Median, Mean, and Variance Stability of a Process Under Temporally Correlated Stochastic Feedback

Roy S. Smith, *Fellow, IEEE*, and Bassam Bamieh, *Fellow, IEEE*

***Abstract*—Stochasticity in feedback gains leads to heavy-tailed state distributions in which the median, mean, and variance have different stability properties. These properties are characterised for a scalar system with temporally correlated stochastic feedback. Necessary and sufficient conditions are obtained for the generic case of log-normal feedback gain distributions. Temporal correlation is modeled by a multivariable Gaussian distribution of the logarithm of the feedback gain values. This correlation has no effect on the stability of the median of the state, but does influence the stability of both the mean and the variance of the state. Examples illustrate both stabilising and destabilising correlations.**

***Index Terms*—Stochastic systems, stability of linear systems, time-varying systems.**

I. Introduction

DYNAMICAL systems with stochastic coefficients arise in many physical models which are either imprecisely known or operate in inherently stochastic environments. Stochasticity in parameters is fundamentally different from stochastic exogenous inputs that enter additively as forcing functions for example. One way to appreciate the difference is to note that stochastic system parameters are actually stochastic feedback loops inside the system description. As is well known, feedback affects dynamical systems in fundamentally different ways than additive exogenous disturbances. Stochasticity in feedback gains also arises in data-based control methods due to the inevitable noise in the data. Similar dynamics arise in financial systems and statistical mechanics.

Most classical stochastic control theory deals with the case of additive, exogenous inputs. If the dynamics are linear and one is interested only in second order processes, then there is no "closure problem", and covariance analysis is sufficient to quantify the system behaviour. However, if one is interested in quantities other than first and second order moments, or if the exogenous inputs are not Gaussian, then these classical methods are not applicable. In the setting where feedback gains are stochastic, there is yet another phenomenon where even if a linear system's parameters are Gaussian processes, the state evolution involves the products of a stochastic state and a stochastic gain, leading to heavy-tailed state distributions. Thus linear systems with stochastic feedbacks can exhibit a very rich phenomenology of behaviors.

This problem was studied in a control context in [1] and [2]. The proof that the limiting distribution is a heavy-tailed distribution appears in [3]. Some aspects of stability in discrete stochastic feedback systems were investigated in [4]. A more complete analysis from a feedback control point of view was given in our prior work [5].

All of the previous work considered the stochasticity to be independent and identically distributed (i.i.d.) between samples. One reason for imposing this assumption is mathematical tractability. The techniques used in the aforementioned references rely heavily on the i.i.d. assumption for the analysis to be tractable. There is in general no clear way to extend those analysis techniques to cases where stochastic gains are temporally correlated. This issue is not purely academic. There are many applications where one would need to introduce some notion of temporal correlation (correlation times) into the theory. The behaviour of various statistics of interest appear to depend on these temporal correlations. To put this issue in some context, consider the i.i.d. case on the one hand, and the case of time-constant, but randomly selected feedback gains on the other. They can be regarded as two extreme cases of a continuum of temporally correlated gains. The limit of short correlation times is conceptually the i.i.d. case, while the limit of long correlation times is the random constant case. These models are quite important in statistical physics where the long correlation times limit is referred to as the "quenched noise" limit [6], and the short correlation times limit is the so-called "annealed limit". Such models are encountered in the physics of disordered systems, and in particular appear in the famous problem of Anderson localization [7]. This remains an active area of research given the

Manuscript received March 17, 2020; revised June 2, 2020; accepted June 18, 2020. Date of publication July 2, 2020; date of current version July 17, 2020. This work was supported by NSF under Award CMMI-1763064 and Award ECCS-1932777. Recommended by Senior Editor V. Ugrinovskii. *(Corresponding author: Roy S. Smith.)*

Roy S. Smith is with the Automatic Control Laboratory, ETH Zürich, 8092 Zurich, Switzerland (e-mail: rsmith@control.ee.ethz.ch).

Bassam Bamieh is with the Mechanical Engineering Department, University of California at Santa Barbara, Santa Barbara, CA 93106 USA (e-mail: bamieh@engineering.ucsb.edu).

Digital Object Identifier 10.1109/LCSYS.2020.3006726

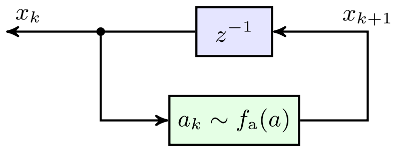

Fig. 1. Stochastic feedback system. At each time instant, k, a new sample $a_k \sim f_{\mathrm{a}}(a)$ is drawn from the distribution.

complexity of the phenomenology that arises in seemingly simple models.

