Rate-limited control of systems with uncertain gain

Victoria Kostina*, Yuval Peres†, Miklós Z. Rácz†, Gireeja Ranade†
*California Institute of Technology, Pasadena, CA
† Microsoft Research, Redmond, WA
vkostina@caltech.edu, peres@microsoft.com, miracz@microsoft.com, giranade@microsoft.com

***Abstract*—Controlling and stabilizing systems involves countering the impact of explicit communication constraints in addition to inherent system model parameter uncertainty and random noise. Here we use an information-theoretic approach to jointly tackle all three issues and understand their interactions. Our main result bounds the minimum communication rate required for the mean-square stability of a system with uncertain system gain. Moreover, our techniques extend to provide a finer characterization of the required rate when specific finite bounds on the second moment of the state are desired.**

1. Introduction

Control systems with physically distributed sensors and actuators are essential to building autonomous agents such as self-driving cars and drones, and control algorithms for such systems must be implemented over communication networks. These systems may be rapidly evolving and difficult to model at the timescale that the control must act at. As a result, system designers must account for information bottlenecks as well as uncertain model parameters as they allocate limited resources in designing the system.

This paper studies the interactions between model uncertainty and rate limitations. We consider the control of the following system:

$$X_{t+1} = A_t X_t + U_t + W_t. \tag{1}$$

Here X_t is the (scalar) state of the system at time t. The system gains $\{A_t\}_{t\geq 0}$ are drawn i.i.d. from a known distribution P_A (having a density) and model the uncertainty the controller has about the system. The additive disturbances $\{W_t\}_{t\geq 0}$ are also i.i.d. from a known distribution P_W. The controller chooses the control U_t causally based on observations $Y_0^t := \{Y_0, \ldots, Y_t\}$ that are transmitted over a rate-limited noiseless channel by an encoder co-located with the system. Our goal is to stabilize the system in a mean-square sense, which thus requires the co-design of an encoder-controller pair.

Without the rate limitation, the system in (1) is mean-square stabilizable if and only if $\sigma_A^2 < 1$, where σ_A^2 is the variance of A_t (see [1]). If the system gain is constant, i.e., $A_t = a$ for every t, then a communication rate of $R > \log|a|$ is necessary and sufficient for mean-square stabilizability (see, e.g., [2], [3], [4], [5]). We simultaneously consider a random system gain and limited communication rate in order to provide a perspective on the parameter uncertainty that is compatible with information-theoretic rate limits.

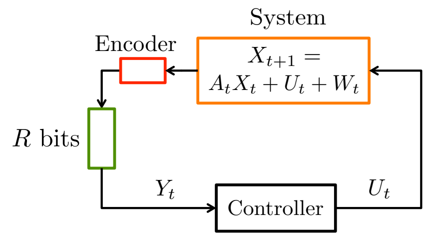

Fig. 1. Rate-limited control: The control U_t can be any function of the observations Y_0^t. The encoder chooses Y_t, an encoding of the state X_t, such that $H(Y_t \mid Y_0^{t-1}) \leq R$. The system gains $\{A_t\}_{t\geq 0}$ are i.i.d. random variables, and so are the additive disturbances $\{W_t\}_{t\geq 0}$.

Our goal is to estimate the critical rate for stability, R^*, defined as the infimum of R for which there exists a variable-rate transmission process $\{Y_t\}_{t\geq 0}$ and a causal control strategy $\{U_t\}_{t\geq 0}$ satisfying the following:

- Transmission rate: $H(Y_t \mid Y_0^{t-1}) \leq R$ for every t;
- Mean-square stability: $\limsup_{t\to\infty} \mathbb{E}[X_t^2] < \infty$.

The following theorem illustrates our results when A_t is Gaussian.

Theorem 1.1. *Assume P_A is Gaussian with mean μ_A and variance $\sigma_A^2 < 1$, and that P_W has a density. Then*

$$R^* \geq \max\left\{ \frac{1}{2}\log\frac{\mu_A^2}{1-\sigma_A^2} - \frac{1}{2}\,,\ \frac{1}{2}\log\left(\mu_A^2 + \sigma_A^2\right) \right\},$$

and further

$$R^* \leq \frac{1}{2}\log\frac{\mu_A^2}{1-\sigma_A^2} + \log\left(1 + \log\left(3 + \frac{\mu_A}{\sqrt{(1-\sigma_A^2)}}\right)\right) + 2 + \log e.$$

The interesting regime for this result is when the system is not self-stabilizing, i.e., when $\mu_A^2 + \sigma_A^2 > 1$. Our primary region of interest is when σ_A^2 goes to 1, in which case the rate goes to infinity, even though μ_A is bounded. In particular, when $1 - \sigma_A^2 = \epsilon \downarrow 0$, and $\mu_A = O(1)$, we obtain that

$$\frac{1}{2}\log\frac{1}{\epsilon} - O(1) \leq R^* \leq \frac{1}{2}\log\frac{1}{\epsilon} + \log\log\frac{1}{\epsilon} + O(1), \quad (2)$$

thus the critical rate is $\frac{1}{2}\log\frac{1}{\epsilon}$ up to first order in $\frac{1}{\epsilon}$.

The second term in the lower bound of Thm. 1.1 is effective in the regime where σ_A^2 is close to zero. Note that when $\sigma_A^2 = 0$, this reduces to the known $R^* \geq \log|a|$ bound from the data-rate theorems.

A more precise version of this theorem is presented in Sec. 3, together with further results for more general distributions P_A and P_W. We believe that it should be possible to close the additive gap in (2) above. This gap might be connected to the tightness of Shannon's lower bound [6], which is known to be tight in the low-distortion regime but can be loose in the high-distortion regime [7], [8], [9], [10], [11].

Finally, while mean-square stability captures the limits of stabilizability, it is desirable to understand how much extra rate is required on top of this absolute minimum in order to achieve a particular bound on the second moment of the state. Our techniques extend to provide refined bounds on the extra rate required to achieve the tighter bound. These results are also provided in Sec. 3.

A. *Proof techniques*

Our proof of the converse uses Shannon's rate-distortion lower bound to bound the expected distortion of the state. The technique allows us to avoid a precise characterization of the complex distributions arising at each step. Furthermore, at each time t, we provide side information to the controller in the form of either the precise previous state of the system (i.e., X_{t-1}) or the random gain of the previous step (i.e., A_{t-1}), and these two cases give the two parts of the converse bound.

The achievable scheme uses a simple uniform quantization to achieve the bound.

