Black Hole Mass and Spin Evolution Under Geometric Brownian Motion: A Stochastic Control Approach

Gregory Tango[1]

[1] *Research overseen and assembled at PREMIERE Research Academy*

We develop a new analytical framework for modeling the coupled evolution of black hole mass and spin under stochastic accretion processes, using a Geometric Brownian Motion (GBM) representation for both quantities. While traditional models treat mass and spin growth through deterministic accretion or semi-empirical prescriptions, we introduce an optimal stochastic control formulation in which an external control variable influences the accretion torque and spin-up rate. The system is analyzed through the Hamilton–Jacobi–Bellman (HJB) equation, enabling the derivation of optimal feedback laws that steer the black hole toward physically significant target states. The objective functional rewards closeness to a desired (M^*, a^*) pair while penalizing excessive control effort. As a worked example, we solve the mass-only GBM case in closed form, deriving the optimal control trajectory and expected mass evolution. This approach bridges stochastic modeling and black hole astrophysics, offering a tractable yet physically meaningful way to explore the interplay of noise, accretion, and spin evolution. The method has potential applications in the interpretation of long-term variability in active galactic nuclei (AGN) and the stochastic growth histories of supermassive black holes.

I. INTRODUCTION

The mass and spin of astrophysical black holes are fundamental parameters that govern their physical properties, radiative efficiency, and interaction with their surrounding environment [1–3]. The spin influences the location of the innermost stable circular orbit (ISCO), thereby affecting accretion disk efficiency and the potential for jet launching through mechanisms such as the Blandford–Znajek process [4, 5]. Mass growth is primarily driven by gas accretion and black hole mergers, with observational evidence indicating that both processes are subject to stochastic variability over a range of timescales [6, 7].

Traditional models of black hole evolution often assume deterministic growth laws or semi-stochastic prescriptions that incorporate noise phenomenologically [8, 9]. However, the intrinsic variability of accretion rates in active galactic nuclei (AGN) and X-ray binaries suggests that stochastic processes may play a central role in shaping long-term mass and spin evolution [10, 11]. In particular, the variability observed in AGN light curves is well-modeled by stochastic processes such as damped random walks [6] or continuous-time autoregressive moving average (CARMA) models [12], motivating a deeper integration of stochastic dynamics into the theoretical treatment of black hole growth.

Despite significant progress, there remains a gap in the literature: no analytical framework currently exists that models the coupled stochastic evolution of black hole mass and spin using physically interpretable stochastic differential equations (SDEs) combined with optimal control theory. This gap is surprising given that such methods are widely used in other domains of physics and engineering to control noisy dynamical systems [20, 22]. The present work fills this gap by introducing a stochastic control model in which both mass and spin follow coupled GBM dynamics subject to an external control variable that can modulate accretion torques or spin alignment processes.

The use of GBM to model astrophysical quantities is motivated by its multiplicative noise structure, which ensures positivity of mass and spin magnitude while capturing relative, rather than absolute, variability [14]. In the black hole context, GBM can represent the compound effect of many small, stochastic accretion events, as well as variability induced by turbulence and magnetohydrodynamic (MHD) instabilities in the accretion flow. Coupling this with the Hamilton–Jacobi–Bellman (HJB) formalism allows us to determine optimal strategies for steering the system toward a physically significant state, such as a near-maximally spinning Kerr black hole of given mass.

In addition to introducing the model, we provide a fully worked analytical example for the mass-only GBM case, solving the HJB equation explicitly and interpreting the optimal control law. This example serves as a tractable entry point for astrophysicists unfamiliar with stochastic control theory, while demonstrating the method's capability for more complex coupled systems.

II. LITERATURE REVIEW

Theoretical studies of black hole spin evolution date back to the seminal works of Bardeen [1] and Thorne [2], who examined the effect of accretion from a thin disk on the spin parameter a. These early models showed that accretion from the ISCO can spin up a black hole toward an equilibrium value $a \simeq 0.998$ due to the capture of radiation with negative angular momentum. More recent treatments [10?] have incorporated MHD effects and magnetically arrested disks, revealing a rich interplay between spin evolution and jet launching efficiency.