Our prior work [5] considered the i.i.d. stochastic feedback case, which was analysed using techniques that do not easily extend to the temporally correlated gains case. In this letter, we present one model of temporal correlations that appears to be tractable. We are able to quantify the behaviour of variances, means and medians of stochastic gains with lognormal distributions and a special type of temporal correlation.

A. Notation

For a random variable a, $a \sim f_{\mathrm{a}}(a)$ denotes that it is drawn from a probability distribution with a density function $f_{\mathrm{a}}(a)$. The expected value of a is $\mathcal{E}\big[a\big] = \mu_a$ and its variance is σ_a^2. The normal distribution of mean μ and variance σ^2 is denoted by $\mathcal{N}(\mu, \sigma^2)$. Log-normal distributions are denoted by $\mathcal{LN}$. If a symmetric matrix X is positive definite this is denoted by $X \succ 0$. The i, j^{th} component of X is denoted by $[X]_{i,j}$. The N-length column vector of ones is denoted by $\mathbf{1}_N$.

II. Model and Problem Formulation

The plant is a scalar system evolving with the dynamics,

$$x_{k+1} = a_k x_k, \quad k = 0, 1, \ldots. \tag{1}$$

The state variable a_k is, at each time step, drawn from a distribution $a_k \sim f_{\mathrm{a}}(a)$. This can be viewed as the stochastic feedback structure given in Figure 1. It is a simple matter to consider more general first order systems with stochastic feedback. It can also be equivalently formulated as multiplicative noise in a first order feedback system.

This is the same problem formulation as that considered in a prior paper [5] in the i.i.d. setting, where a very general class of distributions $f_{\mathrm{a}}(a)$ on positive support was considered. In the current work the theoretical results are derived under the assumption that the feedback is drawn from a log-normal distribution, $f_{\mathrm{a}}(a)$. The log-normal distribution is defined by exponentiating a normally distributed random variable. If $\alpha \sim \mathcal{N}(\mu_\alpha, \sigma_\alpha^2)$, then the random variable defined by $a = \mathrm{e}^{\alpha}$ is log-normally distributed, $a \sim \mathcal{LN}$. The key feature here is that a log-normal distribution is closed under both multiplication and raising to a power.

The state of the system (1) after N time-steps is

$$x_N = x_0 \prod_{k=0}^{N-1} a_k, \tag{2}$$

and without loss of generality it will be assumed that $x_0 = 1$. As x_N is the product of random variables, it is also a random variable. Furthermore, as $a_k \sim \mathcal{LN}$, the state x_N is also a log-normal random variable, $x_N \sim \mathcal{LN}$ with mean and variance denoted by μ_{x_N} and $\sigma_{x_N}^2$ respectively.

The question of stability can be posed in terms of a variety of statistical quantities and this letter is concerned with three definitions of stability:

a) Median Stability: The system is *median stable* if,

$$\lim_{N \longrightarrow \infty} \operatorname{median}(x_N) = 0.$$

b) Mean Stability: The system is *mean stable* if,

$$\lim_{N \longrightarrow \infty} \mathcal{E}\big[x_N\big] = 0.$$

c) Variance Stability: The system is *variance stable* if,

$$\lim_{N \longrightarrow \infty} \mathcal{E}\Big[\big(x_N - \mathcal{E}\big[x_N\big]\big)^2\Big] = 0.$$

Necessary and sufficient conditions for all three stability conditions were derived in [5], under the assumption that the a_k were i.i.d. random variables. The current work restricts the consideration of a_k to log-normal variables, but removes the independence assumption. In contrast to the i.i.d. case, both the mean and variance stability depend strongly on the temporal correlation. However the stability of the median is unchanged by correlation of the a_k random variable.

A. Analysis Framework

The analysis of this system will proceed by taking the obvious step of working with the logarithm of a_k and x_k. By defining,

$$\alpha_k = \ln(a_k), \quad \text{and} \quad \zeta_k = \ln(x_k),$$

the dynamics in (2) can be written instead as,

$$\zeta_N = \zeta_0 + \sum_{k=0}^{N-1} \alpha_k. \tag{3}$$

The definition of α_k as the logarithm of the random variable $a_k \sim \mathcal{LN}$ implies that $\alpha_k \sim \mathcal{N}$. As all of the α_k are normally distributed (and correlated in a manner detailed in Section II-C), and $\zeta_0 = \ln(x_0) = 0$ by assumption, ζ_N is also normally distributed.

The approach to be taken here is to characterise the correlation between the stochastic feedback variables a_k, in terms of a correlation between the logarithmically transformed random variables α_k. As all of the logarithmically transformed variables are characterised in terms of sums of normal variables, the stability conditions can be considered instead in terms of the ζ_k variables. Doing this requires more detail on the relationship between the normal and log-normal distributions.

To facilitate readability we will use Greek symbols for variables in the logarithmic domain (for example ζ_k and α_k) and the Latin alphabet for variables in the state and feedback domain (for example x_k and a_k).

