Chinese Physics Letters

Chin. Phys. Lett. **43** 100002 (2026)

LETTER

DOI: 10.1088/0256-307X/43/10/100002
CSTR: 32039.14.0256-307X.43.10.100002

Exact Finite Response of Higher-Order Statistics in Nonequilibrium Networks

Ruicheng Bao[1,*] and Shiling Liang(梁师翎)[2,3,4,*]

[1]Department of Physics, Graduate School of Science, The University of Tokyo, Hongo, Bunkyo-ku, Tokyo 113-0033, Japan
[2]Center for Systems Biology Dresden, 01307 Dresden, Germany
[3]Max Planck Institute for the Physics of Complex Systems, 01187 Dresden, Germany
[4]Max Planck Institute of Molecular Cell Biology and Genetics, 01307 Dresden, Germany

*Corresponding authors. Email: ruicheng@g.ecc.u-tokyo.ac.jp; shiling@pks.mpg.de

We develop an exact finite-perturbation response theory for the higher-order statistics of general state observables in nonequilibrium Markov networks. For local perturbations, the moment generating function obeys an exact nonlinear response identity in the perturbation strength, controlled by a single kinetic parameter α. This induces an exact nonlinear map for the cumulant generating function and closed expressions for finite responses of all cumulants through partial Bell polynomials. As a central application, we obtain a complete theory of variance control, including sharp saturation limits, a classification of monotone-suppression and sign-reversal regimes, and mean-first-passage-time design principles for optimal noise suppression. The same framework yields exact formulas for operational and symmetrized response resolution along the accessible control branch. A four-state gene-regulation model illustrates how local kinetic control suppresses the occupancy noise of a transcriptionally competent promoter state.

Introduction

Response theory describes how observables change when a system is perturbed. Near thermal equilibrium and for weak perturbations, the fluctuation–dissipation theorem provides a general relation between linear response and spontaneous correlations [1,2]. Far from equilibrium or for finite parameter changes, this relation is no longer valid. However, these conditions are common in driven physical, chemical, and biological systems, including molecular transport [3], biochemical sensing [4−6], adaptation [7], gene regulation [8], and general biochemical processes [9], where sustained fluxes require continuous energy input and experimentally relevant controls are often not infinitesimal [6,7,10−12].

Over the past two decades, nonequilibrium response theory has developed rapidly along two complementary directions. Trajectory-based approaches [13−17] analyze how path probabilities change under a perturbation and typically involve some dynamical quantities, such as the so-called dynamical activity or frenesy [18,19]. In recent years, this viewpoint has led to universal response identities and bounds for nonequilibrium and time-dependent processes [20−30]. Another direction focuses on state observables in nonequilibrium steady states and formulates response directly in terms of stationary probabilities, network topology, and thermodynamic constraints, avoiding incorporating explicit dynamical quantities, thereby making it in principle more experimentally accessible [31−35]. In particular, reference [35] developed a framework for the response of mean values of state and current observables to arbitrarily strong perturbations, extending earlier static-response relations formulated for infinitesimal perturbations to finite control. Experimentally, the two approaches require very different data. By construction, any trajectory-space response compares path probabilities and therefore requires pathwise kinetic information, appearing as dynamical activity or frenesy, traffic and escape-rate contributions, or an equivalent likelihood-ratio observable. Quantitative evaluation generally requires time-resolved data and substantial knowledge of the transition dynamics, making it difficult for snapshot-only or partially observed experiments. By contrast, the finite stationary framework of Ref. [35] can be calibrated from steady-state measurements under controlled perturbations, without reconstructing pathwise dynamical activity. However, it has so far addressed only mean observables. This leaves open how finite perturbations change higher-order fluctuations. Variance determines noise and reliability, while higher cumulants describe asymmetry, intermittency, and non-Gaussian tails. These quantities can be as important as the mean in biochemical sensing, gene regulation, and stochastic transport.

Here, we close this gap by establishing an exact and universal finite-response theory for the higher-order steady-state statistics of general state observables. We derive exact nonlinear response identities for the moment generating function (MGF) and cumulant generating function (CGF) and closed finite-control formulas for cumulants of all orders under arbitrarily strong local perturbations. We then specialize this framework to noise control. The variance is a concave quadratic in the universal gain g, which yields sharp saturation limits, a complete classification of monotone-suppression and sign-reversal regimes, and mean-first-passage-time (MFPT) design principles for optimal control. We further analyze response resolution along the accessible control branch and illustrate the theory with a minimal gene-regulation model, where local kinetic control suppresses the occupancy noise of a transcriptionally competent promoter state.

