SIAM J. CONTROL OPTIM.
Vol. 54, No. 2, pp. 662–689

SEPARATED DESIGN OF ENCODER AND CONTROLLER FOR NETWORKED LINEAR QUADRATIC OPTIMAL CONTROL*

MABEN RABI†, CHITHRUPA RAMESH‡, AND KARL H. JOHANSSON§

Abstract. For a networked control system, we consider the problem of encoder and controller design. We study a discrete-time linear plant with a finite horizon performance cost, comprising a quadratic function of the states and controls, and an additive communication cost. We study separation in design of the encoder and controller, along with related closed-loop properties such as the dual effect and certainty equivalence. The encoder outputs are quantized samples, but our results also apply to two other formats for encoder outputs: real-valued samples at event-triggered times, and real-valued samples over additive noise channels. If the controller and encoder are dynamic, then we show that the performance cost is minimized by a separated design: the controls are updated at each time instant as per a certainty equivalence law, and the encoder is chosen to minimize an aggregate quadratic distortion of the estimation error. This separation is shown to hold even though a dual effect is present in the closed-loop system. We also show that this separated design need not be optimal when the controller or encoder are to be chosen from within restricted classes.

Key words. networked control, linear quadratic optimal control, separation, dual effect, certainty equivalence, quantized control

AMS subject classifications. 93E20, 93E24, 94A29

DOI. 10.1137/14M0970987

1. Introduction.

We consider discrete-time sequential decision problems for a control loop that has a communication bottleneck between the sensor and the controller (Figure 1). The design problem is to choose in concert an encoder and a controller. The encoder maps the sensor's raw data into a causal sequence of channel inputs. Depending on the channel model adopted, the encoder performs either sequential quantization, sampling, or analog companding. The controller maps channel outputs into a causal sequence of control inputs to the plant. Such two-agent problems are generally hard because the information pattern is nonclassical, as the controller has less information than the sensor [45]. This gives scope for the controller to exploit any dual effect present in the loop, even when the plant is linear [14]. These two-agent problems are at the simpler end of a range of design problems arising in networked control systems [11, 3, 21, 1]. Naturally, one seeks formulations of these design problems as stochastic optimization problems whose solutions are tractable in some suitable sense.

The classical partially observed linear quadratic Gaussian (LQG) optimal control problem is a one-agent decision problem [46]. Given a linear, Gauss–Markov plant, one is asked for a causal controller, as a function of noisy linear measurements of the state, to minimize a quadratic cost function of states and controls. This problem has a simple and explicit solution, where the optimal controller "separates" into two policies;

*Received by the editors May 30, 2014; accepted for publication (in revised form) November 25, 2015; published electronically March 24, 2016. C.R. and K.H.J. greatfully acknowledge support from the Swedish Research Council and the Knut and Alice Wallenberg Foundation.
http://www.siam.org/journals/sicon/54-2/M97098.html
†School of Electrical Engineering and Computing, Indian Institute of Technology, Mandi, Himachal Pradesh 175005, India (mabenr@gmail.com).
‡Institut für Automatik, ETH Zürich, Zürich, Switzerland (rameshc@control.ee.ethz.ch).
§Automatic Control, KTH Royal Institute of Technology, Stockholm, SE 100 44, Sweden (kallej@kth.se).

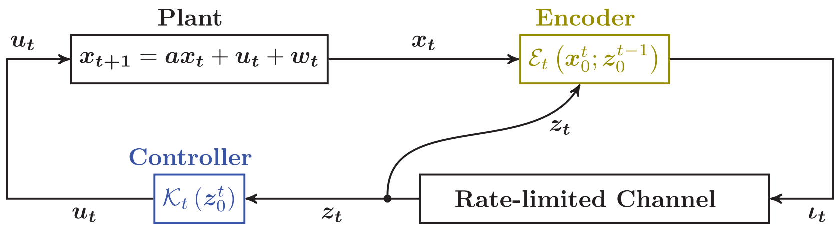

FIG. 1. *Control over a rate-limited channel with perfect feedback.*

one to generate a minimum mean-squared error (MMSE) estimate of the state from the noisy measurements, and the other to control the fully observed Gauss–Markov process corresponding to the estimate. A networked version of this problem is the following two-agent LQG optimal control problem [10]. Given a linear Gauss–Markov plant and a channel model, one is asked for an encoder and controller to minimize a performance cost which is a sum of a communication cost and a quadratic cost on states and controls. The communication cost is charged on decisions at the encoder, which are chosen to satisfy constraints imposed by the channel model. No causal encoding or control policies are, in general, excluded from consideration.

As in the one-agent version, a certain "separated" design is optimal, as has been suggested in various settings since the 1960s [24, 36, 16, 5, 29, 39, 27, 47, 31, 6, 30, 49]. To be precise, the following combination is optimal: certainty equivalence controls with an MMSE estimate of the state, and an encoder that minimizes a distortion for state estimation at the controller. The distortion is the average of a sum of squared estimation errors with time-varying coefficients depending on the coefficients of the performance cost. This separation is different from that obtained in the classical LQG problem, but it is still due to a linear evolution of the state, and the statistical independence of noises from all other current and past variables. As in the classical one-agent version [37, 35], the random variables need not be Gaussian.

1.1. Previous works. In the long history of the two-agent networked LQG problem, different channel models have been treated, leading to different types of encoders. We find in these works that the encoder is either a quantizer, an analog time-dependent compander, or an event-based sampler.

When a discrete alphabet channel is treated, the encoder is a time-dependent quantizer. Quantized control has been explored since the sixties, and structural results for this problem have seen spirited discussions over the years [24, 28, 16]. This problem was revisited by Borkar and Mitter [10] in recent years, setting off a new wave of interest. Surveys can be found in [31, 18]. For an additive noise channel, the encoder is a time-dependent, possibly nonlinear, compander. The corresponding networked LQG problem has been studied in [5], and more recently in [17, 19]. Analog channels with channel use restrictions lead to an encoder being an event-triggered sampler [2]. The networked LQG problem for event-triggered sampling is studied in [30].

The above papers suggested separated designs for the two-agent LQG problem with dynamic encoder and controller, and certainty equivalence controls. This is despite other results [13, 15], confirming the dual effect in the two-agent networked control problem. Thus, there can be an incentive to the controller to influence the estimation error, and yet the optimal controller chooses to ignore this incentive. Furthermore, for the two-agent LQG problem with event-triggered sampling, and with zero-order hold control between samples, Rabi and Johansson [33] showed through numerical computations that it is suboptimal to apply controls affine in the MMSE estimate. The optimal controls are instead nonlinear functions of the received samples. Thus, the literature does not tell us when separation holds, and when it does not, for the general class of two-agent problems.

1.2. Our contributions. We make three main contributions. First, we show that for the combination of a linear plant and nonlinear encoder, the dual effect is present. This confirms the results of Curry [13] and Feng and Loparo [15], by establishing through a counterexample that there is a dual effect in the closed-loop system. In fact, each of the three models we allow for the channel endow the loop with the dual effect. The dual role of the controller lies in reducing the estimation error in the future, using the predicted statistics of the future state and knowledge of the encoding policy. Due to this dual role, we show that, in general, separated designs need not be optimal for linear plants with nonlinear measurements, even with independent and identically distributed (IID) Gaussian noise and quadratic costs. Examples 5 and 6 show instances where the dual effect matters. Example 3 shows how the dual effect in the two-agent networked LQ problem renders useless the techniques that work for the classical, single-agent, partially observed LQ problem. These examples illustrate the insufficiency of arguments offered in [24, 36, 16, 5, 29, 39, 27, 47, 31, 6, 30, 49] for the optimality of separation and certainty equivalent controls.

Our second contribution is a proof for separation in one specific design problem. We prove that for the dynamic encoder-controller design problem, it is optimal to apply separation and certainty equivalence. A key instrument in our proof is the class of "controls-forgetting encoders" (introduced in section 4.2), which we show to be optimal despite it being a strict subset of the general class of state-based encoders. We also notice that the result holds under a variety of communication costs. For example, it holds even when the encoder is an analog compander with hard amplitude limits. Our proof does not require the dual effect to be absent. Hence, there is no contradiction with the fact that separation and certainty equivalence are not optimal for other design problems concerning the same plant-sensor combination. Our work also provides a direct insight into explaining separation or the lack of it, in the form of a property of the optimal cost-to-go function (Example 4 in section 6). Furthermore, we show that when this property does not hold separation is no longer optimal.

Our third contribution points out some important subtleties that arise when dynamic policies are involved. We explicitly demonstrate that with dynamic encoders for LQ optimal control, one cannot extend and apply a result of Bar-Shalom and Tse [7], which mandates absence of a dual effect for certainty equivalence to be optimal. The classical notion of a dual effect was introduced for static measurement policies, and the dual role of the controls has been motivated through the notion of a probing incentive [14]. We ask if the concept of probing applies unchanged for dynamic measurement policies and point out some subtleties in answering this question.

In recent years, there has been a resurgence of interest in problems related to dynamic and decentralized decision making in stochastic control. Old problems and results have been reexamined and reinterpreted to find new insights and develop new methods, such as the common information approach [26, 32]. Others, such as [23], have sought to reinterpret the proof techniques used in [4]. Following in the path of [44], many new counterexamples have been identified that show optimality of nonlinear strategies for control problems under nonclassical information patterns [25, 50]. Similarly, drawing from the many works on two-agent networked LQG problems [13, 15, 10, 31, 18], we have sought to understand why a structural simplification can be found in some dynamic decision problems, despite the nonclassical information pattern and the consequent presence of a dual effect.

1.3. Outline. The remainder of the paper is organized as follows. In section 2, we present a basic problem formulation, pertaining to encoder and controller design for data-rate limited channels. In section 3, we discuss the notion of a dual effect and certainty equivalence, and present a counterexample to establish that there is a dual effect in the considered networked control system. In section 4, we present a proof for separation in the two-agent networked LQ problem. We also indicate extensions to other channel models, including for event-triggered sampling and additive noise channels. In section 6, we present a number of examples to illustrate that, in general, separation does not hold for constrained design problems, followed by the conclusions in section 7.

2. Problem formulation.

In this section, we describe a version of the two-agent networked LQG problem, corresponding to a rate-limited channel model. We consider an instantaneous, error-free, discrete alphabet channel and the logarithm of the size of the alphabet is the bit rate. A control system that uses such a channel to communicate between its sensor and controller is depicted in Figure 1, and comprises four blocks. Each of these blocks, along with the performance cost, are described below, followed by a description of the design problems under consideration.

2.1. Plant. The plant state process $\{x_t\}$ is scalar, and its evolution law is linear:

$$x_{t+1} = ax_t + u_t + w_t \tag{2.1}$$

for $0 \leq t \leq T$. Here $\{u_t\}$ is the controls process, and $\{w_t\}$ is the plant noise process, which is a sequence of independent random variables with constant variance σ_w^2 and zero means. The initial state x_0 has a distribution with mean $\overline{x}_0$ and variance σ_0^2. At any time t, the noise w_t is independent of all state, control, channel input, and channel output data up to and including time t. We assume that the state process is perfectly observed by the sensor.

2.2. Performance cost. The performance cost is a sum of the quadratic cost on states and controls, and a communication cost charged on encoder decisions:

$$J = \mathbb{E}\left[x_{T+1}^2 + p\sum_{i=1}^{T} x_i^2 + q\sum_{i=0}^{T} u_i^2\right] + J^{\text{Comm}}, \tag{2.2}$$

where $p > 0$ and $q > 0$ are suitably chosen scalar weights for the squares of the states and controls, respectively. The communication cost J^{Comm} is an average quantity that depends on the encoding and control policies, and the channel model adopted.

2.3. Channel model. The channel model refers to an input-output description of the communication link from the sensor to the controller. We denote the channel input at time t by ι_t, the corresponding output by z_t, and the encoding map generating ι_t by $\mathcal{E}_t$. In Figure 1, we consider an ideal, discrete alphabet channel that faithfully reproduces inputs, and thus, $\iota_t \equiv z_t$. The encoder's job is to pick, at every time t, the encoding map $\mathcal{E}_t$ producing a channel output letter from the preassigned finite alphabet $z_t \in \{1, \ldots, N\}$, where the nonnegative integer N is the preassigned size of the channel alphabet. Since the alphabet is fixed, we have a hard data-rate constraint at every time. Hence there is no explicit cost attached to communication, so $J^{\text{Comm}} \equiv 0$ in this case. Our results also extend to other channel models that permit the data rate or energy needed for each transmission to be chosen causally by the encoder.

2.4. Controller. The control signal u_t is real valued and is to be computed by a causal policy based on the sequence of channel outputs. The controller has perfect memory, and thus remembers all of its past actions, and the causal sequence of channel outputs. Thus, in general, at every time t the controller's map takes the form

$$\mathcal{K}_t \;:\; \left\{t, \{z_i\}_0^t, \{u_i\}_0^{t-1}\right\} \mapsto u_t.$$

2.5. Encoder. At all times, the encoder knows the entire set of control policies employed by the controller and the statistical parameters of the plant. With this prestored knowledge, the encoder works as a causal quantizer mapping the sequence of plant outputs. Thus, the encoder's map takes the form

$$\mathcal{E}_t \;:\; \left\{t, \{x_i\}_0^t, \{z_i\}_0^{t-1}, \{\mathcal{K}_i(\cdot)\}_0^{t-1}\right\} \mapsto \; z_t.$$

Notice that we do not allow the encoder to directly view the sequence of inputs to the plant. This subtle point plays an important role in the discussion in section 7.

2.6. Design problems. For a given information pattern, different design spaces may arise due to engineering heuristics, hardware or software limitations, etc. Any such design space is a subset of the set of all admissible encoder and controller pairs. We identify three design problems, each associated with its own design space. For these design problems, an adopted channel model can be either the one described in section 2.3, or any of the models from section 5. First, we pose a single-agent design problem which has a classical information pattern.

Design problem 1 (controller-only design). *For the linear plant* (2.1), *the adopted channel model, and a given admissible set of encoding policies*

$$\{\mathcal{E}_t^{\dagger}(\,\cdot\,; \{z_i\}_0^t, \{u_i\}_0^{t-1})\}_0^T,$$

the controller-only design problem asks for a causal sequence of control policies $\{\mathcal{K}_t\}_0^T$ *to minimize the performance cost* (2.2).

Next we pose a design problem where the design space is the largest possible non-randomized set of admissible encoder-controller pairs. We consider every causally time-dependent encoder and controller. In other words, for this type of design problem, regardless of the choices one makes for channel and communication cost, at any time the controller may pick the control input using all channel outputs up till then.

Design problem 2 (dynamic encoder-controller design). *For the linear plant* (2.1) *and the adopted channel model, the dynamic encoder-controller design problem requires one to pick causal sequences of encoding and control policies* $\{\mathcal{E}_t\}_0^T, \{\mathcal{K}_t\}_0^T$ *to minimize the performance cost* (2.2).