B. Normal and Log-Normal Distributions

The monotonic variable relationship

$$\zeta_k = \ln(x_k),$$

will transform the median (and any percentile statistic) directly as,

$$\begin{aligned}\mathrm{Prob}\{x_k < z\} &= \mathrm{Prob}\{\ln(x_k) < \ln(z)\}\\ &= \mathrm{Prob}\{\zeta_k < \ln(z)\}.\end{aligned}$$

Therefore,

$$\mathrm{median}(\zeta_N) = \ln(\mathrm{median}(x_N))$$

and

$$\mathrm{median}(x_N) = \mathrm{e}^{\mathrm{median}(\zeta_N)}. \tag{4}$$

However the relationship between the means and variances is more complicated [8].

$$\mu_\alpha = \ln\left(\frac{\mu_a}{\sqrt{1+\frac{\sigma_a^2}{\mu_a^2}}}\right), \tag{5}$$

$$\sigma_\alpha^2 = \ln\left(1+\frac{\sigma_a^2}{\mu_a^2}\right), \tag{6}$$

and

$$\mu_a = \mathrm{e}^{\mu_\alpha+\sigma_\alpha^2/2}, \tag{7}$$

$$\sigma_a^2 = \left(\mathrm{e}^{\sigma_\alpha^2}-1\right)\left(\mathrm{e}^{2\mu_\alpha+\sigma_\alpha^2}\right). \tag{8}$$

These relationships will be used to transform results from the ζ_N domain to the x_N domain and vice-versa.

C. Temporal Correlation Model

There are a variety of ways in which we can characterise the temporal correlation between the random variables, $a_k \sim \mathcal{LN}$. The approach taken here is to characterise the correlation in terms of the normally distributed $\alpha_k \sim \mathcal{N}$ variables. We model the correlation between the α_k random variables as the output of a linear time-invariant filter (denoted by Γ) driven by an i.i.d. Gaussian noise sequence. This implies that the vector

$$\alpha = \left[\alpha_0 \cdots \alpha_{N-1}\right]^T \tag{9}$$

is modeled as coming from a multivariable Gaussian distribution. This modeling choice gives a relatively straightforward characterisation of the median, mean, and variance stability conditions.

The log-domain filter Γ is specified by its pulse response, $\gamma_\tau,\ \tau = 0, 1, \ldots, \infty$. The model for each α_k is therefore,

$$\alpha_k = \sum_{\tau=0}^{\infty} \gamma_\tau \nu_{k-\tau}, \quad \nu_{k-\tau} \sim \mathcal{N}(\mu_\nu, \sigma_\nu^2). \tag{10}$$

We assume (without loss of generality) that $\gamma_0 = 1$ and that Γ is stable. Note that (10) is the asymptotic limit, implying that the filter generating the α_k is already in equilibrium at time-step $k = 0$.

The appropriate choice of μ_ν and σ_ν^2 can be determined by calculating μ_{α_k} and $\sigma_{\alpha_k}^2$.

$$\mu_{\alpha_k} = \mathcal{E}\left[\sum_{\tau=0}^{\infty}\gamma_\tau \nu_{k-\tau}\right] = \sum_{\tau=0}^{\infty}\gamma_\tau \mathcal{E}\left[\nu_{k-\tau}\right] = \Gamma(\mathrm{e}^{j0})\mu_\nu,$$

where $\Gamma(\mathrm{e}^{j0})$ denotes the zero frequency gain of the filter Γ. By choosing

$$\mu_\nu = \mu_\alpha / \Gamma(\mathrm{e}^{j0}),$$

the correlation model gives the appropriate value of μ_α. To select σ_ν^2 we evaluate the variance of α_k. As the variance scales with the square of the $\mathcal{H}_2$ norm,

$$\sigma_{\alpha_k}^2 = \sum_{i=0}^{\infty}|\gamma_i|^2\sigma_\nu^2 = \|\Gamma\|_2^2\sigma_\nu^2. \tag{11}$$

Therefore selecting

$$\sigma_\nu^2 = \sigma_\alpha^2/\|\Gamma\|_2^2, \tag{12}$$

will give the value of σ_α^2 required for the correlation model.

A similar calculation, using shifted indices, will give the covariances,

$$\mathcal{E}\left[\alpha_i\alpha_j\right] = \left(\sum_{m=0}^{\infty}\gamma_m\gamma_{m+j-i}\right)\sigma_\nu^2. \tag{13}$$

The above results can also be expressed using a multivariable Gaussian formulation for the length-N vector α defined in (9), giving $\alpha \sim \mathcal{N}(\mu_{\alpha_N}, \Sigma_N)$, where the diagonal components of the covariance matrix, Σ_N, are given by (12), and the off-diagonal components are given by (13).

