forall \lambda_t \in \mathbb{R}^n, t = 1 : \nu - 1, \quad (17)$$

which minimizes the right hand side of (16). Plug (17) into (16), use the property of (4e) and in view of (3), (8), the quadratic, first order and constant terms regarding χ_t equates between the left and right hand sides. Thus the ansatz is indeed a solution to the HJBE.

So far, we have shown $4) \implies 1) \implies 2) \implies 4)$, i.e., the equivalence among 1), 2) and 4). To complete the proof, we now show the equivalence between 2) and 3).

First, we prove $2) \implies 3)$, and start with noting that if (3) holds, then (5a) holds trivially. Next, with $\xi_t = \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t$, we know from the above argument that due to the kernel containment (4e), \bar{H}_t can be expressed as

$$\begin{aligned} \bar{H}_t &= \begin{bmatrix} \bar{\mathfrak{R}}_t & \bar{\mathfrak{S}}_t & \bar{g}_t \\ \bar{\mathfrak{S}}_t^T & \bar{\mathfrak{S}}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{\mathfrak{S}}_t & \bar{\mathfrak{S}}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t \\ \bar{g}_t^T & \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{\mathfrak{S}}_t & \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t \end{bmatrix} \\ &= \begin{bmatrix} \bar{\mathfrak{R}}_t \\ \bar{\mathfrak{S}}_t^T \\ \bar{g}_t^T \end{bmatrix} \bar{\mathfrak{R}}_t^\dagger \begin{bmatrix} \bar{\mathfrak{R}}_t & \bar{\mathfrak{S}}_t & \bar{g}_t \end{bmatrix}, \quad t = 1 : \nu - 1. \end{aligned}$$

By (4d), $\bar{\mathfrak{R}}_t^\dagger \succeq 0$ and hence $\bar{H}_t \succeq 0$, i.e., (5b) holds. On the other hand, by the rank property of Schur complement [19, p. 43], it holds that

$$\text{rank}(\bar{H}_t) = \text{rank}(\bar{\mathfrak{R}}_t) + \text{rank}(\bar{H}_t \setminus \bar{\mathfrak{R}}_t).$$

Now, observe that

$$\begin{aligned} \bar{H}_t \setminus \bar{\mathfrak{R}}_t &= \begin{bmatrix} \bar{\mathfrak{S}}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{\mathfrak{S}}_t & \bar{\mathfrak{S}}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t \\ \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{\mathfrak{S}}_t & \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t \end{bmatrix} - \begin{bmatrix} \bar{\mathfrak{S}}_t^T \\ \bar{g}_t^T \end{bmatrix} \bar{\mathfrak{R}}_t^\dagger \begin{bmatrix} \bar{\mathfrak{S}}_t & \bar{g}_t \end{bmatrix} \\ &= \begin{bmatrix} \bar{\mathfrak{S}}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{\mathfrak{S}}_t & \bar{\mathfrak{S}}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t \\ \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{\mathfrak{S}}_t & \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t \end{bmatrix} - \begin{bmatrix} \bar{\mathfrak{S}}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{\mathfrak{S}}_t & \bar{\mathfrak{S}}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t \\ \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{\mathfrak{S}}_t & \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t \end{bmatrix} = 0, \end{aligned}$$

and hence $\text{rank}(\bar{H}_t) = \text{rank}(\bar{\mathfrak{R}}_t)$, i.e., (5c) holds.

Now we prove 3) \implies 2). If (5) holds, it follows that $\bar{\mathfrak{R}}_t \succeq 0$, i.e., (4d) holds. By properties of the generalized Schur complement [19, Thm. 1.20 and p. 43], it follows that

$$\begin{aligned} (I - \bar{\mathfrak{R}}_t \bar{\mathfrak{R}}_t^\dagger) \begin{bmatrix} \bar{\mathfrak{S}}_t & \bar{g}_t \end{bmatrix} &= 0, \implies (4e), \\ \bar{H}_t \setminus \bar{\mathfrak{R}}_t &:= \begin{bmatrix} A^T \bar{P}_{t+1} A + \bar{Q} - \bar{P}_t & \bar{\beta}_t \\ & \bar{\xi}_t \end{bmatrix} - \begin{bmatrix} \bar{\mathfrak{S}}_t^T \\ \bar{g}_t^T \end{bmatrix} \bar{\mathfrak{R}}_t^\dagger \begin{bmatrix} \bar{\mathfrak{S}}_t & \bar{g}_t \end{bmatrix} \\ &= \begin{bmatrix} A^T \bar{P}_{t+1} A + \bar{Q} - \bar{P}_t & \bar{\beta}_t \\ & \bar{\xi}_t \end{bmatrix} - \begin{bmatrix} \bar{\mathfrak{S}}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{\mathfrak{S}}_t & \bar{\mathfrak{S}}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t \\ \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{\mathfrak{S}}_t & \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t \end{bmatrix} \succeq 0, \\ \text{rank}(\bar{H}_t) &= \text{rank}(\bar{\mathfrak{R}}_t) + \text{rank}(\bar{H}_t \setminus \bar{\mathfrak{R}}_t). \end{aligned}$$

Since $\text{rank}(\bar{H}_t) = \text{rank}(\bar{\mathfrak{R}}_t)$, we must have that $\bar{H}_t \setminus \bar{\mathfrak{R}}_t = 0$, and hence (3) follows. This finishes the proof of the entire theorem. \square

Before we continue, we make a few remarks related to Theorem 3.1 and the proof.

Remark 3.2 *From the above proof, it is interesting to note that in the equivalence between point 2) and 3) in Theorem 3.1, we have that (5a) and (5c) are equivalent to that the Riccati recursions in (3) are satisfied, and that (5b), i.e., that $\bar{H}_t \succeq 0$ holds, is equivalent to (4d) and (4e).*

Remark 3.3 *By extending the state space to $\tilde{\mathbf{x}}_t = [\mathbf{x}_t^T, 1]^T$ and with state space matrices given by*

$$\tilde{A} = \begin{bmatrix} A & d \\ \mathbf{0}^T & 1 \end{bmatrix}, \quad \tilde{B} = \begin{bmatrix} B \\ 0 \end{bmatrix}, \quad (18a)$$

and cost matrices given by

$$\tilde{Q} = \begin{bmatrix} \bar{Q} & \bar{q} \\ \bar{q}^T & 0 \end{bmatrix}, \quad \tilde{R} = \bar{R}, \quad (18b)$$

we can indeed rewrite the forward problem (1) as a “stan-

dard” LQ problem (albeit possibly indefinite):

$$\begin{aligned} \min_{\mathbf{x}_{1:\nu}, \mathbf{u}_{1:\nu}} J_N &:= \mathbb{E}_{\mathbf{w}_{\nu-N+1:\nu-1}} \left[\frac{1}{2} \tilde{\mathbf{x}}_\nu^T \tilde{Q} \tilde{\mathbf{x}}_\nu \right. \\ &\quad \left. + \sum_{t=\nu-N+1}^{\nu-1} \left[\frac{1}{2} \tilde{\mathbf{x}}_t^T \tilde{Q} \tilde{\mathbf{x}}_t + \frac{1}{2} \mathbf{u}_t^T \tilde{R} \mathbf{u}_t \right] \right] \quad (19a) \\ \text{s.t. } \tilde{\mathbf{x}}_{t+1} &= \tilde{A} \tilde{\mathbf{x}}_t + \tilde{B} \mathbf{u}_t + \tilde{\mathbf{w}}_t, \\ &\quad t = \nu - N + 1, \dots, \nu - 1, \quad (19b) \\ \tilde{\mathbf{x}}_{t+1} &= \tilde{\mathbf{x}}_t, \quad t = 1, \dots, \nu - N \quad (19c) \\ \tilde{\mathbf{x}}_1 &= \tilde{\bar{\mathbf{x}}}, \quad (19d) \\ \mathbf{u}_1 = \dots = \mathbf{u}_{\nu-N} &= \mathbf{0}, \quad (19e) \end{aligned}$$

with initial value $\tilde{\bar{\mathbf{x}}} = [\bar{\mathbf{x}}^T, 1]^T$, and where $\tilde{\mathbf{w}}_t = [\mathbf{w}_t^T, 0]^T$. However, note that (19) is not a classic LQR in the sense that (\tilde{A}, \tilde{B}) is not controllable. Moreover, $\tilde{\mathbf{x}}$ does not satisfy Assumption 2.6. Furthermore, extending (19) to time-varying d_t or the tracking problem, where $q_t = Qx_t^r$ and x_t^r is the reference signal, is not possible without specifying the problem structure (18). Therefore, existing IOC results cannot be directly applied to (19).

3.2 Analysis of the closed-loop system matrices

In view of (9), it seems like there might be infinitely many choices of control signals that are optimal. However, as we shall see, under the assumptions we impose the optimal control signal is *unique* for the considered forward problem (1).

Proposition 3.4 *Let $\bar{R} \succ 0$. Under Assumptions 2.2, 2.3 and 2.6, $(\bar{Q}, \bar{q}) \in \mathcal{F}(\bar{R})$ is equivalent to $\bar{\mathfrak{R}}_t \succ 0$, $t = 1 : \nu - 1$.*

PROOF. The implication “ \Leftarrow ” follows from Theorem 3.1. To prove the implications “ \Rightarrow ”: by Theorem 3.1, since $(\bar{Q}, \bar{q}) \in \mathcal{F}(\bar{R})$, there exist matrices and vectors that fulfil (3) and (4). In particular, since A is invertible, by (4a) and (4b) we have that

$$\bar{\mathfrak{R}}_t = B^T \bar{P}_{t+1} B + \bar{R} = B^T \bar{P}_{t+1} A A^{-1} B + \bar{R} = \bar{\mathfrak{S}}_t A^{-1} B + \bar{R},$$

for $t = 1 : \nu - 1$. Now, by (4e) we have that $\ker(\bar{\mathfrak{R}}_t) \subset \ker(\bar{\mathfrak{S}}_t^T)$ holds for $t = 1 : \nu - 1$. In particular, this means that

$$z \in \ker(\bar{\mathfrak{R}}_t) \implies \begin{cases} \bar{\mathfrak{R}}_t z = 0 \\ \bar{\mathfrak{S}}_t^T z = 0 \end{cases} \implies \begin{cases} -\bar{\mathfrak{S}}_t A^{-1} B z = \bar{R} z \\ z^T \bar{\mathfrak{S}}_t = 0. \end{cases}$$

This in turn means that for any $z \in \ker(\bar{\mathfrak{R}}_t)$,

$$z^T \bar{R} z = -z^T \bar{\mathfrak{S}}_t A^{-1} B z = 0,$$

and since $\bar{R} \succ 0$ this means that $z = 0$. Therefore, the only vector in $\ker(\bar{\mathfrak{R}}_t)$ is the zero-vector, and since $\bar{\mathfrak{R}}_t$ is

positive semi-definite (see (4d)), this implies that $\bar{\mathfrak{R}}_t$ is in fact strictly positive definite for $t = 1 : \nu - 1$. \square

Corollary 3.5 *Under the assumptions in Proposition 3.4, the optimal control signal for problem (1) takes the form*

$$\mathbf{u}_t = -\bar{\mathfrak{R}}_t^{-1}(\bar{\mathfrak{S}}_t \mathbf{x}_t + \bar{g}_t), \quad t = \nu - N + 1 : \nu - 1 \quad (20a)$$

$$\mathbf{u}_t = 0, \quad t = 1 : \nu - N. \quad (20b)$$

Remark 3.6 *Even if \bar{R} is not strictly positive definite, by the proof of Proposition 3.4 we can still partially characterize $\ker(\bar{\mathfrak{R}}_t)$. In particular, $z \in \ker(\bar{\mathfrak{R}}_t)$ implies that $z^T \bar{R}z = 0$. If \bar{R} is full rank, this is the neutral subspace in the indefinite inner product space defined by \bar{R} ; see, e.g., [15, Chp. 2].*

By Corollary 3.5, under the conditions in Proposition 3.4 there is a unique solution to the forward optimal control problem (1). In this case, the system's behavior (after it starts applying control) is determined by $\tilde{\mathbf{x}}_{t+1} = \tilde{A}_{cl}(t; \bar{Q}, \bar{q})\tilde{\mathbf{x}}_t + \tilde{\mathbf{w}}_t$, for $t = \nu - N + 1 : \nu$, where $\tilde{A}_{cl}(t; \bar{Q}, \bar{q})$ is the closed-loop system matrix at time t for the extended state-space model (see Remark 3.3). In particular,

$$\tilde{A}_{cl}(t; \bar{Q}, \bar{q}) = \begin{bmatrix} A - B\bar{\mathfrak{R}}_t^{-1}\bar{\mathfrak{S}}_t & d - B\bar{\mathfrak{R}}_t^{-1}\bar{g}_t \\ \mathbf{0}^T & 1 \end{bmatrix}, \quad (21)$$

with $\bar{\mathfrak{R}}_t$, $\bar{\mathfrak{S}}_t$, and \bar{g}_t as in (4), and hence it implicitly depends on (\bar{Q}, \bar{q}) in the objective function (1a). This means that the (conditional) distribution of the agent's optimal trajectory $\mathbb{P}(\mathbf{x}_{1:\nu} \mid \mathbf{N} = N, \bar{\mathbf{x}} = \bar{x})$ and optimal control $\mathbb{P}(\mathbf{u}_{1:\nu-1} \mid \mathbf{N} = N, \bar{\mathbf{x}} = \bar{x})$ are implicitly given by solving (1). Moreover, under mild regularity conditions on the probability distribution of $(\bar{\mathbf{x}}, \mathbf{N})$, the formulation in (1) then defines joint probability distributions for $(\mathbf{x}_{1:\nu}, \mathbf{N}, \bar{\mathbf{x}})$ and $(\mathbf{u}_{1:\nu-1}, \mathbf{N}, \bar{\mathbf{x}})$ (cf. [39, Prop. 1.15] or [21, Thm. 5.3]). Before we continue analyzing the identifiability, we present the following corollary that is useful in the analysis to come.