Setup

Consider an irreducible continuous time Markov chain on N states with generator W (column-sum-zero convention: $W_{ij} \geq 0$ for $i \neq j$, $W_{jj} = -\sum_{i\neq j} W_{ij}$) and unique steady state $\boldsymbol{\pi}$ satisfying $W\boldsymbol{\pi} = \mathbf{0}$. For a state observable f, the steady-state average is $\langle f \rangle = \sum_i f_i \pi_i$. We denote the directed traffic by $\phi_{mn} = W_{mn}\pi_n$ and the MFPT from state j to state i by $\langle t_{ij} \rangle$. Throughout, the two-index rate and traffic subscripts are destination-first: W_{mn} and ϕ_{mn} refer to the transition $n \to m$. Transition rates are parametrized as $W_{ij} = \mathrm{e}^{-A_{ij}}$ with $A_{ij} = B_{ij} - E_j - F_{ij}/2$, where $B_{ij} = B_{ji}$ are symmetric barriers, E_j are vertex energies, and

Received 25 June 2026; accepted 27 August 2026; published 20 October 2026

© 2026 Chinese Physical Society and IOP Publishing Ltd. All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Check for updates

$F_{ij} = -F_{ji}$ are driving forces [20,31,35]. The three archetypal local perturbations are: A-type (single-rate): $A_{mn} \to A_{mn} + \Delta$, changing W_{mn} only; B-type (edge-symmetric): $B_{mn} \to B_{mn} + \Delta$, changing both W_{mn} and W_{nm}; and E-type (vertex): $E_m \to E_m + \Delta$, changing all exit rates from state m.

For any state observable f and any of the three perturbation types ($X \in \{A_{mn}, B_{mn}, E_m\}$), we established in Ref. [35] the exact finite-response identity

$$\Delta\langle f\rangle = g(\Delta;\alpha_X)\,\partial_X\langle f\rangle, \tag{1}$$

where $\Delta\langle f\rangle := \langle f\rangle' - \langle f\rangle$ and the universal gain is

$$g(\Delta;\alpha) = \frac{e^{\Delta}-1}{\alpha + (1-\alpha)e^{\Delta}}. \tag{2}$$

The kinetic parameter $\alpha \in [0,1]$ depends only on the perturbation type,

$$\alpha_X = \begin{cases} \phi_{mn}\langle t_{nm}\rangle, & X = A_{mn}, \\ \phi_{mn}\langle t_{nm}\rangle + \phi_{nm}\langle t_{mn}\rangle, & X = B_{mn}, \\ \pi_m, & X = E_m, \end{cases} \tag{3}$$

and, for generic $\alpha \in (0,1)$, the two strong-control limits are

$$g_{-\infty} \equiv \lim_{\Delta\to-\infty} g(\Delta;\alpha) = \frac{-1}{\alpha}, \quad g_{\infty} \equiv \lim_{\Delta\to\infty} g(\Delta;\alpha) = \frac{1}{1-\alpha}. \tag{4}$$

Since

$$\frac{dg}{d\Delta} = \frac{e^{\Delta}}{[\alpha + (1-\alpha)e^{\Delta}]^2} > 0,$$

the full finite-control response curve is parameterized by the compact interval $g \in [g_{-\infty}, g_{\infty}]$. Below we keep the $\Delta > 0$ branch when discussing monotonicity with respect to the physical control parameter, but the globally optimal noise suppression over the full response curve requires comparing both strong-control endpoints. The bound $\alpha \in [0,1]$ follows from the MFPT-traffic inequality of Ref. [35] for A- and B-type perturbations and from the normalization $\pi_m \in [0,1]$ for E-type. Physically, α measures how strongly the perturbed local kinetics is embedded in the global steady-state flow: $\alpha \to 0$ when bypasses dominate, while $\alpha \to 1$ in gateway or bottleneck limits.

Exact Nonlinear Response for the CGF and Cumulant Hierarchy

Let O be a state observable and define its MGF as

$$\mathcal{Z}(\lambda) \equiv \langle e^{\lambda O}\rangle = \sum_i e^{\lambda O_i}\pi_i. \tag{5}$$

Because $e^{\lambda O}$ is itself a state observable for every fixed λ, equation (1) applies for all λ simultaneously,

$$\mathcal{Z}_{\Delta}(\lambda) = \mathcal{Z}_0(\lambda) + g(\Delta;\alpha_X)\,\partial_X\mathcal{Z}(\lambda). \tag{6}$$

Connected fluctuations are encoded in the CGF [36] $\mathcal{K}(\lambda) = \ln\mathcal{Z}(\lambda) = \sum_{n\geq1}\kappa_n\lambda^n/n!$. Combining $\mathcal{Z} = e^{\mathcal{K}}$ with Eq. (6) yields the exact nonlinear response map

$$\Delta\mathcal{K}(\lambda) = \ln\left[1 + g(\Delta;\alpha)\,\partial_X\mathcal{K}(\lambda)\right]. \tag{7}$$

Writing $\partial_X\mathcal{K}(\lambda) = \sum_{n\geq1} u_n\lambda^n/n!$ with $u_n \equiv \partial_X\kappa_n$ and expanding the logarithm generates the exact Bell-polynomial hierarchy [37,38]

$$\Delta\kappa_n = \sum_{p=1}^{n}(-1)^{p+1}(p-1)!\,g^p B_{n,p}(u_1,\ldots,u_{n-p+1}). \tag{8}$$

For example, $B_{n,1} = u_n$, $B_{2,2} = u_1^2$, $B_{3,2} = 3u_1u_2$, $B_{3,3} = u_1^3$, $B_{4,2} = 4u_1u_3 + 3u_2^2$, $B_{4,3} = 6u_1^2u_2$, and $B_{4,4} = u_1^4$; substituting these coefficients into Eq. (8) gives the formulas below. The coefficient extraction is shown step by step in Section III of the Supplemental Materials (SM). The first four orders are

$$\Delta\kappa_1 = g\,u_1, \tag{9a}$$

$$\Delta\kappa_2 = g\,u_2 - g^2u_1^2, \tag{9b}$$

$$\Delta\kappa_3 = g\,u_3 - 3g^2u_1u_2 + 2g^3u_1^3, \tag{9c}$$

$$\Delta\kappa_4 = g\,u_4 - g^2(4u_1u_3 + 3u_2^2) + 12g^3u_1^2u_2 - 6g^4u_1^4. \tag{9d}$$

Thus strong perturbations do not merely rescale fluctuations. Instead, they mix cumulant orders in a universal and exactly controlled way.