B. *Related work*

Our setup is inspired by the uncertainty threshold principle [1] and the extensive work to understand data-rate theorems [3], [5]. The results in [12, Ch. 2] also motivated this exploration, since they suggest the possibility of a unifying information-theoretic perspective that can quantify the joint impact of parameter uncertainty and rate constraints. We aim to develop an understanding that can mesh with our understanding of the impact of multiplicative noise on the actuation channel [13] and on the observation channel [14], [15].

This informational perspective complements the explorations of parameter uncertainty in the robust control literature [16], [17]. Our interest in this paper is in understanding how large uncertainty on the system interacts with an explicit data-rate constraint.

The rate-limited control of a system with uncertain system gain was first considered by Martins et al. [18]. Their converse bound gives a result of the flavor that if the system is stabilizable then $R > \frac{1}{2}\log(\mu_A^2 + \sigma_A^2)$, a result that we also obtain. This is of interest when σ_A^2 is small, but is not tight when σ_A^2 goes to 1. Phat et al. [19] also consider rate-limited control with uncertain parameters from a robust control perspective. Their setup differs from ours in that it is not stochastic, and only considers bounded support for the uncertainty on the parameters. With this they provide a uniform quantization scheme that can stabilize the system.

Recently, Okano and Ishii made progress on understanding rate-limited control of uncertain systems from a worst-case perspective [20], [21], [22]. However, they also bound the support of the parameter uncertainty and do not consider additive noise in their model. The achievable scheme in [22] proposes a non-uniform optimal quantizer for their problem that uses bins with logarithmically decreasing lengths, with the bins closest to zero being the largest. However, this cannot work in the setting where both A_t and X_t can have unbounded support, as is the case in our work. We use a uniform quantizer with variable-rate transmission since we assume that the support of X_t, A_t, W_t might be unbounded.

Beyond mean-square stability our work also provides a bound on the impact of additive noise in the system and calculates the extra rate required to ensure that the second moment of the system state is bounded to within a specified cost. This builds on work in [23], and is similar to bounds of Nair et al. [5] for rate-limited systems without parameter uncertainty. It is also closely related to the idea of sequential rate-distortion [24].

2. Problem Setup

We consider the scalar control system in (1) with system state X_t at time t. The initial state X_0 is random, and has finite second moment, i.e., $\sigma_{X_0}^2 < \infty$. The system gains $\{A_t\}_{t\geq 0}$ are drawn i.i.d. from a distribution P_A, with mean μ_A and variance σ_A^2. The W_t are drawn i.i.d. from P_W with mean 0 and variance σ_W^2, and independently from everything else in the system. Both P_A and P_W are known to the encoder and the controller.

At each time t, the encoder observes the system state X_t and transmits message Y_t (that can depend on all previous messages Y_0^{t-1} and states X_0^t) over a noiseless rate-limited channel to the controller. We allow $Y_t \in \{0,1\}^\star := \cup_{n\geq 0}\{0,1\}^n$, and let $\ell(Y_t)$ represent the length of Y_t. The rate-limit we assume is on the entropy of the transmitted message: we must have $H(Y_t \mid Y_0^{t-1}) \leq R$. The controller takes action U_t at time t based on all the observations it has received till that point, Y_0^t. Such an encoder-controller pair with control strategy $\{U_t\}_{t\geq 0}$ is said to have a *rate-limit of* R.

The controller's objective is to ensure mean-square stability of the system as defined below.

Definition 2.1. *We say that the system* (1) *is* mean-square stabilizable *if there exists a causal encoder-controller pair such that* $\sup_t \mathbb{E}[X_t^2] < \infty$.

In addition to mean-square stabilizability, we also define a more refined notion of stability, when a precise bound on the second moment is desired.

Definition 2.2. *We say that the system* (1) *is* mean-square stabilizable with cost D, *if there exists a causal encoder-controller pair such that* $\limsup_t \mathbb{E}[X_t^2] \leq D$.

Throughout the paper all logarithms are base 2.

3. Results

We are now ready to present our results, starting with the converse, and then showing that this bound is essentially tight. Recall from the uncertainty threshold principle that if $\sigma_A^2 \geq 1$, then it is impossible to stabilize the system in a mean-square sense. This is due to the fact that the multiplicative system gain acts on the state of the system after the control has been applied.[1] On the other hand, if $\mu_A^2 + \sigma_A^2 < 1$, then the system (1) is trivially stabilizable by applying no control ($U_t = 0$ for every t). We thus assume in the following that $\sigma_A^2 < 1$ and $\mu_A^2 + \sigma_A^2 \geq 1$ to focus on the interesting cases.

A. Characterizing mean-square stability

We start by showing the limits of stabilizability when the distribution P_A has a density p_A. For a random variable X having a density on $\mathbb{R}$, we recall that the entropy-power of X is given by $N(X) = \frac{1}{2\pi e}2^{2h(X)}$, where $h(X)$ is the differential entropy of X.

[1]One can show that (1) implies $\mathbb{E}[X_{t+1}^2] \geq \sigma_A^2\mathbb{E}[X_t^2] + \sigma_W^2$, so when $\sigma_A^2 > 1$, then $\mathbb{E}[X_t^2]$ grows exponentially in t. When $\sigma_A^2 = 1$, we have $\mathbb{E}[X_t^2] \geq t\sigma_W^2$, which is unbounded in t when $\sigma_W^2 > 0$. We ignore the knife-edge case of $\sigma_A^2 = 1$ and $\sigma_W^2 = 0$.

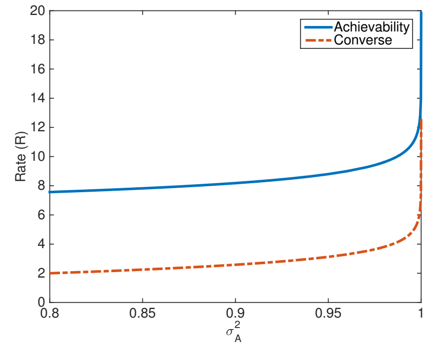

Fig. 2. This figure plots the rate bounds from Thm. 3.1 and Thm. 3.2 as a function of σ_A^2, assuming P_A is Gaussian with $\mu_A = 2$ and also $\sigma_W^2 = 1$. The achievability is the solid blue line that plots (4) and the dotted red line is the converse from (3).