Corollary 3.7 *Under the assumptions in Proposition 3.4, given any $R \succ 0$, for any $(Q, q) \in \mathcal{F}(R)$, let $\{P_t\}_{t=1}^\nu$ be the solution to (3) that corresponds to Q . Accordingly, let \mathfrak{R}_t , \mathfrak{S}_t and g_t be defined as in Theorem 3.1, for $t = 1 : \nu - 1$. Then the matrix*

$$A_{cl}(t; Q) := A - B\mathfrak{R}_t^{-1}\mathfrak{S}_t,$$

as well as the matrix $\tilde{A}_{cl}(t; Q, q)$ in (21), are invertible for all $t = 1 : \nu - 1$.

PROOF. The fact that the matrix $A_{cl}(t; Q)$ is invertible for $t = 1 : \nu - 1$ follows by an argument similar to

the proof of [49, Thm. 2.1], since $R \succ 0$ and $\mathfrak{R}_t \succ 0$ holds for $t = 1 : \nu - 1$. To show that $\tilde{A}_{cl}(t; Q, q)$ has full rank, we simply note that it has the upper block-triangular form (21), and since $A_{cl}(t; Q)$ has full rank so does $\tilde{A}_{cl}(t; Q, q)$. \square

4 Identifiability analysis and persistent excitation

Next, we investigate the inverse problem of recovering (\bar{Q}, \bar{q}) from observations of optimal trajectories to problem (1). Throughout the rest we will therefore, unless explicitly stated otherwise, assume that $\bar{R} = I$ and $(\bar{Q}, \bar{q}) \in \mathcal{F}(\bar{R} = I)$. We start by considering the identifiability of the problem.

To this end, first note that from the analysis in Section 3, for any parameters $(Q, q) \in \mathcal{F}(I)$, the agent's behavior is completely determined by the time-varying closed-loop system matrices $\tilde{A}_{cl}(t; Q, q)$ in (21). In the spirit of [28, Sec. 5.1], we can thus see the sequence of closed-loop system matrices $\{\tilde{A}_{cl}(t; Q, q)\}_{t=1}^{\nu-1}$ as the model structure. Therefore, the fundamental question for identifiability is if there exist two different sets of parameters (Q, q) and (Q', q') such that $\tilde{A}_{cl}(t; Q, q) = \tilde{A}_{cl}(t; Q', q')$ for all $t = 1 : \nu - 1$.

Proposition 4.1 (Identifiability) *Under Assumptions 2.2, 2.3, 2.4 and 2.6, given $(Q, q), (Q', q') \in \mathcal{F}(I)$, if $\tilde{A}_{cl}(t; Q, q) = \tilde{A}_{cl}(t; Q', q')$ for all $t = 1 : \nu - 1$, then $(Q, q) = (Q', q')$.*

PROOF. Assume that $\tilde{A}_{cl}(t; Q, q) = \tilde{A}_{cl}(t; Q', q')$ for $t = 1 : \nu - 1$, and let $(P_k, \eta_k, \mathfrak{S}_k, \mathfrak{R}_k, g_k)_{k=\nu-N+1}^{\nu-1}$ and $(P'_k, \eta'_k, \mathfrak{S}'_k, \mathfrak{R}'_k, g'_k)_{k=\nu-N+1}^{\nu-1}$ be the solutions to (3) and (4) for (Q, q) and (Q', q') , respectively. Moreover, let $Q' = Q + \Delta Q$, $q' = q + \Delta q$, $P'_t = P_t + \Delta P_t$, $\eta'_t = \eta_t + \Delta \eta_t$, $\mathfrak{S}'_t = \mathfrak{S}_t + \Delta \mathfrak{S}_t$, $\mathfrak{R}'_t = \mathfrak{R}_t + \Delta \mathfrak{R}_t$, and $g'_t = g_t + \Delta g_t$. Since $\tilde{A}_{cl}(t; Q, q) = \tilde{A}_{cl}(t; Q', q')$, for $t = 1 : \nu - 1$, it follows that $A_{cl}(t; Q) = A_{cl}(t; Q')$, for $t = 1 : \nu - 1$. Then following the line of arguments in [49, Thm. 2.1], we can conclude that $\Delta Q = 0$. This in turn implies that ΔP_t , $\Delta \mathfrak{R}_t$ and $\Delta \mathfrak{S}_t$ are all zero (cf. (3) and (4)).

Next, since $\tilde{A}_{cl}(t; Q, q) = \tilde{A}_{cl}(t; Q', q')$ and $\Delta \mathfrak{R}_t = 0$, for $t = 1 : \nu - 1$, it follows that $d - B\mathfrak{R}_t^{-1}g_t = d - B\mathfrak{R}_t^{-1}g'_t$ for $t = 1 : \nu - 1$. Therefore, it holds that $B\mathfrak{R}_t^{-1}\Delta g_t = 0$ for $t = 1 : \nu - 1$. Since B is full column rank by assumption, and since $\mathfrak{R}_t^{-1} \succ 0$ by Proposition 3.4, we must have that $\Delta g_t = 0$ for $t = 1 : \nu - 1$. In view of (4c), we therefore have that

$$B^T \Delta \eta_{t+1} = \Delta g_t = 0, \quad t = 1 : \nu - 1.$$

Next, in view of (3c) and (3d), this in turn implies that

$$\Delta\eta_\nu = \Delta q, \quad \Delta\eta_t = A^T \Delta\eta_{t+1} + \Delta q, \quad t = 1 : \nu - 1.$$

Thus, we have

$$\Delta\eta_\nu = \Delta q, \quad B^T \Delta\eta_\nu = B^T \Delta q = 0, \quad (22a)$$

$$\begin{aligned} B^T \Delta\eta_{\nu-1} &= B^T (A^T \Delta\eta_\nu + \Delta q) \\ &= B^T A^T \Delta q + \underbrace{B^T \Delta q}_{=0} = 0. \end{aligned} \quad (22b)$$

Furthermore,

$$\begin{aligned} B^T \Delta\eta_{\nu-2} &= B^T (A^T \Delta\eta_{\nu-1} + \Delta q) \\ &= B^T ((A^T)^2 \Delta\eta_\nu + A^T \Delta q + \Delta q) \\ &= B^T (A^T)^2 \Delta q + \underbrace{B^T A^T \Delta q + B^T \Delta q}_{=0} = 0, \end{aligned} \quad (22c)$$

$$\begin{aligned} \dots \\ B^T \eta_1 &= B^T (A^T)^{\nu-1} \Delta q \\ &+ \underbrace{B^T (A^T)^{\nu-2} \Delta q + \dots + B^T \Delta q}_{=0} = 0. \end{aligned} \quad (22d)$$

Writing (22) in a compact form gives

$$\underbrace{\begin{bmatrix} B^T \\ B^T A^T \\ \vdots \\ B^T (A^T)^{\nu-1} \end{bmatrix}}_{\Gamma} \Delta q = 0.$$

Since (A, B) is controllable by Assumption 2.2, and $\nu \geq n + 1$ by Assumption 2.4, the matrix Γ has full column rank, and hence $\Delta q = 0$. The fact that $\Delta Q = 0, \Delta q = 0$ implies that $Q = Q', q = q'$. \square

This means that the parameters (Q, q) that characterizes the closed-loop system matrices are identifiable. Moreover, in view of (1), we can see $(\bar{\mathbf{x}}, \mathbf{N})$ as the ‘‘input’’ of the model and $\mathbf{y}_{1:\nu}$ as the ‘‘output’’. To this end, in order to uniquely identify the parameters (\bar{Q}, \bar{q}) , the ‘‘input’’ $(\bar{\mathbf{x}}, \mathbf{N})$ needs to be ‘‘persistently exciting’’ [28, Sec. 5.1]. Notably, Assumption 2.4 gives the persistent excitation assumption regarding \mathbf{N} . Moreover, Assumption 2.6 turns out to give a persistent excitation condition for the initial value $\bar{\mathbf{x}}$. In fact, we have the following result, the proof of which we defer to the appendix.

Lemma 4.2 *Let $(\bar{\mathbf{x}}, \mathbf{N})$ be as in Assumption 2.6. Then, for all $N \in \{2, \dots, \nu\}$ such that $\mathbb{P}(\mathbf{N} = N) > 0$, $\text{cov}_{\bar{\mathbf{x}}|\mathbf{N}=N}(\bar{\mathbf{x}}, \bar{\mathbf{x}}) \succ 0$.*

PROOF. See Appendix. \square

This result can now be used to prove the following Lemma, which is useful in the IOC algorithm construction to come. Similarly, the proof of this Lemma is also deferred to the appendix.

Lemma 4.3 (Persistent excitation) *Suppose that $(\bar{Q}, \bar{q}) \in \mathcal{F}(I)$ and let $\tilde{\mathbf{x}}_t := [\mathbf{x}_t^T, 1]^T$. Under Assumptions 2.2, 2.3, 2.4, and 2.6, it holds that $\mathbb{E}_{\mathbf{x}_t|\mathbf{N}=\nu}[\tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^T] \succ 0$, and $\mathbb{E}[\|\tilde{\mathbf{x}}_t\|^2] < \infty$ for all $t = 1 : \nu$.*

PROOF. See Appendix. \square

5 The IOC algorithm

In this section, we construct the IOC algorithm for general linear-quadratic systems with different time-horizon lengths. In particular, we show that the algorithm is statistically consistent, i.e., that it converges in probability to the true underlying parameter. For the sake of brevity, in some of the following we sometimes use the notation (\cdot) for the arguments of some functions.

In order to construct the IOC algorithm, we further make the following assumption.

Assumption 5.1 (Bounded parameters) *The parameter tuple (\bar{Q}, \bar{q}) that governs the agents tracking behaviour lies in the compact set*

$$\mathbb{G}(\varphi) := \left\{ (\bar{Q} \in \mathbb{S}^n, \bar{q} \in \mathbb{R}^n) \mid \left\| \begin{bmatrix} \bar{Q} & \bar{q} \\ \bar{q} & 0 \end{bmatrix} \right\|_F \leq \varphi \right\},$$

for some (potentially unknown) $0 < \varphi < \infty$.

This assumption is mild, since when we solve the corresponding inverse problem in practice, we can always set φ arbitrary large if we have no prior knowledge on the norm bound of the parameters.

5.1 Construction and empirical approximation

To this end, the algorithm is constructed based on the necessary and sufficient optimality conditions in Theorem 3.1. More precisely, we are interested in finding the best $(Q^*, q^*) \in \mathcal{F}(I)$ that suits the collected optimal state trajectory observations $\{y_t^i\}_{t=1}^\nu$ of the agent’s trials $i = 1 : M$. Thus, the IOC algorithm will be built upon an optimization problem which is constructed so that it has a unique optimal solution (Q^*, q^*) which is the ‘‘true’’ $(\bar{Q}, \bar{q}) \in \mathcal{F}(I)$.

First, we construct the objective function for the optimization problem. Given a realization of the planning horizon N , an optimal control signal μ_t^* that is optimal to (Q, q) (assuming that $(Q, q) \in \mathcal{F}(I)$) and an optimal control signal $\bar{\mu}_t$ that is optimal to (\bar{Q}, \bar{q}) , it holds for all state $\chi_t \in \mathbb{R}^n$ that

$$\begin{aligned} 0 &= \frac{1}{2} \chi_t^T Q \chi_t + q^T \chi_t + \frac{1}{2} \|\mu_t^*\|^2 + \mathbb{E}_{\mathbf{w}_t} [V_{t+1}(A\chi_t + B\mu_t^* + d + \mathbf{w}_t)] - V_t(\chi_t) \\ &= \min_{\mu_t} \left\{ \frac{1}{2} \chi_t^T Q \chi_t + q^T \chi_t + \frac{1}{2} \|\mu_t\|^2 \right. \\ &\quad \left. + \mathbb{E}_{\mathbf{w}_t} [V_{t+1}(A\chi_t + B\mu_t + d + \mathbf{w}_t)] \right\} - V_t(\chi_t) \\ &\leq \frac{1}{2} \chi_t^T Q \chi_t + q^T \chi_t + \frac{1}{2} \|\bar{\mu}_t\|^2 + \mathbb{E}_{\mathbf{w}_t} [V_{t+1}(A\chi_t + B\bar{\mu}_t + d + \mathbf{w}_t)] - V_t(\chi_t), \end{aligned}$$

where the inequality follows since $\bar{\mu}_t$ is not necessarily optimal to $(Q, q, \bar{R} = I)$, and where $V_t(\cdot)$ has the form (7), and $P_{t:t+1}$, $\eta_{t:t+1}$, and $\gamma_{t:t+1}$ in $V_t(\cdot)$ are determined by (Q, q) via (3). Moreover, seen intuitively from the other perspective, for given χ_t and $\bar{\mu}_t$, we expect the inequality to hold unless we plug in $(Q, q, \bar{R} = I)$ which renders the state χ_t and control $\bar{\mu}_t$ optimal. We hence define the ‘‘violation’’ of HJBE at time step t by

$$\begin{aligned} \psi_{t,N}(Q, q; \chi_t, \mu_t) &:= \\ &\mathbb{E}_{\mathbf{w}_t} [V_{t+1}(A\chi_t + B\mu_t + d + \mathbf{w}_t)] + \frac{1}{2} \chi_t^T Q \chi_t + q^T \chi_t \\ &\quad + \frac{1}{2} \|\mu_t\|_2^2 - V_t(\chi_t), \quad t = \nu - N + 1 : \nu - 1, \end{aligned}$$

since we expect that $\psi_{t,N}(Q, q; \chi_t, \mu_t) \geq 0$. The latter will be formally proved in Theorem 5.2.