Multiple Local Perturbations

For q independently parametrized local controls $X_a \to X_a + \Delta_a$, repeated use of Eq. (7) gives

$$\mathcal{K}^{(q)}(\lambda) - \mathcal{K}^{(0)}(\lambda) = \sum_{a=1}^{q}\ln[1 + g_a\mathcal{U}_a^{(a-1)}(\lambda)]. \tag{10}$$

Here $\mathcal{K}^{(a)}$ is the CGF after the ath change, while $g_a = g(\Delta_a;\alpha_a^{(a-1)})$ and $\mathcal{U}_a^{(a-1)} = \partial_{X_a}\mathcal{K}^{(a-1)}$ are evaluated in the preceding intermediate state. Although the summands depend on the chosen ordering, their total is the unique endpoint response. Thus, the exact single-control identities remain sequential building blocks at finite strength. More details are given in Section IV of the SM.

Exact Variance Control and Noise Suppression

Among higher-order statistics, the variance is the standard measure of fluctuations and directly quantifies the noise and reliability of a steady-state readout. We therefore specialize to variance control, which allows us to derive design principles for optimal noise suppression. The variance obeys the exact nonlinear law

$$\kappa_2(\Delta) = \kappa_2 + g\,u_2 - g^2u_1^2, \quad g \equiv g(\Delta;\alpha). \tag{11}$$

On the positive branch $g \in [0, g_{\infty}]$, the strongest attainable control is the saturation value

$$\begin{aligned} \kappa_2^{\infty} &= \kappa_2 + g_{\infty}u_2 - g_{\infty}^2u_1^2, \\ \mathcal{S}_{\infty} &\equiv \kappa_2 - \kappa_2^{\infty} = g_{\infty}(g_{\infty}u_1^2 - u_2). \end{aligned} \tag{12}$$

Hence the maximum noise suppression is determined completely by three unperturbed quantities (α, u_1, u_2). The opposite strong-control limit on the negative branch is

$$\begin{aligned} \kappa_2^{-\infty} &= \kappa_2 + g_{-\infty}u_2 - g_{-\infty}^2u_1^2, \\ \mathcal{S}_{-\infty} &\equiv \kappa_2 - \kappa_2^{-\infty} = g_{-\infty}(g_{-\infty}u_1^2 - u_2). \end{aligned} \tag{13}$$

Since $\kappa_2(g) = \kappa_2 + g\,u_2 - g^2u_1^2$ is a concave quadratic in g, the globally optimal noise suppression over the full accessible response curve is always attained at one of the two strong-control endpoints,

$$\kappa_2^{\mathrm{opt}} = \min\{\kappa_2^{-\infty}, \kappa_2^{\infty}\}, \quad \mathcal{S}_{\mathrm{opt}} = \max\{\mathcal{S}_{-\infty}, \mathcal{S}_{\infty}\}. \tag{14}$$

Equivalently, the positive endpoint is better than the negative one if and only if

$$\kappa_2^{\infty} \leq \kappa_2^{-\infty} \iff (2\alpha - 1)u_1^2 \geq \alpha(1-\alpha)u_2. \tag{15}$$

Hence the branchwise analysis below still applies exactly as written to the $\Delta > 0$ branch, while Eq. (14) completes it into a statement about the full response curve.

The term u_2 is the direct linear noise susceptibility, whereas $-g^2u_1^2$ is a universal nonlinear inheritance from the mean response. Noise suppression is therefore possible even when the linear response predicts initial noise amplification.

The control parameter g_{∞} has a clear physical meaning. For an A-type perturbation of the transition $n \to m$, let $\tau_{mn} \equiv 1/\phi_{mn}$ be the mean

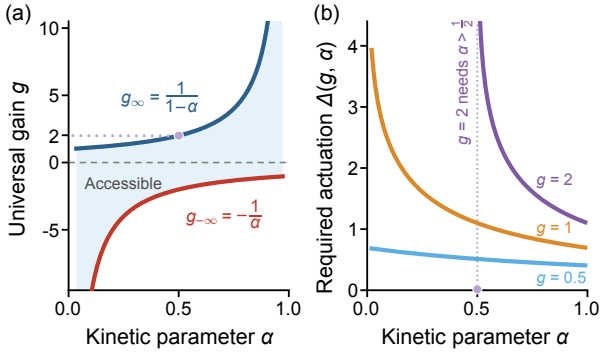


Fig. 1. Kinetic leverage and actuation cost are both set by α. (a) For $0 < \alpha < 1$, the closure of the accessible universal-gain interval is $g \in [-1/\alpha, 1/(1-\alpha)]$ (shaded); its endpoints are approached as $\Delta \to -\infty$ and $\Delta \to +\infty$, respectively, while $g = 0$ at $\Delta = 0$. The interval widens as $\alpha \to 0$ (negative branch) and $\alpha \to 1$ (positive branch). (b) For any fixed target gain g, there exists a threshold (e.g., finite actuation at $g = 2$ requires $\alpha > 1/2$, while $\alpha = 1/2$ is reached only at the $\Delta \to \infty$ endpoint), beyond which the required actuator strength $\Delta(g,\alpha)$ decreases with α. Local controls on kinetic gateways or highly occupied vertices are therefore both stronger and cheaper.