D. Temporal Correlation in the Log-Normal Domain

The filtered i.i.d. noise model for correlation (in (10)) that generates the vector α with correlated components has an equivalent representation in the original problem domain for generating the temporally correlated a vector. This can be determined by transforming (10) to the a domain and gives,

$$a_k = \prod_{\tau=0}^{\infty}(v_{k-\tau})^{\gamma_\tau},$$

where the i.i.d. log-normal random variable $v_{k-\tau}$ is defined by

$$v_{k-\tau} = \mathrm{e}^{\nu_{k-\tau}}, \quad \nu \sim \mathcal{N}(\mu_\nu, \sigma_\nu^2).$$

The mean and variance of $v \sim \mathcal{LN}$ are given by applying the mapping in (7) and (8) to the mean and variance of $\nu \sim \mathcal{N}(\mu_\nu, \sigma_\nu^2)$.

This form of temporal correlation model is certainly not standard, but it does give a log-normal distribution a_k from a product of time shifted i.i.d. log-normal random variables raised to appropriate powers.

TABLE I
STOCHASTIC FEEDBACK GAIN STABILITY CONDITIONS FOR i.i.d. LOG-NORMAL DISTRIBUTIONS $a \sim \mathcal{LN}$

Stability property	in terms of $f_a(a)$ distribution	in terms of $f_\alpha(\alpha)$ distribution
median(x_N)	$\mu_a^2 - \sigma_a^2/\mu_a^2 < 1$	$\mu_\alpha < 0$
mean(x_N)	$\mu_a < 1$	$\mu_\alpha + \sigma_\alpha^2/2 < 0$
variance(x_N)	$\mu_a^2 + \sigma_a^2 < 1$	$\mu_\alpha + \sigma_\alpha^2 < 0$

III. Stability Results

The median, mean and variance stability conditions for this system with i.i.d. stochastic feedback were derived in [5] and are summarised here for comparison. Table I illustrates the three stability conditions. Each of these stability criteria will now be considered for the temporal correlation model introduced in Section II-C.

The basis of the results is the calculation of the mean and variance of

$$\alpha_N = \sum_{k=0}^{N-1} \alpha_k.$$

The evolution of the mean of α_N is unaffected by the correlation between the α_k random variables,

$$\begin{aligned}\mathcal{E}\big[\alpha_N\big] &= \mathcal{E}\left[\sum_{k=0}^{N-1}\sum_{i=0}^{\infty}\gamma_i\nu_{k-i}\right] = \sum_{k=0}^{N-1}\sum_{i=0}^{\infty}\gamma_i\mathcal{E}\big[\nu_{k-i}\big]\\ &= N\Gamma(\mathrm{e}^{j0})\mu_\nu = N\mu_\alpha. \end{aligned} \tag{14}$$

To derive the variance of α_N note that

$$\alpha_N = \begin{bmatrix}1 & \cdots & 1\end{bmatrix}\alpha,$$

where $\alpha \in \mathcal{R}^N$ is the multivariable Gaussian vector defined in (9). The variance of α_N is therefore given by,

$$\text{variance}(\alpha_N) = \text{variance}(\mathbf{1}_N^T\alpha) = \mathbf{1}_N^T\Sigma_N\mathbf{1}_N. \tag{15}$$

The median stability condition depends only on $\mathcal{E}\big[\alpha_N\big]$ as stated formally in the following.

Theorem 1: The stochastic feedback system specified in (1) is *median stable* under correlated log-normal feedback if and only if, $\mu_\alpha < 0$.

Proof: As a is multivariable log-normal, α is multivariable normal and α_N is normal, and so α_N has median equal to its mean. Therefore, from (4) and (14),

$$\text{median}(x_N) = \mathrm{e}^{\text{median}(\alpha_N)} = \mathrm{e}^{\text{mean}(\alpha_N)} = \mathrm{e}^{N\mu_\alpha},$$

which goes to zero as $N \longrightarrow \infty$ if and only if $\mu_\alpha < 0$. ■

The mean and variance of x_N depend on both the mean and variance of α_N. Although the mean of α_N is unchanged by correlation, the variance is not. We therefore expect the conditions for mean and variance stability to depend on the correlation.