Now, plugging (7) and (8) in to the above equation, we have

$$\begin{aligned} \psi_{t,N}(Q, q; \chi_t, \mu_t) &= \mathbb{E}_{\mathbf{w}_t} \left[\frac{1}{2} (A\chi_t + B\mu_t + d + \mathbf{w}_t)^T P_{t+1} \right. \\ &\quad \left. \times (A\chi_t + B\mu_t + d + \mathbf{w}_t) + \eta_{t+1}^T (A\chi_t + B\mu_t + d + \mathbf{w}_t) \right] \\ &\quad + \frac{1}{2} \chi_t^T Q \chi_t + q^T \chi_t + \frac{1}{2} \|\mu_t\|^2 - \frac{1}{2} \chi_t^T P_t \chi_t - \eta_t^T \chi_t \\ &\quad + \frac{1}{2} g_t^T \mathfrak{R}_t^\dagger g_t - \frac{1}{2} d^T P_{t+1} d - \eta_t^T d - \frac{1}{2} \text{tr}(P_{t+1} \Sigma_w). \quad (23) \end{aligned}$$

Given a realization of the planning horizon N , let $\mathbf{x}_{\nu-N+1:\nu}$ and $\mathbf{u}_{\nu-N+1:\nu-1}$ be the optimal trajectory and control. We let $\chi_t = \mathbf{x}_t$ and $\mu_t = \mathbf{u}_t$, and take the expectation of $\psi_{t,N}(Q, q; \mathbf{x}_t, \mathbf{u}_t)$ with respect to

$\mathbf{x}_t | \mathbf{N} = N$. In view of (1b), this gives

$$\begin{aligned} &\mathbb{E}_{\mathbf{x}_t | \mathbf{N} = N} [\psi_{t,N}(Q, q; \mathbf{x}_t, \mathbf{u}_t)] = \\ &\mathbb{E}_{\mathbf{x}_t | \mathbf{N} = N} \left[\mathbb{E}_{\mathbf{w}_t} \left[\frac{1}{2} \underbrace{(A\mathbf{x}_t + B\mathbf{u}_t + d + \mathbf{w}_t)}_{\mathbf{x}_{t+1}} P_{t+1} \right. \right. \\ &\quad \left. \left. \times \underbrace{(A\mathbf{x}_t + B\mathbf{u}_t + d + \mathbf{w}_t)}_{\mathbf{x}_{t+1}} + \eta_{t+1}^T \underbrace{(A\mathbf{x}_t + B\mathbf{u}_t + d + \mathbf{w}_t)}_{\mathbf{x}_{t+1}} \right] \right] \\ &\quad + \frac{1}{2} \mathbf{x}_t^T Q \mathbf{x}_t + q^T \mathbf{x}_t + \frac{1}{2} \|\mathbf{u}_t\|^2 - \frac{1}{2} \mathbf{x}_t^T P_t \mathbf{x}_t - \eta_t^T \mathbf{x}_t \\ &\quad + \frac{1}{2} g_t^T \mathfrak{R}_t^\dagger g_t - \frac{1}{2} d^T P_{t+1} d - \eta_t^T d - \frac{1}{2} \text{tr}(P_{t+1} \Sigma_w) \\ &= \mathbb{E}_{\mathbf{x}_{t+1} | \mathbf{N} = N} \left[\frac{1}{2} \mathbf{x}_{t+1}^T P_{t+1} \mathbf{x}_{t+1} + \eta_{t+1}^T \mathbf{x}_{t+1} \right] \\ &\quad + \mathbb{E}_{\mathbf{x}_t | \mathbf{N} = N} \left[\frac{1}{2} \mathbf{x}_t^T Q \mathbf{x}_t + q^T \mathbf{x}_t + \frac{1}{2} \|\mathbf{u}_t\|^2 - \frac{1}{2} \mathbf{x}_t^T P_t \mathbf{x}_t - \eta_t^T \mathbf{x}_t \right. \\ &\quad \left. + \frac{1}{2} g_t^T \mathfrak{R}_t^\dagger g_t - \frac{1}{2} d^T P_{t+1} d - \eta_t^T d - \frac{1}{2} \text{tr}(P_{t+1} \Sigma_w) \right]. \quad (24) \end{aligned}$$

However, the above expression is constructed based on \mathbf{x}_t , while the observations are in terms of \mathbf{y}_t . To rewrite it in terms of \mathbf{y}_t , first we simply add and subtract some terms in the expression above:

$$\begin{aligned} &\mathbb{E}_{\mathbf{x}_t | \mathbf{N} = N} [\psi_{t,N}(Q, q; \mathbf{x}_t, \mathbf{u}_t)] = \mathbb{E}_{\mathbf{x}_t | \mathbf{N} = N} [\psi_{t,N}(Q, q; \mathbf{x}_t, \mathbf{u}_t)] \\ &\quad + \frac{1}{2} \text{tr}(P_{t+1} \Sigma_v) - \frac{1}{2} \text{tr}(P_t \Sigma_v) + \frac{1}{2} \text{tr}(Q \Sigma_v) \\ &\quad - \frac{1}{2} \text{tr}(P_{t+1} \Sigma_v) + \frac{1}{2} \text{tr}(P_t \Sigma_v) - \frac{1}{2} \text{tr}(Q \Sigma_v). \quad (25) \end{aligned}$$

On the other hand, by Assumption 2.3, $\{\mathbf{v}_t\}_{t=1}^\infty$ are independent of any other stochastic elements. Using the cyclic permutation property of the matrix trace operator, we know that

$$\begin{aligned} &\mathbb{E}_{\mathbf{v}_t} [\mathbf{v}_t^T P_{t+1} \mathbf{v}_t] = \mathbb{E}_{\mathbf{v}_t} [\text{tr}(\mathbf{v}_t^T P_{t+1} \mathbf{v}_t)] = \mathbb{E}_{\mathbf{v}_t} [\text{tr}(P_{t+1} \mathbf{v}_t \mathbf{v}_t^T)] \\ &= \text{tr}(P_{t+1} \Sigma_v). \end{aligned}$$

Similarly, we also have $\mathbb{E}_{\mathbf{v}_t} [\mathbf{v}_t^T P_t \mathbf{v}_t] = \text{tr}(P_t \Sigma_v)$, $\mathbb{E}_{\mathbf{v}_t} [\mathbf{v}_t^T Q \mathbf{v}_t] = \text{tr}(Q \Sigma_v)$, $d^T P_{t+1} d = \text{tr}(P_{t+1} d d^T)$. In view of (8), (2) and the fact that $\mathbb{E}_{\mathbf{w}_t} [\mathbf{w}_t] = 0$, $\mathbb{E}_{\mathbf{v}_t} [\mathbf{v}_t] = 0$, using (24) and (25) we can rewrite

$\mathbb{E}_{\mathbf{x}_t|\mathbf{N}=\mathbf{N}}[\psi_{t,N}(Q, q; \mathbf{x}_t, \mathbf{u}_t)]$ as

$$\begin{aligned}
& \mathbb{E}_{\mathbf{x}_t|\mathbf{N}=\mathbf{N}}[\psi_{t,N}(Q, q; \mathbf{x}_t, \mathbf{u}_t)] \\
&= \mathbb{E}_{\mathbf{v}_{t:t+1}} \left[\mathbb{E}_{\mathbf{x}_{t+1}|\mathbf{N}=\mathbf{N}} \left[\frac{1}{2} \underbrace{(\mathbf{x}_{t+1} + \mathbf{v}_{t+1})^T}_{\mathbf{y}_{t+1}^T} \right. \right. \\
&\quad \times \left. \underbrace{P_{t+1}(\mathbf{x}_{t+1} + \mathbf{v}_{t+1}) + \eta_{t+1}^T(\mathbf{x}_{t+1} + \mathbf{v}_{t+1})}_{\mathbf{y}_{t+1}} \right] \\
&\quad + \mathbb{E}_{\mathbf{x}_t|\mathbf{N}=\mathbf{N}} \left[\frac{1}{2} \underbrace{(\mathbf{x}_t + \mathbf{v}_t)^T}_{\mathbf{y}_t^T} Q \underbrace{(\mathbf{x}_t + \mathbf{v}_t)}_{\mathbf{y}_t} + q^T \underbrace{(\mathbf{x}_t + \mathbf{v}_t)}_{\mathbf{y}_t} \right. \\
&\quad \left. + \frac{1}{2} \|\mathbf{u}_t\|^2 - \frac{1}{2} \underbrace{(\mathbf{x}_t + \mathbf{v}_t)^T}_{\mathbf{y}_t^T} P_t \underbrace{(\mathbf{x}_t + \mathbf{v}_t)}_{\mathbf{y}_t} - \eta_t^T \underbrace{(\mathbf{x}_t + \mathbf{v}_t)}_{\mathbf{y}_t} \right] \\
&\quad + \frac{1}{2} g_t^T \mathfrak{R}_t^\dagger g_t - \frac{1}{2} \text{tr}(P_{t+1} d d^T) - \eta_t^T d - \frac{1}{2} \text{tr}(P_{t+1} \Sigma_w) \\
&\quad - \frac{1}{2} \text{tr}(P_{t+1} \Sigma_v) + \frac{1}{2} \text{tr}(P_t \Sigma_v) - \frac{1}{2} \text{tr}(Q \Sigma_v) \Big] \\
&= \mathbb{E}_{\mathbf{y}_{t+1}|\mathbf{N}=\mathbf{N}} \left[\frac{1}{2} \mathbf{y}_{t+1}^T P_{t+1} \mathbf{y}_{t+1} + \eta_{t+1}^T \mathbf{y}_{t+1} \right] + \mathbb{E}_{\mathbf{x}_t|\mathbf{N}=\mathbf{N}} \left[\frac{1}{2} \|\mathbf{u}_t\|^2 \right] \\
&\quad + \mathbb{E}_{\mathbf{y}_t|\mathbf{N}=\mathbf{N}} \left[\frac{1}{2} \mathbf{y}_t^T Q \mathbf{y}_t + q^T \mathbf{y}_t - \frac{1}{2} \mathbf{y}_t^T P_t \mathbf{y}_t - \eta_t^T \mathbf{y}_t \right] \\
&\quad + \frac{1}{2} g_t^T \mathfrak{R}_t^\dagger g_t - \frac{1}{2} \text{tr}(P_{t+1} d d^T) - \eta_t^T d - \frac{1}{2} \text{tr}(P_{t+1} \Sigma_w) \\
&\quad - \frac{1}{2} \text{tr}(P_{t+1} \Sigma_v) + \frac{1}{2} \text{tr}(P_t \Sigma_v) - \frac{1}{2} \text{tr}(Q \Sigma_v) \\
&=: \mathbb{E}_{\mathbf{y}_{t:t+1}|\mathbf{N}=\mathbf{N}} [\tilde{\psi}_{t,N}(Q, q, P_{t:t+1}, \eta_t, \xi_t; \mathbf{y}_{t:t+1})] \\
&\quad + \mathbb{E}_{\mathbf{x}_t|\mathbf{N}=\mathbf{N}} \left[\frac{1}{2} \|\mathbf{u}_t\|^2 \right],
\end{aligned}$$

where we introduce $\xi_t := g_t^T \mathfrak{R}_t^\dagger g_t$. We construct the objective function $\Psi(Q, q, P_{1:\nu}, \eta_{1:\nu}, \xi_{1:\nu-1})$ by summing up the above equation from $t = \nu - N + 1$ to $\nu - 1$, but excluding the terms $\mathbb{E}_{\mathbf{x}_t|\mathbf{N}=\mathbf{N}}[\frac{1}{2}\|\mathbf{u}_t\|^2]$ which are constants, and taking the expectation over \mathbf{N} . In particular,

$$\Psi(\cdot) := \sum_{N=2}^{\nu} \mathbb{P}(\mathbf{N} = N) \mathbb{E}_{\mathbf{y}_{\nu-N+1:\nu}|\mathbf{N}=\mathbf{N}} [\tilde{\psi}_N(\cdot)], \quad (26)$$

where

$$\begin{aligned}
\tilde{\psi}_N(\cdot) &= \sum_{t=\nu-N+1}^{\nu-1} \tilde{\psi}_{t,N}(\cdot) \\
&= \frac{1}{2} \mathbf{y}_\nu^T P_\nu \mathbf{y}_\nu + \eta_\nu^T \mathbf{y}_\nu - \frac{1}{2} \mathbf{y}_{\nu-N+1}^T P_{\nu-N+1} \mathbf{y}_{\nu-N+1} \\
&\quad - \eta_{\nu-N+1}^T \mathbf{y}_{\nu-N+1} - \frac{1}{2} \text{tr}(P_\nu \Sigma_v) + \frac{1}{2} \text{tr}(P_{\nu-N+1} \Sigma_v) \\
&\quad + \sum_{t=\nu-N+1}^{\nu-1} \left(\frac{1}{2} \xi_t - \frac{1}{2} \text{tr}(P_{t+1} d d^T) - \eta_{t+1}^T d + \frac{1}{2} \mathbf{y}_t^T Q \mathbf{y}_t \right. \\
&\quad \left. + q^T \mathbf{y}_t - \frac{1}{2} \text{tr}(P_{t+1} \Sigma_w) - \frac{1}{2} \text{tr}(Q \Sigma_v) \right). \quad (27)
\end{aligned}$$

The objective function (26) can be rewritten as a joint expectation over $\mathbf{y}_{1:\nu}$ and \mathbf{N} . However, we find the form in (26) more useful both in analysis and in implementation.

The objective function (26) is constructed with the idea that $\Psi(\cdot) + \sum_{t=1}^{\nu-1} \mathbb{E}_{\mathbf{x}_t, \mathbf{N}} [\frac{1}{2} \|\mathbf{u}_t\|^2]$, where the latter is the discarded constant, should be bounded from below by 0 for all $(Q, q) \in \mathcal{F}(I)$. Therefore, ideally we would consider the problem of finding the point (Q^*, q^*) that minimizes (26) subject to (3) and (4). From Theorem 3.1, we know that (3) and (4) are equivalent to (5) and $\xi_t = g_t^T \mathfrak{R}_t^\dagger g_t$ for $t = 1 : \nu - 1$. However, while (26) is linear and hence a convex function, neither the constraint (5) nor the constraint $\xi_t = g_t^T \mathfrak{R}_t^\dagger g_t$ for $t = 1 : \nu - 1$ are convex. Even so, note that the only nonconvex part in (5) is (5c). We therefore consider the relaxed convex problem obtained by removing the constraints (5c) and $\xi_t = g_t^T \mathfrak{R}_t^\dagger g_t$ for $t = 1 : \nu - 1$.³ In particular, the optimization problem for IOC reads

$$\begin{aligned}
& \min_{(Q, q) \in \mathbb{G}(\varphi)} \Psi(Q, q, P_{1:\nu}, \eta_{1:\nu}, \xi_{1:\nu-1}) \\
& \quad \{P_t \in \mathbb{S}_+^n(\varphi)\}_{t=1:\nu}, \\
& \quad \{\eta_t \in \mathcal{R}_\varphi^1(0)\}_{t=1:\nu}, \\
& \quad \{\xi_t \in \mathcal{R}_\varphi^1(0)\}_{t=1:\nu-1} \\
& \text{s.t.} \quad P_\nu = Q, \quad (28a) \\
& \quad \eta_\nu = q \quad (28b) \\
& \quad H_t \succeq 0, \quad t = 1 : \nu - 1, \quad (28c)
\end{aligned}$$

where H_t has the same form as (5b). In Section 5.2, we prove that the unique optimal solution to this optimization problem is indeed (\bar{Q}, \bar{q}) .