interval between successive uses of that transition. Let $\tau_{mn}^{\mathrm{edge}}$ be the mean time, starting from state n, until the next use of the directed edge $n \to m$; it includes any intervening excursions from and returns to state n. The renewal MFPT decomposition of Ref. [35] is then $\tau_{mn} = \langle t_{nm} \rangle + \tau_{mn}^{\mathrm{edge}}$. Equation (3) then reads $\alpha_{\mathrm{A}} \equiv \alpha_{A_{mn}} = \langle t_{nm} \rangle / \tau_{mn}$, so that

$$g_{\infty}^{\mathrm{A}} = \frac{1}{1-\alpha_{\mathrm{A}}} = \frac{\tau_{mn}}{\tau_{mn}^{\mathrm{edge}}}, \quad g_{-\infty}^{\mathrm{A}} = -\frac{1}{\alpha_{\mathrm{A}}} = -\frac{\tau_{mn}}{\langle t_{nm} \rangle}. \tag{16}$$

Large g_{∞} therefore identifies kinetic gateways: most of the inter-event time is spent in the return first-passage excursion rather than in the subsequent edge-event waiting time. The two strong-control endpoints compare the same inter-event time with these two complementary renewal contributions. For B-type perturbations, α_{B} adds the analogous forward and backward excursion fractions and thus quantifies the two-way bottleneck character of the edge; for E-type perturbations, $g_{\infty}^{\mathrm{E}} = 1/(1-\pi_m)$, so frequently occupied vertices are the strongest local levers. More generally, the perturbation needed to realize any target gain $g < g_{\infty}$ is

$$\Delta(g,\alpha) = \ln\left[\frac{1+\alpha g}{1-(1-\alpha)g}\right], \tag{17}$$

which decreases with α for fixed g. Thus, gateway edges or highly occupied vertices are not only more powerful. They are also cheaper to actuate at a fixed target response. Figure 1 visualizes this leverage principle directly. Equation (11) yields the exact noise-control regimes. If $u_2 \le 0$, then $\partial_g \kappa_2 = u_2 - 2g u_1^2 \le 0$ for all $g > 0$: the variance is monotonically nonincreasing; except for the trivial case $u_1 = u_2 = 0$, it is suppressed monotonically, and the optimum is the saturation value (12). If $0 < u_2 < g_{\infty} u_1^2$, the variance is initially amplified but eventually suppressed. For $u_1 \neq 0$, the maximal-noise point and the sign-reversal point are

$$g_{\mathrm{pk}} = \frac{u_2}{2u_1^2}, \quad g_{\mathrm{sr}} = \frac{u_2}{u_1^2}, \quad \Delta_{\mathrm{pk/sr}} = \Delta(g_{\mathrm{pk/sr}}, \alpha), \tag{18}$$

with peak value $\kappa_2^{\mathrm{pk}} = \kappa_2 + u_2^2/(4u_1^2)$. If $u_1 = 0$, the nonlinear inheritance term vanishes and Eq. (11) reduces to the purely linear form $\kappa_2(\Delta) = \kappa_2 + g u_2$. At g_{pk} the direct linear amplification u_2 is exactly balanced by the nonlinear inheritance $2g u_1^2$; at g_{sr} the total noise shift vanishes. The remarkable point is that g_{pk} and g_{sr} depend only on the ratio u_2/u_1^2: kinetics enters only through whether these targets are accessible and how costly they are. Indeed, suppression on the positive branch is possible iff $g_{\infty} > g_{\mathrm{sr}}$, equivalently

$$\alpha > \alpha_{\mathrm{c}} \equiv 1 - \frac{u_1^2}{u_2} \quad (u_2 > 0), \tag{19}$$