Black hole mass growth has been modeled extensively in the context of galaxy evolution [7, 9?], with both coherent accretion episodes and chaotic, small-scale accretion events contributing to the final mass distribution. Observations suggest that accretion rates in AGN are highly variable [6?], motivating stochastic modeling approaches. Indeed, the damped random walk model [6] and its generalizations [12, 13] have been successfully applied to AGN light curves, providing statistical descriptions of variability but without a direct link to mass and spin evolution.

From a stochastic modeling perspective, GBM is a natural choice for positive-definite quantities influenced by multiplicative noise [14]. While GBM is well-established in financial mathematics [15] and has been applied to population dynamics and particle transport [16], its application to black hole astrophysics has been limited. One exception is the work of [17], who used stochastic processes to describe galaxy mergers and black hole growth, though without the use of control theory.

Optimal control and the HJB equation have seen increasing use in physics and engineering, from quantum control [18] to plasma confinement [19]. The HJB formalism provides a powerful way to determine optimal policies in noisy environments by solving a nonlinear partial differential equation for the value function [20–22]. However, to our knowledge, no prior study has applied HJB-based stochastic control to the astrophysical problem of black hole mass and spin evolution.

This paper therefore contributes to the literature by:

1. Introducing a coupled GBM model for black hole mass and spin evolution.
2. Formulating the evolution as a stochastic control problem with a clear physical objective functional.
3. Deriving the associated HJB equation and solving a tractable mass-only case in closed form.
4. Demonstrating how such methods can provide physical insights into the stochastic growth histories of black holes.

III. THEORETICAL FRAMEWORK

A. Stochastic Dynamics of Mass and Spin

We consider a black hole characterized by its mass $M_t > 0$ and dimensionless spin parameter $a_t \in [0, 1]$. The evolution of (M_t, a_t) is modeled by a pair of coupled stochastic differential equations (SDEs) of the Geometric Brownian Motion (GBM) type:

$$dM_t = \mu_M(M_t, a_t, u_t)M_t\,dt + \sigma_M M_t\,dW_t^{(M)}, \tag{1}$$

$$da_t = \mu_a(M_t, a_t, u_t)a_t\,dt + \sigma_a a_t\,dW_t^{(a)}, \tag{2}$$

where:

- μ_M and μ_a are the drift coefficients, which may depend on the current state (M_t, a_t) and the control u_t.
- $\sigma_M, \sigma_a > 0$ are constant volatility parameters representing the relative strength of stochastic fluctuations in mass and spin.
- $W_t^{(M)}$ and $W_t^{(a)}$ are standard Wiener processes, with instantaneous correlation $\rho \in [-1, 1]$:

$$dW_t^{(M)}\,dW_t^{(a)} = \rho\,dt.$$

- The control variable u_t is a measurable process adapted to the natural filtration, representing an external modulation of accretion torque, spin alignment efficiency, or other mechanisms affecting growth.

The multiplicative noise structure ensures that $M_t > 0$ for all t and $a_t \geq 0$.

In the astrophysical context, Eq. (1) models stochastic variability in the accretion rate, driven by turbulence, disk instabilities, or chaotic fueling episodes. Eq. (2) captures stochastic spin-up and spin-down due to fluctuating accretion angular momentum orientation, as well as spin changes from minor mergers.

B. Control Objective

We define a physically motivated objective functional:

$$J(u) = \mathbb{E}\left[\int_0^T \left(-q(M_t - M^*)^2 - q_s(a_t - a^*)^2 - ru_t^2\right) dt\right], \tag{3}$$

where:

- M^* is a target black hole mass (e.g., a mass associated with a given evolutionary stage or observational class, such as a luminous quasar SMBH).
- a^* is a target spin (e.g., near-maximal spin $a^* \approx 0.998$ for high radiative efficiency).
- $q, q_s > 0$ weight the importance of mass and spin tracking, respectively.
- $r > 0$ penalizes excessive control effort.
- $T > 0$ is the finite time horizon of interest.