Nevertheless, the distribution of $\bar{\mathbf{x}}, \{\mathbf{w}_t\}, \{\mathbf{v}_t\}$ and \mathbf{N} are usually not a priori known in practice, and hence the distribution of \mathbf{y}_t and \mathbf{N} are not known. Therefore, it is not possible to calculate the objective function (26) explicitly, and hence we cannot solve (28) directly. But since we have the optimal state trajectory observations $\{\mathbf{y}_t^i\}_{t=1}^{\nu}$ of the agent's trials, i.e., realizations of I.I.D. random processes $\{\mathbf{y}_t^i\}_{t=1}^{\nu}$ for $i = 1 : M$, we can instead empirically estimate the objective function. To this end, let M_N denote the number of observations which has a planning horizon of N time steps. Clearly, $\sum_{N=2}^{\nu} M_N = M$. Then, for each value of N , the expectation in (26) is ap-

³ Another possibility would be to substitute $\xi_t = g_t^T \mathfrak{R}_t^\dagger g_t$ into the matrix (5b), which would (also) give a convex problem; see [31, 32].

proximated by the empirical mean as

$$\begin{aligned} & \mathbb{E}_{\mathbf{y}_{\nu-N+1:\nu} | \mathbf{N}=\mathbf{N}} [\tilde{\psi}_N(\cdot)] \approx \\ & \frac{1}{M_N} \sum_{i_N=1}^{M_N} \left[\frac{1}{2} \mathbf{y}_{\nu}^{i_N T} P_{\nu} \mathbf{y}_{\nu}^{i_N} + \eta_{\nu}^T \mathbf{y}_{\nu}^{i_N} - \frac{1}{2} \mathbf{y}_{\nu-N+1}^{i_N T} P_{\nu-N+1} \mathbf{y}_{\nu-N+1}^{i_N} \right. \\ & - \eta_{\nu-N+1}^T \mathbf{y}_{\nu-N+1}^{i_N} - \frac{1}{2} \text{tr}(P_{\nu} \Sigma_v) + \frac{1}{2} \text{tr}(P_{\nu-N+1} \Sigma_v) \\ & + \sum_{t=\nu-N+1}^{\nu-1} \left(\frac{1}{2} \xi_t - \frac{1}{2} \text{tr}(P_{t+1} d d^T) - \eta_{t+1}^T d + \frac{1}{2} \mathbf{y}_t^{i_N T} Q \mathbf{y}_t^{i_N} \right. \\ & \left. \left. + q^T \mathbf{y}_t^{i_N} - \frac{1}{2} \text{tr}(P_{t+1} \Sigma_w) - \frac{1}{2} \text{tr}(Q \Sigma_v) \right) \right]. \end{aligned}$$

On the other hand, approximating the expectation over \mathbf{N} is the same as estimating the probabilities $\mathbb{P}(\mathbf{N} = N)$ using the empirical estimates M_N/M . This, together with the above expression, gives that

$$\begin{aligned} & \Psi(Q, q, P_{1:\nu}, \eta_{1:\nu}, \xi_{1:\nu-1}) \approx \\ & \Psi_E^{\mathbf{y}}(Q, q, P_{1:\nu}, \eta_{1:\nu}, \xi_{1:\nu-1}) = \frac{1}{M} \sum_{N=2}^{\nu} \sum_{i_N=1}^{M_N} \left[\frac{1}{2} \mathbf{y}_{\nu}^{i_N T} P_{\nu} \mathbf{y}_{\nu}^{i_N} \right. \\ & + \eta_{\nu}^T \mathbf{y}_{\nu}^{i_N} - \frac{1}{2} \mathbf{y}_{\nu-N+1}^{i_N T} P_{\nu-N+1} \mathbf{y}_{\nu-N+1}^{i_N} - \eta_{\nu-N+1}^T \mathbf{y}_{\nu-N+1}^{i_N} \\ & - \frac{1}{2} \text{tr}(P_{\nu} \Sigma_v) + \frac{1}{2} \text{tr}(P_{\nu-N+1} \Sigma_v) + \sum_{t=\nu-N+1}^{\nu-1} \left(\frac{1}{2} \xi_t \right. \\ & - \frac{1}{2} \text{tr}(P_{t+1} d d^T) - \eta_{t+1}^T d + \frac{1}{2} \mathbf{y}_t^{i_N T} Q \mathbf{y}_t^{i_N} + q^T \mathbf{y}_t^{i_N} \\ & \left. \left. - \frac{1}{2} \text{tr}(P_{t+1} \Sigma_w) - \frac{1}{2} \text{tr}(Q \Sigma_v) \right) \right]. \quad (29) \end{aligned}$$

We therefore consider the estimator

$$\begin{aligned} & \min_{\substack{(Q, q) \in \mathbb{G}(\varphi) \\ \{P_t \in \mathbb{S}_+^n(\varphi)\}_{t=1:\nu}, \\ \{\eta_t \in \mathcal{B}_{\varphi}^n(0)\}_{t=1:\nu}, \\ \{\xi_t \in \mathcal{B}_{\varphi}^1(0)\}_{t=1:\nu-1}}} \Psi_E^{\mathbf{y}}(Q, q, P_{1:\nu}, \eta_{1:\nu}, \xi_{1:\nu-1}) \\ & \text{s.t. (28a)–(28c) hold.} \quad (30) \end{aligned}$$

In practice, an estimate is obtained by solving (30) for a given realization $\{\mathbf{y}_{1:\nu}^i\}_{i=1}^M$ of $\{\mathbf{y}_{1:\nu}^i\}_{i=1}^M$. We will use the notation $\Psi_E^{\mathbf{y}}(\cdot)|_{\mathbf{y}=\mathbf{y}}$ to denote the objective function at the given realization.

5.2 Statistical consistency analysis

In this section, we analyze the statistical consistency of the IOC algorithm. To proceed, we first show that the optimization problem (28) is well-posed, i.e., the objective function (26) is bounded from below on its feasible domain (28a), (28b) and (28c). In addition, we show the “true” (\bar{Q}, \bar{q}) is actually the unique global minimizer.

Theorem 5.2 *Let $(\bar{Q}, \bar{q}) \in \mathcal{F}(I)$ be the “true” parameters of the stochastic linear-quadratic control problem (1) that governs the agent, and let $\mathbf{x}_{1:\nu}$, $\mathbf{u}_{1:\nu}$, and $\mathbf{y}_{1:\nu}$ be distributed accordingly. Under Assumptions 2.2, 2.3, 2.4, 2.6, and 5.1, for any feasible solution $(Q, q, P_{1:\nu}, \eta_{1:\nu}, \xi_{1:\nu-1})$ of the optimization problem (28), the objective function (26) is bounded from below by $-\sum_{t=1}^{\nu-1} \mathbb{E}_{\mathbf{x}_t, \mathbf{N}} [\frac{1}{2} \|\mathbf{u}_t\|^2]$. Moreover, for φ that is large enough, $(\bar{Q}, \bar{q}, \bar{P}_{1:\nu}, \bar{\eta}_{1:\nu}, \bar{\xi}_{1:\nu-1})$ is the unique globally optimal solution achieving the lower bound, where $\bar{P}_{1:\nu}, \bar{\eta}_{1:\nu}$ are generated by (3) and $\bar{\xi}_t = \bar{g}_t^T \bar{\mathfrak{R}}_t^{\dagger} \bar{g}_t$, $t = 1 : \nu - 1$.*

PROOF. See Appendix. \square

Having shown that the optimization problem (28) has (\bar{Q}, \bar{q}) as unique globally optimal solution, next we turn to the estimator (30). We show that it is statistically consistent, but to this end we first have the following Lemmas.

Lemma 5.3 (Boundedness of estimator) *The feasible region in problem (30) is compact. Moreover, for any realization, the cost function $\Psi_E^{\mathbf{y}}(Q, q, P_{1:\nu}, \eta_{1:\nu}, \xi_{1:\nu-1})|_{\mathbf{y}=\mathbf{y}}$ is bounded on the feasible region, and the optimization problem (30) is convex and admits an optimal solution.*

PROOF. It has already been established that the feasible region is convex, see the arguments in Section 5.1. Moreover, the feasible region is bounded by construction. Next, for any realization, $\Psi_E^{\mathbf{y}}(Q, q, P_{1:\nu}, \eta_{1:\nu}, \xi_{1:\nu-1})|_{\mathbf{y}=\mathbf{y}}$ is a linear function, and hence (30) is a convex problem. Moreover, since the cost function is linear, it is bounded on the compact feasible domain. Finally, by Weierstrass’ theorem (see, e.g., [7, Prop. A.8]) the optimization problem (30) admits an optimal solution. \square

Lemma 5.4 (Uniform law of large numbers) *For large enough φ and under Assumptions 2.2, 2.3, 2.4, 2.6, and 5.1, the optimal value*

$$\begin{aligned} & \sup_{\substack{(Q, q) \in \mathbb{G}(\varphi) \\ \{P_t \in \mathbb{S}_+^n(\varphi)\}_{t=1:\nu}, \\ \{\eta_t \in \mathcal{B}_{\varphi}^n(0)\}_{t=1:\nu}, \\ \{\xi_t \in \mathcal{B}_{\varphi}^1(0)\}_{t=1:\nu-1}}} |\Psi_E^{\mathbf{y}}(\cdot) - \Psi(\cdot)| \\ & \text{s.t. (28a)–(28c) hold.} \end{aligned}$$

converges to 0 almost surely as $M \rightarrow \infty$.

PROOF. The proof follows along the lines of [47, Proof of Lem. 4.2]. First, the argument inside the expectation in $\Psi(\cdot)$ can be bounded from above by an integrable function of the random variables, but which is independent of the parameters $(Q, q, P_{1:\nu}, \eta_{1:\nu}, \xi_{1:\nu-1})$. This can

be done using bounds from Lemmas 4.3 and 5.3. Using this bound in terms of an integrable function, the result follows from [20, Thm. 2]. \square

Theorem 5.5 (Statistical consistency) *For large enough φ and under Assumptions 2.2, 2.3, 2.4, 2.6, and 5.1, given a realization of M trajectories, let $(Q^M, q^M, P_{1:\nu}^M, \eta_{1:\nu}^M, \xi_{1:\nu-1}^M)$ be a corresponding optimal solution to (30). Then $Q^M \xrightarrow{p} \bar{Q}$ and $q^M \xrightarrow{p} \bar{q}$ as $M \rightarrow \infty$.*

PROOF. The result follows by verifying the conditions in [43, Thm. 5.7]. In particular, since (30) is convex, $(Q^M, q_{1:\nu}^M, P_{1:\nu}^M, \eta_{1:\nu}^M, \xi_{1:\nu-1}^M)$ is a globally optimal solution. This means that

$$\begin{aligned} & \Psi_E^y(Q^M, q^M, P_{1:\nu}^M, \eta_{1:\nu}^M, \xi_{1:\nu-1}^M) |_{\mathbf{y}=y} \\ & \leq \Psi_E^y(\bar{Q}, \bar{q}, \bar{P}_{1:\nu}, \bar{\eta}_{1:\nu}, \bar{\xi}_{1:\nu-1}) |_{\mathbf{y}=y}. \end{aligned}$$

Moreover, since convergence almost surely implies convergence in probability [21, Lem. 3.2], Lemma 5.4 implies that

$$\begin{aligned} & \sup_{\substack{(Q,q) \in \mathbb{G}(\varphi) \\ \{P_t \in \mathbb{S}_+^n(\varphi)\}_{t=1:\nu} \\ \{\eta_t \in \mathbb{R}^n(0)\}_{t=1:\nu} \\ \{\xi_t \in \mathbb{R}^1(0)\}_{t=1:\nu-1}}} |\Psi_E^y(\cdot) - \Psi(\cdot)| \\ & \text{s.t.} \quad (28a)-(28c) \text{ hold.} \end{aligned}$$

converges to 0 in probability as $M \rightarrow \infty$. Finally, the fact that the feasible region to (28) and (30) is compact (see Lemma 5.3), and that (28) has a unique optimal solution (see Theorem 5.2), by [43, p. 46] the last condition also holds. Hence, the result follows. \square

5.3 On implementation and the computational complexity of the estimator

To get a point estimate from the estimator (30), the data (i.e., the observed trajectories) are used in the optimization problem (30). This problem can be solved using any appropriate method for solving the convex optimization problem. In fact, even for modest sized problems, one can use many off-the-shelf convex optimization solvers. For example, problem (30) can be implemented, almost exactly as stated, in frameworks for disciplined convex programming, including YALMIP [29], CVX [16, 17], and CVXPY [1, 11].