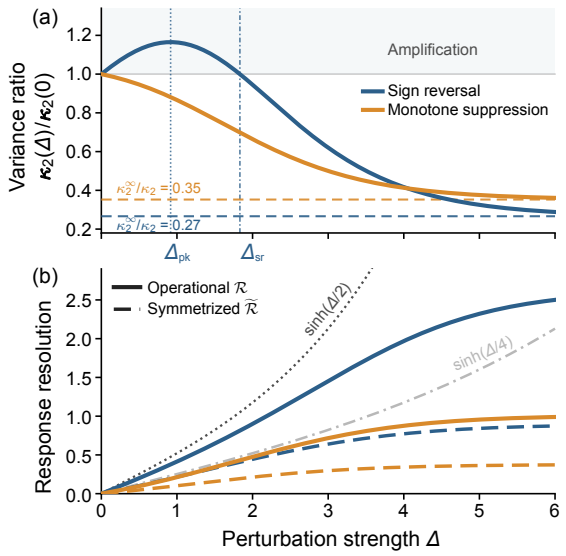


Fig. 2. Exact noise suppression and response resolution. (a) The variance ratio $\kappa_2(\Delta)/\kappa_2(0)$ for two different observables in the same six-state model shows sign-reversal and monotone-suppression behaviours. Blue dotted and dash-dotted vertical lines mark the maximal-noise point and the sign-reversal point predicted by Eq. (18); horizontal dashed lines mark the exact $\Delta \to \infty$ variance-ratio endpoints 0.266 (blue) and 0.353 (orange). Both panels show $0 \le \Delta \le 6$, corresponding for $\alpha = 5/7$ to $0 \le g \le 3.470$ within the full positive branch $0 \le g < g_{\infty} = 3.5$; $g = g_{\infty}$ is reached only as $\Delta \to \infty$. (b) The operational response resolution $\mathcal{R}(\Delta)$ and the symmetrized response resolution $\widetilde{\mathcal{R}}(\Delta)$ for the same two observables of panel (a), which are bounded by $\sinh(\Delta/2)$ and $\sinh(\Delta/4)$, respectively. Even when the variance is nonmonotonic, the corresponding response resolution increases monotonically and approaches a sharp saturation value.

with the obvious understanding that $\alpha_{\mathrm{c}} < 0$ means suppression is already accessible for all $\alpha \in [0,1]$. For A-type control this becomes the MFPT condition $\tau_{mn}/\tau_{mn}^{\mathrm{edge}} > u_2/u_1^2$. If instead $g_{\infty} \le g_{\mathrm{sr}}$, the positive branch never suppresses the variance: it may still be monotone-enhancing or show a finite-Δ turnover, but it remains above the baseline noise level throughout. Figure 2 shows both the sign-reversal and monotone-suppression scenarios.

The saturation value itself carries a design message. Whenever suppression is feasible one has

$$\frac{\partial \mathcal{S}_{\infty}}{\partial g_{\infty}} = 2 g_{\infty} u_1^2 - u_2 > 0, \tag{20}$$

so stronger kinetic leverage always improves the optimal attainable suppression on a suppressing branch. Moreover, for a desired suppression amount $\mathcal{S} > 0$ and generic $u_1 \neq 0$, the required gain solves

$$g_*(\mathcal{S}) = \frac{u_2 + \sqrt{u_2^2 + 4u_1^2 \mathcal{S}}}{2u_1^2}, \quad \Delta_*(\mathcal{S}) = \Delta(g_*(\mathcal{S}), \alpha), \tag{21}$$

while for the special case $u_1 = 0$ one simply has $g_* = \mathcal{S}/(-u_2)$. Equation (21) makes the MFPT-based design principle explicit on the $\Delta > 0$ branch: large α reduces the actuation cost for any target suppression and can be the difference between an inaccessible and an accessible control objective. For the full response curve, the globally optimal suppression is still fixed simply by Eq. (14), i.e., by comparing the two strong-control endpoints.

A useful linear benchmark follows from applying the mean-response bound of Ref. [35] to the nonnegative observable $Q_i = (O_i - \kappa_1)^2$, with $\kappa_1 = \langle O \rangle$ the unperturbed mean held fixed under ∂_X:

$$\left|\partial_X \ln \kappa_2\right| \le 1, \quad X = A_{mn}, B_{mn}, E_m, \tag{22}$$

with the stronger B-type bound $|\partial_{B_{mn}} \ln \kappa_2| \le \tanh(F_{c \ni mn}^{\max}/4)$. Integrated to finite Δ, these give only loose exponential envelopes. The true nonlinear limit is Eq. (12).

Response Resolution

Suppressing noise is only useful if the signal—the shift in the mean—is not suppressed even faster. We therefore quantify how reliably a finite perturbation can be distinguished from baseline against its own noise, the response resolution. We use the operational ratio defined in Ref. [35],

$$\mathcal{R}(\Delta) \equiv \frac{|\Delta\kappa_1|}{\sqrt{\kappa_2(\Delta)}} = \frac{|g\,u_1|}{\sqrt{\kappa_2 + g u_2 - g^2 u_1^2}}, \tag{23}$$

which obeys the finite-control bound,

$$\mathcal{R}(\Delta) \le \begin{cases} \sinh(|\Delta|/2), & X = A, E, \\ \sinh(|\Delta|/2)\tanh(F^{\max}_{c\ni mn}/4), & X = B. \end{cases} \tag{24}$$

Here $F^{\max}_{c\ni mn}$ denotes the largest cycle affinity among cycles that pass through the perturbed edge. A complementary symmetrized ratio compares the signal with the mean noise scale of the initial and final steady states,

$$\widetilde{\mathcal{R}}(\Delta) \equiv \frac{|\Delta\kappa_1|}{\sqrt{\kappa_2} + \sqrt{\kappa_2(\Delta)}} = \frac{|g\,u_1|}{\sqrt{\kappa_2} + \sqrt{\kappa_2 + g u_2 - g^2 u_1^2}}. \tag{25}$$

A midpoint-splitting argument then yields a novel exact finite-control bound

$$\widetilde{\mathcal{R}}(\Delta) \le \begin{cases} \sinh(|\Delta|/4), & X = A, E, \\ \sinh(|\Delta|/4)\tanh(F^{\max}_{c\ni mn}/4), & X = B, \end{cases} \tag{26}$$

whose proof is given in the SM. The dashed curves in Fig. 2 show this symmetrized quantity, while the gray dotted and dash-dotted curves show the bounds $\sinh(\Delta/2)$ and $\sinh(\Delta/4)$, respectively.