Next we pose a design problem where the controller and encoder must respect a restriction on selecting the control signals or encoding maps. At every time, the control values must be chosen from a restricted set $\mathcal{U}$, such as the interval $(-1, 1)$ or the finite set $\{-1, 0, 1\}$. Likewise, the encoding maps have to be chosen from within restricted sets. For example, the encoding maps may be constrained to consist of two quantization cells $(-\infty, \theta), (\theta, \infty)$, where the encoder threshold θ must be chosen from a restricted set Θ, say the interval $(-5, 5)$. Subject to these constraints, the controller and encoder policies are still to be dynamically chosen.

Design problem 3 (constrained encoder-controller design). *For the linear plant* (2.1), *and the adopted channel model, the constrained encoder-controller design problem requires one to pick causal sequences of encoding and control policies*

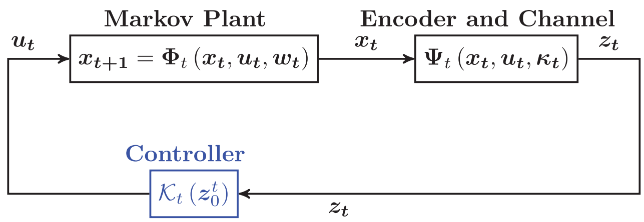

FIG. 2. *Setup for definitions in section* 3.

$\{\mathcal{E}_t\}_0^T, \{\mathcal{K}_t\}_0^T$, *subject to the constraints represented by* $\theta \in \Theta$ *and* $u_k \in \mathcal{U}$, *to minimize the performance cost* (2.2).

For all three design problems presented above, we assume the existence of measurable policies minimizing the associated costs. We avoid investigating the necessary technical qualifications except to say that if need be, one may allow randomized polices, or even reject the class of merely measurable policies in favor of the class of universally measurable policies [9].

3. Dual effect and certainty equivalence. We begin by presenting a definition of dual effect [14] and certainty equivalence [22]. We then present an example to establish that there is a dual effect of the controls in the networked control system introduced in section 2.

3.1. Dual effect. In a feedback control loop, the dual effect is an effect that the controller may see in the rest of the loop. When it is present, the control laws affect not just the first moment, but also second, third, and higher central moments of the controller's nonlinear filter for the state. Below, we state this formally for a controlled Markov process with partial observations available to the controller:

$$x_{t+1} = \Phi_t\big(x_t, u_t, w_t\big), \qquad z_t = \Psi_t\big(x_t, u_t, \kappa_t\big), \tag{3.1}$$

where the sequences $\{x_t\}$ and $\{u_t\}$ are the real-valued plant state and control processes, respectively; see Figure 2. The sequence $\{z_t\}$ is the observation process and the sequences $\{w_t\}$ and $\{\kappa_t\}$ are the plant noise and observation noise processes, respectively. Assume that all the primitive random variables are defined on a suitable probability triple $[\Omega, \mathcal{F}, \mathcal{P}]$. Now consider two arbitrary admissible sets of control policies, $\{\mathcal{K}(t,\cdot)\}, \{\widetilde{\mathcal{K}}(t,\cdot)\}$. Once we pick one such set of control policies, this together with the measure $\mathcal{P}$ defines the states, observations, and controls as random processes. The choice of policies fixes their statistics. We can advertise this relationship by (1) specifying random variables, x_t, for example, in the form $x_t(\omega;\mathcal{K})$, (2) specifying a filtration, for example, the one generated by the z-process as $\mathcal{F}^{\mathcal{K},z}$, or (3) specifying an expected value of a functional, $\mathbb{E}[F_t]$, for example, in the form

$$\mathbb{E}_{\mathcal{P},\mathcal{K}}\left[F_t\Big(t, \{x_i(\omega;\mathcal{K})\}_0^t, \{z_i(\omega;\mathcal{K})\}_0^t, \{u_i(\omega;\mathcal{K})\}_0^t\Big)\right],$$

where ω stands for any element of the sample space of the primitive random variables. To minimize the notational burden, we advertise the dependence on the set of control policies only as needed. We now define the dual effect by defining its absence.

DEFINITION 3.1 (dual effect). *The networked control system in Figure* 2 *is said to have* no dual effect of second order *if*

1. *for any two sets* $\mathcal{K}, \widetilde{\mathcal{K}}$ *of admissible control policies, and*
2. *for any two time instants* t, s,

we have $\mathcal{F}_t^{\mathcal{K},z} = \mathcal{F}_t^{\widetilde{\mathcal{K}},z}$ *for every* t*, and for any given event* $X \in \mathcal{F}_t^{\mathcal{K},z}$*,*

$$\mathbb{E}_{\mathcal{P},\mathcal{K}}\left[\Big(x_t(\omega;\mathcal{K}) - \mathbb{E}_{\mathcal{P},\mathcal{K}}\big[x_t(\omega;\mathcal{K})\big|\big\{z_i(\omega;\mathcal{K})\big\}_0^s, \omega\in X\big]\Big)^2\Big|\big\{z_i(\omega;\mathcal{K})\big\}_0^s, \omega \in X\right]$$
$$= \mathbb{E}_{\mathcal{P},\widetilde{\mathcal{K}}}\left[\Big(x_t(\omega;\widetilde{\mathcal{K}}) - \mathbb{E}_{\mathcal{P},\widetilde{\mathcal{K}}}\big[x_t(\omega;\widetilde{\mathcal{K}})\big|\big\{z_i(\omega;\widetilde{\mathcal{K}})\big\}_0^s, \omega\in X\big]\Big)^2\Big|\big\{z_i(\omega;\widetilde{\mathcal{K}})\big\}_0^s, \omega \in X\right].$$

Thus, we require equality of the two sets of covariances of filtering, prediction, and smoothing errors, corresponding to any two choices of control strategies. In the definition above, by choosing one set of control policies, say $\widetilde{\mathcal{K}}$ as resulting in $u_t = 0$, for all t, we obtain the definition of Bar-Shalom and Tse [7].

3.2. Certainty equivalence. For the controlled Markov process (3.1), consider

$$J^{\text{general}} = \mathbb{E}\left[L\Big(\{x_i\}_1^{T-1}, \{u_i\}_0^T\Big)\right]$$

to be the objective function, where L is a given deterministic, nonnegative cost function. Imagine that a muse could at time t supply to the controller the exact values of all primitive random variables by informing the controller of the exact element ω of the sample space Ω. With such complete and acausal information, the controller could, in principle, solve the deterministic optimization problem

$$\inf_u J_t\left(u;\omega\right) = \inf_u L\Big(\{x_i\left(\omega\right)\}_0^T, \{u_i\left(\omega\right)\}_0^{t-1}, u, \{u_i\left(\omega\right)\}_{t+1}^T\Big).$$

Let $u_t^*(\omega)$ be an optimal control law for this deterministic optimization problem. We now state the definition of certainty equivalence from van der Water and Willems [40].

DEFINITION 3.2. *A certainty equivalence control law for the plant* (2.1) *with the performance cost* (2.2) *has the form*

$$\mathbb{E}\left[u_t^*\left(\omega\right)\Big|\{z_i\left(\omega\right)\}_0^t, \{u_i\left(\omega\right)\}_0^{t-1}\right].$$

Clearly, this law is causal. Notice also that its form is tied to the performance cost, and to the statistics of the state and observation processes. It is possible for certainty equivalence control laws to be nonlinear, and such laws can be optimal even when separated designs may not be. For linear plants, they can sometimes be linear or affine, as indicated by the following proposition from [40] adapted to our problem.

LEMMA 3.3 (affine certainty equivalence laws for linear plants). *For the plant* (3.1)*, with* $\Phi_t = ax_t + u_t + w_t$*, and the quadratic performance cost* (2.2) *with* $J^{Comm} = 0$*, the following are certainty equivalence laws:*

$$u_t^{CE} = -\,k_t^{CE}\Big(a\cdot\mathbb{E}\left[x_t\Big|\{z_i\}_0^t, \{u_i\}_0^{t-1}\right] + \mathbb{E}\left[w_t\Big|\{z_i\}_0^t, \{u_i\}_0^{t-1}\right]\Big),$$

where $k_i^{CE} = \beta_{i+1}/(q+\beta_{i+1})$*,* $\alpha_i = \beta_{i+1}+\alpha_{i+1}$*,* $\beta_i = p + a^2 q\beta_{i+1}/(q+\beta_{i+1})$*,* $\alpha_{T+1} = 0$*, and* $\beta_{T+1} = 1$*.*

DEFINITION 3.4 (certainty equivalence property). *The certainty equivalence property holds for a stochastic control problem if it is optimal to apply the certainty equivalence control law.*

For the stochastic control problem described in Lemma 3.3, with nonlinear measurements that do not result in a dual effect of the controls, Bar-Shalom and Tse [7] showed that the certainty equivalence property holds.

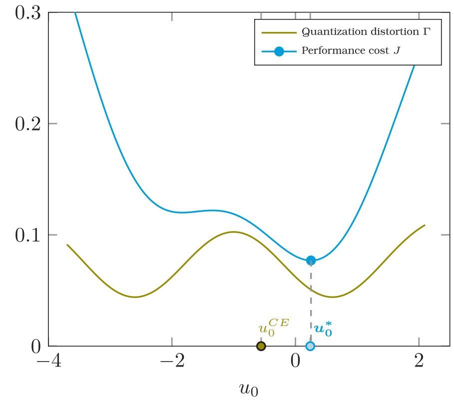

FIG. 3. *Plot for Example* 1.

We now consider a simple example, and show that there is a dual effect of the control signal in the closed-loop system presented in section 2.

Example 1. For the plant (2.1), let $a = 1$, $x_0 = 2$, and $\sigma_0 = 0$. Let this information be known to the encoder and the controller, which simply means that $z_0 = x_0$. Let the variance $\sigma_w^2 = 0.7^2$. For the objective function, let the horizon end $T = 1$, and let $p = q = 0.01$. Let the channel alphabet be the discrete set $\{1, 2, 3\}$. For the given threshold $\theta = 1.6$, let the encoder at $t = 1$ be

$$\xi_1(x_1) = \begin{cases} 1 & \text{if } x_1 \in (-\infty, -\theta), \\ 2 & \text{if } x_1 \in (-\theta, \theta), \\ 3 & \text{if } x_1 \in (\theta, +\infty). \end{cases} \tag{3.2}$$

The optimal control law at $t = 1$ is $u_1 = -a\,\widehat{x}_{1|1}/(q+1)$, where $\widehat{x}_{1|1} = \mathbb{E}[x_1 | x_0, u_0, z_1]$. Using the encoding policy ξ_1 and the optimal control signal u_1, the performance cost with $J^{\text{Comm}} = 0$ can be written as a function of the control at $t = 0$:

$$J(u_0) = \sigma_w^2 + qu_0^2 + \left(p + \frac{qa^2}{q+1}\right)\mathbb{E}\left[x_1^2 \,|x_0, u_0\right] + \overbrace{\frac{a^2}{q+1}\mathbb{E}\left[\left(x_1 - \widehat{x}_{1|1}\right)^2 \Big| x_0, u_0, z_1\right]}^{\triangleq\, \Gamma},$$

where Γ is the quantization distortion, which is thus proportional to the conditional variance of the controller's MMSE in estimating x_1. Notice that Γ is a function of u_0, thus resulting in a dual effect of the control signal in the plant-encoder-channel combination. Figure 3 shows how the quantization distortion Γ depends on u_0. The total cost J is also plotted. The optimal value u_0^* is shown to be different from the certainty equivalent control u_0^{CE}. ▯

4. Dynamic encoder-controller design. In this section we solve the dynamic encoder-controller design problem (Design problem 2). We work out the details for the discrete alphabet channel with the fixed alphabet size N. We begin by examining a known structural property of optimal encoders. This states that it is optimal for the encoder to apply a quantizer on the state x_t, with the shape of the quantizer depending only on past quantizer outputs. Next, we present a structural property for encoders called *controls forgetting*, which leads to separation. Finally, we show that the optimal encoder for Design problem 2 is indeed controls forgetting and it leads to separation and certainty equivalence.

4.1. Known structural properties of optimal encoders. Let us now formulate the encoder's Markov decision problem. Fix the control policies to be $u_t = \mathcal{K}_t^\dagger(\{z_i\}_0^t)$, where for each time t, $\mathcal{K}_t^\dagger$ is a prescribed admissible law. Then the optimization problem reduces to one of picking encoding policies. This is a single-agent, sequential decision problem, and hence one with a classical information pattern. The action space for this decision problem is the infinite dimensional function space of discrete-valued encoders. At time t, the encoder takes as input the current and previous states, all previous outputs, and all previous encoding maps. For convenience, we can view this encoding map as a function of only the current state but with the rest of the inputs considered as parameters determining the form of this function. Thus, without loss of generality the encoder can be described as the function $\xi_t(\cdot) : \mathbb{R} \to \{1, \ldots, N\}$ having x_t as its argument with its shape determined by $(\{x_i\}_0^{t-1}\{z_i\}_0^{t-1}\{\xi_i(\cdot)\}_0^{t-1})$. Hence the action space at times t can be described as

$$\Big\{\xi\left(\cdot\right) : \mathbb{R} \to \{1, \ldots, N\}\,,\ \text{Borel measurable}\Big\}.$$

Identifying encoders as decisions to be picked is not enough, as the signal x_t need not be Markov. We utilize the following property, proved by Striebel [38] and Varaiya and Walrand [42]: for every design problem we set up in section 2, the signals x_t, $\{z_i\}_0^t$, $\{\xi_i(\cdot)\}_0^{t-1}$ form sufficient statistics for the encoding decision at time t. Hence, at every time t, performance is not degraded by the encoder choosing to quantize just x_t instead of quantizing the entire waveform $\{x_0, \ldots, x_t\}$. The shape of the quantizer is allowed to vary with past encoder shapes, past encoder outputs, and past control inputs. But given the sufficient statistics, the encoder can forget the data $\{x_0, \ldots, x_{t-1}\}$.

Denote by $\mathcal{D}_{t^-}^{\text{con}}$ the data at the controller just after it has read the channel output z_t and just before it has generated the control value u_t. Similarly denote by $\mathcal{D}_{t^+}^{\text{con}}$ the data at the controller just after it has generated the control value u_t. Then

$$\mathcal{D}_{t^-}^{\text{con}} = \Big\{\{z_i\}_0^t, \{\xi_i\left(\cdot\right)\}_0^t, \{u_i\}_0^{t-1}\Big\}, \quad \mathcal{D}_{t^+}^{\text{con}} = \Big\{\mathcal{D}_{t^-}^{\text{con}}, u_t\Big\} = \Big\{\{z_i\}_0^t, \{\xi_i\left(\cdot\right)\}_0^t, \{u_i\}_0^t\Big\},$$

and let $\widehat{x}_{t|t} \triangleq \mathbb{E}[x_t|\mathcal{D}_{t^-}^{\text{con}}]$.