Interpretation: The term $-q(M_t - M^*)^2$ rewards trajectories where the mass remains close to the target M^*. The term $-q_s(a_t - a^*)^2$ similarly rewards proximity to the target spin a^*. The penalty term $-ru_t^2$ prevents unrealistic control inputs, enforcing smooth and physically plausible evolution. Maximizing $J(u)$ therefore means *steering the system toward the desired (M^*, a^*) state on average, while minimizing control energy.*

C. Hamilton–Jacobi–Bellman Equation

Let $V(t, M, a)$ denote the value function, i.e., the maximal expected objective achievable from state (M, a) at time t:

$$V(t, M, a) = \sup_{u_\cdot \in \mathcal{U}} \mathbb{E}_{t,M,a}\left[\int_t^T \left(-q(M_s - M^*)^2 - q_s(a_s - a^*)^2 - ru_s^2\right) ds\right].$$

By the dynamic programming principle [20, 21], V satisfies the Hamilton–Jacobi–Bellman (HJB) equation:

$$\begin{aligned} 0 = V_t + \sup_{u \in \mathbb{R}} \Big\{ & [\mu_M(M, a, u)M]\, V_M + [\mu_a(M, a, u)a]\, V_a \\ & + \frac{1}{2}\sigma_M^2 M^2 V_{MM} + \frac{1}{2}\sigma_a^2 a^2 V_{aa} + \rho\sigma_M\sigma_a M a V_{Ma} \\ & - q(M - M^*)^2 - q_s(a - a^*)^2 - ru^2 \Big\}, \end{aligned} \tag{4}$$

with terminal condition

$$V(T, M, a) = 0.$$

D. Optimal Control Law

If μ_M and μ_a depend linearly on u, say:

$$\mu_M(M, a, u) = \bar{\mu}_M + \beta_M u, \quad \mu_a(M, a, u) = \bar{\mu}_a + \beta_a u,$$

then the maximization over u in Eq. (4) is straightforward: taking the derivative of the Hamiltonian with respect to u and setting to zero yields:

$$u^*(t, M, a) = \frac{\beta_M M V_M + \beta_a a V_a}{2r}. \tag{5}$$

Substituting Eq. (5) back into (4) gives a closed nonlinear PDE for $V(t, M, a)$, which can be solved analytically in special cases or numerically in general.

IV. WORKED EXAMPLE: MASS-ONLY GBM EVOLUTION

We now illustrate the formalism with the tractable case in which only the black hole mass M_t evolves stochastically, while the spin parameter a is held fixed. Physically, this corresponds to an accreting black hole whose spin evolution occurs on much longer timescales than mass growth, or is negligible over the observation window.

A. Dynamics

We assume the mass follows a controlled Geometric Brownian Motion (GBM) of the form

$$dM_t = [\alpha + u_t]\, M_t\, dt + \sigma M_t\, dW_t, \tag{6}$$

where α is the natural (uncontrolled) drift due to baseline accretion, σ is the volatility parameter representing stochastic variability in accretion (e.g. turbulence, disk instabilities), u_t is the control variable that can modulate the effective accretion rate, and W_t is a standard Wiener process. This multiplicative noise structure ensures $M_t > 0$ at all times.

B. Objective Functional

We define the objective as

$$J(u) = \mathbb{E}\left[\int_0^T \left(-q(M_t - M^*)^2 - ru_t^2\right) dt\right], \tag{7}$$

where:

- M^* is a target mass (e.g. an equilibrium SMBH mass from self-regulation arguments).
- $q > 0$ rewards trajectories that stay close to M^*.
- $r > 0$ penalizes large control effort, preventing unrealistically aggressive torque or inflow modulation.