On the positive control branch $g \in [0, g_\infty]$, the gain is itself strictly monotone, $\mathrm{d}g/\mathrm{d}\Delta = \mathrm{e}^{\Delta}/[\alpha + (1-\alpha)\mathrm{e}^{\Delta}]^2 > 0$, so monotonicity in g and in Δ are equivalent. For the operational ratio one finds (details in SM)

$$\frac{\mathrm{d}\mathcal{R}}{\mathrm{d}g} = \frac{|u_1|\,(2\kappa_2 + g u_2)}{2\,\kappa_2(\Delta)^{3/2}} \ge 0, \quad g > 0, \tag{27}$$

where nonnegativity follows from $2\kappa_2 + g u_2 = \kappa_2 + \kappa_2(\Delta) + g^2 u_1^2 > 0$, with strict inequality whenever $u_1 \neq 0$. Thus $\mathcal{R}(\Delta)$ is monotone nondecreasing along the accessible branch even in the sign-reversal regime where the variance itself is nonmonotonic. For the symmetrized ratio,

$$\frac{\mathrm{d}\widetilde{\mathcal{R}}}{\mathrm{d}g} = \frac{|u_1|}{\left[\sqrt{\kappa_2} + \sqrt{\kappa_2(\Delta)}\right]^2}\left[\sqrt{\kappa_2} + \frac{2\kappa_2 + g u_2}{2\sqrt{\kappa_2(\Delta)}}\right] \ge 0, \tag{28}$$

with the same nonnegativity mechanism, again strict whenever $u_1 \neq 0$. Hence both response–resolution measures are monotone nondecreasing with control strength on the accessible branch. Their exact saturation values are obtained by setting $g \to g_\infty$: $\mathcal{R}_\infty = g_\infty |u_1| / \sqrt{\kappa_2^\infty}$ and $\widetilde{\mathcal{R}}_\infty = g_\infty |u_1| / (\sqrt{\kappa_2} + \sqrt{\kappa_2^\infty})$.

The higher cumulants follow the same bounded-gain logic. Every $\Delta\kappa_n(g)$ is a degree-n polynomial fixed by Eq. (8), and therefore has exact strong-control values obtained by setting $g \to g_{-\infty}$ or $g \to g_\infty$. Skewness, kurtosis, and more general shape descriptors can thus be optimized by solving finite-degree polynomial equations inside the compact interval $[0, g_\infty]$; representative examples are shown in the SM.

Numerical Illustrations

All panels in Fig. 2 are computed exactly for a six-state two-basin nonequilibrium network. Basin A ($\{0,1,2\}$) and basin B ($\{3,4,5\}$) each support a biased three-state cycle (forward rate 2.0, backward rate 0.5) and are coupled by a weak asymmetric gateway edge $2 \to 3$ (rate 0.05) and $3 \to 2$ (rate 0.02). The perturbed parameter is the A-type gateway rate on $2 \to 3$, for which the exact kinetic parameter is $\alpha = 5/7$ and hence $g_\infty = 7/2$. The observable is $O_i = i$ for the sign-reversal example in Fig. 2, and the three-level readout $O = (0,0,1,0,1,2)$ for the monotone-suppression curve in Fig. 2. All stationary distributions, MFPTs, and cumulants are obtained by direct linear-algebra calculations.

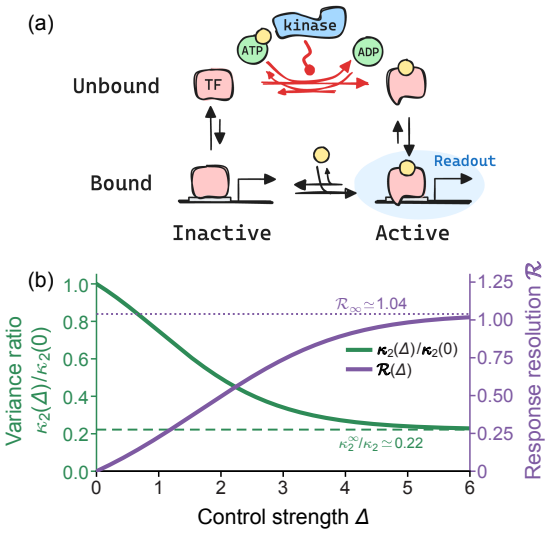


Fig. 3. Noise control in gene regulation. (a) A minimal four-state model for an active transcription-factor/promoter module, with a kinase acting as the B-type control on the $0 \leftrightarrow 1$ activating edge. (b) Exact finite-control response of the transcriptionally competent promoter occupancy $O = (0,0,0,1)$ under B-type control. The plotted $0 \le \Delta \le 6$ maps to $0 \le g \le 2.573$ within the full branch $0 \le g < g_\infty \simeq 2.59$; the dashed and dotted lines are the exact $\Delta \to \infty$ endpoints, corresponding in the common-scaling regime to $[K]'/[K] \to 0$. The green solid curve shows $\kappa_2(\Delta)/\kappa_2(0)$ and the green dashed line its strong-control limit; the purple solid curve shows $\mathcal{R}(\Delta)$ and the purple dotted line its saturation value $\mathcal{R}_\infty$. The output variance is monotonically suppressed and saturates at $\kappa_2^\infty/\kappa_2 \simeq 0.22$, while the operational response resolution grows monotonically to $\mathcal{R}_\infty \simeq 1.04$.