The problem we consider has two decision makers that jointly minimize a given cost function. The information available to these decision makers is not the same, and neither is the information available to each agent a subset of the information available to the agent downstream in the feedback loop. Thus, the information pattern here is neither classical nor nested. We apply the *common information approach*[1] to our problem. This approach allows a designer to treat a problem with multiple decision makers as a classical control problem with a single decision maker that has access to partial state information. When applied to our setup, this approach leads to the following structural result at the encoder: the encoding policy $\xi_t\left(\cdot\right)$ is selected based on the information available to the controller at the previous time instant, namely, $\mathcal{D}_{(t-1)^+}^{\text{con}}$. At times t^- and t^+, the data $\mathcal{D}_{(t-1)^-}^{\text{con}}$ and $\mathcal{D}_{(t-1)^+}^{\text{con}}$, respectively, comprise the common information in this problem. The encoding map $\xi_t\left(\cdot\right)$ is applied

[1]This approach can be traced back to a conjecture by Witsenhausen in [45] regarding a structure for the optimal control laws in the n-step delayed sharing information pattern problem. This conjecture was shown to be true by Varaiya and Walrand in [41] for $n = 1$ and false for $n > 1$. Our terminology is derived from [26, 32], where this conjecture has been resolved and developed into a sequential approach for problems with nonclassical information patterns.

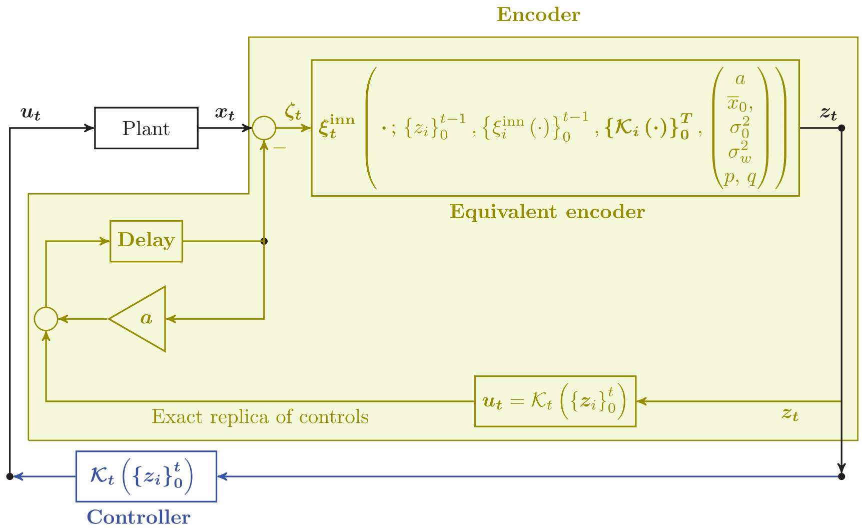

FIG. 4. *The block diagram of Figure 1 with innovation encoding.*

to the state x_t, which is private information available to the encoder. A similar approach has been used before for problems of quantized control [12, 43, 48].

4.2. Controls-forgetting encoders and separation. We now present a structural property of encoders which ensures separation in design. Recall the plant (2.1) and cost (2.2), and define ζ_t as the control free part of the state $\zeta_0 = x_0$, $\zeta_{i+1} = x_{i+1} - \sum_{j=0}^{i} a^{i-1} u_j$ for $i \geq 0$. At the encoder, the change of variables

$$\left(x_t, \{z_i\}_0^{t-1}; \{\mathcal{K}_i(\cdot)\}_0^T\right) \longmapsto \left(\zeta_t, \{z_i\}_0^{t-1}; \{\mathcal{K}_i(\cdot)\}_0^T\right) \tag{4.1}$$

is causal and causally invertible. Hence the statistics $(\zeta_t, \{z_i\}_0^{t-1}; \{\mathcal{K}_i(\cdot)\}_0^T)$ are also sufficient statistics at the encoder.

DEFINITION 4.1 (innovation encoder [10]). *We say an admissible encoder is an "innovation" encoder if its inputs are* $(\zeta_t, \{z_i\}_0^{t-1}; \{\mathcal{K}_i(\cdot)\}_0^T)$ *and its output is* z_t.

The networked control system in Figure 1, redrawn with an innovation encoder, is shown in Figure 4. Here, the control free part of the state is not affected by the control policies, but obeys the recursion $\zeta_{t+1} = a\zeta_t + w_t$. For any sequence of causal encoders, one can find an equivalent sequence of innovation encoders such that when these two sets operate on the same sequence of plant outputs, they produce two sequences of channel inputs that are equal with probability one. Hence, if for a plant and channel, the dual effect is present in a certain class of causal encoders, then the dual effect is also present in the equivalent class of innovation encoders [15]. This is what the following example illustrates.

Example 2 (dual effect in a loop with fixed innovation encoder). Consider Example 1, but with the encoder replaced by an innovation encoder. For the given threshold $\theta = 1.6$, let the encoder at time $t = 1$ be the following innovation encoder:

$$\xi_1^{\text{inn}}(\zeta_1) = \begin{cases} 1 & \text{if } a\zeta_1 + \mathcal{K}_0(z_0) \in (-\infty, -\theta), \\ 2 & \text{if } a\zeta_1 + \mathcal{K}_0(z_0) \in (-\theta, \theta), \\ 3 & \text{if } a\zeta_1 + \mathcal{K}_0(z_0) \in (\theta, +\infty). \end{cases} \tag{4.2}$$

Notice from (3.2) and (4.2) that this innovation encoder ξ_t^{inn} is equivalent to the causal encoder ξ_t of Example 1. For the same applied control policy $\mathcal{K}_0$, and for the same realizations of primitive random variables, we get $\xi_1^{\text{inn}}(\zeta_1(\omega)) = \xi_1(x_1(\omega))$. The results in Figure 3 apply here also because for an event $X \in \mathcal{F}^{(x_0,z_1)}$, we have

$$\mathbb{P}\left[x_1 \in X \left| x_0,\, z_1 = \xi_1^{\text{inn}}(\zeta_1)\right.\right] = \mathbb{P}\left[x_1 \in X \left| x_0,\, z_1 = \xi_1(x_1)\right.\right]. \qquad \square$$

The encoder (quantizer) in the loop causes the dual effect. Furthermore, the encoder's presence renders useless the techniques that worked in the case of the classical, single-agent, partially observed LQ control problem. The next example illustrates this.

Example 3. We examine a scalar system as it evolves from time step 0 to time step 1. We have $x_0 \sim \mathcal{N}(\mu_0, \sigma_0)$,

$$x_1 = x_0 + u_0 + w_0,$$

where w_0 is the process noise variable which is independent of x_0, u_0, and $w_0 \sim \mathcal{N}(0, \sigma_w)$. We adopt the specific quantizing strategy given below (on the left in the form of an encoder for x_t, and on the right, in the equivalent, innovation form):

$$\xi_0(x_0) = \begin{cases} -1 & \text{if } x_0 \le 0, \\ +1 & \text{if } x_0 > 0, \end{cases} \qquad \xi_0^{\text{inn}}(\zeta_0) = \begin{cases} -1 & \text{if } \zeta_0 \le 0, \\ +1 & \text{if } \zeta_0 > 0, \end{cases}$$

$$\xi_1(x_1) = \begin{cases} -1 & \text{if } x_1 \le 0, \\ +1 & \text{if } x_1 > 0, \end{cases} \qquad \xi_1^{\text{inn}}(\zeta_1) = \begin{cases} -1 & \text{if } \zeta_1 \le -u_0(z_0), \\ +1 & \text{if } \zeta_1 > -u_0(z_0). \end{cases}$$

Since the encoder at time 0 is binary, the general control law at time 0 has the form

$$u_0(z_0) = \begin{cases} \alpha & \text{if } z_0 = -1, \\ \beta & \text{if } z_0 = +1, \end{cases}$$

where α, β are arbitrary real numbers. The process $\widehat{x}_{t|t}$ is fully observed at the controller. We have $\widehat{x}_{0|0} = \mathbb{E}[x_0|z_0]$, and as noted in [49], one can write

$$\widehat{x}_{1|1} = \widehat{x}_{0|0} + u_0 + \overline{w}_0, \tag{4.3}$$

where the noiselike random variable $\overline{w}_0$ is given by $\overline{w}_0 \triangleq \mathbb{E}\left[x_1 \left| z_0, z_1\right.\right] - \mathbb{E}\left[x_1 \left| z_0\right.\right]$. Then one can treat the problem as the control of the fully observed process $\widehat{x}_{t|t}$ to minimize the given cost, which can be rewritten as the following sum of two terms:

$$J = \mathbb{E}\left[\widehat{x}_{1|1}^2 + p \cdot \widehat{x}_{0|0}^2 + p \cdot \left(x_0 - \widehat{x}_{0|0}\right)^2 + q u_0^2\right] + \mathbb{E}\left[\left(x_1 - \widehat{x}_{1|1}\right)^2\right]. \tag{4.4}$$

For the single-agent partially observed LQ control problem, transforming the problem to the form in (4.4) results in two special things: (1) the random process $\{\overline{w}_t\}$ is statistically independent of the control process $\{u_t\}$ and of the "state" process $\{\widehat{x}_{t|t}\}$, and (2) because the dual effect is absent, the second term on the right-hand side (RHS) of (4.4) does not vary with $\{u_t\}$. Therefore, by considering $\{\widehat{x}_{t|t}\}$ as the process to be controlled, we get a single-agent, fully observed LQ control problem.

In the two-agent problems considered in this paper, neither of the above-mentioned special things may happen. For this specific example, we show in Figure 5 how the second moments of $\overline{w}_0$ and $x_1 - \widehat{x}_{1|1}$ vary with u_0. Hence, the dual effect persists in the two-agent networked problem with innovation encoding. $\square$

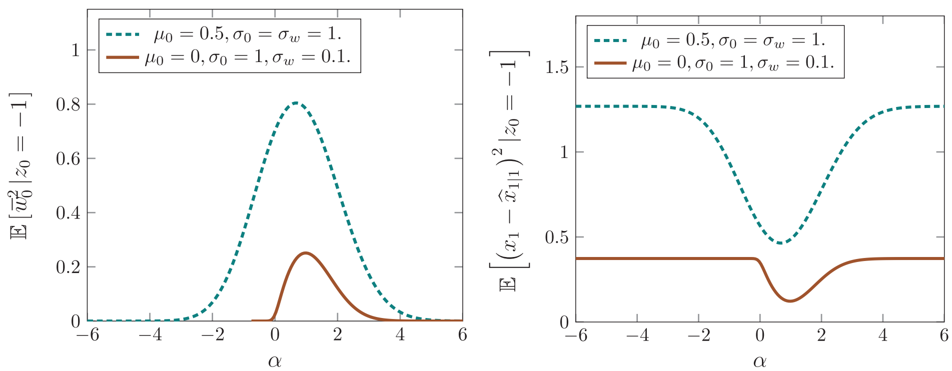

FIG. 5. *Plots for Example* 3.

Next we define a class of encoders for which at prescribed times t, the statistics of $\overline{w}_t$, $x_{t+1} - \widehat{x}_{t+1|t+1}$ are independent of the control u_t.

DEFINITION 4.2 (controls-forgetting encoder). *Denote by* $\rho^{\zeta}_{\tau|\tau-1}(\cdot)$ *the conditional density of* ζ_τ *given the data* $\mathcal{D}^{\text{con}}_{(\tau-1)^-}$. *An admissible encoding strategy is controls forgetting from time* τ *if it takes the form*

$$\xi_t^{CF,\tau}\left(x_t; \mathcal{D}^{\text{con}}_{(t-1)^-}\right) = \begin{cases} \xi_t^{\dagger}\left(x_t; \mathcal{D}^{\text{con}}_{(t-1)^-}\right) & \text{if } t \le \tau, \\ \epsilon_t\left(\zeta_t;\, \rho^{\zeta}_{\tau|\tau-1}(\cdot),\, \{z_i\}_{\tau}^{t-1}, \{\epsilon_i(\cdot)\}_{\tau}^{t-1}\right) & \text{if } t \ge \tau+1, \end{cases}$$

where (1) $\xi_t^{\dagger}(\cdot; \mathcal{D}^{\text{con}}_{(t-1)^-})$ *is any admissible policy for encoding at time* t, (2) *for* $t \ge \tau+1$ *the policies* $\epsilon_t(\cdot; \rho^{\zeta}_{\tau|\tau-1}(\cdot), \{z_i\}_{\tau}^{t-1}, \{\epsilon_i(\cdot)\}_{\tau}^{t-1})$ *are adapted to the data*

$$\mathcal{D}^{CF,\tau}_{(t-1)^+} = \left(\rho^{\zeta}_{\tau|\tau-1}(\cdot), \{z_i\}_{\tau}^{t-1}, \{\epsilon_i(\cdot)\}_{\tau}^{t-1}\right) \subset \mathcal{D}^{\text{con}}_{(t-1)^+} \quad \textit{for } t \ge \tau,$$

and (3) *for fixed values of the data* $\mathcal{D}^{CF,\tau}_{(t-1)^+}$, *the map* $\epsilon_t(\cdot)$ *produces the same output regardless of both the controls* $\{u_i\}_{\tau}^{t}$ *and the control policies* $\{\mathcal{K}_i(\cdot)\}_{t+1}^{T}$.

Clearly such controls-forgetting encoders exist. For example, consider a set of encoders that quantize in sequence $\zeta_{\tau+1}, \ldots, \zeta_T$ to minimize the estimation distortion $\sum_{i=\tau+1}^{T} \mathbb{E}[(\zeta_i - \widehat{\zeta}_{i|i})^2]$, where $\widehat{\zeta}_{i|i} = \mathbb{E}[\zeta_i | \mathcal{D}^{CF,\tau}_{(i-1)^+}]$. Let the nonnegative function $\psi(\cdot)$ represent some notion of cost. For example, $\psi(x) := x^2$.

LEMMA 4.3 (distortions incurred by controls-forgetting encoders also forget controls). *Fix the time* $t = \tau$ *and the distortion measure* ψ. *If the encoder is controls forgetting from time* τ, *then for times* $i \ge \tau+1$, *the distortions* $\mathbb{E}[\psi(x_i - \widehat{x}_{i|i}) | \mathcal{D}^{\text{con}}_{i^-}]$ *are statistically independent of the partial set of controls* $\{u_i\}_{i=\tau}^{T}$.

Proof. The unconditional statistics of $\{\zeta_t\}$ are independent of the entire control waveform, no matter what the encoder is. For times $i \ge \tau+1$ and for sets $X \in \mathcal{F}_i^z$, $\mathbb{P}[\zeta_i \in X | \mathcal{D}^{\text{con}}_{i^-}]$ is independent of $\{u_i\}_{i=\tau}^{T}$ because the encoding maps ξ_i are controls forgetting from time τ. Since $\zeta_t - \widehat{\zeta}_{t|t} = x_t - \widehat{x}_{t|t}$, for all t, the lemma follows. ∎

DEFINITION 4.4 (controls affine from time τ). *A controller affine from time τ takes the following form:*

$$\mathcal{K}_i^{\text{mult},\,\tau}\left(\mathcal{D}_{i^-}^{\text{con}}\right) = \begin{cases} u_i^{\dagger} & \text{if } i < \tau, \\ u_i^{\text{aff}} = k_i \widehat{x}_{i|i} + d_i & \text{if } i \geq \tau, \end{cases} \tag{4.5}$$

where the controls $u_i^{\dagger}$ are generated by an admissible strategy $\{\mathcal{K}_i^{\dagger}(\cdot)\}_{i=0}^T$, the controls u_i^{aff} are generated by an affine strategy $\{\mathcal{K}_i^{\text{aff}}(\cdot)\}_{i=0}^T$, with the gains $\{k_i\}_0^T$ and offsets $\{d_i\}_0^T$ computed offline, and $\widehat{x}_{i|i} = \mathbb{E}[x_i|\{y_j\}_{j=0}^i]$.