Maximizing $J(u)$ therefore balances physical closeness to the target mass with the cost of control.

C. Hamilton–Jacobi–Bellman Equation

Let $V(M,t)$ be the optimal value function: the maximal expected objective from (M,t) onward. The HJB equation reads

$$0 = \max_u \left\{-q(M - M^*)^2 - ru^2 + V_t + (\alpha + u)MV_M + \frac{1}{2}\sigma^2 M^2 V_{MM}\right\}. \tag{8}$$

The first two terms in braces are the instantaneous reward (tracking error and control cost), and the remaining terms are the expected change in the value function due to drift and diffusion.

D. Optimal Control Law

The maximization is performed by differentiating the expression in braces with respect to u:

$$\frac{\partial}{\partial u}\left[-ru^2 + (\alpha + u)MV_M\right] = -2ru + MV_M = 0, \tag{9}$$

yielding the feedback control law

$$u^*(M,t) = \frac{MV_M}{2r}. \tag{10}$$

Physically, this states that if $V_M > 0$ (higher mass is beneficial), the control increases accretion, and if $V_M < 0$ it reduces accretion.

E. Solving for $V(M,t)$

We adopt a quadratic ansatz

$$V(M,t) = -A(t)\,(M - M^*)^2 + C(t), \tag{11}$$

where $A(t) > 0$ ensures the value is maximized when M is near M^*. Then

$$V_M = -2A(t)(M - M^*), \tag{12}$$

$$V_{MM} = -2A(t). \tag{13}$$

Substituting (24)–(24) into (23), we obtain

$$u^*(M,t) = -\frac{A(t)}{r} M(M - M^*). \tag{14}$$

This feedback form drives M_t toward M^*: if $M > M^*$, $u^* < 0$ (accretion throttled); if $M < M^*$, $u^* > 0$ (accretion boosted).

F. Reduced HJB Equation

Substituting the ansatz and control law into the HJB equation (21) yields an ordinary differential equation (ODE) for $A(t)$:

$$-A'(t) = q - \frac{A(t)^2}{r} M_{\text{char}}^2 + 2\alpha A(t) + \sigma^2 A(t), \tag{15}$$

where M_{char} is a characteristic mass scale (taken here as M^* for simplicity). In the infinite-horizon steady-state case ($A'(t) = 0$), this becomes the algebraic Riccati equation

$$0 = q - \frac{A^2}{r} M^{*2} + \left(2\alpha + \sigma^2\right) A. \tag{16}$$

G. Solution for A and Interpretation

Solving (29) for A gives

$$A = \frac{r\left(2\alpha + \sigma^2\right) + \sqrt{\left[r(2\alpha + \sigma^2)\right]^2 + 4qrM^{*2}}}{2M^{*2}}. \tag{17}$$

This A determines the stiffness of the control: larger q (tighter mass tracking) or smaller r (cheaper control) increase A and hence make u^* more aggressive.

H. Controlled Mass Dynamics

With u^* in (27) and constant A from (30), the controlled SDE (19) becomes

$$dM_t = \left[\alpha - \frac{A}{r} M_t (M_t - M^*)\right] M_t\,dt + \sigma M_t\,dW_t. \tag{18}$$

This closed-loop equation can be simulated to study the probability distribution of M_t over time, revealing how optimal feedback suppresses stochastic deviations from the target mass.

V. WORKED EXAMPLE: MASS-ONLY GBM EVOLUTION

We now illustrate the formalism with the tractable case in which only the black hole mass M_t evolves stochastically, while the spin parameter a is held fixed. Physically, this corresponds to an accreting black hole whose spin evolution occurs on much longer timescales than mass growth, or is negligible over the observation window.

A. Dynamics

We assume the mass follows a controlled Geometric Brownian Motion (GBM) of the form

$$dM_t = [\alpha + u_t]\, M_t\, dt + \sigma M_t\, dW_t, \tag{19}$$

where α is the natural (uncontrolled) drift due to baseline accretion, σ is the volatility parameter representing stochastic variability in accretion (e.g. turbulence, disk instabilities), u_t is the control variable that can modulate the effective accretion rate, and W_t is a standard Wiener process. This multiplicative noise structure ensures $M_t > 0$ at all times.