Theorem 3.1. *Assume that P_A has a density p_A, that P_W has a density p_W, and that $\sigma_A^2 < 1$. If the system* (1) *is mean-square stabilizable using a causal encoder-controller pair with rate-limit R, then*

$$R \geq \max\left\{\frac{1}{2}\log\frac{\mu_A^2 N(A)}{1-\sigma_A^2}, \frac{1}{2}\log(\mu_A^2+\sigma_A^2)\right\}. \quad (3)$$

This converse result is the combination of two bounds stated together. The assumption above that P_A has a density is only used to obtain the $\frac{1}{2}\log\frac{\mu_A^2 N(A)}{1-\sigma_A^2}$ bound, while the assumption that P_W has a density is only used to obtain the other bound of $\frac{1}{2}\log(\mu_A^2+\sigma_A^2)$.

The converse result of Theorem 3.1 is complemented by the following sufficient condition, which is achieved using a simple uniform quantization scheme.

Theorem 3.2. *Assume that $\sigma_A^2 < 1$. If*

$$\begin{aligned} R \geq {} & \frac{1}{2}\log\frac{\mu_A^2}{1-\sigma_A^2} \\ & + \log\left(1+\log\left(3+\frac{\mu_A}{\sqrt{(1-\sigma_A^2)}}\right)\right) + 2 + \log e, \quad (4) \end{aligned}$$

then the system (1) *can be mean-square stabilized using a causal encoder-controller pair with rate-limit R.*

When P_A is Gaussian, then $N(A) = \sigma_A^2$ and Theorem 3.1 and Theorem 3.2 give Theorem 1.1.

Both the necessary and sufficient bounds on the rate blow up as $\frac{1}{2}\log\frac{1}{1-\sigma_A^2}$ as $\sigma_A^2 \nearrow 1$, as seen in Fig. 2.

B. Characterizing mean-square stability with finite cost

By considering only the effect of the uncertain system gain acting on the additive noise, one can show (see Section 4-B) that for the system (1), under any causal encoder-controller pair, we must have:

$$\liminf_{t\to\infty} \mathbb{E}[X_t^2] \geq \frac{\sigma_W^2}{1-\sigma_A^2}. \tag{5}$$

For any $D > \frac{\sigma_W^2}{1-\sigma_A^2}$, we characterize the critical rate R required to mean-square stabilize with cost D the system (1) (recall Def. 2.2). These results bound how the required rate blows up as $D \searrow \frac{\sigma_W^2}{1-\sigma_A^2}$.

Theorem 3.3. *Assume that P_A has a density p_A, that P_W has a density p_W, and that $\sigma_A^2 < 1$. If the system* (1) *is mean-square stabilizable with cost D using a causal encoder-controller pair with rate-limit R, then*

$$\begin{aligned} R \geq \max\Bigg\{ &\frac{1}{2}\log\frac{\mu_A^2 N(A)}{1-\sigma_A^2}, \frac{1}{2}\log(\mu_A^2+\sigma_A^2), \\ &\frac{1}{2}\log\left(\frac{\mu_A^2}{1-\sigma_A^2} \times \frac{D\cdot N(A)+N(W)}{D-\frac{\sigma_W^2}{1-\sigma_A^2}}\right)\Bigg\}. \end{aligned} \tag{6}$$

Theorem 3.4. *For any $D > \frac{\sigma_W^2}{1-\sigma_A^2}$, the system* (1) *can be mean-square stabilized with cost D. This is achievable using an encoder-controller pair with rate-limit R for any*

$$R > \frac{1}{2}\log\frac{\mu_A^2}{1-\sigma_A^2} + \log M_D, \tag{7}$$

where

$$\begin{aligned} M_D = e \times &\left(\frac{3\sqrt{1-\sigma_A^2}}{\mu_A} + \sqrt{\frac{D+\sigma_{X_0}^2}{D-\frac{\sigma_W^2}{1-\sigma_A^2}}}\right) \\ &\times \left(1+\log\left(3+\sqrt{\frac{\mu_A^2}{1-\sigma_A^2}}\sqrt{\frac{D+\sigma_{X_0}^2}{D-\frac{\sigma_W^2}{1-\sigma_A^2}}}\right)\right). \end{aligned}$$

The third term in (6) represents an extra rate penalty that must be paid to contain the second moment of the system below D. The $\log M_D$ term in (7) plays a similar role for the quantization scheme. As $D \searrow \frac{\sigma_W^2}{1-\sigma_A^2}$, both terms blow up as $-\frac{1}{2}\log\left(D-\frac{\sigma_W^2}{1-\sigma_A^2}\right)$. This is plotted in Fig. 3, along with a simulation of the achievable strategy.

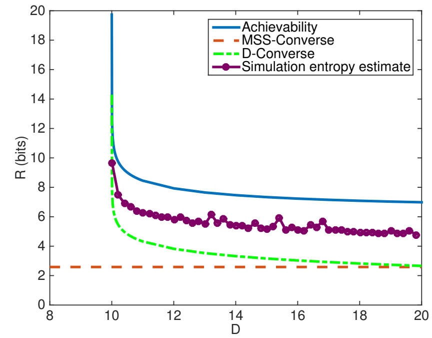

Fig. 3. This figure plots the rate bounds from Thm. 3.3 and Thm. 3.4 as a function of D, assuming P_A and P_W are Gaussian with $\mu_A = 2$, $\sigma_A^2 = .9$, and $\mu_W = 0$, $\sigma_W^2 = 1$, respectively. The achievability is the solid blue line that plots (7) and the dotted lines plot the converse bound for cost D (purple) and for mean-squared stability from Thm. 5.1 (red). The achievability and converse bounds blow up as $D \searrow \frac{\sigma_W^2}{1-\sigma_A^2} = 10$. The purple line with points demarcated plots an estimate of the encoder output entropy by simulating the proposed control strategy. This estimate was computed using a normal approximation of the state distribution at steady state.