4.3. Preliminary lemmas. The main result ahead is Theorem 4.9 that states that the solution to Design problem 2 is to apply a separated design and certainty equivalence controls. Now we do some groundwork towards proving that result.

Given an admissible encoder, the partial control waveform $\{u_j\}_{j=i}^T$ affects only the cost-to-go from time j: $\mathbb{E}[x_{T+1}^2] + \sum_{j=i}^T \mathbb{E}[px_j^2 + qu_j^2]$. In the classical single agent LQ problem, the "prescribed encoder" is simply the linear observation process with prescribed signal-to-noise ratios (SNRs). There, this cost-to-go can be expressed as a quadratic function of $\{u_j\}_{j=i}^T, \{x_j\}_{j=i}^T$, and $\{\widehat{x}_{j|j}\}_{j=i}^T$. But in our two agent LQ problem, because of the dual effect, the cost-to-go may have a nonquadratic dependence on the controls $\{u_j\}_{j=i}^T$. However we show that by restricting to controls-forgetting encoders and affine controls, the cost-to-go does get a quadratic dependence on controls. We use this reasoning and backward dynamic programming to show that for time $t = i$, the following conclusions fall out:

- it is optimal at time $t = i$ to apply as control a linear function of $\widehat{x}_{i|i}$, and
- it is optimal at time $t = i$ to apply an encoding map that is controls forgetting from time $i - 1$.

LEMMA 4.5 (optimal control at time $t = T$). *It is optimal to apply the linear law $u_T^* = -a\,\widehat{x}_{1|1}/(q+1)$. And $V_T^*(\mathcal{D}_{T^-}^{\text{con}}) = \min_{u_T} \mathbb{E}[x_{T+1}^2 + qu_T^2|\mathcal{D}_{T^-}^{\text{con}}]$, the optimum cost-to-go, can be written as the expected value of a quadratic form in x_T and $\widehat{x}_{T|T}$.*

Proof. One is given $\mathcal{D}_{T^-}^{\text{con}}$, and is asked to pick u_T to minimize the cost-to-go

$$\begin{aligned} V_T\left(u_T; \mathcal{D}_{T^-}^{\text{con}}\right) &= \mathbb{E}\left[x_{T+1}^2 + qu_T^2 \,|\mathcal{D}_{T^-}^{\text{con}}\right] \\ &= \sigma_w^2 + \mathbb{E}\left[a^2\, x_T^2 + 2\,a\,x_T u_T + (1+q)\,u_T^2 \,|\mathcal{D}_{T^-}^{\text{con}}\right] \\ &= \sigma_w^2 + \frac{a^2}{1+q}\mathbb{E}\left[q\,x_T^2 + \left(x_T - \widehat{x}_{\,T|T}\right)^2 \middle| \mathcal{D}_{T^-}^{\text{con}}\right] \\ &\quad + (1+q)\left(u_T - \frac{a\widehat{x}_{\,T|T}}{1+q}\right)^2. \end{aligned}$$

This proves the lemma. ∎

LEMMA 4.6 (optimal ξ_i for separated, quadratic cost-to-go). *Fix the time $t = i$. Consider the dynamic encoder-controller design problem (Design problem 2) for the linear plant (2.1), and the performance cost (2.2). Suppose that we apply an admissible controller $\widetilde{\mathcal{K}}$ along with an encoder $\xi_t^{CF,i}$ that is controls forgetting from time i. Furthermore, suppose that the partial sets of policies $\{\xi_{i+1}^{CF,i}(\cdot), \ldots, \xi_T^{CF,i}(\cdot)\}$, $\{\widetilde{\mathcal{K}}_i(\cdot), \widetilde{\mathcal{K}}_{i+1}(\cdot), \ \ldots, \widetilde{\mathcal{K}}_T(\cdot)\}$ are chosen such that the following properties hold:*

1. *The cost-to-go at time i takes the separated form*

$$\mathbb{E}\left[x_{T+1}^2 + p\sum_{j=i}^{T} x_i^2 + q\sum_{j=i}^{T} u_i^2 \middle| \mathcal{D}_{i^+}^{\mathrm{con}} \right] = \mathbb{E}\left[J_i^{\mathrm{con}}(u_i, x_i) \middle| \mathcal{D}_{i^+}^{\mathrm{con}} \right] + \mathbb{E}\left[\Gamma_{i+1} \middle| \mathcal{D}_{i^+}^{\mathrm{con}} \right],$$

where, $J_i^{\mathrm{con}}(u_i, x_i) = \overline{\alpha} + \alpha\,\sigma_w^2 + \overline{\beta}\,x_i + \widetilde{\beta}\,x_i^2 + \overline{\nu}\,\widehat{x}_{i|i} + \widehat{\nu}\,x_i\widehat{x}_{i|i} + \widetilde{\nu}\,\widehat{x}_{i|i}^2$, and the term Γ_{i+1} is a weighted sum of future distortions that depends only on the random sequence $\{x_j - \widehat{x}_{j|j}\}_{j=i+1}^T$.
2. *The coefficients of the quadratic J_i^{con} may depend on the partial set of control policies $\{\widetilde{\mathcal{K}}_j(\cdot)\}_{j=i}^T$ but not on the partial set of encoding maps $\{\xi_j^{CF,i}(\cdot)\}_{j=i+1}^T$.*
3. *The term Γ_{i+1} depends on the encoding maps $\{\xi_j^{CF,i}(\cdot)\}_{j=i+1}^T$ but not on the partial set of control policies $\{\widetilde{\mathcal{K}}_j(\cdot)\}_{j=i}^T$.*

Then, it is optimal to apply an encoding map at time $t = i$ that does not depend on the data: $(u_{i-1}, \{\widetilde{\mathcal{K}}_j(\cdot)\}_{j=i}^T)$. It also follows that the shapes of the encoding maps $\{\xi_j^{CF,i}(\cdot)\}_{j=i+1}^T$ and their performance do not depend on the control u_{i-1}.

Proof. The proof exploits three facts: first, the special form of $J_i^{\mathrm{con}}(u_i, x_i)$ makes the encoder's performance cost at time i a sum of a quadratic distortion between x_i and $\widehat{x}_{i|i}$, and a term gathering distortions at later times. Second, the minimum of the sum distortion depends only on the intrinsic shape of the conditional density $\rho_{i|i-1}(\cdot)$ and not on its mean. Third, these facts and the controls-forgetting nature of later encoding maps allows the encoder to "ignore" the control u_{i-1}. We now start by writing the cost-to-go as

$$\begin{aligned}
\mathbb{E}\left[J_i^{\mathrm{con}} + \Gamma_{i+1} \middle| \mathcal{D}_{i^+}^{\mathrm{con}} \right] &= \mathbb{E}\left[\overline{\alpha} + \alpha\,\sigma_w^2 + \overline{\beta}\,x_i + \widetilde{\beta}\,x_i^2 \middle| \mathcal{D}_{i^+}^{\mathrm{con}} \right] \\
&\quad + \mathbb{E}\left[\overline{\nu}\,\widehat{x}_{i|i} + \widehat{\nu}\,x_i\,\widehat{x}_{i|i} + \widetilde{\nu}\,\widehat{x}_{i|i}^2 \middle| \mathcal{D}_{i^+}^{\mathrm{con}} \right] + \mathbb{E}\left[\Gamma_{i+1} \middle| \mathcal{D}_{i^+}^{\mathrm{con}} \right] \\
&= \overline{\alpha} + \alpha\,\sigma_w^2 + \mathbb{E}\left[\left(\overline{\beta} + \overline{\nu}\right) x_i + \left(\widehat{\nu} + \widetilde{\nu} + \widetilde{\beta}\right) x_i^2 \middle| \mathcal{D}_{i^+}^{\mathrm{con}} \right] \\
&\quad - (\widehat{\nu} + \widetilde{\nu})\,\mathbb{E}\left[x_i^2 - \widehat{x}_{i|i}^2 \middle| \mathcal{D}_{i^+}^{\mathrm{con}} \right] + \mathbb{E}\left[\Gamma_{i+1} \middle| \mathcal{D}_{i^+}^{\mathrm{con}} \right] \\
&= \overline{\alpha} + \alpha\,\sigma_w^2 + \mathbb{E}\left[\left(\overline{\beta} + \overline{\nu}\right) x_i + \left(\widehat{\nu} + \widetilde{\nu} + \widetilde{\beta}\right) x_i^2 \middle| \mathcal{D}_{i^+}^{\mathrm{con}} \right] \\
&\quad - (\widehat{\nu} + \widetilde{\nu})\,\mathbb{E}\left[\left(x_i - \widehat{x}_{i|i}\right)^2 \middle| \mathcal{D}_{i^+}^{\mathrm{con}} \right] + \mathbb{E}\left[\Gamma_{i+1} \middle| \mathcal{D}_{i^+}^{\mathrm{con}} \right]. \qquad (4.6)
\end{aligned}$$

Given $\mathcal{D}_{(i-1)^+}^{\mathrm{con}}$ the part of the RHS that depends on the encoding map $\xi_i(\cdot)$ is

$$-(\widehat{\nu} + \widetilde{\nu})\,\mathbb{E}\left[\left(x_i - \widehat{x}_{i|i}\right)^2 \middle| \mathcal{D}_{i^+}^{\mathrm{con}} \right] + \mathbb{E}\left[\Gamma_{i+1} \middle| \mathcal{D}_{i^+}^{\mathrm{con}} \right].$$

Notice that the first term is the scaled quantization error variance of the quantizer $\xi_i(\cdot)$. This reduction of the encoder's performance cost to a sum of current and future quantization distortions is possible because the term $J_i^{\mathrm{con}}(u_i, x_i)$ has been assumed to be quadratic in x_i and $\widehat{x}_{i|i}$. The reduced performance cost of the encoder is a function only of the quantizer $\xi_i(\cdot)$ and the conditional density $\rho_{i|i-1}(x|\mathcal{D}_{(i-1)^-}^{\mathrm{con}})$. Indeed, given the data $\mathcal{D}_{(i-1)^-}^{\mathrm{con}}$ and probabilities $\pi_{i|i-1}(x_i, \Delta) =$

$\mathbb{P}[x_i \in \Delta | \mathcal{D}^{\text{con}}_{(i-1)^-}, u_{i-1}]$, this cost is the following average:

$$\Gamma_i\left(\xi_i(\cdot);\mathcal{D}^{\text{con}}_{(i-1)^+}\right) = \sum_{\text{cells } \Delta} \pi_{i|i-1}(x_i,\Delta)\cdot\mathbb{E}\left[\Gamma_{i+1}\left(\mathcal{D}^{\text{con}}_{i^-}\right)\middle|\mathcal{D}^{\text{con}}_{(i-1)^+}, x_i\in\Delta\right] + \sum_{\text{cells } \Delta} \pi_{i|i-1}(x_i,\Delta)\cdot\lambda\cdot\mathbb{E}\left[\left(x_i-\widehat{x}_{i|i}\right)^2\middle|\mathcal{D}^{\text{con}}_{(i-1)^+}, x_i\in\Delta\right],$$

where $\lambda = -(\widehat{\nu}+\widetilde{\nu})$. The cost Γ_i does depend on both $\xi_i(\cdot)$ and u_i, but for given data $\mathcal{D}^{\text{con}}_{(i-1)^-}$ and control u_{i-1}, the minimum of Γ_i over all admissible quantizers $\xi_i(\cdot)$ may possibly depend on $\mathcal{D}^{\text{con}}_{(i-1)^-}$ but not on the control u_{i-1}. To see this consider two arbitrary possible values $u, \widetilde{u}$ for u_{i-1}. Suppose that one is given the quantizer

$$\xi(x) = \begin{cases} 1 & \text{if } x\in(-\infty,\delta_1), \\ 2 & \text{if } x\in(\delta_1,\delta_2), \\ \vdots & \vdots \\ N & \text{if } x\in(\delta_{N-1},+\infty), \end{cases}$$

meant for quantizing a random variable with the density $\rho_{i|i-1}(x|\mathcal{D}^{\text{con}}_{(i-1)^-}, u_{i-1}=u)$. Consider the quantizer $\widetilde{\xi}$ constructed by taking each cell $\Delta = (\underline{\delta}, \overline{\delta})$ in ξ, and generating a new cell $\widetilde{\Delta} = (\underline{\delta}-u+\widetilde{u},\ \overline{\delta}-u+\widetilde{u})$, and stipulating that the new quantizer $\widetilde{\xi}$ assigns to the cell $\widetilde{\Delta}$ the same channel input that the quantizer ξ assigns to Δ.

Because of the linear evolution, $x_i = ax_{i-1}+u_{i-1}+w_{i-1}$, and because the random variable w_{i-1} is independent of the data $\mathcal{D}^{\text{con}}_{(i-1)^+}$, we have the convolution relations

$$\rho(x) = \rho_{i|i-1}\left(\frac{\cdot-u}{a}\right)\circledast\rho_w(\cdot)\bigg|_x \text{ and}$$
$$\widetilde{\rho}(x) = \rho_{i|i-1}\left(\frac{\cdot-\widetilde{u}}{a}\right)\circledast\rho_w(\cdot)\bigg|_x,$$

leading to the following symmetry w.r.t. translations:

$$\rho_{i|i-1}\left(x-u\middle|\mathcal{D}^{\text{con}}_{(i-1)^-}, u_{i-1}=u\right) = \rho_{i|i-1}\left(x-\widetilde{u}\middle|\mathcal{D}^{\text{con}}_{(i-1)^-}, u_{i-1}=\widetilde{u}\right). \tag{4.7}$$

Then we get the following equalities for each pair of cells $\Delta, \widetilde{\Delta}$, with $\widetilde{x}_{i|i} = x_i - \widehat{x}_{i|i}$,

$$\mathbb{P}\left[x_i\in\Delta\middle|\mathcal{D}^{\text{con}}_{(i-1)^-}, u_{i-1}=u\right] = \mathbb{P}\left[x_i\in\widetilde{\Delta}\middle|\mathcal{D}^{\text{con}}_{(i-1)^-}, u_{i-1}=\widetilde{u}\right],$$
$$\Gamma_{i+1}\left(\mathcal{D}^{\text{con}}_{(i-1)^-}, u_{i-1}=u, x_i\in\Delta\right) = \Gamma_{i+1}\left(\mathcal{D}^{\text{con}}_{(i-1)^-}, u_{i-1}=\widetilde{u}, x_i\in\widetilde{\Delta}\right), \text{ and}$$
$$\mathbb{E}\left[\widetilde{x}^2_{i|i}\middle|\mathcal{D}^{\text{con}}_{(i-1)^-}, u_{i-1}=u, x_i\in\Delta\right] = \mathbb{E}\left[\widetilde{x}^2_{i|i}\middle|\mathcal{D}^{\text{con}}_{(i-1)^-}, u_{i-1}=\widetilde{u}, x_i\in\widetilde{\Delta}\right].$$

Then the performance of any quantizer ξ designed for $u_{i-1}=u$ can be matched by $\widetilde{\xi}$ for $u_{i-1}=\widetilde{u}$, and vice versa. Hence, we can conclude that for any $u, \widetilde{u}$,

$$\inf_{\xi}\ \Gamma_i\left(\xi(\cdot);\mathcal{D}^{\text{con}}_{(i-1)^-}, u_{i-1}=u,\right) = \inf_{\xi}\ \Gamma_i\left(\xi(\cdot);\mathcal{D}^{\text{con}}_{(i-1)^-}, u_{i-1}=\widetilde{u},\right).$$

Notice that this optimal encoder is controls forgetting from time $i-1$. ∎

As the optimal control u_T^* is a linear function on $\widehat{x}_{T|T}$, the encoder ξ_T begets a performance cost that is quadratic in x_T, $\widehat{x}_{T|T}$. Then the above lemma renders the optimal encoding map ξ_T^* to be controls forgetting from time $T-1$. This reduction also holds at earlier times.