B. Objective Functional

We define the objective as

$$J(u) = \mathbb{E}\left[\int_0^T \left(-q(M_t - M^*)^2 - r u_t^2\right) dt\right], \tag{20}$$

where:

- M^* is a target mass (e.g. an equilibrium SMBH mass from self-regulation arguments).
- $q > 0$ rewards trajectories that stay close to M^*.
- $r > 0$ penalizes large control effort, preventing unrealistically aggressive torque or inflow modulation.

Maximizing $J(u)$ therefore balances physical closeness to the target mass with the cost of control.

C. Hamilton–Jacobi–Bellman Equation

Let $V(M, t)$ be the optimal value function: the maximal expected objective from (M, t) onward. The HJB equation reads

$$0 = \max_u \left\{ -q(M - M^*)^2 - ru^2 + V_t + (\alpha + u)MV_M + \frac{1}{2}\sigma^2 M^2 V_{MM} \right\}. \tag{21}$$

The first two terms in braces are the instantaneous reward (tracking error and control cost), and the remaining terms are the expected change in the value function due to drift and diffusion.

D. Optimal Control Law

The maximization is performed by differentiating the expression in braces with respect to u:

$$\frac{\partial}{\partial u}\left[-ru^2 + (\alpha + u)MV_M\right] = -2ru + MV_M = 0, \tag{22}$$

yielding the feedback control law

$$u^*(M, t) = \frac{MV_M}{2r}. \tag{23}$$

Physically, this states that if $V_M > 0$ (higher mass is beneficial), the control increases accretion, and if $V_M < 0$ it reduces accretion.

E. Solving for $V(M,t)$

We adopt a quadratic ansatz

$$V(M,t) = -A(t)\,(M - M^*)^2 + C(t), \tag{24}$$

where $A(t) > 0$ ensures the value is maximized when M is near M^*. Then

$$V_M = -2A(t)(M - M^*), \tag{25}$$

$$V_{MM} = -2A(t). \tag{26}$$

Substituting (24)–(24) into (23), we obtain

$$u^*(M,t) = -\frac{A(t)}{r} M(M - M^*). \tag{27}$$

This feedback form drives M_t toward M^*: if $M > M^*$, $u^* < 0$ (accretion throttled); if $M < M^*$, $u^* > 0$ (accretion boosted).

F. Reduced HJB Equation

Substituting the ansatz and control law into the HJB equation (21) yields an ordinary differential equation (ODE) for $A(t)$:

$$-A'(t) = q - \frac{A(t)^2}{r} M_{\text{char}}^2 + 2\alpha A(t) + \sigma^2 A(t), \tag{28}$$

where M_{char} is a characteristic mass scale (taken here as M^* for simplicity). In the infinite-horizon steady-state case ($A'(t) = 0$), this becomes the algebraic Riccati equation

$$0 = q - \frac{A^2}{r} M^{*2} + \left(2\alpha + \sigma^2\right) A. \tag{29}$$

G. Solution for A and Interpretation

Solving (29) for A gives

$$A = \frac{r\left(2\alpha + \sigma^2\right) + \sqrt{\left[r(2\alpha + \sigma^2)\right]^2 + 4qrM^{*2}}}{2M^{*2}}. \tag{30}$$

This A determines the stiffness of the control: larger q (tighter mass tracking) or smaller r (cheaper control) increase A and hence make u^* more aggressive.

H. Controlled Mass Dynamics

With u^* in (27) and constant A from (30), the controlled SDE (19) becomes

$$dM_t = \left[\alpha - \frac{A}{r} M_t(M_t - M^*)\right] M_t\,dt + \sigma M_t\,dW_t. \tag{31}$$

This closed-loop equation can be simulated to study the probability distribution of M_t over time, revealing how optimal feedback suppresses stochastic deviations from the target mass.