The only difference between an implementation in a framework for disciplined convex programming and the stated problem in (30), is that the cost function can be rewritten to make the implementation more efficient. To this end, observe that for any $Z \in \mathbb{S}^n$ and any $a \in \mathbb{R}^n$,

$a^T Z a = \text{tr}(Z a a^T)$. This means that the objective function (29) can be rewritten as

$$\begin{aligned} \Psi_E^y(Q, q, P_{1:\nu}, \eta_{1:\nu}, \xi_{1:\nu-1}) &= \frac{1}{M} \sum_{N=2}^{\nu} \left[\frac{1}{2} \text{tr}(P_{\nu} \mathbf{Y}_{\nu}^{(N)}) + \eta_{\nu}^T \mathbf{y}_{\nu}^{(N)} \right. \\ & - \frac{1}{2} \text{tr}(P_{\nu-N+1} \mathbf{Y}_{\nu-N+1}^{(N)}) - \eta_{\nu-N+1}^T \mathbf{y}_{\nu-N+1}^{(N)} - \frac{M_N}{2} \text{tr}(P_{\nu} \Sigma_{\nu}) \\ & + \frac{M_N}{2} \text{tr}(P_{\nu-N+1} \Sigma_{\nu}) + \sum_{t=\nu-N+1}^{\nu-1} \left(\frac{M_N}{2} \xi_t - \frac{M_N}{2} \text{tr}(P_{t+1} d d^T) \right. \\ & - M_N \eta_{t+1}^T d + \frac{1}{2} \text{tr}(Q \mathbf{Y}_t^{(N)}) + q^T \mathbf{y}_t^{(N)} \\ & \left. \left. - \frac{M_N}{2} \text{tr}(P_{t+1} \Sigma_w) - \frac{M_N}{2} \text{tr}(Q \Sigma_w) \right) \right], \end{aligned}$$

where $\mathbf{y}_t^{(N)} = \sum_{i_N=1}^{M_N} \mathbf{y}_t^{i_N}$ and $\mathbf{Y}_t^{(N)} = \sum_{i_N=1}^{M_N} \mathbf{y}_t^{i_N} (\mathbf{y}_t^{i_N})^T$.

Note that $\mathbf{y}_t^{(N)}$ and $\mathbf{Y}_t^{(N)}$ are collecting all the samples at time-point t from trajectories with the same planning horizon length N , and that these can be pre-computed from the data before assembling the optimization problem (30). Moreover, the sizes of $\mathbf{y}_t^{(N)}$ and $\mathbf{Y}_t^{(N)}$ only depend on the dimension of the state space, n . It means that the size of the optimization problem does not grow with the amount of data collected.

More specifically, since $Q \in \mathbb{S}^n$, $q \in \mathbb{R}^n$, $\{P_t \in \mathbb{S}_+^n\}_{t=1:\nu}$, $\{\eta_t \in \mathbb{R}^n\}_{t=1:\nu}$, and $\{\xi_t \in \mathbb{R}\}_{t=1:\nu-1}$, the number of variables in the problem is $n(n+1)/2 + n + \nu n(n+1)/2 + \nu n + \nu$. Moreover, the LMI constraints in (28c) are ν symmetric matrices of size $(m+n+1) \times (m+n+1)$. This means that, e.g., $n = 12$, $m = 4$, and $\nu = 80$ gives a problem with a total of 7370 scalar variables and 80 LMI constraints of size 17×17 . As we demonstrate in Section 6.1, this can be handled by standard off-the-shelf convex optimization solvers.

6 Numerical examples

In this section, we present two numerical examples. The first example, in Section 6.1, illustrates that the problem (30) can be solved efficiently with off-the-shelf convex optimization solvers. The second example, in Section 6.2, applies the developed methodology to a non-zero sum pursuit-evasion game, where the pursuer models the evaders objective function using collected data. In both examples, the problem is solved on a MacBook Pro with Apple M1 eight-core CPU and 16GB of RAM, and the implementation is done using YALMIP [29] in Matlab and solved by MOSEK [4].

6.1 Demonstration of performance for a system with both modest size and planning horizon

To illustrate the performance of the method, we generate a system with modest size and modest planning horizon length. In particular, to ensure that Assumption 2.2

holds, we generate continuous-time matrices $\hat{A} \in \mathbb{R}^{12 \times 12}$ and $\hat{B} \in \mathbb{R}^{12 \times 4}$ in controllable canonical form

$$\hat{A} = \begin{bmatrix} & & & I_4 \\ & & & \\ & & I_4 & \\ & & & \\ & & & I_4 \\ a_1 I_4 & a_2 I_4 & a_3 I_4 & a_4 I_4 \end{bmatrix}, \quad \hat{B} = \begin{bmatrix} \mathbf{0}_4 \\ \mathbf{0}_4 \\ \mathbf{0}_4 \\ I_4 \end{bmatrix}.$$

We sample the coefficients a_i , $i = 1 : 4$ from a standard normal distribution $\mathcal{N}(0, 1)$. Next, we discretize the system by letting $A = e^{\hat{A}\Delta t}$ and $B = \int_0^{\Delta t} e^{\hat{A}t} dt \hat{B}$, using the sampling period $\Delta t = 0.1$. We choose \tilde{Q} to be the Hermitian part of a randomly drawn matrix with shifted eigenvalues so that the smallest eigenvalue is -0.1 .⁴ Namely, we let $G' = (G + G^T)/2$ and $\tilde{Q} = G' - (\sigma_{\min}(G') + 0.1)I$, where $\sigma_{\min}(\cdot)$ is the smallest eigenvalue of a matrix and where $G \in \mathbb{R}^{12 \times 12}$ and elements are randomly drawn from $\mathcal{N}(0, 1)$. We set $\nu = 80$, and verify that the conditions in point 2) in Theorem 3.1 hold, i.e., that $\tilde{\mathfrak{R}}_t \succ 0$ for $t = 1 : 79$. The process noise \mathbf{w}_t and measurement noise \mathbf{v}_t are drawn from multivariate normal distribution $\mathcal{N}(0, \Sigma_w)$ and $\mathcal{N}(0, \Sigma_v)$, respectively, with covariance matrices that are randomly generated from a Wishart distribution of degree 12, i.e., with the same degrees of freedom as the dimension of the state. Moreover, the Wishart distribution used to draw the covariance matrices has itself a random covariance of $0.01GG^T$, where each element in $G \in \mathbb{R}^{12 \times 12}$ was drawn from a standard normal distribution. Finally, we generate $M = 5 \times 10^4$ optimal trajectories, with the planning horizon lengths N drawn uniformly from the integers in the interval $[2, 80]$ and with initial value $\bar{\mathbf{x}}$ drawn from $\mathcal{N}(0, 100I_{12})$.

The time to solve the optimization problem (30), as reported by MOSEK, is 4.85 seconds. Moreover, the relative error of the estimate is 0.0347, where the relative error is computed as $\frac{\|\tilde{Q}_{est} - \tilde{Q}\|_F}{\|\tilde{Q}\|_F}$, where \tilde{Q} is defined in (18b) and \tilde{Q}_{est} is defined analogously. This shows that solving the IOC problem for systems of “moderate” size and planning horizon length can be done efficiently with off-the-shelf solvers.

6.2 Identification of cost in non-zero sum pursuit-evasion game

In this section, we demonstrate the performance of the proposed IOC algorithm on a non-zero sum two-dimensional finite-horizon linear-quadratic pursuit-evasion game, cf. [40]. For a more extensive treatment of pursuit-evasion games, see, e.g., [5]. To this end, let $\mathbf{x}_t \in \mathbb{R}^2$ be the distance between the pursuer and the

⁴ We shift the eigenvalues in order to make sure that we get a forward problem that is well-posed, i.e., that $(\tilde{Q}, \bar{q}) \in \mathcal{F}(I)$.

evader, and let $\mathbf{u}_t^p, \mathbf{u}_t^e \in \mathbb{R}^2$ be the control signal of the pursuer and the evader, respectively. In particular, for each realization $(\bar{\mathbf{x}}, N)$ of $(\bar{\mathbf{x}}, N)$, we assume that the evader solves the following problem

$$\min_{\substack{\mathbf{x}_{1:\nu}, \\ \mathbf{u}_{1:\nu}^e}} J_N := \mathbb{E}_{\mathbf{w}_{\nu-N+1:\nu-1}} \left[\frac{1}{2} \mathbf{x}_{\nu}^T Q^e \mathbf{x}_{\nu} + \sum_{t=\nu-N+1}^{\nu-1} \left[\frac{1}{2} \mathbf{x}_t^T Q^e \mathbf{x}_t + \frac{1}{2} \|\mathbf{u}_t^e\|^2 \right] \right] \quad (31a)$$

$$\text{s.t. } \mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t^e + B\mathbf{u}_t^p + \mathbf{w}_t, \quad (31b)$$

$$t = \nu - N + 1 : \nu - 1, \quad (31b)$$

$$\mathbf{x}_{t+1} = \mathbf{x}_t, \quad t = 1 : \nu - N \quad (31c)$$

$$\mathbf{x}_1 = \bar{\mathbf{x}}, \quad (31d)$$

$$\mathbf{u}_1^e = \dots = \mathbf{u}_{\nu-N}^e = 0, \quad (31e)$$

where (A, B) is discretized in the same way as in Sec. 6.1 from the continuous-time dynamics $\dot{\mathbf{x}} = \hat{A}\mathbf{x} + \hat{B}\mathbf{u}^e + \hat{B}\mathbf{u}^p$ using the sampling period $\Delta t = 0.1$, and where $\nu = 20$. Notably, $Q^e \prec 0$. In practice, as a pursuer, Q^e is unknown. In order to gain advantages over the evader and predict its future movements, as a pursuer, one can first use some “trivial” and “dummy” movements $\mathbf{u}_{\nu-N+1:\nu}^p$ that are easy for the evader to predict (i.e., known by the evader) in the first a few rounds of the game. During these rounds, the pursuer collects the evader’s behaviour data and use the proposed IOC algorithm to estimate Q^e . In particular, here we assume the pursuer choose control \mathbf{u}_t^p to be a constant during the data collection phase for convenience. Consequently, the forcing term $d = B\mathbf{u}_t^p$ would be constant for the evader (cf. (1b)).⁵ The pursuer observes the noisy distance (see (2)) between the pursuer and evader, which is the optimal solution to (31).

To simulate this, we choose $\hat{A} = 0$, $\hat{B} = I_2$, $Q^e = -0.1I_2$, and for each time step in the trajectories process noise \mathbf{w}_t and measurement noise \mathbf{v}_t are drawn from multivariate normal distribution $\mathcal{N}(0, \Sigma_w)$ and $\mathcal{N}(0, \Sigma_v)$, respectively, with covariance matrices

$$\Sigma_w \approx 10^{-2} \begin{bmatrix} 1.04 & 0.68 \\ 0.68 & 1.00 \end{bmatrix} \quad \text{and} \quad \Sigma_v \approx 10^{-2} \begin{bmatrix} 2.33 & -2.25 \\ -2.25 & 2.18 \end{bmatrix}.$$

The latter matrices were randomly generated by drawing two elements from a Wishart distribution of degree 2, i.e., with the same degrees of freedom as the dimension of the state. The Wishart distribution is itself generated analogously to the distribution in Section 6.1. As “dummy” movements for the pursuer, we choose $\mathbf{u}_t^p = [-1, -1]^T$, $t = \nu - N + 1 : \nu$, and hence

⁵ In fact, it does not matter what kinds of strategy the pursuer uses in the data collection phase, as long as the evader can foresee it, since, as mentioned in Sec. 2.1, the results still hold for time-varying d_t .

the constant forcing term in the dynamics is given by $d = B[-1, -1]^T$. Finally, the random variable \mathbf{N} is taken to be uniformly distributed on the integers between 2 and $\nu = 20$.

To test the performance of the algorithm, we generate 100 batches of trajectories, where each batch consists of 50000 trajectories. For each batch, we divide the trajectories into groups of size $M = 100 + 100(k - 1)$, for $k = 1, \dots, 500$, where each larger group contains all the trajectories of a smaller group. For each such group of trajectories, we solve the IOC problem (with φ set to 10^6), and this procedure is repeated for all the 100 batches. This means that we obtain estimates $Q_{est}^{\ell, M}$, for $\ell = 1, \dots, 100$ and $M = 100, 200, \dots, 50000$. For each value of M , the relative error $\|Q_{est}^{\ell, M} - Q^e\|_F / \|Q^e\|_F$ is averaged over the batches, and the resulting empirical mean and empirical standard deviation (as a function of M) is shown in Figure 1. From the figure we see that, in line with the statistical consistency proved in Theorem 5.5, both the mean and the standard deviation decreases with increasing M . Moreover, in Figure 1 the logarithm of the mean and the logarithm of the standard deviation appears to be (approximately) affine in $\log(M)$. The figure also shows the corresponding lines obtained by fitting an affine model to each of the two sets of logarithmic data. From this fit, we see that **Mean of relative error** $\approx \mathcal{O}(M^{-0.48})$ and **Standard deviation of relative error** $\approx \mathcal{O}(M^{-0.55})$. We hence suspect that the convergence rate is $\mathcal{O}(M^{-0.5})$, and that $\sqrt{M}(Q_M - \bar{Q})$ is asymptotically normal, just like most M-estimators such as maximum log-likelihood [43, p. 51]. However, a theoretical analysis of this is left for future work.




Fig. 1. Log-log plot of the mean and standard deviation of the relative error of Q_{est} as a function of the number of trajectories. The estimates are obtained using noisy data, as described in Section 6.2.

7 Conclusion

In this work, we have considered the inverse optimal control problem for discrete-time finite-horizon general in-

definite linear-quadratic problems with stochastic planning horizons. We first investigate the necessary and sufficient conditions for when the forward problem is solvable. The identifiability of the corresponding inverse optimal control problem is analyzed and proved. Furthermore, based on the underlying necessary and sufficient condition, we construct the estimator of the inverse optimal control problem as the solution to a convex optimization problem, and prove that the estimator is statistically consistent. The performance of the estimator is illustrated on a numerical example of identifying the evaders cost in non-zero sum pursuit-evasion game.

Acknowledgements

The authors would like to thank the anonymous reviewers for helpful comments that helped improve the paper.