Noise Control in Gene Regulation

Enzymatically controlled interconversion between inactive and active molecular species is a ubiquitous motif in cellular signaling and can act as tunable filters of fluctuations and information transmission [39-41]. It also provides a route for regulating transcription-factor (TF) activity and hence downstream gene expression [42]. Figure 3 illustrates these ideas with a minimal four-state model for an active TF/promoter module: $0 = (I,U)$, $1 = (A,U)$, $2 = (I,B)$, and $3 = (A,B)$, where A and I refer to the active and inactive TFs, and U and B refer to the unbound and bound promoters, respectively.

The system is controlled through a kinase, an enzyme that catalyzes the activation of the unbound TF, converting it from the inactive to the active state. As a catalyst, the kinase lowers the activation barrier of this conversion, modulating the forward and backward rates of the $0 \leftrightarrow 1$ edge by the same factor. This is precisely the B-type perturbation introduced above. In an enzyme-dominated common-scaling regime, in which the two catalyzed rates share the factor $W'/W = [K]'/[K]$, our convention $W'/W = \mathrm{e}^{-\Delta}$ gives $\Delta = -\ln([K]'/[K])$; hence positive Δ corresponds to lowering the effective catalyst abundance. More generally, Δ can be calibrated directly from the measured common rate-rescaling factor. The readout is taken to be the active-bound state (state 3), with the binary form $O = (0,0,0,1)$, in which the active TF enhances the association of RNA polymerase and thereby raises the gene-expression level. If transcription initiation from that state occurs with rate r, then the mean RNA production rate is simply $r\kappa_1$, so controlling κ_1 and κ_2 corresponds directly to controlling the mean and noise of promoter competence. Because this observable is binary, the entire steady-state distribution is Bernoulli and therefore fixed by a single number $p(\Delta) = \kappa_1(\Delta) = \pi_{\mathrm{A,B}}(\Delta)$; in particular,

$\kappa_2(\Delta) = p(\Delta)[1-p(\Delta)]$ is maximal at $p = 1/2$ and decreases as the promoter becomes more decisively OFF or ON. The present control branch drives the system away from half occupation toward a low-noise OFF state.

For the parameter set in Fig. 3, the unperturbed kinetics give $\alpha \simeq 0.614$ and hence $g_\infty \simeq 2.59$. The transcriptionally competent occupancy changes from $p_0 = \pi_{A,B} \simeq 0.259$ to the exact strong-control limit $p_\infty \simeq 0.0446$. Correspondingly, the promoter–noise ratio is strongly reduced, $\kappa_2^\infty/\kappa_2 \simeq 0.222$, so the maximal suppression is about 78%. At the same time the operational response resolution grows monotonically to $\mathcal{R}_\infty \simeq 1.04$, while the complementary symmetrized ratio (not shown) saturates at $\widetilde{\mathcal{R}}_\infty \simeq 0.333$. This example therefore provides a design rule: perturb the activation step with large kinetic leverage α and with a response direction that moves the transcriptionally competent promoter occupancy away from the high-noise midpoint $p = 1/2$. Because the readout is binary, all higher cumulants are simultaneously controlled as well.

Experimental Implementation

All inputs of the theory are measurable. Steady-state snapshots at a few control strengths give the cumulants $\kappa_n(\Delta)$, with Δ known from the measured rate rescaling. A weak control gives the slopes u_n, because $g(\Delta;\alpha) = \Delta$ to first order for every α. One finite strength Δ_* then fixes α from the mean shift alone [35],

$$g_* = \frac{\kappa_1(\Delta_*) - \kappa_1}{u_1}, \quad \alpha = \frac{e^{\Delta_*}}{e^{\Delta_*} - 1} - \frac{1}{g_*}, \tag{29}$$

for $u_1 \neq 0$. Equations (6) and (8) predict the full response of every cumulant; the MFPTs that enter Eq. (3) are never measured. By Eq. (6), the ratio $[\mathcal{Z}_{\Delta_1}(\lambda) - \mathcal{Z}_0(\lambda)]/[\mathcal{Z}_{\Delta_2}(\lambda) - \mathcal{Z}_0(\lambda)]$ is independent of λ, which tests the one-parameter form and determines α from two finite controls alone. Note that the measurement-induced extrinsic noise is removed by the differences themselves: noise independent of the state and the control shifts each measured cumulant by a control-independent amount, which drops out of $\Delta\kappa_n$, and a detector gain or offset cancels in Eq. (29). To implement this framework in the gene-regulation example, we measure the fraction of cells in the transcriptionally competent state at several kinase concentrations.