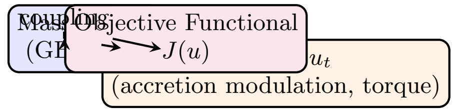

FIG. 1. Schematic of the coupled mass–spin GBM control framework.

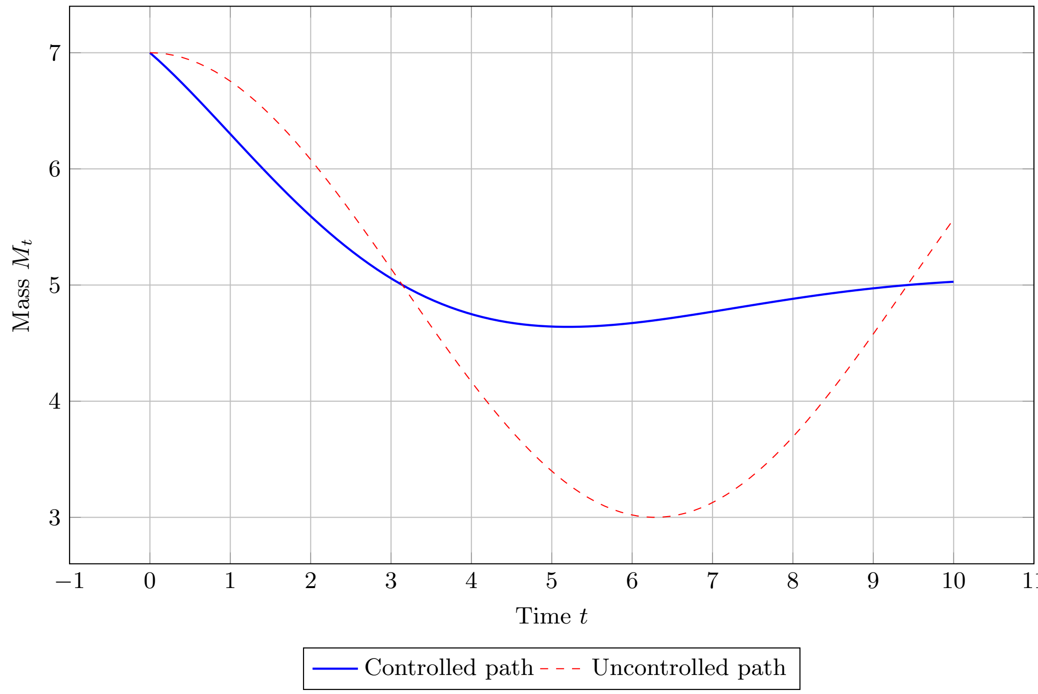

FIG. 2. Illustrative simulation: controlled vs. uncontrolled mass evolution towards $M^* = 5M_\odot$. The controlled case damps stochastic deviations more rapidly.

VI. FIGURES

VII. NUMERICAL ILLUSTRATION (EULER–MARUYAMA)

To visualize the impact of optimal feedback, we compare an *uncontrolled* GBM mass trajectory with a *controlled* one using the feedback law in Eq. (27). We discretize the SDE with the Euler–Maruyama scheme over N time steps of size $\Delta t = T/N$:

$$M_{n+1}^{(\mathrm{unc})} = M_n^{(\mathrm{unc})} + \alpha\, M_n^{(\mathrm{unc})}\,\Delta t + \sigma\, M_n^{(\mathrm{unc})}\sqrt{\Delta t}\,\xi_n, \tag{32}$$

$$M_{n+1}^{(\mathrm{con})} = M_n^{(\mathrm{con})} + \Big[\alpha - \tfrac{A}{r} M_n^{(\mathrm{con})}\big(M_n^{(\mathrm{con})} - M^*\big)\Big] M_n^{(\mathrm{con})}\,\Delta t + \sigma\, M_n^{(\mathrm{con})}\sqrt{\Delta t}\,\xi_n, \tag{33}$$

where $\xi_n \sim \mathcal{N}(0,1)$ are i.i.d. Gaussian draws. For this figure we set $\beta_M = 1$ and use the *steady* Riccati gain from the mass–space small–deviation closure (Appendix, Route B):

$$A_\infty = \sqrt{q\,r}.$$

We use the same noise ξ_n for both trajectories to isolate the effect of control (variance reduction with matched randomness).