To quantify the dependency of the mean of x_N on correlation consider the following.

$$\begin{aligned}\Gamma(\mathrm{e}^{j0})^2 &= \left(\sum_{m=0}^{\infty}\gamma_m\right)^2 = \left(\sum_{m=0}^{\infty}\gamma_m\right)\left(\sum_{j=0}^{\infty}\gamma_j\right)\\ &= \sum_{m=0}^{\infty}\gamma_m\sum_{t=-m}^{\infty}\gamma_{m+t} = \sum_{t=-\infty}^{\infty}\sum_{m=0}^{\infty}\gamma_m\gamma_{m+t}. \end{aligned} \tag{16}$$

Theorem 2: The stochastic feedback system specified in (1) is *mean stable* under correlated log-normal feedback if and only if,

$$\mu_\alpha + \frac{\Gamma(\mathrm{e}^{j0})^2}{2\|\Gamma\|_2^2}\sigma_\alpha^2 < 0.$$

Proof: Consider the ratio,

$$\frac{\text{mean}(x_N)}{\text{mean}(x_{N-1})} = \mathrm{e}^{\ln\left(\frac{\mu_{x_N}}{\mu_{x_{N-1}}}\right)}.$$

The following will show that, in the limit, the exponent is constant. In which case the conditions under which the exponent is negative are equivalent to the mean stability of x_N. To this end consider,

$$\ln\left(\frac{\mu_{x_N}}{\mu_{x_{N-1}}}\right) = \ln\big(\mu_{x_N}\big) - \ln\big(\mu_{x_{N-1}}\big)$$

Observe that

$$\mu_{x_N} = \mathrm{e}^{N\mu_\alpha + \frac{1}{2}\mathbf{1}_N^T\Sigma_N\mathbf{1}_N}. \tag{17}$$

Using the analogous result for $\mu_{x_{N-1}}$ implies that,

$$\begin{aligned}&\ln\left(\frac{\mu_{x_N}}{\mu_{x_{N-1}}}\right)\\ &\quad = \mu_\alpha + \frac{1}{2}\left(\mathbf{1}_N^T\Sigma_N\mathbf{1}_N - \begin{bmatrix}\mathbf{1}_{N-1}^T 0\end{bmatrix}\Sigma_N\begin{bmatrix}\mathbf{1}_{N-1}\\ 0\end{bmatrix}\right).\end{aligned}$$

The term in parentheses is simply the sum of the elements N^{th} row of Σ_N and of the elements of the N^{th} column of Σ_N, less the $\Sigma_{N,N}$ element. This is therefore equal to

$$\mu_\alpha + \frac{1}{2}\left(\sum_{l=1}^{N}[\Sigma_N]_{N,l} + \sum_{n=1}^{N-1}[\Sigma_N]_{n,N}\right).$$

From (13) this can be expressed in terms of the pulse response coefficients of Γ as

$$\mu_\alpha + \frac{\sigma_\nu^2}{2}\left(\sum_{l=1}^{N}\sum_{m=0}^{\infty}\gamma_m\gamma_{m+l-N} + \sum_{n=1}^{N-1}\sum_{m=0}^{\infty}\gamma_m\gamma_{m+N-n}\right).$$

Now define an index $t = l-N$ for the first sum and $t = N-n$ for the second. Note that t ranges from $-N+1$ to 0 in the first summation pair and from 1 to $N-1$ in the second. These contiguous ranges can be combined to give,

$$\mu_\alpha + \frac{\sigma_\nu^2}{2}\left(\sum_{t=-N+1}^{N-1}\sum_{m=0}^{\infty}\gamma_m\gamma_{m+t}\right).$$

In the limit as $N \longrightarrow \infty$, equation (16) shows that the term in parentheses converges to $\Gamma(\mathrm{e}^{j0})^2$. Substituting (12) gives the required condition. ■

Observe that in the uncorrelated case $\Sigma_N = \sigma_\nu^2 I = \sigma_\alpha^2 I$, and $\Gamma(\mathrm{e}^{j0})^2/\|\Gamma\|_2^2 = 1$. In this case Theorem 2 is equivalent to the condition for mean stability given in Table 1. The mean of x_N for any finite N can be calculated directly from (17). This calculation is used in the illustrative examples in Section IV.

The condition for variance stability can also be stated in terms of the mean and variance of $f_\alpha(\alpha)$ and the gains associated with the correlation filter Γ.

Theorem 3: The stochastic feedback system specified in (1) is *variance stable* under correlated log-normal feedback if and only if,

$$\mu_\alpha + \frac{\Gamma(\mathrm{e}^{j0})^2}{\|\Gamma\|_2^2}\sigma_\alpha^2 < 0.$$

Proof: From (8), the variance of x_N is

$$\sigma_{x_N}^2 = \left(\mathrm{e}^{\sigma_{\alpha_N}^2} - 1\right)\left(\mathrm{e}^{2\mu_{\alpha_N} + \sigma_{\alpha_N}^2}\right).$$

From (14), $\mu_{\alpha_N} = N\mu_\alpha$. The argument in the proof of Theorem 2 shows that the asymptotic variance of α_N with respect to N is,

$$\lim_{N\longrightarrow\infty}\sigma_{\alpha_N}^2 = N\frac{\Gamma(\mathrm{e}^{j0})^2}{\|\Gamma\|_2^2}\sigma_\alpha^2.$$