4. Key observations

A. The evolution of the second moment

For any control $\{U_t\}_{t\geq 0}$ we have that

$$\begin{aligned} \mathbb{E}[X_{t+1}^2] &= \mathbb{E}[(A_tX_t+U_t+W_t)^2] \\ &= \mathbb{E}[(A_tX_t-\mu_AX_t+\mu_AX_t+U_t+W_t)^2] \\ &= \mathbb{E}[((A_t-\mu_A)X_t+\mu_AX_t+U_t)^2]+\mathbb{E}[W_t^2], \end{aligned} \tag{8}$$

since W_t is independent of all the other terms in the expectation. Now we note that $(A_t-\mu_A)X_t$ and $(\mu_AX_t+U_t)$ in (8) are orthogonal:

$$\begin{aligned} &\mathbb{E}[(A_t-\mu_A)X_t\cdot(\mu_AX_t+U_t)] \\ &\qquad = \mathbb{E}[(A_t-\mu_A)]\cdot\mathbb{E}[X_t(\mu_AX_t+U_t)] = 0, \end{aligned}$$

since A_t is independent of X_t and U_t. We can rewrite the expectation from (8) as:

$$\begin{aligned} \mathbb{E}[X_{t+1}^2] &= \mathbb{E}[(A_t-\mu_A)^2X_t^2]+\mathbb{E}[(\mu_AX_t+U_t)^2]+\sigma_W^2 \\ &= \sigma_A^2\mathbb{E}[X_t^2]+\mu_A^2\mathbb{E}[(X_t+U_t)^2]+\sigma_W^2, \end{aligned} \tag{9}$$

where, without loss of generality, we can scale the control U_t by $\frac{1}{\mu_A}$ in (9).

The expression (9) makes it evident that the controller only needs to focus on minimizing the term $\mathbb{E}[(X_t+U_t)^2]$; a scaling penalty of σ_A^2 from the first term is inevitable.

B. The minimum cost

Next we show that (5) is true. From (9), we have that for every t and for any control strategy $\{U_t\}_{t\geq 0}$,

$$\begin{aligned}\mathbb{E}[X_{t+1}^2] &= \sigma_A^2\mathbb{E}[X_t^2] + \mu_A^2\mathbb{E}[(X_t+U_t)^2] + \mathbb{E}[W_t^2]\\ &\geq \sigma_A^2\mathbb{E}[X_t^2] + \sigma_W^2.\end{aligned}$$

Since $\sigma_A^2 < 1$, recursion gives:

$$\mathbb{E}[X_{t+1}^2] \geq \sigma_A^{2t}\mathbb{E}[X_0^2] + \frac{1-\sigma_A^{2t}}{1-\sigma_A^2}\sigma_W^2.$$

Taking the limit $t\to\infty$ gives (5).

5. Converse proofs

We first tackle the converse for mean-square stability and then discuss mean-square stability with cost D.

A. Mean-square stability converse

The mean-square stability converse bound comes from identifying the information bottleneck that limits the controller. When σ_A^2 is large, we provide the controller with the system state for free to obtain a bound. However, when σ_A^2 is small, the primary control challenge comes from the uncertainty in the state itself, and we provide the random gain A_{t-1} as side information to the controller. These two bounds give the converse and we prove each term separately, starting with the former.

Before we prove the theorems in this section, we introduce some notation and previous results.

Definition 5.1. *For a random variable X, the* mean-square distortion-rate function *(or simply* distortion*) with rate R is defined as*

$$\mathbb{D}_R(X) := \inf_{Z:I(X;Z)\leq R} \mathbb{E}[(X+Z)^2],$$

where $I(X;Z)$ denotes the mutual information between X and Z. Similarly, the conditional distortion *of X given Y is defined as*

$$\mathbb{D}_R(X\mid Y) := \inf_{Z:I(X;Z|Y)\leq R} \mathbb{E}[(X+Z)^2].$$

Since $I(X_t;Y_t\mid Y_0^{t-1}) \leq H(Y_t\mid Y_0^{t-1})$, we have $I(X_t;Y_t\mid Y_0^{t-1})\leq R$, given our assumptions on the rate. Hence the definitions of distortion capture the second-moment error that is necessarily made when describing a random variable under the communication constraints.

Shannon's lower bound [6] states that for a continuous random variable X,

$$\mathbb{D}_R(X) \geq N(X)\cdot 2^{-2R}. \tag{10}$$

For a Gaussian X the inequality (10) holds with equality.

The entropy-power inequality [25] states that for independent random variables X and Y,

$$N(X+Y)\geq N(X)+N(Y). \tag{11}$$

Theorem 5.1. *Assume P_A has a density p_A and that $\sigma_A^2 < 1$. If the system* (1) *is mean-square stabilizable using a causal encoder-controller pair with rate-limit R, then*

$$R\geq \frac{1}{2}\log\frac{\mu_A^2 N(A)}{1-\sigma_A^2}.$$

Proof. Let $R < \frac{1}{2}\log\frac{\mu_A^2N(A)}{1-\sigma_A^2}$. Choose $\beta > 0$ such that $1+\beta \leq \sigma_A^2 + \frac{N(A)\mu_A^2 2^{-2R}}{1+\beta}$. Such a β exists because of the assumption on R. We show by induction that $\mathbb{E}[X_t^2] > c\cdot(1+\beta)^t$ for all $t\geq 0$, for some positive constant c. For the base case $t=0$, this inequality holds by choosing c small enough.

We will show that any control satisfies

$$\begin{aligned}\mathbb{E}[X_{t+1}^2] \geq \sigma_A^2\mathbb{E}[X_t^2] &+ N(A)\mu_A^2 2^{-2R}\mathbb{E}[X_{t-1}^2]\\ &+ N(W)\mu_A^2 2^{-2R} + \sigma_W^2.\end{aligned} \tag{12}$$

Once we have (12), we can apply the induction hypothesis, i.e., $\mathbb{E}[X_t^2] > c\ (1+\beta)^t$, to get:

$$\begin{aligned}\mathbb{E}[X_{t+1}^2] &\geq \sigma_A^2 c\ (1+\beta)^t + N(A)\mu_A^2 2^{-2R} c\ (1+\beta)^{t-1}\\ &\geq c\ (1+\beta)^{t+1},\end{aligned}$$

where the second inequality follows from the definition of β. The statement then follows by taking the lim inf as $t\to\infty$.