LEMMA 4.7 (encoder separation for affine controls). *If the two conditions hold,* (A) *for any admissible control strategy, an admissible encoder strategy minimizing the performance cost* (2.2) *exists, and* (B) *the control strategy is affine from time* τ, *then the following two conclusions hold:* (a) *an encoder that is controls forgetting from time* τ *minimizes the partial LQ cost*

$$\mathbb{E}\left[x_{T+1}^2 + p\sum_{i=\tau+1}^{T} x_i^2 + q\sum_{i=\tau}^{T} u_i^2 \,\middle|\, \mathcal{D}_{\tau^+}^{\text{con}} \right],$$

and (b) *the shapes of the minimizing encoding maps from time* τ *and their performance are independent of the data,* $\{\, u_{\tau-1}^{\dagger},\ \{k_i\}_{i=\tau}^{T},\ \{d_i\}_{i=\tau}^{T} \}$.

Proof. We prove this by mathematical induction. For a given control strategy, define

$$W_T = \mathbb{E}\left[x_{T+1}^2 + px_T^2 + qu_T^2 \,\middle|\, \mathcal{D}_{(T-1)^+}^{\text{con}} \right], \qquad W_T^* = \inf_{\xi_T(\cdot)} W_T$$

$$W_i = \mathbb{E}\left[px_i^2 + qu_i^2 \,\middle|\, \mathcal{D}_{(i-1)^+}^{\text{con}} \right] + \mathbb{E}\left[W_{i+1}^* \left(\mathcal{D}_{i^+}^{\text{con}}\right) \,\middle|\, \mathcal{D}_{(i-1)^+}^{\text{con}}, \xi_i(\cdot) \right], \quad W_i^* = \inf_{\xi_i(\cdot)} W_i.$$

Induction hypothesis for time i: for some time $t = i$ such that $\tau \leq i < T$, we have the following three assumptions: (1) for every $j \geq i+1$, the optimal value function $W_j^*(\mathcal{D}_{(j-1)^-}^{\text{con}})$ takes the form

$$\alpha_j\sigma_w^2 + \overline{\alpha}_j + \overline{\beta}_j\,\widehat{x}_{j|j} + \mathbb{E}\left[\widetilde{\beta}_j {x_j}^2 + \widetilde{\lambda}_j\left(x_j - \widehat{x}_{j|j}\right)^2 + \widetilde{\Gamma}_{j+1}^*\left(\mathcal{D}_{(j+1)^-}^{\text{con}}\right) \,\middle|\, \mathcal{D}_{j^-}^{\text{con}} \right],$$

where the $\alpha_j, \overline{\alpha}_j, \widetilde{\beta}_j, \overline{\beta}_j, \widetilde{\lambda}_j$ are known nonnegative real numbers for $j \geq i+1$, (2) for each such j, the nonnegative function $\widetilde{\Gamma}_{j+1}^*(\mathcal{D}_{j^-}^{\text{con}})$ is assumed to be independent of the partial waveform $\{u_j, u_{j+1}, \ldots, u_T\}$, and (3) the optimal partial set of encoding maps $\{\xi_j^*(\cdot)\}_{i+1}^{T}$ is a set that is controls forgetting from time i. We will now show if this hypothesis holds for time i, then it holds for time $i-1$. Assuming that the partial set of optimal encoding maps $\{\xi_j^*(\cdot)\}_{i+1}^{T}$ is employed, we get

$$\begin{aligned} W_i &= \mathbb{E}\left[p\,x_i^2 + q\,u_i^2 \,\middle|\, \mathcal{D}_{(i-1)^+}^{\text{con}} \right] + \mathbb{E}\left[W_{i+1}^*\left(\mathcal{D}_{i^+}^{\text{con}}\right) \,\middle|\, \mathcal{D}_{(i-1)^+}^{\text{con}}, \xi_i(\cdot) \right], \\ &= p\,\mathbb{E}\left[x_i^2 \,\middle|\, \mathcal{D}_{(i-1)^+}^{\text{con}} \right] + q\,\mathbb{E}\left[u_i^2 \,\middle|\, \mathcal{D}_{(i-1)^+}^{\text{con}} \right] + \alpha_{i+1}\sigma_w^2 + \overline{\alpha}_{i+1} \\ &\quad + \widetilde{\beta}_{i+1}\mathbb{E}\left[{x_{i+1}}^2 \,\middle|\, \mathcal{D}_{(i+1)^-}^{\text{con}} \right] + \overline{\beta}_{i+1}\mathbb{E}\left[x_{i+1} \,\middle|\, \mathcal{D}_{(i+1)^-}^{\text{con}} \right] + \mathbb{E}\left[\widetilde{\Gamma}_{i+1}^*\left(\mathcal{D}_{(i+1)^-}^{\text{con}}\right) \,\middle|\, \mathcal{D}_{i^-}^{\text{con}} \right]. \end{aligned}$$

We can rewrite the above expression as

$$\begin{aligned} W_i = \alpha_i\sigma_w^2 + \overline{\alpha}_i + \widetilde{\beta}_i\mathbb{E}\left[x_i^2 \,|\mathcal{D}_{i^-}^{\text{con}} \right] + \overline{\beta}_i\mathbb{E}\left[x_i \,|\mathcal{D}_{i^-}^{\text{con}} \right] \\ + \mathbb{E}\left[\widetilde{\Gamma}_{i+1}^*\left(\mathcal{D}_{(i+1)^-}^{\text{con}}\right) \,\middle|\, \mathcal{D}_{i^-}^{\text{con}} \right] + \widetilde{\lambda}_i\mathbb{E}\left[\left(x_i - \widehat{x}_{i|i}\right)^2 |\mathcal{D}_{i^-}^{\text{con}} \right], \end{aligned}$$

where the coefficients $\alpha_i = \alpha_{i+1} + \widetilde{\beta}_{i+1}$, $\overline{\alpha}_i = \overline{\alpha}_{i+1} + \widetilde{\beta}_{i+1}d_i^2 + q\,d_i^2 + \overline{\beta}_{i+1}d_i$, $\overline{\beta}_i = 2(q\,k_i d_i + a\widetilde{\beta}_{i+1}d_i + \widetilde{\beta}_{i+1}k_i d_i)$, $\widetilde{\beta}_i = p_i + a^2\widetilde{\beta}_{i+1} + k_i^2\widetilde{\beta}_{i+1} + 2ak_i\widetilde{\beta}_{i+1} + q\,k_i^2\widetilde{\beta}_{i+1}$, and $\widetilde{\lambda}_i = q\,k_i^2 + k_i^2\widetilde{\beta}_{i+1} + 2ak_i\widetilde{\beta}_{i+1}$.

We have thus $W_i = \mathbb{E}[\text{A quadratic in } x_i, \widehat{x}_{i|i}] + \mathbb{E}[\text{ future distortions }]$. This and the fact that the encoder is controls forgetting from time $t = i$ meet the requirements of Lemma 4.6. Then we get the optimal encoding map ξ_i^* to be controls forgetting from time $t = i - 1$, and

$$\widetilde{\Gamma}_i = \min_{\xi} \mathbb{E}\left[\widetilde{\Gamma}_{i+1}^*\left(\mathcal{D}^{\text{con}}_{(i+1)^-}\right)\middle|\,\mathcal{D}^{\text{con}}_{i^-}\right] + \widetilde{\lambda}_i\,\mathbb{E}\left[\left(x_i - \widehat{x}_{i|i}\right)^2|\mathcal{D}^{\text{con}}_{i^-}\right]$$

is independent of the partial set of controls $\{u_j\}_{j=i-1}^{T}$. From this it follows that the induction hypothesis is also true for time $i - 1$. ∎

LEMMA 4.8 (certainty equivalence controls for controls-forgetting encoders). *If the encoder is controls forgetting from time τ, then the partial LQ cost*

$$\mathbb{E}\left[x_{T+1}^2 + p\sum_{i=\tau+1}^{T} x_i^2 + q\sum_{i=\tau}^{T} u_i^2\middle|\,\mathcal{D}^{\text{con}}_{\tau^-}\right]$$

is minimized by the control law for $i \geq \tau$, $u_i^* = k_i^*\,\widehat{x}_{i|i}$, where $k_i^* = a\beta_{i+1}/(q + \beta_{i+1})$, $\beta_i = p + a^2 q\beta_{i+1}/(q + \beta_{i+1})$, and, $\beta_{T+1} = 1$.

Proof. Define the cost-to-go at time $t = T - 1$ as $V_{T-1} = \mathbb{E}[W_T(\epsilon_T(\cdot); \mathcal{D}^{\text{con}}_{(T-1)^+})]$. Because of Lemma 4.5,

$$V_{T-1} = \sigma_w^2 + \left(p + \frac{a^2 q}{q+1}\right)\mathbb{E}\left[x_T^2\,\middle|\,\mathcal{D}^{\text{con}}_{(T-1)^-}, u_{T-1}\right] + \mathbb{E}\left[\left(x_T - \widehat{x}_{T|T}\right)^2\middle|\,\mathcal{D}^{\text{con}}_{T^-}\right].$$

Because the encoder is controls forgetting from time τ, the last term, which is the distortion due to the encoder ξ_T, is independent of the partial set of controls $\{u_i\}_{i=\tau+1}^{T}$. Hence the only part of V_{T-1} that depends on the control u_{T-1} is the quadratic

$$\begin{aligned} &q\,u^2 + \left(p + \frac{a^2 q}{q+1}\right)\mathbb{E}\left[x_T^2\,\middle|\,\mathcal{D}^{\text{con}}_{(T-1)^-}, u_{T-1}\right] \\ &\quad = q\,u^2 + \left(p + \frac{a^2 q}{q+1}\right)\Big\{a^2\mathbb{E}\left[x_T^2\,\middle|\,\mathcal{D}^{\text{con}}_{(T-1)^-}, u_{T-1}\right] \\ &\qquad\qquad + 2a\,\widehat{x}_{T-1|T-1}\,u_{T-1} + u_{T-1}^2 + \sigma_w^2\Big\}. \end{aligned}$$

Hence $u_{T-1}^* = -a(p + \frac{a^2 q}{q+1})\widehat{x}_{T-1|T-1}/(q + p + \frac{a^2 q}{q+1})$, and the resulting value function

$$\begin{aligned} V_{T-1}^* = &\left(1 + p + \frac{a^2 q}{q+1}\right)\sigma_w^2 + \frac{a^2 q\left(p + \frac{a^2 q}{q+1}\right)}{q + p + \frac{a^2 q}{q+1}}\mathbb{E}\left[x_T^2\,\middle|\,\mathcal{D}^{\text{con}}_{(T-1)^-}, u_{T-1}\right] \\ &+ \frac{a^2\left(p + \frac{a^2 q}{q+1}\right)}{q + p + \frac{a^2 q}{q+1}}\mathbb{E}\left[\left(x_{T-1} - \widehat{x}_{T-1|T-1}\right)^2\middle|\,\mathcal{D}^{\text{con}}_{(T-2)^+}\right] \\ &+ \mathbb{E}\left[\left(x_T - \widehat{x}_{T|T}\right)^2\middle|\,\mathcal{D}^{\text{con}}_{T^-}\right]. \end{aligned}$$

Repeating this procedure backwards in time, we get for times $i \geq \tau$, the optimal control laws as $u_i^* = -k_i^*\,\widehat{x}_{i|i}$. ∎

4.4. Main theorem. Lemma 4.7 implies that for a preassigned controller affine from time zero, there exist optimal encoding maps that are controls forgetting from time zero. Lemma 4.8 is complementary and implies that for a preassigned encoder that is controls forgetting from time zero, the optimal control laws have linear forms. For Design problem 2, an optimal pair of strategies has a similar simplified structure. This pair consists of a controls-forgetting encoder and control laws linear in $\hat{x}_{i|i}$. In general, this controls-forgetting encoder does not minimize the aggregate squared estimation error. The goal it accomplishes is slightly different. It minimizes a sum of state estimation errors with the time-varying weights λ_i. Details follow.

THEOREM 4.9 (optimality of separation and certainty equivalence). *For Design problem* 2, *with the discrete alphabet channel of constant alphabet size, the quadratic performance cost* (2.2) *is minimized by applying the linear control laws*

$$u_t^* = -k_t^* \, \widehat{x}_{\,t|t} \tag{4.8}$$

in combination with the following encoder which is controls forgetting from time 0*:*

$$\epsilon_t^* \left(\zeta_t \, ; \{z_i\}_0^{t-1} , \{\epsilon_i \, (\cdot)\}_0^{t-1} \right) = \underset{\epsilon(\cdot)}{\arg\inf} \; \Gamma_i \Big(\epsilon \, (\cdot) \; ; \; \{z_i\}_0^{t-1} , \{\epsilon_i \, (\cdot)\}_0^{t-1} \Big), \tag{4.9}$$

where $k_i^* = a\beta_{i+1}/(q+\beta_{i+1}), \beta_i = p + a^2 q \beta_{i+1}/(q+\beta_{i+1}), \beta_{T+1} = 1,$ *and,*

$$\begin{aligned}
\Gamma_t &= \lambda_t \cdot \mathbb{E}\left[\left(\zeta_t - \widehat{\zeta}_{t|t} \right)^2 \middle| \epsilon_t \, (\cdot) \, , \mathcal{D}^{\text{con}}_{(t-1)^+} \right] + \mathbb{E}\left[\Gamma^*_{t+1} \left(\overline{x}_0, \sigma_0^2, \{z_i\}_0^t \, , \{\epsilon_i \, (\cdot)\}_0^t \right) \right], \\
\Gamma_T &= \mathbb{E}\left[\left(\zeta_T - \widehat{\zeta}_{T|T} \right)^2 \middle| \epsilon_T \, (\cdot) \, , \overline{x}_0, \sigma_0^2, \{z_i\}_0^{T-1} \, , \{\epsilon_i \, (\cdot)\}_0^{T-1} \right], \\
\Gamma_t^* &= \inf_{\epsilon(\cdot)} \Gamma_t \, (\epsilon) \, ,
\end{aligned}$$

where $\lambda_i = a^2 \beta_{i+1}^2/(q+\beta_{i+1})$*. Also, this control law is a certainty equivalence law.*

Proof. We start with Lemma 4.5, and then repeatedly apply, in sequence, Lemmas 4.7 and 4.8. This proves optimality. Lemma 3.3 implies that the control laws of (4.8) are indeed certainty equivalence control laws as per [40]. □

The optimal controller splits into a least squares estimator computing $\widehat{x}_{t|t}$ and a time-dependent gain. Computing $\widehat{x}_{t|t}$ is intrinsically hard because quantization is a nonlinear operation. If one ignores this computational burden, then, at least formally, the optimal controller resembles that for the classical LQG problem.