A Deferred proofs

PROOF OF PROPOSITION 2.5. We show the contraposition of the statement, i.e., that if $(\bar{Q}, \bar{q}, \bar{R})$ is such that there exists an initial value $\bar{x} \in \mathbb{R}^n$ and a time-horizon length $N \in \{2, \dots, \nu\}$ so that the optimal control problem (1) is unbounded below, then there exists an initial value \bar{x}' so that (1) is unbounded from below for planning horizon length ν . To this end, consider planning horizon $\mathbf{N} = \nu$. By Theorem 3.1, if $(\bar{Q}, \bar{q}) \notin \mathcal{F}(\bar{R})$ then there exists an N such that (4d) or (4e) does not hold for $t = \nu - N + 1$. Splitting the summation in the objective function as

$$J_\nu = \mathbb{E}_{\mathbf{w}_{1:\nu-1}} \left[\frac{1}{2} \mathbf{x}_\nu^T \bar{Q} \mathbf{x}_\nu + \bar{q}^T \mathbf{x}_\nu + \sum_{t=\nu-N+1}^{\nu-1} \left[\frac{1}{2} \mathbf{x}_t^T \bar{Q} \mathbf{x}_t + \bar{q}^T \mathbf{x}_t + \frac{1}{2} \mathbf{u}_t^T \bar{R} \mathbf{u}_t \right] + \sum_{t=1}^{\nu-N} \left[\frac{1}{2} \mathbf{x}_t^T \bar{Q} \mathbf{x}_t + \bar{q}^T \mathbf{x}_t + \frac{1}{2} \mathbf{u}_t^T \bar{R} \mathbf{u}_t \right] \right],$$

and following along the lines of the proof of “1) \implies 2)” in Theorem 3.1, similar to (15) we get the following inequality

$$J_\nu \leq \mathbb{E}_{\mathbf{w}_{1:\nu-1}} \left[\frac{1}{2} \mathbf{u}_{\nu-N+1}^T \bar{\mathfrak{R}}_{\nu-N+1} \mathbf{u}_{\nu-N+1} + \mathbf{x}_{\nu-N+1}^T \bar{\mathfrak{S}}_{\nu-N+1}^T \mathbf{u}_{\nu-N+1} + \bar{g}_{\nu-N+1}^T \mathbf{u}_{\nu-N+1} + \mathbf{x}_{\nu-N+1}^T \bar{\mathfrak{S}}_{\nu-N+1}^T \bar{\mathfrak{R}}_{\nu-N+1}^\dagger \bar{\mathfrak{S}}_{\nu-N+1} \mathbf{x}_{\nu-N+1} + \bar{g}_{\nu-N+1}^T \bar{\mathfrak{R}}_{\nu-N+1}^\dagger \bar{\mathfrak{S}}_{\nu-N+1} \mathbf{x}_{\nu-N+1} + \tau(\mathbf{x}_{\nu-N+1}) + \sum_{t=1}^{\nu-N} \left[\frac{1}{2} \mathbf{x}_t^T \bar{Q} \mathbf{x}_t + \bar{q}^T \mathbf{x}_t + \frac{1}{2} \mathbf{u}_t^T \bar{R} \mathbf{u}_t \right] \right].$$

Moreover, from the same proof we know that we can select $\mathbf{x}_{\nu-N+1}$ in order to make the terms outside of the last summation unbounded from below. Now, by invertability of A , we can select $\mathbf{u}_t = 0$ and $\mathbf{x}_t = A^{-1}(\mathbf{x}_{t+1} - d - \mathbf{w}_t)$ for $t = 1 : \nu - N$. For any value of $\mathbf{x}_{\nu-N+1}$, this gives an initial condition and a sequence of states and controls that fulfill the constraints (1b)-(1e) (note that (1c) and (1e) are vacuous since $\mathbf{N} = \nu$). Moreover it is easily seen that J_ν is bounded from above by an expression similar to the one in the previous proof, but containing an additional constant $\tilde{\tau}(\mathbf{x}_{\nu-N+1})$. Following a logic similar to the reminder of the proof of “1) \implies 2)” in Theorem 3.1, the result follows. \square

PROOF OF LEMMA 4.2. Let $N \in \{2, \dots, \nu\}$ be such that $\mathbb{P}(\mathbf{N} = N) > 0$. The matrix $\text{cov}_{\bar{\mathbf{x}}|\mathbf{N}=N}(\bar{\mathbf{x}}, \bar{\mathbf{x}})$ is symmetric. Let $\text{cov}_{\bar{\mathbf{x}}|\mathbf{N}=N}(\bar{\mathbf{x}}, \bar{\mathbf{x}}) = \sum_{i=1}^n \lambda_i v_i$ be an orthonormal eigen-decomposition of the matrix. This means that we can write $\bar{\mathbf{x}} = \sum_{i=1}^n \alpha_i v_i$ for some real-valued random variables $\alpha_i, i = 1, \dots, n$.

Assume that $\text{cov}_{\bar{\mathbf{x}}|\mathbf{N}=N}(\bar{\mathbf{x}}, \bar{\mathbf{x}})$ is not (strictly) positive definite. Then at least one eigenvalue is zero; without loss of generality, let $\lambda_1 = 0$. Then we have that

$$\begin{aligned} 0 &= v_1^T \text{cov}_{\bar{\mathbf{x}}|\mathbf{N}=N}(\bar{\mathbf{x}}, \bar{\mathbf{x}}) v_1 \\ &= \mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N} (v_1^T \bar{\mathbf{x}} \bar{\mathbf{x}}^T v_1) - \mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N} (v_1^T \bar{\mathbf{x}}) \mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N} (\bar{\mathbf{x}}^T v_1) \\ &= \mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N} (\alpha_1^2) - \mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N} (\alpha_1) \mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N} (\alpha_1), \end{aligned}$$

and thus that $(\mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N}(\alpha_1))^2 = \mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N}(\alpha_1^2)$. By Jensen’s inequality [6, Prop. 9.24], we know that $(\mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N}(\alpha_1))^2 \leq \mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N}(\alpha_1^2)$, and by following the proof of [6, Prop. 9.24], we see that equality holds if and only if α_1 is constant a.s. (otherwise, the inequality from [6, Thm. 9.23] is strict at some points). To this end, let $\alpha_1 = c$ a.s. for some constant c . If $c = 0$, then for $\chi = v_1$, Assumption 2.6 does not hold. If $c \neq 0$, then the probability mass of $\bar{\mathbf{x}}$ is located on a hyperplane defined by $\alpha_1 = c$, which does not pass through the origin. In this case, let $\chi = v_2$ and note that for $\epsilon < c$ we have that $\mathbb{P}(\bar{\mathbf{x}} \in \mathcal{B}_\epsilon(\rho v_2)) = 0$ for all ρ , hence violating Assumption 2.6. Therefore, we must have $\text{cov}_{\bar{\mathbf{x}}|\mathbf{N}=N}(\bar{\mathbf{x}}, \bar{\mathbf{x}}) \succ 0$. \square

PROOF OF LEMMA 4.3. Note that

$$\begin{aligned} \text{cov}_{\bar{\mathbf{x}}|\mathbf{N}=N}(\bar{\mathbf{x}}, \bar{\mathbf{x}}) &= \mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N} [\bar{\mathbf{x}} \bar{\mathbf{x}}^T] - \mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N} [\bar{\mathbf{x}}] \mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N} [\bar{\mathbf{x}}]^T \\ &= \underbrace{\mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N} \left[\begin{bmatrix} \bar{\mathbf{x}} \bar{\mathbf{x}}^T & \bar{\mathbf{x}} \\ \bar{\mathbf{x}}^T & 1 \end{bmatrix} \right]}_{\mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N} [\tilde{\bar{\mathbf{x}}} \tilde{\bar{\mathbf{x}}}^T]} \setminus 1. \end{aligned}$$

Hence by Lemma 4.2 and [19, Thm. 1.12], we know that for all N such that $\mathbb{P}(\mathbf{N} = N) > 0$, $\mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N} [\tilde{\bar{\mathbf{x}}} \tilde{\bar{\mathbf{x}}}^T] \succ 0$. On the other hand, by Assumption 2.3, \mathbf{w}_t is uncorrelated with the noiseless $\mathbf{z}_t := A \mathbf{x}_t + B \mathbf{u}_t + d$, for $t = \nu - N + 1 : \nu - 1$, and for such t it thus hold that

$$\begin{aligned} &\mathbb{E}_{\mathbf{x}_{t+1}|\mathbf{N}=N} [\tilde{\mathbf{x}}_{t+1} \tilde{\mathbf{x}}_{t+1}^T] \\ &= \mathbb{E}_{\mathbf{x}_{t+1}|\mathbf{N}=N} \left[\begin{bmatrix} \mathbf{z}_t + \mathbf{w}_t \\ 1 \end{bmatrix} \begin{bmatrix} \mathbf{z}_t^T + \mathbf{w}_t^T & 1 \end{bmatrix} \right] \\ &= \mathbb{E}_{\mathbf{x}_t|\mathbf{N}=N} [\tilde{\mathbf{z}}_t \tilde{\mathbf{z}}_t^T] + \begin{bmatrix} \Sigma_w & 0 \\ 0 & 0 \end{bmatrix} \\ &= \tilde{A}_{cl}(t; \bar{Q}, \bar{q}) \mathbb{E}_{\mathbf{x}_t|\mathbf{N}=N} [\tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^T] \tilde{A}_{cl}(t; \bar{Q}, \bar{q})^T + \begin{bmatrix} \Sigma_w & 0 \\ 0 & 0 \end{bmatrix}, \end{aligned}$$

where $\tilde{\mathbf{z}}_t = \begin{bmatrix} \mathbf{z}_t^T & 1 \end{bmatrix}^T$. In particular, note that

$$\begin{aligned} &\mathbb{E}_{\mathbf{x}_{\nu-N+2}|\mathbf{N}=N} [\tilde{\mathbf{x}}_{\nu-N+2} \tilde{\mathbf{x}}_{\nu-N+2}^T] \\ &= \mathbb{E}_{\mathbf{x}_{\nu-N+1}|\mathbf{N}=N} [\tilde{\mathbf{z}}_{\nu-N+1} \tilde{\mathbf{z}}_{\nu-N+1}^T] + \begin{bmatrix} \Sigma_w & 0 \\ 0 & 0 \end{bmatrix} \\ &= \tilde{A}_{cl}(\nu - N + 1; \bar{Q}, \bar{q}) \underbrace{\mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N} [\tilde{\bar{\mathbf{x}}} \tilde{\bar{\mathbf{x}}}^T]}_{\succ 0} \tilde{A}_{cl}(\nu - N + 1; \bar{Q}, \bar{q})^T \\ &\quad + \underbrace{\begin{bmatrix} \Sigma_w & 0 \\ 0 & 0 \end{bmatrix}}_{\succeq 0} \succ 0, \end{aligned} \tag{A.1}$$

since $\tilde{A}_{cl}(t; \bar{Q}, \bar{q})$ is invertible for all $t = 1 : \nu - 1$ and since positive definiteness is invariant under congruence. By induction, we thus have $\mathbb{E}_{\mathbf{x}_t|\mathbf{N}=N} [\tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^T] \succ 0$ for all N such that $\mathbb{P}(\mathbf{N} = N) > 0$, and in particular thus for $\mathbf{N} = \nu$ by Assumption 2.4.

Now we show $\mathbb{E}[\|\tilde{\mathbf{x}}_t\|^2] < \infty$. First, note that $\mathbb{E}[\|\tilde{\bar{\mathbf{x}}}\|^2] = \mathbb{E}[\|\bar{\mathbf{x}}\|^2] + 1$. By Assumption 2.6 we have $\mathbb{E}[\|\bar{\mathbf{x}}\|^2] < \infty$, and hence $\mathbb{E}[\|\tilde{\bar{\mathbf{x}}}\|^2] < \infty$. In addition, by taking trace on both sides of (A.1), moving the trace inside the expectation, rearranging terms, and using Cauchy-Schwartz inequality, we have that

$$\begin{aligned} &\mathbb{E}_{\mathbf{x}_{\nu-N+2}|\mathbf{N}=N} [\|\tilde{\mathbf{x}}_{\nu-N+2}\|^2] \\ &\leq \mathbb{E}_{\bar{\mathbf{x}}|\mathbf{N}=N} [\|\tilde{\bar{\mathbf{x}}}\|^2] \cdot \|\tilde{A}_{cl}(\nu - N + 1; \bar{Q}, \bar{q})\|_F^2 + \text{tr}(\Sigma_w). \end{aligned}$$

Using an induction argument similar to the one above, we thus have that $\mathbb{E}_{\mathbf{x}_t|\mathbf{N}=N} [\|\tilde{\mathbf{x}}_t\|^2] < \infty$ for all $t = \nu -$

$N + 1 : \nu$ and all N such that $\mathbb{P}(\mathbf{N} = N) > 0$. Finally,

$$\mathbb{E}[\|\tilde{\mathbf{x}}_t\|^2] = \sum_{N=1}^{\nu} \mathbb{P}(\mathbf{N} = N) \mathbb{E}_{\tilde{\mathbf{x}}_t | \mathbf{N}=N} [\|\tilde{\mathbf{x}}_t\|^2] < \infty,$$

which proves the lemma. \square

PROOF OF THEOREM 5.2 As can be seen from the construction of the objective function in Section 5.1, and in view of (1e),

$$\begin{aligned} & \Psi(Q, q, P_{1:\nu}, \eta_{1:\nu}, \xi_{1:\nu-1}) + \sum_{t=1}^{\nu-1} \mathbb{E}_{\mathbf{x}_t, \mathbf{N}} \left[\frac{1}{2} \|\mathbf{u}_t\|^2 \right] \\ &= \sum_{N=2}^{\nu} \mathbb{P}(\mathbf{N} = N) \mathbb{E}_{\mathbf{y}_{\nu-N+1:\nu} | \mathbf{N}=N} [\tilde{\psi}_N(\cdot)] \\ &+ \sum_{N=2}^{\nu} \mathbb{P}(\mathbf{N} = N) \sum_{t=\nu-N+1}^{\nu-1} \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} \left[\frac{1}{2} \|\mathbf{u}_t\|^2 \right] \\ &= \sum_{N=2}^{\nu} \mathbb{P}(\mathbf{N} = N) \sum_{t=\nu-N+1}^{\nu-1} \left(\mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} [\psi_{t,N}(Q, q; \mathbf{x}_t, \mathbf{u}_t) + \frac{1}{2} \|\mathbf{u}_t\|^2] \right) \end{aligned}$$