Now consider the following mutual information, which we can bound by R.

$$I(X_t;Y_t\mid Y_0^{t-1}) \leq H(Y_t\mid Y_0^{t-1}) \leq R.$$

To show (12), based on (9), we focus on the term

$$\inf_{\substack{U_t=f(Y_0^t):\\ I(X_t;Y_t|Y_0^{t-1})\leq R}} \mathbb{E}[(X_t+U_t)^2], \tag{13}$$

where the infimum is over all U_t that are functions of Y_0^t. We lower bound this by considering the distortion when the controller is provided with extra information about the precise value of X_{t-1}. By providing X_{t-1} as side-information, we effectively allow the encoder to focus on communicating uncertainty about A_{t-1}.

$$\begin{aligned}
&\inf_{\substack{U_t=f(Y_0^t):\\ I(X_t;Y_t|Y_0^{t-1})\leq R}} \mathbb{E}[(X_t+U_t)^2] \\
&\quad= \inf_{\substack{U_t=f(Y_0^t):\\ I(X_t;Y_t|Y_0^{t-1})\leq R}} \mathbb{E}[\mathbb{E}[(X_t+U_t)^2 \mid X_{t-1}]] \\
&\quad\geq \mathbb{E}[\inf_{\substack{U_t=f(Y_0^t):\\ I(X_t;Y_t|Y_0^{t-1})\leq R}} \mathbb{E}[(X_t+U_t)^2 \mid X_{t-1}]].
\end{aligned}$$

Moving the infimum inside only decreases the value of the expectation. The inner expectation is over A_{t-1} and W_{t-1} and the outer expectation is over X_{t-1}. Expanding the terms inside, the quantity within the outer expectation equals

$$\inf_{\substack{U_t=f(Y_0^t):\\ I(X_t;Y_t|Y_0^{t-1})\leq R}} \mathbb{E}[(A_{t-1}X_{t-1}+U_{t-1}+W_{t-1}+U_t)^2 \mid X_{t-1}]$$

Since $U_t = f(Y_0^t)$ and U_{t-1} is a function of Y_0^{t-1}, we can remove U_{t-1} from the infimization to get that (13) is lower-bounded by

$$\mathbb{E}[\inf_{\substack{U_t=f(Y_0^t):\\ I(X_t;Y_t|Y_0^{t-1})\leq R}} \mathbb{E}[(A_{t-1}X_{t-1}+W_{t-1}+U_t)^2 \mid X_{t-1}]].$$

Now, similar to [23], we apply Shannon's lower bound from (10) to get:

$$\begin{aligned}
&\mathbb{E}[\inf_{\substack{U_t=f(Y_0^t):\\ I(X_t;Y_t|Y_0^{t-1})\leq R}} \mathbb{E}[(A_{t-1}X_{t-1}+W_{t-1}+U_t)^2 \mid X_{t-1}]] \\
&\quad\geq \mathbb{E}[N(A_{t-1}X_{t-1}+W_{t-1}|X_{t-1}) \cdot 2^{-2R}].
\end{aligned}$$

We can now apply the entropy-power inequality (11) (since W_{t-1} is independent of everything else) to get:

$$\begin{aligned}
&\mathbb{E}[N(A_{t-1}X_{t-1}+W_{t-1}|X_{t-1}) \cdot 2^{-2R}] \\
&\quad\geq \mathbb{E}[N(A_{t-1}X_{t-1}|X_{t-1})]2^{-2R} + N(W_{t-1})2^{-2R}
\end{aligned}$$

Using $N(AX \mid X) = X^2 N(A)$ gives:

$$\begin{aligned}
&\inf_{\substack{U_t=f(Y_0^t):\\ I(X_t;Y_t|Y_0^{t-1})\leq R}} \mathbb{E}[(X_t+U_t)^2] \\
&\quad\geq \left\{N(A)\mathbb{E}[X_{t-1}^2] + N(W)\right\} \cdot 2^{-2R}.
\end{aligned} \tag{14}$$

Plugging (14) into (9) gives us (12). □

Theorem 5.2. *Assume that P_W has a density p_W and $\sigma_A^2 < 1$. If (1) is mean-square stabilizable using a causal encoder-controller pair with rate-limit R, then*

$$R \geq \frac{1}{2}\log\left(\mu_A^2+\sigma_A^2\right).$$

The proof uses the distortion-rate function through Lemma 5.3, which bounds the ability of the controller to reduce the second moment. This lemma generalizes a corresponding result of [23].

Lemma 5.3. *For every $t \geq 1$ and $R, S \geq 0$ we have*

$$\begin{aligned}
&\inf_{\substack{U_{t-1}=f(Y_0^{t-1}):\\ I(X_{t-1};Y_{t-1}|Y_0^{t-2})\leq R}} \mathbb{D}_S(X_t \mid Y_0^{t-1}) \\
&\quad\geq \left(\mu_A^2+\sigma_A^2\right)\mathbb{D}_{R+S}(X_{t-1} \mid Y_0^{t-2}).
\end{aligned} \tag{15}$$

The proof of this lemma is deferred to the Appendix A.

Proof of Thm. 5.2. We follow a dynamic programming strategy and first consider optimization over the last time step. Consider

$$\begin{aligned}
&\inf_{\substack{U_t=f(Y_0^t):\\ I(X_t;Y_t|Y_0^{t-1})\leq R}} \mathbb{E}[X_{t+1}^2] = \sigma_A^2\mathbb{E}[X_t^2] \\
&\quad+ \mu_A^2 \inf_{\substack{U_t=f(Y_0^t):\\ I(X_t;Y_t|Y_0^{t-1})\leq R}} \left\{\mathbb{E}[(X_t+U_t)^2]\right\} + \sigma_W^2,
\end{aligned}$$

where we use (9) to expand the terms. Since U_t is a function of Y_0^t, the middle term above is equal to the distortion function of X_t at rate R given Y_0^{t-1}, i.e.,

$$\inf_{\substack{U_t=f(Y_0^t):\\ I(X_t;Y_t|Y_0^{t-1})\leq R}} \mathbb{E}[(X_t+U_t)^2] = \mathbb{D}_R(X_t \mid Y_0^{t-1}).$$

Thus we have that

$$\inf_{\substack{U_t=f(Y_0^t):\\ I(X_t;Y_t|Y_0^{t-1})\leq R}} \mathbb{E}[X_{t+1}^2] \geq \mu_A^2\mathbb{D}_R(X_t \mid Y_0^{t-1}). \tag{16}$$

Now consider the next step in the dynamic program and consider the infimum over U_{t-1}. Lemma 5.3 (with $S = R$) implies that

$$\begin{aligned}
&\inf_{\substack{U_{t-1}=f(Y_0^{t-1}):\\ I(X_{t-1};Y_{t-1}|Y_0^{t-2})\leq R}} \mathbb{D}_R(X_t \mid Y_0^{t-1}) \\
&\quad\geq (\mu_A^2+\sigma_A^2)\mathbb{D}_{2R}(X_{t-1} \mid Y_0^{t-2}).
\end{aligned} \tag{17}$$