Note that in general the sequence of weights $\{\lambda_i\}_0^T$ depends on the parameters of the performance cost, including the control penalty coefficient q. In the two special cases, (1) the coefficients $q = 0$, $p = 1$, or (2) the quantity $p + a^2 q - q > 0$ and $p + a^2 q - q + \sqrt{(p + a^2 q - q)^2 + 4pq} = 2$, it turns out that the weights $\beta_i \equiv 1 \, \forall i$, and hence the weights $\lambda_i \equiv a^2/(q+1) \, \forall i$. Thus in these two cases, optimal encoders ignore the parameters of the performance cost and simply minimize the usual aggregate squared error in state estimation.

4.5. Extension to the multivariable case. Theorem 4.9 can be extended to situations where the state, control, and noise signals are vectors, as well as where the objective function (2.2) includes cross terms involving the state and control. We can also extend to the case where the sensor has access only to partial and noisy observations of the state. To carry out these extensions, we need no more than the standard arguments of LQG control. Below, we mention only the key steps corresponding to the lemmas of section 4.3. Consider a partially observed, linear multivariable plant,

$$x_{t+1} = Ax_t + Bu_t + Ew_t, \quad y_t = Cx_t + Dv_t, \tag{4.10}$$

where the state $x_t \in \mathbb{R}^n$, the control $u_t \in \mathbb{R}^m$, the output $y_t \in \mathbb{R}^p$, the process noise $w_t \in \mathbb{R}^{l_1}$, and the measurement noise $v_t \in \mathbb{R}^{l_2}$. Let the two noise sequences $\{w_t\}, \{v_t\}$ be IID sequences that are mutually independent of each other.

For any matrix M, let M' denote its transpose. For all times, let $\mathbb{E}[w_t w_t'] = \Sigma_w$ and $\mathbb{E}[v_t v_t'] = \Sigma_v$. Let the performance objective be defined as

$$J_{\text{general}} = \mathbb{E}\left[x_{T+1}' \, P_{T+1} \, x_{T+1}\right] + \sum_{i=0}^{T} \mathbb{E}\left[\begin{pmatrix} x_i' & u_i' \end{pmatrix} \begin{bmatrix} P & R' \\ R & Q \end{bmatrix} \begin{pmatrix} x_i \\ u_i \end{pmatrix}\right], \tag{4.11}$$

where P and $\left[\begin{smallmatrix} P & R' \\ R & Q \end{smallmatrix}\right]$ are symmetric and positive semidefinite, and Q is symmetric and positive definite. It is easy to see that an extension of Lemma 4.5 holds for the multivariable case. Precisely, the optimal control law at the terminal decision time T is $u_T^* = -K_T \, \widehat{x}_{T|T}$, where $K_T = -(B'P_{T+1}B + Q)^{-1}(R + B'P_{T+1}A)$ and the optimal cost-to-go $V_T^*(\mathcal{D}_{T^-}^{\text{con}}) = \text{trace}(E\Sigma_w E') + \mathbb{E}[(x_T - \widehat{x}_{T|T})' M_T (x_T - \widehat{x}_{T|T}) | \mathcal{D}_{T^-}^{\text{con}}] + \mathbb{E}[x_T'(P + A'P_{T+1}A - M_T)x_T | \mathcal{D}_{T^-}^{\text{con}}]$, where the matrix $M_T = K_T'(B'P_{T+1}B + Q)K_T$. Next, the following generalization of Lemma 4.6 can be proved.

LEMMA 4.10 (multivariable version of Lemma 4.6). *Assume the hypothesis of* Lemma 4.6 *but for the multivariable, partially observed plant* (4.10), *the objective function* (4.11), *and the following new definition of* J_i^{con}:

$$J_i^{con} = \alpha + \beta' \, x_i + \widehat{\beta}' \, \widehat{x}_{i|i} + \begin{pmatrix} x_i' & \widehat{x}_{i|i}' \end{pmatrix} \begin{bmatrix} \widehat{P} & \widehat{R}' \\ \widehat{R} & \widehat{Q} \end{bmatrix} \begin{pmatrix} x_i \\ \widehat{x}_{i|i} \end{pmatrix}.$$

Then, it is optimal to apply an encoding map at time $t = i$ that does not depend on the data: $(u_{i-1}, \{\widetilde{\mathcal{K}}_j(\cdot)\}_i^T)$. It also follows that the shapes of the encoding maps $\{\xi_j^{CF,i}(\cdot)\}_{i+1}^T$ and their performance do not depend on the control u_{i-1}.

Proof Sketch: We can rewrite the part of the cost-to-go that depends on the control u_{i-1}. As in (4.6), it is possible to rewrite this in such a way that the only dependence on $\widehat{x}_{i|i}$ is through a quadratic form of the estimation error $x_i - \widehat{x}_{i|i}$:

$$\begin{aligned} \mathbb{E}\left[J_i^{\text{con}} + \Gamma_{i+1} \middle| \mathcal{D}_{i^+}^{\text{con}} \right] = \alpha + \mathbb{E}\left[\left(\beta + \widehat{\beta}\right)' x_i + x_i' \left(\widehat{P} + \widehat{R} + \widehat{R}' + \widehat{Q}\right) x_i \middle| \mathcal{D}_{i^+}^{\text{con}} \right] \\ + \mathbb{E}\left[\left(x_i - \widehat{x}_{i|i}\right)' \left(\widehat{R} + \widehat{R}' + \widehat{Q}\right) \left(x_i - \widehat{x}_{i|i}\right) + \Gamma_{i+1} \middle| \mathcal{D}_{i^+}^{\text{con}} \right]. \end{aligned}$$

The part of the RHS that depends on ξ_i is

$$\mathbb{E}\left[\Gamma_{i+1} \middle| \mathcal{D}_{i^+}^{\text{con}} \right] + \mathbb{E}\left[\left(x_i - \widehat{x}_{i|i}\right)' \left(\widehat{R} + \widehat{R}' + \widehat{Q}\right) \left(x_i - \widehat{x}_{i|i}\right) \middle| \mathcal{D}_{i^+}^{\text{con}} \right].$$

The minimum of this quantity over different ξ_i will be independent of u_{i-1} if the density $\rho_{i|i-1}$ is symmetric w.r.t. translations in the control. If the matrix A is invertible, then $\rho_{Ax_i|z_0^{i-1}}(x) = \rho_{i-1|i-1}(A^{-1}x)$. Let $u, \widetilde{u}$ be two possible values for u_{i-1}. Then

$$\rho_{i|i-1}(x) = \rho_{Ax_i|z_0^{i-1}}(\cdot - Bu) \circledast \rho_{Ew}(\cdot)\Big|_x \text{ and}$$

$$\widetilde{\rho}_{i|i-1}(x) = \rho_{Ax_i|z_0^{i-1}}(\cdot - B\widetilde{u}) \circledast \rho_{Ew}(\cdot)\Big|_x.$$

If the following three conditions hold, (1) the matrix A is invertible, (2) the conditional density $\rho_{i-1|i-1}$ is a "well-behaved" function, for example, a function of bounded variation, and (3) the noise random variables w_i, v_i have well-behaved densities, then it is straightforward to deduce the following symmetry w.r.t. translations:

$$\rho_{\,i|i-1}\left(x-u\,\middle|\,\mathcal{D}^{\mathrm{con}}_{(i-1)^-}, u_{i-1}=u\right)=\rho_{\,i|i-1}\left(x-\widetilde{u}\,\middle|\,\mathcal{D}^{\mathrm{con}}_{(i-1)^-}, u_{i-1}=\widetilde{u}\right).$$

If the matrix A is not invertible, or if any of the relevant densities have Dirac delta functions, then too, this symmetry property holds. Proving that needs some slightly more delicate arguments. The rest is similar to the proof of Lemma 4.6. □

The remaining lemmas of section 4.3 are straightforward to generalize to the multivariable case. Moreover, our results clearly extend also to the case of deterministic, time-varying coefficients of the plant dynamics and of the objective function.

5. Dynamic designs for other models of channels. The results of section 4.4 for Design problem 2 apply to the case where the quantizer word lengths at different times are deterministic but time varying. In this case, the communication cost J^{Comm} in (2.2) may take a positive functional form depending on the channel model. Our results also extend to other channel models, all coming from within three broad classes of messaging a sequence of real numbers: (1) quantized messaging, (2) unquantized but irregular, event-triggered sampling, and (3) unquantized messaging corrupted by additive channel noise. For each of these channel models, we find that the dynamic LQ design problem gets a separated optimal solution despite the existence of a dual effect in the corresponding networked control system. To obtain this design simplification, we assume that at all times, the channel output is perfectly visible to the encoder. Thus in each one of our channel models, there is an ideal, delay-free feedback channel copying the actual channel outputs back to the encoder.

To show these extensions, we need to find the appropriate versions of Lemma 4.6. Once this is done, all the steps in the proofs for Lemmas 4.7–4.8 and Theorem 4.9 can be repeated. We derived these versions of Lemma 4.6 in [34] for three different channel models. We show that for each of the channel models in [34], an encoder that is controls forgetting from time zero will be optimal in combination with the certainty equivalence control laws of (4.8). Below we outline such developments when the channel is real valued with added noise, for the scalar, fully observed plant (2.1) and the objective function (2.2). For discussions on the other channel models, see [34].

5.1. Messaging over a noisy linear channel. The considered channel model (see Figure 6) is a generalization of the classical additive white Gaussian noise channel, where we let the channel noise be colored and non-Gaussian. This channel accepts real valued inputs ι_t and delivers outputs z_t with noise added. For $0 \leq t \leq T$, the channel output $z_t = \iota_t + \chi_t$, where the channel noise process $\{\chi_i\}$ is IID with mean zero and variance $\sigma_\chi^2 < \infty$. At time t, the noise χ_t is independent of the state, control,

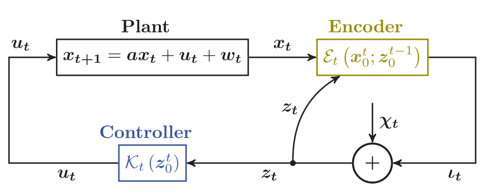

FIG. 6. *Block diagram for additive white noise channel.*

and process noise up to and including time t. For this style of messaging, we describe a model that allows the encoder to choose the SNR for each message. The model will also specify costs incurred for choosing message SNRs.

Let the even function $\phi(\cdot)$ increase with increasing magnitude of argument, and let $\phi(0) = 0$. An example is $\phi(\iota) = \iota^2$. Let $\overline{\iota}$ denote an upper limit on inputs to the channel. Then the communication cost incurred at a time t can be described as

$$\varphi_t = \begin{cases} \phi(\iota_t) & \text{if } |\iota_t| \leq \overline{\iota}, \\ +\infty & \text{if } |\iota_t| > \overline{\iota}. \end{cases}$$

Let $\mathfrak{p} \leq \phi(\overline{\iota})$ denote an upper limit on the average power of channel inputs over the entire horizon. We define the communication cost from time t to the end as

$$J^{\text{Comm}}_{[t,T]} = \begin{cases} m \cdot \mathbb{E}\left[\sum_{j=t}^{T} \varphi\left(\iota_j\right)\right] & \text{if } \sum_{j=0}^{T} \varphi\left(\iota_j\right) \leq \mathfrak{p} \cdot (T+1)\,, \\ +\infty & \text{if } \sum_{j=0}^{T} \varphi\left(\iota_j\right) > \mathfrak{p} \cdot (T+1)\,, \end{cases} \tag{5.1}$$

where m is a fixed nonnegative scalar. Then the cost $J^{\text{Comm}} = J^{\text{Comm}}_{[0,T]}$.

It is straightforward to see that $\{x_t,\, \{\iota_j\}_0^{t-1}, \{\xi_j\}_0^{t-1}, \{z_j\}_0^{t-1}\}$ are sufficient statistics at the encoder. As with quantized and event-triggered messaging, there is scope for the dual effect since the encoding map may be nonlinear. Clearly there is no dual effect introduced in the case where (1) the upper limit on inputs is removed, and (2) the encoder implements an affine encoder. But in general, there is scope for introducing the dual effect. If the encoder implements the quadratic encoder $\xi_t^{\text{quadratic}} = \eta x_t^2$, then there is a second-order dual effect. Another example of an admissible encoder that introduces a dual effect is the piecewise-constant encoder:

$$\xi_t = \begin{cases} -\overline{\iota} & \text{if } x_t \in (-\infty, -\theta)\,, \\ 0 & \text{if } x_t \in (-\theta, +\theta)\,, \\ \overline{\iota} & \text{if } x_t \in (+\theta, -\infty)\,, \end{cases}$$

where the threshold θ is fixed. In fact, this encoder has nearly the same input-output behavior as the encoders considered in Examples 1 and 2. Hence, it is easy to setup an example with an additive noise channel such that the dual effect is present. When there is a finite, hard limit $\overline{\iota}$ on amplitudes of channel inputs, then the dual effect is present for any encoder other than the trivial one, $\xi_t \equiv$ constant. As with other types of messaging, we can show that even though the dual effect is present, the dynamic encoder-controller problem has a separated solution and certainty equivalence controls are optimal.

LEMMA 5.1 (controls-forgetting compander optimal for affine controls). *Fix time $t = i$ and apply control laws affine from time i. Suppose that for all times $j > i$ the optimal encoding policies $\mathcal{E}_j^*(\cdot)$ and their performances are independent of the partial control waveform $\{u_i, \ldots, u_T\}$. Then, for all times $j > i-1$ the optimal encoding policies $\mathcal{E}_j^*(\cdot)$ and their performances are independent of the slightly longer waveform $\{u_{i-1}, u_i, \ldots, u_T\}$.*

Proof. We carry out two steps. First we show that because the cost-to-go is quadratic, the quantizer's objective at time i is to minimize a sum Γ_i of current and future estimation distortions. Second we show that the minimum of this sum distortion is independent of the control u_{i-1}. Thus the optimal encoder becomes controls forgetting from time $i-1$. ◻

The main result for communication over a noisy linear channel is then as follows.