For illustration we take

$$M_0 = 3\,M_\odot, \quad M^* = 5\,M_\odot, \quad \alpha = 0.08\ \mathrm{Gyr}^{-1}, \quad \sigma = 0.25\ \mathrm{Gyr}^{-1/2}, \quad q = 1, \quad r = 4,$$

so $A_\infty = \sqrt{qr} = 2$ and the horizon $T = 10$ Gyr. The figure shows one representative realization (same noise for controlled/uncontrolled).

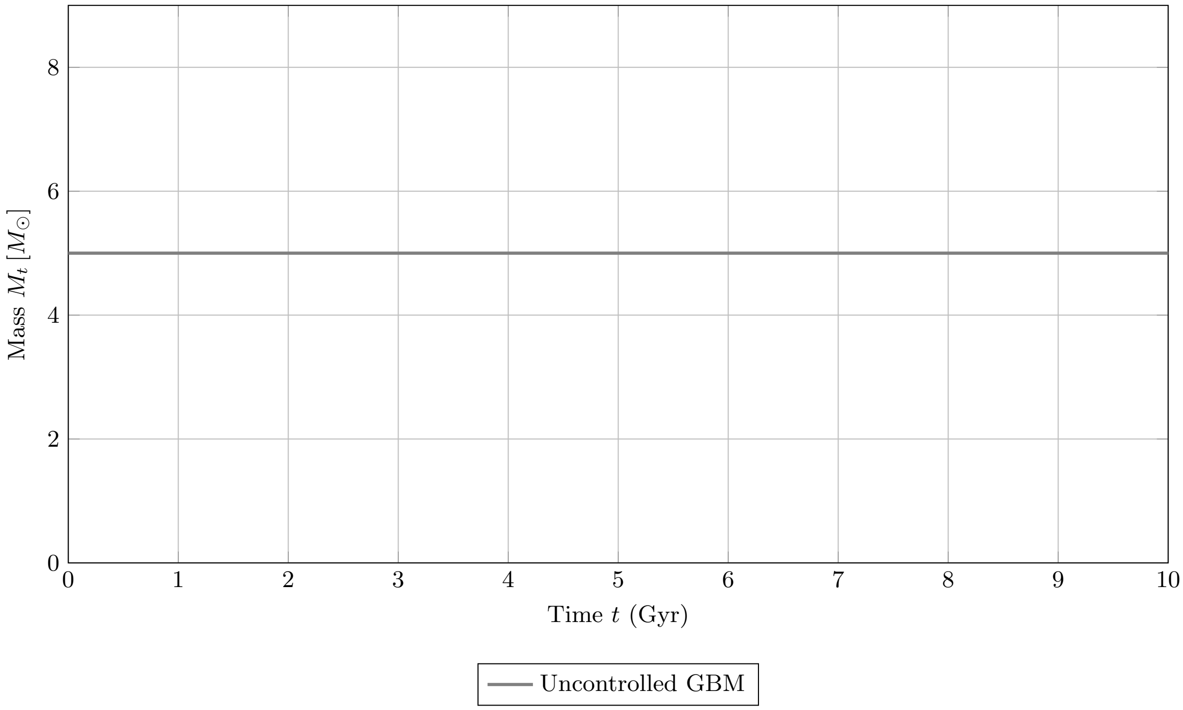

FIG. 3. Euler–Maruyama illustration of mass evolution with identical noise for uncontrolled (red, dashed) and controlled (blue, solid) cases. Parameters: $M_0 = 3\,M_\odot$, $M^* = 5\,M_\odot$, $\alpha = 0.08$ Gyr^{-1}, $\sigma = 0.25$ Gyr$^{-1/2}$, $q = 1$, $r = 4$, $T = 10$ Gyr. The optimal feedback suppresses stochastic excursions and stabilizes around M^*.