TABLE II
SIMULATION EXAMPLE CASE STUDY CONFIGURATIONS

Label	Filter	$\Gamma(e^{j0})^2$	μ_α	σ_α^2	μ_a
Case 1a	i.i.d.	—	-0.1	0.3	1.05
1b	Γ_1	0.09	-0.1	0.3	1.05
Case 2a	i.i.d.	—	-0.1	0.1	0.95
2b	Γ_2	7.29	-0.1	0.1	0.95

Therefore

$$\lim_{N\longrightarrow\infty} \sigma_{x_N}^2 = \left(e^{N\sigma_\alpha^2\Gamma(e^{j0})^2/\|\Gamma\|_2^2}-1\right)\left(e^{2N\mu_\alpha+N\sigma_\alpha^2\Gamma(e^{j0})^2/\|\Gamma\|_2^2}\right)$$
$$= e^{2N(\mu_\alpha+\sigma_\alpha^2\Gamma(e^{j0})^2/\|\Gamma\|_2^2)} - e^{2N(\mu_\alpha+\sigma_\alpha^2\Gamma(e^{j0})^2/2\|\Gamma\|_2^2)}.$$

The exponent of the first term is larger than the that of the second and so determines whether or not the limit goes to zero. Convergence to zero as $N \longrightarrow \infty$ occurs if and only if the exponent of the first term is negative. ■

IV. ILLUSTRATIVE EXAMPLES

Two examples of correlated stochastic feedback are used to illustrate the effects of temporal correlation. As introduced in Section III, the correlation is modeled in the logarithmic variable domain as a causal LTI filter driven by a normally distributed i.i.d. random variable. We present two case examples and for each compare the effect of correlated α_k variables with independent α_k variables. The details of the two cases are summarized in Table II.

For simplicity we compare two FIR filters, specified here by their pulse responses,

$$\Gamma_1 : \gamma_1(k) = \begin{bmatrix}1.0 & -0.9 & 0.2 & -0.6 & 0 & \cdots\end{bmatrix}$$

and

$$\Gamma_2 : \gamma_2(k) = \begin{bmatrix}1.0 & 0.9 & 0.2 & 0.6 & 0 & \cdots\end{bmatrix}.$$

Each has the same 2-norm scaling, $\|\Gamma\|_2^2 = 2.21$, but differ in their zero frequency gains (see Table II).

Figure 2 illustrates multiple sample trajectories generated by (1) with $a_k = e^{\alpha_k}$ generated with the mean and variance given for Case 1a in Table 2. As

$$\mu_a = e^{\mu_\alpha+\sigma_\alpha^2/2} = 1.05 > 1,$$

the mean of x_N (blue dashed line) grows exponentially. However $\mu_\alpha = -0.1 < 0$ and so the median of x_N (magenta dashed line) decays to zero. This is a case where the system has a stable median and an unstable mean. The simulations are run for up to $N = 50$ and it is clear that the sample estimate of the mean (red solid line) deteriorates as N increases. This is because the stable median implies that the probability of a sample trajectory exceeding the mean decays exponentially to zero as N increases (see [5] for details).

Figure 3 illustrates Case 1b where the mean and variance of α_k are the same as Case 1a, but the α_k are now correlated. From Theorem 2 the low value of $\Gamma_1(e^{j0})$ leads to mean stability. In this case both the median and mean decay to zero with different exponents.

The differing stability cases can also be examined by looking at the evolution of the probability density functions of

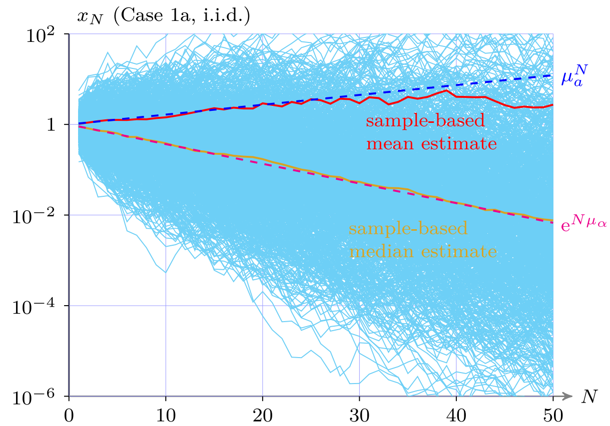

Fig. 2. Simulations (1,000 realisations) of Case 1a trajectories. The random variable $a_k \sim \mathcal{LN}$ is i.i.d. The theoretical and sample medians, as well as the theoretical and sample means, are shown. The median is stable and the mean is unstable.