Combining (17) and (16) we get that

$$\begin{aligned}
&\inf_{\substack{U_{t-1}=f(Y_0^{t-1}):\\ I(X_{t-1};Y_{t-1}|Y_0^{t-2})\leq R}} \inf_{\substack{U_t=f(Y_0^t):\\ I(X_t;Y_t|Y_0^{t-1})\leq R}} \mathbb{E}[X_{t+1}^2] \\
&\quad\geq \mu_A^2(\mu_A^2+\sigma_A^2)\mathbb{D}_{2R}(X_{t-1} \mid Y_0^{t-2}).
\end{aligned} \tag{18}$$

Subsequent applications of Lemma 5.3 give that

$$\inf_{U_0^t} \mathbb{E}[X_{t+1}^2] \geq \mu_A^2(\mu_A^2+\sigma_A^2)^t\mathbb{D}_{(t+1)R}(X_0), \tag{19}$$

where the infimum is over all causal rate-limited control strategies U_0^t. By Shannon's Lower Bound from (10),

$$\mathbb{D}_{(t+1)R}(X_0) \geq N(X_0) \cdot 2^{-2(t+1)R},$$

and since $(\mu_A^2 + \sigma_A^2) \cdot 2^{-2R} > 1$ by assumption, we get from (19) that $\liminf_{t\to\infty} \mathbb{E}[X_t^2] = \infty$. □

Thm. 5.1 and Thm. 5.2 together imply Thm. 3.1.

B. Mean-square stability with cost D converse

Proof of Thm. 3.3. We build on the proofs of Thm. 5.1 and Thm. 5.2, which provide the first two terms in the max of (6). However, since now we aim to have a specific finite bound on the second moment of the state, we need to incorporate the incremental impact of the additive noise at each time step. This leads to the final term in the bound, which depends on both the desired cost D and the variance of the additive noise σ_W^2.

Recall the recursion (12) for the second moment of the state. We can solve this recursion (see Appendix B) to obtain that

$$\liminf_{t\to\infty} \mathbb{E}[X_t^2] \quad \geq \quad \frac{N(W)\mu_A^2 2^{-2R} + \sigma_W^2}{1 - \sigma_A^2 - N(A)\mu_A^2 2^{-2R}}. \quad (20)$$

From the assumption on R we know that the right hand side of (20) is at least D, which completes the proof. □

6. Achievability proof

In this section we analyze the uniform quantization scheme to obtain Theorems 3.2 and 3.4. Thm. 3.2 follows directly from the bound in Thm. 3.4, since as $D \to \infty$, M_D converges to

$$e \times \left(1 + \frac{3\sqrt{1-\sigma_A^2}}{\mu_A}\right) \times \left(1 + \log\left(3 + \frac{\mu_A}{\sqrt{1-\sigma_A^2}}\right)\right),$$

and furthermore $\mu_A^2 + \sigma_A^2 \geq 1$, so the expression in the first bracket above is at most 4.

Proof of Thm. 3.4. The encoder-controller strategy uses a uniform quantization scheme with bins of size 2δ, where δ is given by:

$$\delta^2 = \frac{D(1 - \sigma_A^2) - \sigma_W^2}{\mu_A^2}. \quad (21)$$

As illustrated in Fig. 4, the quantization bins are $[(2k-1)\delta, (2k+1)\delta)$, $k \in \mathbb{Z}$, and the quantization points are the midpoints of the bins, i.e., $\{2k\delta, k \in \mathbb{Z}\}$. Let $\widehat{X}_t$ denote the midpoint of the bin that X_t lands in, which can be obtained from X_t as follows:

$$\widehat{X}_t = \mathrm{sgn}(X_t) \left\lfloor \frac{|X_t| + \delta}{2\delta} \right\rfloor \cdot 2\delta, \quad (22)$$

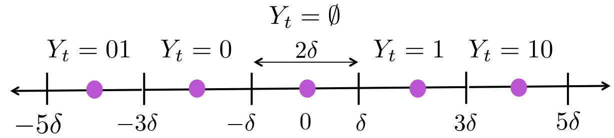

Fig. 4. Uniform quantization scheme with bins of length 2δ.

since if $X_t \in [(2k-1)\delta, (2k+1)\delta)$, then $|k| = \lfloor \frac{|X_t|+\delta}{2\delta} \rfloor$. Since the bins have length 2δ, we have that $|\widehat{X}_t - X_t| \leq \delta$. The encoder transmits Y_t, a variable length binary encoding of $\widehat{X}_t$, to the controller: $Y_t = \emptyset$ if $\widehat{X}_t = 0$, $Y_t = 0$ if $\widehat{X}_t = 2\delta$, $Y_t = 1$ if $\widehat{X}_t = -2\delta$, $Y_t = 00$ if $\widehat{X}_t = 4\delta$, and so on.

Based on the received binary string Y_t, the controller applies the control $U_t = -\mu_A \widehat{X}_t$. Using (9) (note the scaling in U_t) we can write the evolution of the second moment of the state as follows:

$$\mathbb{E}[X_t^2] = \sigma_A^2 \mathbb{E}[X_{t-1}^2] + \mu_A^2 \mathbb{E}[(X_{t-1} - \widehat{X}_{t-1})^2] + \sigma_W^2.$$

Since $|\widehat{X}_t - X_t| \leq \delta$ for every t, we can recursively bound the second moment:

$$\begin{aligned} \mathbb{E}[X_t^2] &\leq \sigma_A^2 \mathbb{E}[X_{t-1}^2] + \mu_A^2\delta^2 + \sigma_W^2 \\ &\leq \sigma_A^2 \left[\sigma_A^2 \mathbb{E}[X_{t-2}^2] + \mu_A^2\delta^2 + \sigma_W^2\right] + \mu_A^2\delta^2 + \sigma_W^2 \\ &\leq \sigma_A^{2t}\sigma_{X_0}^2 + \frac{1-\sigma_A^{2t}}{1-\sigma_A^2}(\mu_A^2\delta^2 + \sigma_W^2). \end{aligned} \quad (23)$$

Since $\sigma_A^2 < 1$ we thus have that

$$\limsup_{t\to\infty} \mathbb{E}[X_t^2] \leq \frac{\mu_A^2\delta^2 + \sigma_W^2}{1 - \sigma_A^2} = D,$$

where the last equality follows from the definition of δ in (21). Thus this encoder-controller scheme stabilizes the system with cost at most D.