Now, by the definition of $\psi_{t,N}(Q, q; \mathbf{x}_t, \mathbf{u}_t)$ in (23), as in (24) we have that

$$\begin{aligned} & \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} [\psi_{t,N}(Q, q; \mathbf{x}_t, \mathbf{u}_t)] \\ &= \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} \left[\mathbb{E}_{\mathbf{w}_t} \left[\frac{1}{2} (\mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t + d + \mathbf{w}_t)^T P_{t+1} \right. \right. \\ &\quad \times (\mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t + d + \mathbf{w}_t) + \eta_{t+1}^T (\mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t + d + \mathbf{w}_t) \\ &\quad + \frac{1}{2} \mathbf{x}_t^T Q \mathbf{x}_t + q^T \mathbf{x}_t + \frac{1}{2} \|\mathbf{u}_t\|^2 - \frac{1}{2} \mathbf{x}_t^T P_t \mathbf{x}_t - \eta_t^T \mathbf{x}_t \\ &\quad \left. \left. + \frac{1}{2} g_t^T \mathfrak{R}_t^\dagger g_t - \frac{1}{2} d^T P_{t+1} d - \eta_t^T d - \frac{1}{2} \text{tr}(P_{t+1} \Sigma_w) \right] \right]. \end{aligned}$$

By opening the parenthesis and computing the expectation with respect to \mathbf{w}_t , and using Assumption 2.3, we have that

$$\begin{aligned} & \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} [\psi_{t,N}(Q, q; \mathbf{x}_t, \mathbf{u}_t)] \\ &= \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} \left[\frac{1}{2} (\mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t + d)^T P_{t+1} (\mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t + d) \right. \\ &\quad + \eta_{t+1}^T (\mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t + d) - \frac{1}{2} \mathbf{x}_t^T P_t \mathbf{x}_t - \eta_t^T \mathbf{x}_t + \frac{1}{2} \xi_t \\ &\quad \left. + \frac{1}{2} \mathbf{x}_t^T Q \mathbf{x}_t + q^T \mathbf{x}_t \right] \\ &= \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} \left[\frac{1}{2} \begin{bmatrix} \mathbf{u}_t \\ \mathbf{x}_t \\ 1 \end{bmatrix} \begin{bmatrix} \mathbf{u}_t^T & \mathbf{x}_t^T & 1 \end{bmatrix} H_t - \frac{1}{2} \|\mathbf{u}_t\|^2 \right], \quad (\text{A.2}) \end{aligned}$$

where H_t has the form (5b). On the other hand, since $(Q, q, P_{1:\nu}, \eta_{1:\nu}, \xi_{1:\nu-1})$ is feasible, by the constraint (28c) we have that $H_t \succeq 0$. Therefore, it holds that

$$\begin{aligned} \Psi(\cdot) &= \sum_{N=2}^{\nu} \mathbb{P}(\mathbf{N} = N) \sum_{t=\nu-N+1}^{\nu-1} \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} \left[\right. \\ &\quad \left. \frac{1}{2} \begin{bmatrix} \mathbf{u}_t^T & \mathbf{x}_t^T & 1 \end{bmatrix} H_t \begin{bmatrix} \mathbf{u}_t \\ \mathbf{x}_t \\ 1 \end{bmatrix} - \frac{1}{2} \|\mathbf{u}_t\|^2 \right] \quad (\text{A.3}) \\ &\geq - \sum_{N=2}^{\nu} \mathbb{P}(\mathbf{N} = N) \sum_{t=\nu-N+1}^{\nu-1} \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} \left[\frac{1}{2} \|\mathbf{u}_t\|^2 \right] \\ &= - \sum_{t=1}^{\nu-1} \mathbb{E}_{\mathbf{x}_t, \mathbf{N}} \left[\frac{1}{2} \|\mathbf{u}_t\|^2 \right]. \end{aligned}$$

This proves the first part of the theorem.

Next, we show that the lower bound is actually attained by $(\bar{Q}, \bar{q}, \bar{P}_{1:\nu}, \bar{\eta}_{1:\nu}, \bar{\xi}_{1:\nu-1})$. By using Theorem 3.1 we have that the true underlying \bar{Q} and \bar{q} , together with corresponding solution $\{\bar{P}_t \in \mathbb{S}^n\}_{t=1:\nu}$ and $\{\bar{\eta}_t \in \mathbb{R}^n\}_{t=1:\nu}$ to the Riccati recursions (3), and with $\bar{\xi}_t = \bar{g}_t^T \bar{\mathfrak{R}}_t^\dagger \bar{g}_t$ for $t = 1 : \nu - 1$, is a feasible solution to the optimization problem, if φ is large enough. For this feasible solution $(\bar{Q}, \bar{q}, \bar{P}_{1:\nu}, \bar{\eta}_{1:\nu}, \bar{\xi}_{1:\nu-1})$, we can decompose the corresponding \bar{H}_t as in (13) and in this case, together with (20), (A.2) can be written as

$$\begin{aligned} & \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} [\psi_{t,N}(Q, q; \mathbf{x}_t, \mathbf{u}_t)] \\ &= \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} \left[\frac{1}{2} \begin{bmatrix} \mathbf{u}_t^T & \mathbf{x}_t^T & 1 \end{bmatrix} \bar{H}_t \begin{bmatrix} \mathbf{u}_t \\ \mathbf{x}_t \\ 1 \end{bmatrix} - \frac{1}{2} \|\mathbf{u}_t\|^2 \right] \\ &= \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} \left[\frac{1}{2} \begin{bmatrix} \mathbf{u}_t^T & \mathbf{x}_t^T & 1 \end{bmatrix} \begin{bmatrix} \bar{\mathfrak{R}}_t \\ \bar{\mathfrak{S}}_t^T \\ \bar{g}_t^T \end{bmatrix} \bar{\mathfrak{R}}_t^\dagger \begin{bmatrix} \bar{\mathfrak{R}}_t & \bar{\mathfrak{S}}_t^T & \bar{g}_t^T \end{bmatrix} \begin{bmatrix} \mathbf{u}_t \\ \mathbf{x}_t^T \\ 1 \end{bmatrix} \right. \\ &\quad \left. - \frac{1}{2} \|\mathbf{u}_t\|^2 \right] \\ &= \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} \left[\underbrace{\frac{1}{2} \|(\bar{\mathfrak{R}}_t)^{\frac{1}{2}} (\bar{\mathfrak{R}}_t \mathbf{u}_t + \bar{\mathfrak{S}}_t \mathbf{x}_t + \bar{g}_t)\|^2}_{=0} - \frac{1}{2} \|\mathbf{u}_t\|^2 \right]. \end{aligned}$$

This shows that the lower bound for the objective function $\Psi(\cdot)$ is attained by $(\bar{Q}, \bar{q}, \bar{P}_{1:\nu}, \bar{\eta}_{1:\nu}, \bar{\xi}_{1:\nu-1})$.

Finally, we show that the ‘‘true’’ $(\bar{Q}, \bar{q}, \bar{P}_{1:\nu}, \bar{\eta}_{1:\nu}, \bar{\xi}_{1:\nu-1})$ is actually the unique global optimizer to (28). To this end, let $(Q^*, q^*, P_{1:N}^*, \eta_{1:N}^*, \xi_{1:N-1}^*)$ be an optimal solution to (28), and let us also use \star to denote other vectors

and matrices obtained using this optimal solution. Since the solution is optimal, it must be feasible, which implies that $H_t^* \succeq 0, \forall t = 1 : \nu - 1$. Hence it follows that $\mathfrak{R}_t^* \succeq 0$, $\ker(\mathfrak{R}_t^*) \subset [\ker(\mathfrak{S}_t^{*T}) \cap \ker(g_t^{*T})]$ and $H_t^* \mathfrak{R}_t^* \succeq 0$, see [19, Thm. 1.20, p. 43]. In view of the above ‘‘kernel containment’’ and (A.3), the optimal objective value can be further rewritten as

$$\begin{aligned} \Psi(Q^*, q^*, P_{1:\nu}^*, \eta_{1:\nu}^*, \xi_{1:\nu-1}^*) &= \sum_{N=2}^{\nu} \mathbb{P}(\mathbf{N} = N) \sum_{t=\nu-N+1}^{\nu-1} \\ &\mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} \left[\frac{1}{2} \begin{bmatrix} \mathbf{u}_t^T & \mathbf{x}_t^T & 1 \end{bmatrix} \begin{bmatrix} I & & \\ \mathfrak{S}_t^{*T} \mathfrak{R}_t^{*\dagger} & I & \\ g_t^{*T} \mathfrak{R}_t^{*\dagger} & & I \end{bmatrix} \right. \\ &\left. \begin{bmatrix} \mathfrak{R}_t^* & & \\ A^T P_{t+1}^* A + Q^* - P_t^* - \mathfrak{S}_t^{*T} \mathfrak{R}_t^{*\dagger} \mathfrak{S}_t^* & \beta_t^* - \mathfrak{S}_t^{*T} \mathfrak{R}_t^{*\dagger} g_t^* & \\ \beta_t^{*T} - g_t^{*T} \mathfrak{R}_t^{*\dagger} \mathfrak{S}_t^* & & \xi_t^* - g_t^{*T} \mathfrak{R}_t^{*\dagger} g_t^* \end{bmatrix} \right. \\ &\left. \begin{bmatrix} I & \mathfrak{R}_t^{*\dagger} \mathfrak{S}_t^{*\dagger} & \mathfrak{R}_t^{*\dagger} g_t^{*\dagger} \\ & I & \\ & & I \end{bmatrix} \begin{bmatrix} \mathbf{u}_t \\ \mathbf{x}_t \\ 1 \end{bmatrix} - \frac{1}{2} \|\mathbf{u}_t\|^2 \right] \end{aligned}$$

Recalling the notation $\tilde{\mathbf{x}}_t = [\mathbf{x}_t^T, 1]^T$ and the fact that $\mathfrak{R}_t^* \succeq 0$, we in turn get

$$\begin{aligned} \Psi(\cdot) &= \sum_{N=2}^{\nu} \mathbb{P}(\mathbf{N} = N) \sum_{t=\nu-N+1}^{\nu-1} \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} \left[-\frac{1}{2} \|\mathbf{u}_t\|^2 \right. \\ &+ \frac{1}{2} \begin{bmatrix} \mathbf{u}_t^T & \mathbf{x}_t^T & \mathfrak{S}_t^{*T} \mathfrak{R}_t^{*\dagger} + g_t^{*T} \mathfrak{R}_t^{*\dagger} & \mathbf{x}_t^T & 1 \end{bmatrix} \\ &\times \begin{bmatrix} \mathfrak{R}_t^* & & \\ & H_t^* \setminus \mathfrak{R}_t^* & \\ & & \begin{bmatrix} \mathbf{u}_t + \mathfrak{R}_t^{*\dagger} \mathfrak{S}_t^* \mathbf{x}_t + \mathfrak{R}_t^{*\dagger} g_t^* \\ \mathbf{x}_t \\ 1 \end{bmatrix} \end{bmatrix} \\ &= \sum_{t=1}^{\nu-1} \mathbb{E}_{\mathbf{x}_t, \mathbf{N}} \left[-\frac{1}{2} \|\mathbf{u}_t\|^2 \right] + \sum_{N=2}^{\nu} \mathbb{P}(\mathbf{N} = N) \sum_{t=\nu-N+1}^{\nu-1} \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} \left[\right. \\ &\left. \frac{1}{2} \|(\mathfrak{R}_t^*)^{\frac{1}{2}} (\mathbf{u}_t + \mathfrak{R}_t^{*\dagger} [\mathfrak{S}_t^* g_t^*] \tilde{\mathbf{x}}_t)\|^2 + \frac{1}{2} \text{tr}((H_t^* \setminus \mathfrak{R}_t^*) \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^T) \right], \end{aligned}$$

where $(\cdot)^{\frac{1}{2}}$ denotes the (uniquely defined) positive semi-definite matrix square root (see [18, Thm. 7.2.6]). Note that since $H_t^* \mathfrak{R}_t^* \succeq 0$, all terms except $\sum_{t=1}^{\nu-1} \mathbb{E}_{\mathbf{x}_t, \mathbf{N}} [-\frac{1}{2} \|\mathbf{u}_t\|^2]$ are non-negative. Hence, in order for the lower bound $\sum_{t=1}^{\nu-1} \mathbb{E}_{\mathbf{x}_t, \mathbf{N}} [-\frac{1}{2} \|\mathbf{u}_t\|^2]$ to be attained, we must have that

$$\mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} \left[\|(\mathfrak{R}_t^*)^{\frac{1}{2}} (\mathbf{u}_t + \mathfrak{R}_t^{*\dagger} [\mathfrak{S}_t^* g_t^*] \tilde{\mathbf{x}}_t)\|^2 \right] = 0, \quad (\text{A.4a})$$

$$\begin{aligned} \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} \left[\text{tr}((H_t^* \setminus \mathfrak{R}_t^*) \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^T) \right] &= \text{tr}(H_t^* \setminus \mathfrak{R}_t^*) \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=N} [\tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^T] \\ &= 0, \quad t = \nu - N + 1 : \nu - 1, \quad (\text{A.4b}) \end{aligned}$$

for all N such that $\mathbb{P}(\mathbf{N} = N) > 0$. In particular, by Assumption 2.4 it must be true for $\mathbf{N} = \nu$. From Lemma 4.3, we know that $\mathbb{E}_{\mathbf{x}_t | \mathbf{N}=\nu} [\tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^T] \succ 0$. Thus it follows that, for $N = \nu$, (A.4b) implies that $H_t^* \setminus \mathfrak{R}_t^* = 0$ holds for $t = 1 : \nu - 1$. By using the observation in Remark 3.2, we therefore have that $(Q^*, q^*, P_{1:\nu}^*, \eta_{1:\nu}^*)$ satisfies the generalized Riccati iterations (3).