Discussion

We have developed an exact and universal theoretical framework for analyzing higher-order steady-state fluctuations of state observables in nonequilibrium Markov networks under finite local perturbations. The framework gives an exact nonlinear response identity for the MGF and a closed relation for the CGF, which determines the finite response of every cumulant from unperturbed response coefficients and a single kinetic parameter. For the variance, we obtained the full finite-control law, the strong-control limits, the conditions for monotonic suppression and sign reversal, and universal bounds on response resolution. The MFPT interpretation of the kinetic parameter also identifies gateway transitions and highly occupied states as effective control points for noise suppression.

Acknowledgements

R. B. was supported by JSPS KAKENHI (Grant No. 25KJ0766). S. L. acknowledges financial support from the Max Planck Society.

References

[1] Kubo R 1966 *Rep. Prog. Phys.* **29** 255
[2] Marconi U M B, Puglisi A, Rondoni L, and Vulpiani A 2008 *Phys. Rep.* **461** 111
[3] Phillips R 2020 *The Molecular Switch: Signaling and Allostery* (Princeton University Press)
[4] Govern C C and ten Wolde P R 2014 *Phys. Rev. Lett.* **113** 258102
[5] Kaizu K, De Ronde W, Paijmans J, Takahashi K, Tostevin F, and Ten Wolde P R 2014 *Biophys. J.* **106** 976
[6] ten Wolde P R, Becker N B, Ouldridge T E, and Mugler A 2016 *J. Stat. Phys.* **162** 1395
[7] Lan G, Sartori P, Neumann S, Sourjik V, and Tu Y 2012 *Nat. Phys.* **8** 422
[8] Mahdavi S D, Salmon G L, Daghlian P, Garcia H G, and Phillips R 2024 *Proc. Natl. Acad. Sci. USA* **121** e2411395121
[9] Mallory J D, Kolomeisky A B, and Igoshin O A 2020 *Proc. Natl. Acad. Sci. USA* **117** 8884
[10] Seifert U 2012 *Rep. Prog. Phys.* **75** 126001
[11] Peliti L and Pigolotti S 2021 *Stochastic Thermodynamics: An Introduction* (Princeton University Press)
[12] Falasco G and Esposito M 2025 *Rev. Mod. Phys.* **97** 015002
[13] Harada T and Sasa S 2005 *Phys. Rev. Lett.* **95** 130602
[14] Speck T and Seifert U 2006 *Europhys. Lett.* **74** 391
[15] Baiesi M, Maes C, and Wynants B 2009 *Phys. Rev. Lett.* **103** 010602
[16] Baiesi M and Maes C 2013 *New J. Phys.* **15** 013004
[17] Maes C 2020 *Front. Phys.* **8** 229
[18] Maes C and Van Wieren M H 2006 *Phys. Rev. Lett.* **96** 240601
[19] Maes C 2020 *Phys. Rep.* **850** 1
[20] Dechant A and Sasa S 2020 *Proc. Natl. Acad. Sci. USA* **117** 6430
[21] Aslyamov T, Ptaszyński K, and Esposito M 2025 *Phys. Rev. Lett.* **134** 157101
[22] Ptaszyński K, Aslyamov T, and Esposito M 2024 *Phys. Rev. Lett.* **133** 227101
[23] Kwon E, Chun H M, Park H, and Lee J S 2025 *Phys. Rev. Lett.* **135** 097101
[24] Hasegawa Y 2024 *Phys. Rev. Lett.* **132** 087102
[25] Zheng J and Lu Z 2025 *Phys. Rev. E* **112** L012103
[26] Ptaszyński K, Aslyamov T, and Esposito M 2026 *Phys. Rev. E* **113** 024130
[27] Aslyamov T, Ptaszyński K, and Esposito M 2026 *Phys. Rev. Lett.* **136** 067102
[28] Zheng J and Lu Z 2026 arXiv: 2602.18631 [cond-mat.stat-mech]
[29] Bao R and Liang S 2026 arXiv: 2608.11202 [cond-mat.stat-mech]
[30] Liang S and Bao R 2026 arXiv: 2608.06368 [cond-mat.stat-mech]
[31] Owen J A, Gingrich T R, and Horowitz J M 2020 *Phys. Rev. X* **10** 011066
[32] Owen J A and Horowitz J M 2023 *Nat. Commun.* **14** 1280
[33] Fernandes Martins G and Horowitz J M 2023 *Phys. Rev. E* **108** 044113
[34] Aslyamov T and Esposito M 2024 *Phys. Rev. Lett.* **133** 107103
[35] Bao R and Liang S 2024 arXiv: 2412.19602 [cond-mat.stat-mech]
[36] Touchette H 2009 *Phys. Rep.* **478** 1
[37] Comtet L 1974 *Advanced Combinatorics: The Art of Finite and Infinite Expansions* (Springer Dordrecht)
[38] Riordan J 1968 *Combinatorial Identities* (John Wiley & Sons)
[39] Gomez-Uribe C, Verghese G C, and Mirny L A 2007 *PLoS Comput. Biol.* **3** e246
[40] Hinczewski M and Thirumalai D 2014 *Phys. Rev. X* **4** 041017
[41] Andrews S S, Peria W J, Yu R C, Colman-Lerner A, and Brent R 2016 *Cell Syst.* **3** 444
[42] Whitmarsh A J and Davis R J 2000 *Cell. Mol. Life Sci.* **57** 1172