THEOREM 5.2 (optimality of separation and certainty equivalence for additive noise channel). *For Design problem* 2, *with the additive noise channel, the performance cost* (2.2) *with communication cost* (5.1) *is minimized by applying* $u_t^* = -k_t^* \widehat{x}_{t|t}$ *in combination with the compander, which is controls forgetting from time zero,* $\epsilon_t^*(\zeta_t ; \{z_i\}_0^{t-1}, \{\epsilon_i\}_0^{t-1}) = \arg\inf_\epsilon \Gamma_i(\epsilon ; \{z_i\}_0^{t-1}, \{\epsilon_i\}_0^{t-1})$, *where* $\beta_{T+1} = 1$, $\beta_i = p + a^2 q\beta_{i+1}/(q + \beta_{i+1})$, $\lambda_i = a^2\beta_{i+1}^2/(q + \beta_{i+1})$, $k_i^* = a\beta_{i+1}/(q + \beta_{i+1})$, *and*

$$\Gamma_T = \begin{cases} +\infty & \text{if } \sum_{i=0}^{T} \varphi(\iota_i) > \mathfrak{P}(T+1), \\ \mathbb{E}\left[\left(\zeta_T - \widehat{\zeta}_{T|T}\right)^2 + m\varphi(\iota_T) \,\middle|\, \epsilon_T, \{z_i, \epsilon_i\}_0^{T-1} \right] & \text{otherwise,} \end{cases}$$

$$\Gamma_t = \lambda_t \, \mathbb{E}\left[\left(\zeta_t - \widehat{\zeta}_{t|t}\right)^2 + m\varphi(\iota_t) \,\middle|\, \epsilon_t, \mathcal{D}^{\text{con}}_{(t-1)^+} \right] + \mathbb{E}\left[\Gamma_{t+1}^* \left(\{z_i, \epsilon_i\}_0^t \right) \right],$$

$$\Gamma_t^* = \inf_\epsilon \Gamma_t(\epsilon).$$

Moreover, this control law is a certainty equivalence law.

Proof. Starting with the result of Lemma 4.5 and repeatedly applying Lemmas 5.1 and 4.8 proves optimality. Lemma 3.3 implies that the control laws of (4.8) are indeed certainty equivalence control laws as per [40]. ∎

6. Constrained encoder-controller design. We now use our understanding of the dynamic encoder-controller design problem (Design problem 2) to examine the constrained encoder-controller design problem (Design problem 3). In this section, we show that, in general, separation in design of encoder and controller is not optimal for these design problems. Some of these counterexamples illustrate that the distortion term in the cost-to-go lacks symmetry w.r.t. translations (4.7). Recall that this property was instrumental in ensuring separation in the dynamic encoder-controller design problem (see proof of Lemma 4.6).

6.1. Symmetry w.r.t. translations leads to separation. We present a simple example of a dynamic encoder-controller design problem: the encoder is specified in a parametric form, but the choice of the parameters can be dynamic, with no restrictions on the set of parameters. We show that the optimal controller uses the certainty equivalence law.

Example 4. For the scalar linear plant (2.1), with initial state x_0 given by a zero mean Gaussian with variance σ_x^2, and process noise w_k given by a zero mean Gaussian with finite variance σ_w^2, let the horizon length $T = 2$. Let the cost coefficients p and q remain unspecified. Let the channel alphabet be the discrete set $\{1, 2\}$. The controller receives a quantized version of the state

$$z_k = \begin{cases} 1 & \text{if } x_k \le \delta_k, \\ 2 & \text{otherwise}. \end{cases}$$

The quantizer thresholds δ_0 and δ_1 are to be chosen along with the control signals u_0 and u_1, to jointly minimize the two-step horizon control cost.

We use dynamic programming to find the optimal values for u_1, δ_1 and u_0, δ_0, in the specified order. From Lemma 4.5, we know that u_1^* is given by the certainty equivalence law $u_1^* = -a\,\hat{x}_{1|1}/(q+1)$, where the MMSE estimate of x_1 is given by $\hat{x}_{1|1} = \mathbb{E}[x_1 | \{z_i\}_0^1]$. Then, let us consider the cost-to-go at the previous time step,

$$V_0 = \mathbb{E}\left[a^2(p + a^2)x_0^2 + (q + p + a^2)u_0^2 + 2a(p + a^2)x_0 u_0 - \frac{a^2}{q+1}\hat{x}_{1|1}^2 \,\middle|\, z_0 \right] + \kappa, \tag{6.1}$$

where $\kappa = (1 + p + a^2)\sigma_w^2$. The above cost-to-go is to be minimized by selecting suitable u_0 and δ_1 simultaneously. To do this, we first need to find an expression for $\mathbb{E}[\hat{x}_{1|1}^2|z_0]$. The encoder outputs at times $k = 0, 1$ tell us the quantization cells in which x_0, x_1 lie. We use this information to find an expression for the estimate $\hat{x}_{1|1}$, as shown in the appendix of [34], and write the optimal cost-to-go as

(6.2)

$$V_0^* = \min_{u_0,\delta_1} \left\{ \mathbb{E}\left[a^2(p+a^2)x_0^2 + \overbrace{\left(q+p+a^2\frac{q}{q+1}\right)u_0^2 + 2a\left(p+a^2\frac{q}{q+1}\right)x_0u_0}^{\text{function of } u_0} \middle| z_0 \right] \right.$$
$$\left. \underbrace{-\frac{a^2}{q+1}\frac{\sum_{j=1}^{N}\vartheta^2\left(\frac{\delta_{j-1}-u_0}{\sigma_2}, \frac{\delta_j-u_0}{\sigma_2}\right)}{\mathbb{P}\left[x_0 \in (\theta_{l-1}, \theta_l)\right]}}_{\triangleq \Gamma_1: \quad \text{function of } u_0 \text{ and } \mathcal{E}_1} + (1+p+a^2)\sigma_w^2 \right\},$$

where $\sigma_2^2 = \sigma_w^2 + a^2\sigma_x^2$, and where $z_0 = l$ implies x_0 lies in the lth cell (θ_{l-1}, θ_l), and $z_1 = j$ implies x_1 lies in the jth cell (δ_{j-1}, δ_j), for $1 \leq j \leq N$. The term $\vartheta(\underline{r}, \bar{r})$ in the above equation is given by

(6.3)

$$\vartheta(\underline{r}, \bar{r}) = \left[-a\sigma_x\phi\left(\frac{\theta_l}{\sigma_x}\right)\Phi\left(r\frac{\sigma_2}{\sigma_w} - \theta_l\frac{a}{\sigma_w}\right) - \sigma_2\phi(r)\Phi\left(\frac{\theta_l}{\sigma_1} - r\frac{a\sigma_x}{\sigma_w}\right) \right.$$
$$\left. + a\sigma_x\phi\left(\frac{\theta_{l-1}}{\sigma_x}\right)\Phi\left(r\frac{\sigma_2}{\sigma_w} - \theta_{l-1}\frac{a}{\sigma_w}\right) + \sigma_2\phi(r)\Phi\left(\frac{\theta_{l-1}}{\sigma_1} - r\frac{a\sigma_x}{\sigma_w}\right) \right]_{r=\underline{r}}^{\bar{r}},$$

where $\sigma_1^2 = \sigma_x^2\sigma_w^2/\sigma_2^2$ and $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution.

The quantization distortion term Γ_1 in (6.2) possesses symmetry w.r.t. translations, as defined in (4.7). Thus, for any value of the control signal u_0, the minimum value is given by $\Gamma_1^*(\mathcal{E}_1)$, a term that depends only on the encoder. Then, the cost-to-go w.r.t. the control signal u_0 comprises only the terms in the first row in (6.2). Hence, we obtain separation. Furthermore, the optimal control is given by the certainty equivalence law $u_0^{\text{CE}} = -a(p + a^2\frac{q}{q+1})/(p + q + a^2\frac{q}{q+1})\hat{x}_{0|0}$. Thus, the certainty equivalence property holds for this setup.

We illustrate symmetry w.r.t. translations in Figure 7. For the choice of parameters $a = 1$, $p = 1$, and $q = 1$, we evaluate the quantization distortion term Γ_1 and show that the minimum that this function attains over the range of the quantizer threshold δ_1 is invariant over u_0. To evaluate the cost-to-go, we make an arbitrary choice, $\delta_0 = 0$, for the quantizer threshold at time $k = 0$, and we compute the estimates and probabilities using this choice. ◻

6.2. Optimal constrained encoder. We now impose a restriction on the choice of encoder parameters. The one-bit quantizer that we considered in the previous example selects two semi-infinite intervals as the quantizer cells: $\Delta_1 = (-\infty, \delta_k]$ and $\Delta_2 = (\delta_k, \infty)$. We restrict the choice of the quantizer threshold to a constraint set, such that $\delta_k \in \Theta$. In the following example, we see that separation is not optimal.

Example 5. Consider the same setup as in Example 4, with the restriction that the quantizer threshold be chosen from the interval $\Theta = (-1, 1)$. The quantizer thresholds $\delta_0 \in \Theta$ and $\delta_1 \in \Theta$ are to be chosen along with the control signals u_0 and u_1, to jointly minimize the two-step horizon control cost.

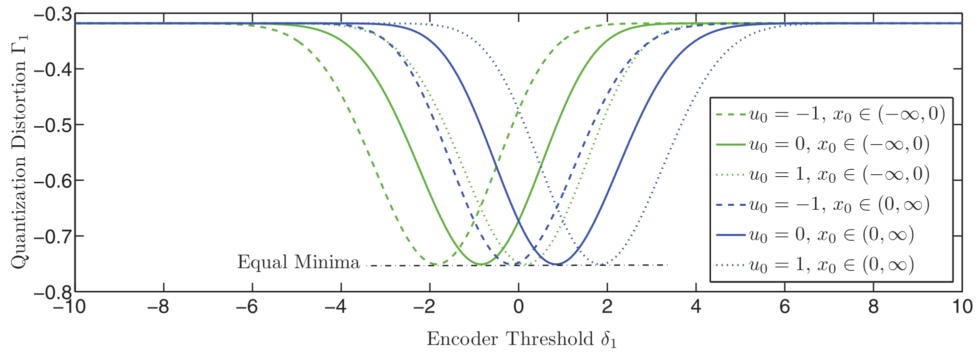

FIG. 7. *Illustration of symmetry w.r.t. translations of the quantization distortion term* Γ_1 *in* (6.2). *Different values of* u_0 *result in the same minimum value for* Γ_1 *thus resulting in separation and certainty equivalence in Example* 4.

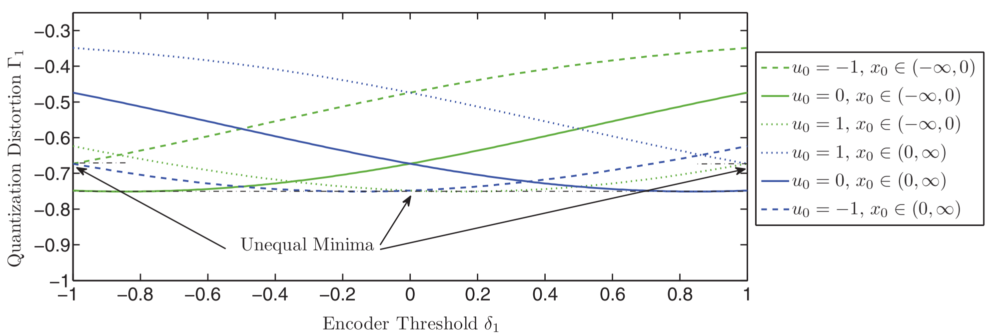

FIG. 8. *This plot illustrates the lack of symmetry w.r.t. translations of* Γ_1*, when the quantizer thresholds are restricted to be chosen from an interval, such as in Example* 5. *Different values of* u_0 *do not result in the same minimum value for* Γ_1 *over the range of* δ_1*, thus resulting in a lack of separation and certainty equivalence.*

We follow the same procedure as before. The optimal control signal u_1 is given by the certainty equivalence law as $u_1^* = u_1^{\mathrm{CE}}$. This gives us the same cost-to-go V_0 from (6.1). For the parameters $a = 1$, $p = 1$, and $q = 1$, we plot Γ_1 over a range of quantizer thresholds $\delta_1 \in \Theta$, for three arbitrary choices of u_0, in Figure 8. By restricting the range of quantizer thresholds to Θ, the curves do not reach their minima at the same δ_1 as in Figure 7. In particular, the minima for $u_0 = -1$, when $x_0 \in (-\infty, 0)$, and $u_0 = 1$, when $x_0 \in (0, \infty)$ are higher than before. Thus, the minimum value of Γ_1 obtained over the range of δ_1 now varies depending on the choice of u_0. Consequently, there is no longer a symmetry w.r.t. translations, and separation is not achieved. Furthermore, the optimal control signal u_0^* must be chosen along with δ_1^* to optimize the entire cost-to-go including the term Γ_1. Thus, u_0^* must minimize a nonquadratic expression in this problem, and cannot be chosen independently of the encoding policy. Hence, separation of the controller and encoder is no longer optimal. $\square$

6.3. Optimal constrained controller. We now remove the restriction on the encoder parameters, and instead impose a restriction on controls. Specifically, all control values must come from a specified discrete set $\mathcal{U}$.

Example 6. Consider the same setup as in Example 4, with the restriction that the control signal be chosen from a discrete set $\mathcal{U} = \{-1, 0, 1\}$. The quantizer thresholds δ_0 and δ_1 are to be chosen along with the control signals $u_0 \in \mathcal{U}$ and $u_1 \in \mathcal{U}$, to jointly minimize the two-step horizon control cost.

The unconstrained minimizer for the cost-to-go at the terminal time is given by the certainty equivalent control u_1^{CE}. The best we can do, given the constraint set $\mathcal{U}$, is to choose the control value from the discrete set $\mathcal{U}$ that results in the lowest cost-to-go. Using this principle, we find the optimal control signal u_1^* to be

$$u_1^* = \begin{cases} -1 & \text{if } \hat{x}_{1|1} \geq (q+1)/2a\ , \\ \ \ 0 & \text{if } (q+1)/2a \geq \hat{x}_{1|1} \geq -(q+1)/2a\ , \\ \ \ 1 & \text{if } \hat{x}_{1|1} \leq -(q+1)/2a\ . \end{cases}$$

The optimality regions are identified by comparing $\min_{u_1 \in \mathcal{U}} V_1(u_1)$ evaluated at each permissible value of u_1, and determining the switching points. The cost-to-go V_0, obtained by averaging over the three different cost-to-go functions obtained at time $k = 1$, is given by

$$\begin{aligned} V_0 = (1+p+a^2)\sigma_w^2 + \mathbb{E}\Big[& a^2(p+a^2)x_0^2 + (q+p+a^2)u_0^2 + 2a(p+a^2)x_0u_0 \\ & + (-2a\hat{x}_{1|1} + q + 1)\mathbb{1}_{\left\{\hat{x}_{1|1} \geq \frac{q+1}{2a}\right\}} \\ & + (2a\hat{x}_{1|1} + q + 1)\mathbb{1}_{\left\{\hat{x}_{1|1} \leq -\frac{q+1}{2a}\right\}} \Big| z_0 \Big]\ . \end{aligned}$$

We denote the terms in the above cost-to-go that directly depend on the choice of the encoder threshold δ_1 as Γ_1^{RC}. Using the expression for $\hat{x}_{1|1}$ and the posterior density for x_1 (derived in the appendix of [34]), we compute Γ_1^{RC} as

$$\begin{aligned} \Gamma_1^{\mathrm{RC}} &= \mathbb{E}\Big[(-2a\hat{x}_{1|1} + q + 1)\mathbb{1}_{\left\{\hat{x}_{1|1} \geq \frac{q+1}{2a}\right\}} + (2a\hat{x}_{1|1} + q + 1)\mathbb{1}_{\left\{\hat{x}_{1|1} \leq -\frac{q+1}{2a}\right\}} \Big| z_0\Big] \\ &= \sum_{j=1}^{N} \frac{\mathbb{P}\left[x_0 \in (\theta_{l-1}, \theta_l), x_1 \in (\delta_{j-1}, \delta_j)\right]}{\mathbb{P}\left[x_0 \in (\theta_{l-1}, \theta_l)\right]} \Big((-2a\hat{x}_{1|1} + q + 1)\mathbb{1}_{\left\{\hat{x}_{1|1} \geq \frac{q+1}{2a}\right\}} \\ &\qquad + (2a\hat{x}_{1|1} + q + 1)\mathbb{1}_{\left\{\hat{x}_{1|1} \leq -\frac{q+1}{2a}\right\}} \Big)\ . \end{aligned}$$

Evaluating the above expression for parameters $a = 1$, $p = 1$, $q = 1$, and some arbitrary choice of quantizer threshold $\delta_0 = 0$, we plot Γ_1^{RC} over a range of quantizer thresholds δ_1, for different choices of u_0 from the discrete set $\mathcal{U}$, in Figure 9. Notice that the minimum values of Γ_1^{RC} obtained over the range of δ_1 vary depending on the choice of u_0. In other in words, there is no symmetry w.r.t. translations. Consequently, a separation in design of the controller and encoder is no longer optimal. □

7. Conclusions. We examined some two-agent networked control problems and found that the dual effect is present, in general. This makes these two-agent decision problems hard, and we do not get the simplifications that are obtainable for the classical single-agent LQ problem. Thus it is not surprising that neither controls-forgetting encoders nor certainty equivalence controls are optimal for two-agent decision problems that have constraints on encoders or controls. However, it is surprising that for the dynamic encoder-controller design problem, it is optimal to apply a controls-forgetting encoder in combination with certainty equivalence controls.