What remains is to bound from above the output entropy of the quantizer at each step by the expression on the right hand side of (7). Our starting point is the following bound:

$$\begin{aligned} H(Y_t \mid Y_0^{t-1}) &\leq H(Y_t) \\ &\leq \mathbb{E}[\ell(Y_t)] + \log \mathbb{E}[\ell(Y_t) + 1] + \log e, \end{aligned} \quad (24)$$

where the second inequality is due to [26]. Thus it remains to bound the expected length of the encoding, $\mathbb{E}[\ell(Y_t)]$.

We observe that

$$\ell(Y_t) = \left\lceil \log\left(1 + \left\lfloor \frac{|X_t| + \delta}{2\delta} \right\rfloor\right) \right\rceil.$$

By removing the floor and the ceiling and adding 1 we do not decrease this quantity, so

$$\begin{aligned}\ell\left(Y_t\right) &\leq \log\left(1+\frac{|X_t|+\delta}{2\delta}\right)+1 \\ &= \log\left(3\delta+|X_t|\right)-\log\delta.\end{aligned}$$

By Jensen's inequality we can bound the expectation of the first term in the line above:

$$\begin{aligned}\mathbb{E}\left[\log\left(3\delta+|X_t|\right)\right] &\leq \log\left(3\delta+\mathbb{E}\,|X_t|\right) \\ &\leq \log\left(3\delta+\sqrt{\mathbb{E}\left[X_t^2\right]}\right).\end{aligned}$$

From (23) we have that $\mathbb{E}\left[X_t^2\right] \leq D+\sigma_{X_0}^2$ for every t, and so, together with the previous two displays, we obtain that

$$\mathbb{E}[\ell(Y_t)] \leq \log\left(3\delta+\sqrt{D+\sigma_{X_0}^2}\right)-\log\left(\delta\right). \quad (25)$$

Putting together the bounds (24) and (25) and using the definition of δ from (21) we obtain that

$$H(Y_t \mid Y_0^{t-1}) \leq \frac{1}{2}\log\frac{\mu_A^2}{1-\sigma_A^2}+\log M_D. \qquad \square$$

7. Discussion

Our main contribution is to provide a unified perspective on communication constraints and parameter uncertainty in control systems. While we determine the critical rate up to first order, the additive gap between the achievability and the converse grows unboundedly as σ_A^2 goes to 1. Closing this gap remains a challenge.

A more challenging question is to discuss stability limits for higher and lower moments of the state in the scalar setting. It is of particular interest to understand the behavior of the logarithm of the state, which captures the typical behavior of the system. The second moment analysis here relies on orthogonality properties of the underlying random variables, and understanding the typical behavior might need a different approach.

Acknowledgments

We thank Philip Chou, Tamás Linder, and Kartik Venkat for useful pointers and helpful discussions.

References

[1] M. Athans, R. Ku, and S. Gershwin, "The uncertainty threshold principle: Some fundamental limitations of optimal decision making under dynamic uncertainty," *IEEE Transactions on Automatic Control*, vol. 22, no. 3, pp. 491–495, 1977.

[2] W. S. Wong and R. W. Brockett, "Systems with Finite Communication Bandwidth Constraints—Part I: State Estimation Problems," *IEEE Transactions on Automatic Control*, vol. 42, no. 9, pp. 1294–1299, 1997.

[3] S. Tatikonda and S. Mitter, "Control under communication constraints," *IEEE Transactions on Automatic Control*, vol. 49, no. 7, pp. 1056–1068, 2004.

[4] G. Nair and R. Evans, "Stabilizability of Stochastic Linear Systems with Finite Feedback Data Rates," *SIAM Journal on Control and Optimization.*, vol. 43, no. 2, pp. 413–436, 2004.

[5] G. N. Nair, F. Fagnani, S. Zampieri, and R. J. Evans, "Feedback control under data rate constraints: An overview," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 108–137, 2007.

[6] C. E. Shannon, "Coding Theorems for a Discrete Source With a Fidelity Criterion," *IRE Nat. Conv. Rec.*, vol. 4, no. 1, pp. 142–163, 1959.

[7] H. Gish and J. Pierce, "Asymptotically Efficient Quantizing," *IEEE Trans. Inf. Theory*, vol. 14, no. 5, pp. 676–683, 1968.

[8] P. A. Chou, T. Lookabaugh, and R. M. Gray, "Entropy-Constrained Vector Quantization," *IEEE Trans. on Acoust., Speech, and Signal Process.*, vol. 37, no. 1, pp. 31–42, 1989.

[9] T. Linder and R. Zamir, "On the Asymptotic Tightness of the Shannon Lower Bound," *IEEE Trans. on Information Theory*, vol. 40, no. 6, pp. 2026–2031, 1994.

[10] T. Koch, "The Shannon Lower Bound is Asymptotically Tight," *IEEE Transactions on Information Theory*, vol. to appear, 2016.

[11] V. Kostina, "Data compression with low distortion and finite blocklength," *arXiv preprint arXiv:1510.02190*, 2015.

[12] G. Ranade, "Active systems with uncertain parameters: an information-theoretic perspective," Ph.D. dissertation, University of California, Berkeley, 2014.

[13] G. Ranade and A. Sahai, "Control Capacity," in *International Symposium on Information Theory (ISIT)*. IEEE, 2015.

[14] ——, "Non-coherence in estimation and control," in *51st Annual Allerton Conf. on Comm., Control, and Comp.*, 2013.

[15] J. Ding, Y. Peres, and G. Ranade, "A tiger by the tail: when multiplicative noise stymies control," in *International Symposium on Information Theory (ISIT)*. IEEE, 2016.

[16] M. Fu and L. Xie, "Quantized feedback control for linear uncertain systems," *International Journal of Robust and Nonlinear Control*, vol. 20, no. 8, pp. 843–857, 2010.

[17] K. Zhou and J. C. Doyle, *Essentials of robust control*, 1998.

[18] N. Martins, M. Dahleh, and N. Elia, "Feedback Stabilization of Uncertain Systems in the Presence of a Direct Link," *IEEE Trans. Autom. Control*, vol. 51, no. 3, pp. 438–447, 2006.

[19] V. N. Phat, J. Jiang, A. V. Savkin, and I. R. Petersen, "Robust stabilization of linear uncertain discrete-time systems via a limited capacity communication channel," *Systems & Control Letters*, vol. 53, no. 5, pp. 347–360, 2004.

[20] K. Okano and H. Ishii, "Data rate limitations for stabilization of uncertain systems," in *51st Conference on Decision and Control (CDC)*. IEEE, 2012, pp. 3286–3291.