Now, to show that the optimal solution to (28) is unique, first assume that $(Q^*, q^*, P_{1:\nu}^*, \eta_{1:\nu}^*)$ is an optimal solution such that $\mathfrak{R}_t^* \succ 0$, for $t = 1 : \nu - 1$. In this case, also $(\mathfrak{R}_t^*)^{\frac{1}{2}} \succ 0$ for $t = 1 : \nu - 1$. By (A.4a), this means that, conditioned on $\mathbf{N} = \nu$, we have $\mathbf{u}_t = -\mathfrak{R}_t^{*\dagger} [\mathfrak{S}_t^* g_t^*] \tilde{\mathbf{x}}_t$ a.s. for $t = 1 : \nu - 1$. But conditioned on $\mathbf{N} = \nu$, we also have that $\mathbf{u}_t = -\bar{\mathfrak{R}}_t^\dagger [\bar{\mathfrak{S}}_t \bar{g}_t] \tilde{\mathbf{x}}_t$. Therefore

$$\mathfrak{R}_t^{*\dagger} [\mathfrak{S}_t^* g_t^*] \tilde{\mathbf{x}}_t = \bar{\mathfrak{R}}_t^\dagger [\bar{\mathfrak{S}}_t \bar{g}_t] \tilde{\mathbf{x}}_t, \text{ a.s.},$$

for $t = 1 : \nu - 1$. Multiplying from the right with $\tilde{\mathbf{x}}_t^T$ and taking expectation $\mathbb{E}_{\mathbf{x}_t | \mathbf{N}=\nu}$ on both sides, we have that

$$\mathfrak{R}_t^{*\dagger} [\mathfrak{S}_t^* g_t^*] \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=\nu} [\tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^T] = \bar{\mathfrak{R}}_t^\dagger [\bar{\mathfrak{S}}_t \bar{g}_t] \mathbb{E}_{\mathbf{x}_t | \mathbf{N}=\nu} [\tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^T],$$

for $t = 1 : \nu - 1$. Using Lemma 4.3, we know that $\mathbb{E}_{\mathbf{x}_t | \mathbf{N}=\nu} [\tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^T] \succ 0$ and hence the matrix is full rank. Therefore, it must hold that

$$\mathfrak{R}_t^{*\dagger} [\mathfrak{S}_t^* g_t^*] = \bar{\mathfrak{R}}_t^\dagger [\bar{\mathfrak{S}}_t \bar{g}_t], \quad t = 1 : \nu - 1,$$

and thus, by (21), that

$$\tilde{A}_{cl}(t; Q^*, q^*) = \tilde{A}_{cl}(t; \bar{Q}, \bar{q}), \quad t = 1 : \nu - 1.$$

By Proposition 4.1, we therefore have that $\bar{Q} = Q^*$ and that $\bar{q} = q^*$. This implies that in the subset of the feasible region (28a)-(28c) where $\mathfrak{R}_t \succ 0$, for $t = 1 : \nu - 1$, it holds that $(\bar{Q}, \bar{q}, \bar{P}_{1:\nu}, \bar{\eta}_{1:\nu}, \bar{\xi}_{1:\nu-1})$ is the unique globally optimal solution to (28). Next, suppose that there exists a minimizer $(Q^*, q^*, P_{1:\nu}^*, \eta_{1:\nu}^*, \xi_{1:\nu-1}^*)$ such that $\mathfrak{R}_t^* \succeq 0$ but not strictly positive definite for some $t \in \{1, \dots, \nu - 1\}$. In particular, this means that $(Q^*, q^*, P_{1:\nu}^*, \eta_{1:\nu}^*, \xi_{1:\nu-1}^*) \neq (\bar{Q}, \bar{q}, \bar{P}_{1:\nu}, \bar{\eta}_{1:\nu}, \bar{\xi}_{1:\nu-1})$. Moreover, since (28) is a convex optimization problem that attains an optimal solution, the set of all optimal solutions is a nonempty convex set [38, Thm. 27.2]. This means that $(\alpha \bar{Q} + (1 - \alpha) Q^*, \alpha \bar{q} + (1 - \alpha) q^*, \{\alpha \bar{P}_t + (1 - \alpha) P_t^*\}_{t=1}^{\nu}, \{\alpha \bar{\eta}_t + (1 - \alpha) \eta_t^*\}_{t=1}^{\nu}, \{\alpha \bar{\xi}_t + (1 - \alpha) \xi_t^*\}_{t=1}^{\nu-1})$ are all optimal for all $\alpha \in [0, 1]$; for each $\alpha \in [0, 1]$ we denote the optimal solution $(Q^\alpha, q^\alpha, P_{1:\nu}^\alpha, \eta_{1:\nu}^\alpha, \xi_{1:\nu-1}^\alpha)$. Since the eigenvalues of $\mathfrak{R}_t, t = 1 : \nu - 1$, depends smoothly on P_t (see (4)), we can select α close enough to 1 so that $(Q^\alpha, q^\alpha, P_{1:\nu}^\alpha, \eta_{1:\nu}^\alpha, \xi_{1:\nu-1}^\alpha)$ will be such that $\mathfrak{R}_t^\alpha \succ 0$ for all $t = 1 : \nu - 1$. However, this contradicts

the fact that $(\bar{Q}, \bar{q}, \bar{P}_{1:\nu}, \bar{\eta}_{1:\nu}, \bar{\xi}_{1:\nu-1})$ is the unique globally optimal solution to (28) with $\Re_t \succ 0$, $t = 1 : \nu - 1$. Therefore, there can be no optimal solution such that \Re_t^* is not (strictly) positive definite for all $t \in \{1, \dots, \nu - 1\}$, and hence $(\bar{Q}, \bar{q}, \bar{P}_{1:\nu}, \bar{\eta}_{1:\nu}, \bar{\xi}_{1:\nu-1})$ is the unique globally optimal solution to (28). \square

References

- [1] Akshay Agrawal, Robin Verschueren, Steven Diamond, and Stephen Boyd. A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1):42–60, 2018.
- [2] R. McNeill Alexander. *Optima for animals*. Princeton University Press, Princeton, NJ, 1996.
- [3] Brian D O Anderson and John B Moore. *Optimal control: linear quadratic methods*. Dover publications, Mineola, NY, 2007.
- [4] MOSEK ApS. *The MOSEK optimization toolbox for MATLAB manual. Version 9.0.*, 2019.
- [5] Tamer Başar and Geert Jan Olsder. *Dynamic non-cooperative game theory*. Academic Press, London, UK, 1982.
- [6] H.H. Bauschke and P.L. Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, Cham, 2nd edition, 2017.
- [7] Dimitri P. Bertsekas. *Nonlinear Programming*. Athena scientific, Belmont, MA, 2nd edition, 1999.
- [8] Dimitri P. Bertsekas. *Dynamic programming and optimal control: Volume 1*. Athena scientific, Belmont, MA, 2nd edition, 2000.
- [9] Stephen Boyd, Laurent El Ghaoui, Eric Feron, and Venkataramanan Balakrishnan. *Linear matrix inequalities in system and control theory*. SIAM, Philadelphia, PA, 1994.
- [10] Shuping Chen, Xunjing Li, and Xun Yu Zhou. Stochastic linear quadratic regulators with indefinite control weight costs. *SIAM Journal on Control and Optimization*, 36(5):1685–1702, 1998.
- [11] Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.
- [12] Heinz W. Engl, Martin Hanke, and Andreas Neubauer. *Regularization of inverse problems*. Kluwer Academic Publishers, Dordrecht, 2000.
- [13] Augusto Ferrante and Lorenzo Ntogramatzidis. A note on finite-horizon LQ problems with indefinite cost. *Automatica*, 52:290–293, 2015.
- [14] Augusto Ferrante and Lorenzo Ntogramatzidis. A discussion on the discrete-time finite-horizon indefinite LQ problem. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 216–220. IEEE, 2016.
- [15] Israel Gohberg, Peter Lancaster, and Leiba Rodman. *Indefinite linear algebra and applications*. Birkhäuser Verlag, Basel, 2005.
- [16] Michael Grant and Stephen Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, pages 95–110. Springer, London, United Kingdom, 2008.
- [17] Michael Grant and Stephen Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, March 2014.
- [18] Roger A. Horn and Charles R. Johnson. *Matrix analysis*. Cambridge university press, New York, NY, 2013.
- [19] Roger A. Horn and Fuzhen Zhang. Basic properties of the Schur complement. In Fuzhen Zhang, editor, *The Schur Complement and Its Applications*, pages 17–46. Springer, Boston, MA, 2005.
- [20] Robert I Jennrich. Asymptotic properties of nonlinear least squares estimators. *The Annals of Mathematical Statistics*, 40(2):633–643, 1969.
- [21] Olav Kallenberg. *Foundations of modern probability*. Springer, 1997.
- [22] Rudolf E. Kalman. When is a linear control system optimal? *Journal of Basic Engineering*, 86(1):51–60, 1964.
- [23] Arezou Keshavarz, Yang Wang, and Stephen Boyd. Imputing a convex objective function. In *2011 IEEE international symposium on intelligent control*, pages 613–619. IEEE, 2011.
- [24] Yibei Li, Yu Yao, and Xiaoming Hu. Continuous-time inverse quadratic optimal control problem. *Automatica*, 117:108977, 2020.
- [25] Yibei Li, Han Zhang, Yu Yao, and Xiaoming Hu. A convex optimization approach to inverse optimal control. In *2018 37th Chinese Control Conference (CCC)*, pages 257–262. IEEE, 2018.
- [26] Bosen Lian, Wenqian Xue, Frank L Lewis, and Tianyou Chai. Robust inverse q-learning for continuous-time linear systems in adversarial environments. *IEEE Transactions on Cybernetics*, 2021.
- [27] Lennart Ljung. *System Identification (2nd Ed.): Theory for the User*. Prentice Hall PTR, USA, 1999.
- [28] Lennart Ljung and Torkel Glad. On global identifiability for arbitrary model parametrizations. *Automatica*, 30(2):265–276, 1994.
- [29] J. Löfberg. YALMIP : A toolbox for modeling and optimization in MATLAB. In *In Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004.
- [30] Andrew Y Ng and Stuart Russell. Algorithms for inverse reinforcement learning. In *Proceeding of the 17th International Conference on Machine Learning*, pages 663–670, 2000.
- [31] Kenneth Nordström. Convexity of the inverse and Moore–Penrose inverse. *Linear Algebra and its Applications*, 434(6):1489–1512, 2011.
- [32] Kenneth Nordström. A note on the convexity of the Moore–Penrose inverse. *Linear Algebra and its Applications*, 538:143–148, 2018.
- [33] M Cody Priess, Richard Conway, Jongeun Choi, John M Popovich, and Clark Radcliffe. Solutions to the inverse LQR problem with application to bi-

- ological systems analysis. *IEEE Transactions on control systems technology*, 23(2):770–777, 2014.
- [34] Mustapha Ait Rami, X Chen, and Xun Yu Zhou. Discrete-time indefinite LQ control with state and control dependent noises. *Journal of Global Optimization*, 23(3):245–265, 2002.
- [35] André CM Ran and Harry L Trentelman. Linear quadratic problems with indefinite cost for discrete time systems. *SIAM journal on matrix analysis and applications*, 14(3):776–797, 1993.
- [36] Karl J. Åström. *Introduction to stochastic control theory*. Dover, Mineola, NY, 2006. Unabridged republication of original published by Academic Press, 1970.
- [37] Robert E Reid, A Kemal Tugcu, and Barry C Mears. An optimal controller arising from minimization of a quadratic performance criterion of indefinite form. *IEEE transactions on automatic control*, 28(10):985–987, 1983.
- [38] R Tyrrell Rockafellar. *Convex analysis*. Princeton university press, Princeton, NJ, 1970.
- [39] Jun Shao. *Mathematical statistics*. Springer, New York, NY, 1999.
- [40] Alan Wilbor Starr and Yu-Chi Ho. Nonzero-sum differential games. *Journal of optimization theory and applications*, 3(3):184–206, 1969.
- [41] Noureddine Toumi, Roland Malhamé, and Jerome Le Ny. A tractable mean field game model for the analysis of crowd evacuation dynamics. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 1020–1025. IEEE, 2020.
- [42] Noureddine Toumi, Roland Malhamé, and Jerome Le Ny. A spatial partitioning based crowd evacuation model. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 5247–5252. IEEE, 2021.
- [43] Adrianus W. van der Vaart. *Asymptotic statistics*. Cambridge university press, Cambridge, United Kingdom, 1998.
- [44] Wenqian Xue, Patrik Kolaric, Jialu Fan, Bosen Lian, Tianyou Chai, and Frank L Lewis. Inverse reinforcement learning in tracking control based on inverse optimal control. *IEEE Transactions on Cybernetics*, 2021.
- [45] Wenqian Xue, Bosen Lian, Jialu Fan, Patrik Kolaric, Tianyou Chai, and Frank L Lewis. Inverse reinforcement q-learning through expert imitation for discrete-time systems. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [46] Chengpu Yu, Yao Li, Hao Fang, and Jie Chen. System identification approach for inverse optimal control of finite-horizon linear quadratic regulators. *Automatica*, 129:109636, 2021.
- [47] Han Zhang and Axel Ringh. Inverse linear-quadratic discrete-time finite-horizon optimal control for indistinguishable homogeneous agents: A convex optimization approach. *Automatica*, 148:110758, 2023.
- [48] Han Zhang, Axel Ringh, Weihang Jiang, Shaoyuan Li, and Xiaoming Hu. Statistically consistent inverse optimal control for linear-quadratic tracking with random time horizon. In *2022 41st Chinese Control Conference (CCC)*, pages 1515–1522, 2022.
- [49] Han Zhang, Jack Umenberger, and Xiaoming Hu. Inverse optimal control for discrete-time finite-horizon linear quadratic regulators. *Automatica*, 110:108593, 2019.
- [50] Xun Yu Zhou and Duan Li. Continuous-time mean-variance portfolio selection: A stochastic lq framework. *Applied Mathematics and Optimization*, 42(1):19–33, 2000.