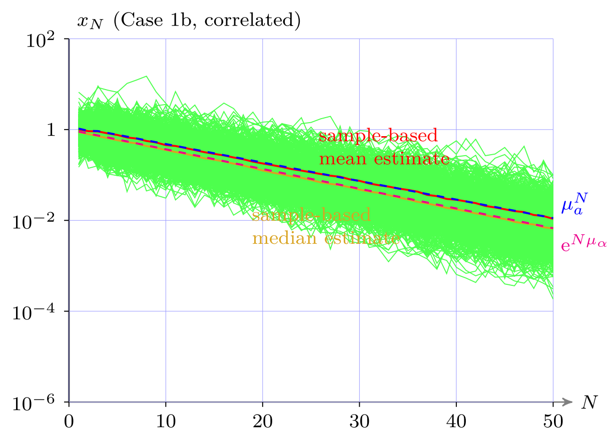

Fig. 3. Simulations of Case 1b trajectories. The correlation is modeled by the Γ_1 filter in the $\alpha_k \sim \mathcal{N}(\mu_\alpha, \sigma_\alpha^2)$ domain. The theoretical and sample medians, as well as the theoretical and sample means, are shown. The median are the mean are both stable.

ζ_N and x_N. Figure 4 illustrates the evolution of these densities for the first $N = 10$ steps in Cases 1a and 1b. For $N = 0$, both the i.i.d. and correlated cases have the same initial density. For both cases the evolution of the mean of ζ_N is identical—correlation does not affect median stability. However correlation modifies the variance of α_N—in this case significantly reducing its growth with N. This is also clear from the spread of sample trajectories illustrated in Figure 3.

The effect of the variance reduction in the ζ_N domain on the mean of x_N is illustrated in Figure 5. All distributions shown are log-normal distributions—the key issue here is the evolution of the median and the mean as a function of N. In the i.i.d. case the mean of x_N (blue $+$) grows with N, while the median (blue circles) decays to zero. In the correlated case the median of x_N (olive circles) is the same as the i.i.d. case, but the mean of x_N (olive $+$) decays to zero. Figures 2 to 5 illustrate that the correlation in Case 1 stabilises the mean of x_N while leaving its median unchanged.

Figure 6 illustrates the evolution of x_N distributions for Case 2. The effect of the correlation in Case 2 is the opposite of that in Case 1; The mean of x_N for i.i.d. feedback gains is stable, and is destabilised by correlated feedback gains.

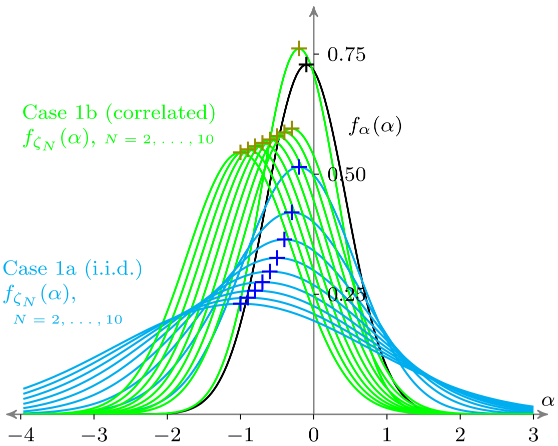

Fig. 4. Evolution of the probability density functions in the α_k domain for Case 1. The black line shows the distribution of the $\alpha \sim \mathcal{N}(\mu_\alpha, \sigma_\alpha^2)$ variable. The cyan plots show the evolution of ζ_N distribution for $N = 2, \ldots, 10$ in the i.i.d. case. The green plots show the evolution of ζ_N distribution in the temporally correlated (using Γ_1) case. Means of each distribution are shown by $+$.

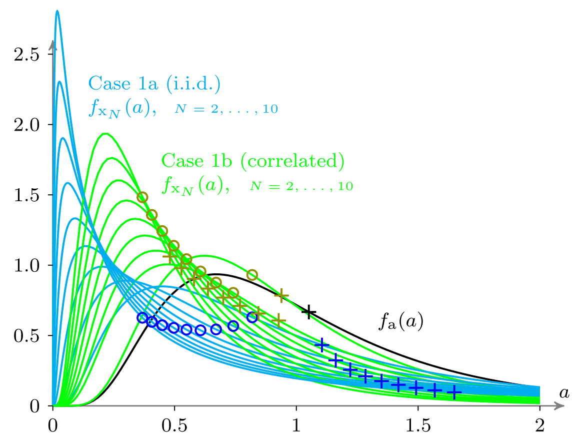

Fig. 5. Evolution of the probability density functions in the a_k domain for Case 1. The black line shows the distribution of the $a \sim \mathcal{LN}$ variable. The cyan plots show the evolution of the x_N distribution for $N = 2, \ldots, 10$ in the i.i.d. case. The green plots show the evolution of the x_N distribution in the correlated (using Γ_1) case. The means of each distribution are shown by $+$ and the medians by circles.