[21] ——, "Data Rate Limitations for Stabilization of Uncertain Systems over Lossy Channels," in *American Control Conference (ACC)*. IEEE, 2012, pp. 1260–1265.

[22] ——, "Minimum data rate for stabilization of linear systems with parametric uncertainties," 2014, arXiv preprint.

[23] V. Kostina and B. Hassibi, "Rate-cost tradeoffs in control," 2016, preprint.

[24] S. Tatikonda, A. Sahai, and S. Mitter, "Control of LQG systems under communication constraints," in *Conference on Decision and Control*, vol. 1. IEEE, 1998, pp. 1165–1170.

[25] T. Cover and J. Thomas, *Elements of information theory*, 2nd ed. Wiley, 2006.

[26] N. Alon and A. Orlitsky, "A lower bound on the expected length of one-to-one codes," *IEEE Transactions on Information Theory*, vol. 40, no. 5, pp. 1670–1672, 1994.

Appendix A
Proof of Lemma 5.3

Lemma 5.3. *For every $t \geq 1$ and $R, S \geq 0$ we have*

$$\inf_{\substack{U_{t-1}=f(Y_0^{t-1}):\\ I(X_{t-1};Y_{t-1}|Y_0^{t-2})\leq R}} \mathbb{D}_S(X_t \mid Y_0^{t-1}) \geq \left(\mu_A^2 + \sigma_A^2\right) \mathbb{D}_{R+S}(X_{t-1} \mid Y_0^{t-2}). \quad (15)$$

Proof. By definition we have

$$\inf_{\substack{U_{t-1}=f(Y_0^{t-1}):\\ I(X_{t-1};Y_{t-1}|Y_0^{t-2})\leq R}} \mathbb{D}_S(X_t \mid Y_0^{t-1}) = \inf_{\substack{U_{t-1}=f(Y_0^{t-1}):\\ I(X_{t-1};Y_{t-1}|Y_0^{t-2})\leq R}} \inf_{Z: I(X_t;Z|Y_0^{t-1})\leq S} \mathbb{E}[(X_t+Z)^2] \quad (26)$$

We can substitute an expansion of X_t and take conditional expectation over A_{t-1} and W_{t-1} to write:

$$\begin{aligned}&\mathbb{E}[(X_t+Z)^2]\\ &= \mathbb{E}[\mathbb{E}[(A_{t-1}X_{t-1} + U_{t-1} + W_{t-1} + Z)^2 \mid A_{t-1}, W_{t-1}]]\\ &= \mathbb{E}\big[A_{t-1}^2 \mathbb{E}\big[\big(X_{t-1} + \tfrac{U_{t-1}+W_{t-1}+Z}{A_{t-1}}\big)^2 \mid A_{t-1}, W_{t-1}\big]\big]\end{aligned}$$

Pulling the infimum inside the expectation can only decrease the quantity of interest, since U_{t-1} and Z can then depend on the values of A_{t-1} and W_{t-1}. Hence, (26) is bounded from below by

$$\mathbb{E}\Big[A_{t-1}^2 \times \inf_{\substack{U_{t-1}=f(Y_0^{t-1}):\\ I(X_{t-1};Y_{t-1}|Y_0^{t-2})\leq R}} \inf_{Z: I(X_t;Z|Y_0^{t-1})\leq S} \mathbb{E}\big[\big(X_{t-1} + \tfrac{U_{t-1}+W_{t-1}+Z}{A_{t-1}}\big)^2 \mid A_{t-1}, W_{t-1}\big]\Big]. \quad (27)$$

We can lower bound (27) by taking a joint infimum over U_{t-1} and Z, where now only the sum of the mutual informations above are bounded by $R+S$. Now, the following is true about the mutual informations in the infima:

$$I(X_{t-1}; Y_{t-1} \mid Y_0^{t-2}) + I(X_t; Z \mid Y_0^{t-1}) = I(X_{t-1}; Y_{t-1}, Z \mid Y_0^{t-2}).$$

Thus we can bound (27) from below by

$$\mathbb{E}\Big[A_{t-1}^2 \times \inf_{\substack{U_{t-1}=f(Y_0^{t-1}),Z:\\ I(X_{t-1};Y_{t-1},Z|Y_0^{t-2})\leq R+S}} \mathbb{E}\big[\big(X_{t-1} + \tfrac{U_{t-1}+W_{t-1}+Z}{A_{t-1}}\big)^2 \mid A_{t-1}, W_{t-1}\big]\Big]. \quad (28)$$

Since we can treat A_{t-1} and W_{t-1} as constants inside the expectation, the infimum inside the expectation is equal to $\mathbb{D}_{R+S}(X_{t-1} \mid Y_0^{t-2})$. Finally, taking expectation over A_{t-1} gives the factor of $(\mu_A^2 + \sigma_A^2)$ and the statement of the lemma. □

Appendix B
Recursions

Lemma B.1. *Assume that the nonnegative sequence $\{z_t\}_{t=0}^{\infty}$ satisfies*

$$z_{t+1} \geq K z_t + L z_{t-1} + M \quad (29)$$

for every $t \geq 1$, where K, L, and M are all nonnegative, and $K + L < 1$. Then

$$\liminf_{t\to\infty} z_t \geq \frac{M}{1-K-L}. \quad (30)$$

Proof. Define the nonnegative sequence $\{v_t\}_{t=0}^{\infty}$ recursively by setting $v_0 := z_0$, $v_1 := z_1$, and for every $t \geq 1$ let

$$v_{t+1} = K v_t + L v_{t-1} + M. \quad (31)$$

Since all the terms are nonnegative, we have that $z_t \geq v_t$ for every $t \geq 0$. Solving the recurrence (31) (for instance using generating functions) we obtain that

$$v_t = \frac{M}{1-K-L} + c_1 \left(\tfrac{K-\sqrt{K^2+4L}}{2}\right)^t + c_2 \left(\tfrac{K+\sqrt{K^2+4L}}{2}\right)^t,$$

for appropriate constants c_1 and c_2. Since $K + L < 1$, we obtain that $\lim_{t\to\infty} v_t = \frac{M}{1-K-L}$, and hence (30) follows. □

In the proof of Thm 3.3 in Sec. 5-B we apply Lemma B.1 with $K = \sigma_A^2$, $L = N(A)\mu_A^2 2^{-2R}$, and $M = N(W)\mu_A^2 2^{-2R} + \sigma_W^2$.