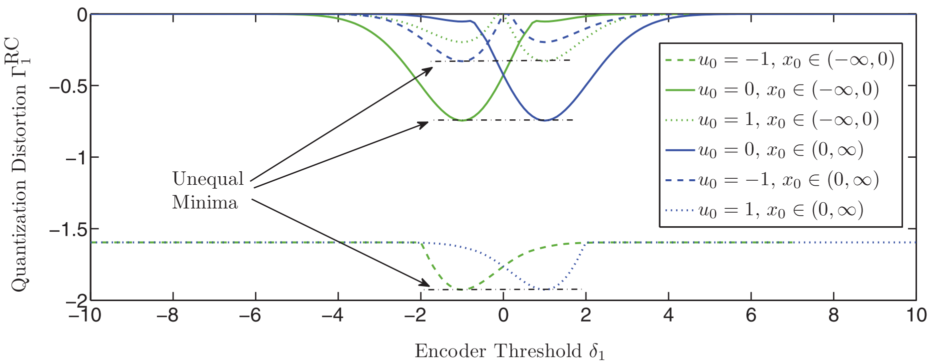

FIG. 9. *This plot illustrates the lack of symmetry w.r.t. translations for Γ_1^{RC}, when the controls are restricted to be chosen from a discrete set $\mathcal{U}$ in Example* 6. *Different values of u_0 do not result in the same minimum value for Γ_1^{RC} over the range of δ_1, thus resulting in the lack of separation and certainty equivalence.*

Our results also have consequences for some one-agent decision problems. For the networked control loop in Figure 1, consider the one-agent decision problem obtained by fixing the encoder and asking for the design of the controller alone. A theorem of Bar-Shalom and Tse [7] applies to the special case where the prescribed encoder is restricted to be memoryless. This is precisely the case where the encoding maps may depend on time, but not on inputs or outputs of the encoder. For such encoders the theorem in [7] says that certainty equivalence controls are optimal *if and only if* there is no dual effect in the loop. Such a neat dichotomy need not hold in cases where the prescribed encoder is dynamic. To see this, consider a two-agent dynamic design problem where it is optimal to apply the pair $(\mathcal{E}^{CF,*}, \mathcal{K}^{CE})$ of controls-forgetting encoder and certainty equivalence controls. Since the encoding strategy $\mathcal{E}^{CF,*}$ is controls forgetting, its individual encoding polices take the form

$$\xi_t^{CF,*}\left(x_t\,;\,\mathcal{D}^{\mathrm{con}}_{(t-1)^-}\right) = \epsilon_t^*\left(\zeta_t\,;\,\{z_i\}_\tau^{t-1}\right).$$

Consider the new encoding strategy $\mathcal{E}^\dagger$ whose individual encoding maps take the form

$$\begin{aligned}\xi_t^\dagger\left(x_t\,;\,\mathcal{D}^{\mathrm{con}}_{(t-1)^-}\right) &= \epsilon_t^*\left(\zeta_t + \sum_{i=0}^{t-1} a^{t-1-i}u_i - \sum_{i=0}^{t-1} a^{t-1-i}u_i^{CE}\,;\,\{z_i\}_\tau^{t-1}\right)\\ &= \epsilon_t^*\left(x_t - \sum_{i=0}^{t-1} a^{t-1-i}u_i^{CE}\,;\,\{z_i\}_\tau^{t-1}\right),\end{aligned}$$

where u_i are the controls generated by the actual controller, and u_i^{CE} are the certainty equivalence controls as per the law (4.8). Since in general $u_i \neq u_i^{CE}$, this encoder is not controls forgetting, which means that the loop gets a second-order dual effect. But notice that the pair $(\mathcal{E}^\dagger, \mathcal{K}^{CE})$ incurs exactly the same performance as the pair $(\mathcal{E}^{CF,*}, \mathcal{K}^{CE})$. Hence for the one-agent decision problem obtained by fixing the encoder to be $\mathcal{E}^\dagger$, it is optimal to apply certainty equivalence controls. This is an example of a one-agent decision problem for a linear plant and quadratic performance costs where certainty equivalence controls are optimal, even though there is a dual effect in the loop. Thus we can conclude that the theorem of Bar-Shalom and Tse cannot generalize to the case of dynamic encoders.

Acknowledgments. Adam Molin first suggested to us that for event-triggered systems, the separated solution may not be optimal for the zero-order hold control design problem, and yet be optimal for the dynamic encoder-controller design problem. Comments from the anonymous reviewers have considerably improved the presentation. To Lei Bao, Mikael Skoglund, Henrik Sandberg, Ashutosh Nayyar, John S. Baras, and Armand Makowski we are grateful for many discussions over the last few years.

REFERENCES

[1] B. R. Andrievsky, A. S. Matveev, and A. L. Fradkov, *Control and estimation under information constraints: Toward a unified theory of control, computation and communications*, Autom. Remote Control, 71 (2010), pp. 572–633.

[2] K. J. Åström and B. Bernhardsson, *Comparison of Riemann and Lebesgue sampling for first order stochastic systems*, in Proceedings of the 41st IEEE conference on Decision and Control, Las Vegas, NV, IEEE, Piscataway, NJ, 2002, pp. 2011–2016.

[3] J. Baillieul and P. J. Antsaklis, *Control and communication challenges in networked real-time systems*, Proceedings of the IEEE, 95 (2007), pp. 9–28.

[4] R. Bansal and T. Başar, *Stochastic teams with nonclassical information revisited: When is an affine law optimal?*, IEEE Trans. Automat. Control, 32 (1987), pp. 554–559.

[5] R. Bansal and T. Başar, *Simultaneous design of measurement and control strategies for stochastic systems with feedback*, Automatica J. IFAC, 25 (1989), pp. 679–694.

[6] L. Bao, M. Skoglund, and K. H. Johansson, *Iterative encoder-controller design for feedback control over noisy channels*, IEEE Trans. Automat. Control, 56 (2011), pp. 265–278.

[7] Y. Bar-Shalom and E. Tse, *Dual effect, certainty equivalence, and separation in stochastic control*, IEEE Trans. Automat. Control, 19 (1974), pp. 494–500.

[8] B. Bernhardsson, *Dual control of a first-order system with two possible gains*, Internat. J. Adapt. Control Signal Process., 3 (1989), pp. 15–22.

[9] D. P. Bertsekas and S. E. Shreve, *Stochastic Optimal Control: The Discrete Time Case*, Math. Sci. Eng. 139, Academic Press, New York, 1978.

[10] V. S. Borkar and S. K. Mitter, *LQG control with communication constraints*, in Communications, Computation, Control, and Signal Processing, A. Paulraj, V. Roychowdhury, and C. D. Schaper, eds., Springer, New York, 1997, pp. 365–373.

[11] V. S. Borkar, S. K. Mitter, and S. Tatikonda, *Markov control problems under communication constraints*, Commun. Inf. Syst., 1 (2001), pp. 15–32.

[12] V. S. Borkar, S. K. Mitter, and S. Tatikonda, *Optimal sequential vector quantization of Markov sources*, SIAM J. Control Optim., 40 (2001), pp. 135–148.

[13] R. E. Curry, *Estimation and Control with Quantized Measurements*, M.I.T. Press, Cambridge, MA, 1970.

[14] A. A. Fel′dbaum, *Dual-control theory.* I, Automat. Remote Control, 21 (1960), pp. 874–880.

[15] X. Feng and K. A. Loparo, *Active probing for information in control systems with quantized state measurements: A minimum entropy approach*, IEEE Trans. Automat. Control, 42 (1997), pp. 216–238.

[16] T. R. Fischer, *Optimal quantized control*, IEEE Trans. Automat. Control, 27 (1982), pp. 996–998.

[17] J. S. Freudenberg, R. H. Middleton, and J. H. Braslavsky, *Minimum variance control over a Gaussian communication channel*, IEEE Trans. Automat. Control, 56 (2011), pp. 1751–1765.

[18] M. Fu, *Lack of separation principle for quantized linear quadratic Gaussian control*, IEEE Trans. Automat. Control, 57 (2012), pp. 2385–2390.

[19] K. Gatsis, A. Ribeiro, and G. J. Pappas, *Optimal power management in wireless control systems*, IEEE Trans. Automat. Control, 59 (2014), pp. 1495–1510.

[20] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*, Kluwer Academic, Norwell, MA, 1991.

[21] G. C. Goodwin, E. I. Silva, and D. E. Quevedo, *A brief introduction to the analysis and design of networked control systems*, in Proceedings of the IEEE conference on Decision and Control, IEEE, Piscataway, NJ, 2008, pp. 1–13.

[22] D. P. Joseph and T. J. Tou, *On linear control theory*, Trans. AIEE, Part II Appl. Ind., 80 (1961), pp. 193–196.

[23] A. A. Kulkarni and T. P. Coleman, *An optimizer's approach to stochastic control problems with nonclassical information structures*, in Proceedings of the IEEE Control and Decision Conference, IEEE, Piscataway, NJ, 2012, pp. 154–159.
[24] R. E. Larson, *Optimum quantization in dynamic systems*, IEEE Trans. Automat. Control, 12 (1967), pp. 162–168.
[25] G. M. Lipsa and N. C. Martins, *Optimal memoryless control in Gaussian noise: A simple counterexample*, Automatica J. IFAC, 47 (2011), pp. 552–558.
[26] A. Mahajan and D. Teneketzis, *Optimal design of sequential real-time communication systems*, IEEE Trans. Inform. Theory, 55 (2009), pp. 5317–5338.
[27] A. S. Matveev and A. V. Savkin, *The problem of LQG optimal control via a limited capacity communication channel*, Systems Control Lett., 53 (2004), pp. 51–64.
[28] R. Marleau and J. Negro, *Comments on "Optimum quantization in dynamic systems"*, IEEE Trans. Automat. Control, 17 (1972), pp. 273–274.
[29] S. K. Mitter, *Control with limited information*, Eur. J. Control, 7 (2001), pp. 122–131.
[30] A. Molin and S. Hirche, *On the optimality of certainty equivalence for event-triggered control systems*, IEEE Trans. Automat. Control, 58 (2013), pp. 470–474.
[31] G. N. Nair, F. Fagnani, S. Zampieri, and R. J. Evans, *Feedback control under data rate constraints: An overview*, Proc. IEEE, 95 (2007), pp. 108–137.
[32] A. Nayyar, *Sequential Decision Making in Decentralized Systems*, Ph.D. thesis, The University of Michigan, Ann Arbor, MI, 2011.
[33] M. Rabi and K. H. Johansson, *Optimal stopping for updating controls*, in Proccedings of the Second International Workshop on Sequential Methods, UTT, Troyes, France, 2009.
[34] M. Rabi, C. Ramesh and K. H. Johansson, *Separated design of encoder and controller for networked linear quadratic optimal control*, preprint, arXiv:1405.0135 [math.OC], 2014.
[35] J. G. Root, *Optimum control of non-Gaussian linear stochastic systems with inaccessible state variables*, SIAM J. Control, 7 (1969), pp. 317–323.
[36] G. Sauer and J. Melsa, *Stochastic control with continuously variable observation costs for a class of discrete nonlinear systems*, IEEE Trans. Automat. Control, 19 (1974), pp. 234–239.
[37] L. Shaw, *Optimum stochastic control*, in Disciplines and Techniques of Systems Control, J. Peschon, ed., Blaisdell, New York, 1965, pp. 125–185.
[38] C. Striebel, *Sufficient statistics in the optimum control of stochastic systems*, J. Math. Anal. Appl., 12 (1965), pp. 576–592.
[39] S. Tatikonda, A. Sahai, and S. K. Mitter, *Stochastic linear control over a communication channel*, IEEE Trans. Automat. Control, 49 (2004), pp. 1549–1561.
[40] H. van de Water and J. C. Willems, *The certainty equivalence property in stochastic control theory*, IEEE Trans. Automat. Control, 26 (1981), pp. 1080–1087.
[41] P. Varaiya and J. C. Walrand, *On delayed sharing patterns*, IEEE Trans. Automat. Control, 23 (1978), pp. 443–445.
[42] P. Varaiya and J. C. Walrand, *Causal coding and control for Markov chains*, Systems Control Lett., 3 (1983), pp. 189–192.
[43] J. C. Walrand and P. Varaiya, *Optimal causal coding-decoding problems*, IEEE Trans. Inform. Theory, IT-29 (1983), pp. 814–820.
[44] H. S. Witsenhausen, *A counterexample in stochastic optimum control*, SIAM J. Control, 6 (1968), pp. 131–147.
[45] H. S. Witsenhausen, *Separation of estimation and control for discrete time systems*, Proc. IEEE, 59 (1971), pp. 1557–1566.
[46] W. M. Wonham, *On the separation theorem of stochastic control*, SIAM J. Control, 6 (1968), pp. 312–326.
[47] W. Wu and A. Arapostathis, *Optimal control of stochastic systems with costly observations - the general Markovian model and the LQG problem*, in Proceedings of the American Control Conference, vol. 1, American Automatic Control Council Evanston, IL, 2005, pp. 294–299.
[48] S. Yüksel, *On optimal causal coding of partially observed Markov sources in single and multiterminal Settings*, IEEE Trans. Inform. Theory, 59 (2013), pp. 424–437.
[49] S. Yüksel, *Jointly optimal LQG quantization and control policies for multi-dimensional systems*, IEEE Trans. Automat. Control, 59 (2014), pp. 1612–1617.
[50] A. Zaidi, S. Yüksel, T. Oechtering, and M. Skoglund, *On optimal policies for control and estimation over Gaussian relay channels*, Automatica J. IFAC, 49 (2013), pp. 2892–2897.