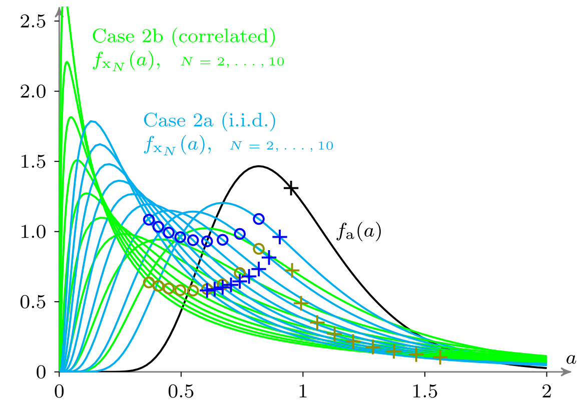

Fig. 6. Evolution of the probability density functions in the x domain for Case 2. In this case the i.i.d. evolution of x_N has a stable median (blue circles) and a stable median (blue $+$). The evolution for the correlated case shows a stable median (olive circles) but an unstable mean (olive $+$).

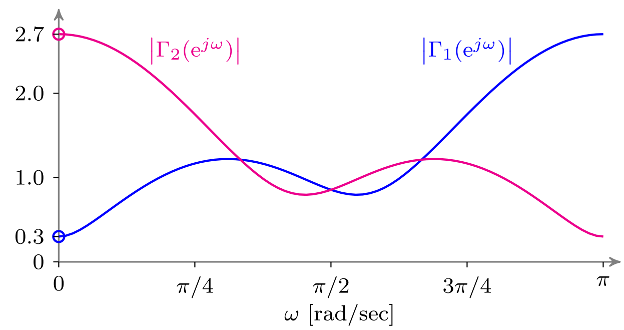

Fig. 7. Frequency domain comparison of the α-domain temporal correlation filters, Γ_1 and Γ_2. Both filters have $\|\Gamma\|_2^2 = 2.21$. The DC gains of each filter are marked by circles.

The differences between the correlation filters is illustrated via their frequency responses in Figure 7. The stability criteria for the mean of x_N is determined by the ratio, $\Gamma(e^{j0})^2/\|\Gamma\|_2^2$. For i.i.d. feedback gains this factor is one. For the correlated feedback in Case 1b the ratio is less than one, whereas in Case 2b it is larger than one.

V. Conclusion and Discussion

We have provided necessary and sufficient conditions for the median, mean, and variance stability of the state of a scalar system under temporally correlated stochastic feedback. The condition for the stability of the median is unaffected by correlation. Correlation is modeled via an LTI filter in the log domain of the feedback gain. The stability of the mean and the variance of the state is determined by properties of this filter. In particular a high pass filter (more precisely a small $\Gamma(e^{j0})^2/\|\Gamma\|_2^2$ ratio) has a stabilising effect on the mean and the variance. In contrast a low-pass filter (high $\Gamma(e^{j0})^2/\|\Gamma\|_2^2$ ratio) has a destabilising effect. The heuristic interpretation is that with a low-pass filter, the destabilising effects of variance persist longer than in the high-pass case, leading to destabilisation with a lower variance. In the high-pass filter case subsequent correlated feedback gains tend to have the opposite sign reducing the destabilising effect of the variance.

References

[1] K. Åström, "On a first-order stochastic differential equation," *Int. J. Control*, vol. 1, no. 4, pp. 301–326, 1965.

[2] J. C. Willems and G. L. Blankenship, "Frequency domain stability criteria for stochastic systems," *IEEE Trans. Autom. Control*, vol. AC-16, no. 4, pp. 292–299, Aug. 1971.

[3] H. Kesten, "Random difference equations and renewal theory for products of random matrices," *Acta Mathematica*, vol. 131, pp. 207–248, Dec. 1973.

[4] M. Milisavljević and E. I. Verriest, "Stability and stabilization of discrete systems with multiplicative noise," in *Proc. Eur. Control Conf.*, 1997, pp. 3503–3508.

[5] R. S. Smith and B. Bamieh, "Stochasticity in feedback loops; great expectations and guaranteed ruin," 2020. [Online]. Available: arXiv:1912.08267.

[6] J. Toner, N. Guttenberg, and Y. Tu, "Swarming in the dirt: Ordered flocks with quenched disorder," *Phys. Rev. Lett.*, vol. 121, no. 24, 2018, Art. no. 248002.

[7] E. Abrahams, *50 Years of Anderson Localization*. Singapore: World Sci., 2010.

[8] D. J. Finney, "On the distribution of a variate whose logarithm is normally distributed," *Suppl. J. Roy. Statist. Soc.*, vol. 7, no. 2, pp. 155–161, 1941.