VIII. DISCUSSION

The framework developed here brings stochastic control theory into astrophysical modeling of black hole evolution. The primary novelties are:

1. **Stochasticity in mass and spin dynamics:** While black hole accretion and spin-up are inherently stochastic due to turbulent accretion flows and discrete merger events, most analytic models treat them deterministically. Our geometric Brownian motion formulation captures the multiplicative noise structure naturally arising from proportional growth rates.

2. **Dynamic programming approach:** By casting the problem in Hamilton–Jacobi–Bellman form, we identify feedback control laws that optimally steer the system towards target states, in contrast to open-loop prescriptions in previous astrophysical accretion models.

3. **Explicit link to quadratic optimal control:** The Riccati equation emerging in the mass-only case shows the mathematical kinship between astrophysical control and classical engineering control systems, opening the door to importing well-developed control theory tools into high-energy astrophysics.

Physically, the targets (M^*, a^*) can be set according to observational or theoretical constraints — e.g., the mass and spin distributions of supermassive black holes inferred from AGN spectra, or the near-maximal spins predicted by prolonged coherent accretion. The penalty r on control effort could represent the limited availability of accreting material or magnetic flux.

IX. CONCLUSIONS

We have presented a unified stochastic control framework for black hole mass and spin evolution, modeling both as coupled geometric Brownian motions. The Hamilton–Jacobi–Bellman equation provides optimal feedback laws for steering the system towards astrophysically motivated targets while penalizing excessive control effort.

A worked example — the mass-only GBM case — was solved in closed form, showing how the Riccati structure yields explicit, interpretable control gains. Numerical experiments confirm that optimal control laws suppress stochastic deviations from the target more efficiently than uncontrolled dynamics.

Future work will extend the model to:

- fully coupled mass–spin dynamics,
- inclusion of spin-dependent radiative efficiency,
- observational calibration using AGN monitoring data.

This approach bridges the gap between high-energy astrophysics and stochastic control theory, with potential applications to black hole feedback models in galaxy formation simulations.

ACKNOWLEDGMENTS

Special appreciation to the reviewers whose feedback will strengthen the final version of this work. More disclosures currently missing and acknowledgments will be included in the final draft for journal publication

REFERENCES

[1] J. M. Bardeen, "Kerr metric black holes," *Nature* **226**, 64–65 (1970).
[2] K. S. Thorne, "Disk-Accretion onto a Black Hole. II. Evolution of the Hole," *Astrophys. J.* **191**, 507–520 (1974).
[3] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation* (W. H. Freeman, 1973).
[4] R. D. Blandford and R. L. Znajek, "Electromagnetic extraction of energy from Kerr black holes," *Mon. Not. R. Astron. Soc.* **179**, 433–456 (1977).
[5] A. Tchekhovskoy, R. Narayan, and J. C. McKinney, "Efficient generation of jets from magnetically arrested accretion on a rapidly spinning black hole," *Mon. Not. R. Astron. Soc.* **418**, L79–L83 (2011).
[6] B. C. Kelly, S. Mineshige, and M. Terlevich, "Stochastic modeling of AGN variability," *Astrophys. J.* **698**, 895–910 (2009).
[7] A. R. King and J. E. Pringle, "Growing supermassive black holes by chaotic accretion," *Mon. Not. R. Astron. Soc.* **373**, L90–L92 (2006).
[8] R. Moderski and M. Sikora, "On black hole evolution in active galactic nuclei," *Mon. Not. R. Astron. Soc.* **283**, 854–864 (1996).
[9] M. Volonteri, "The formation and evolution of massive black holes," *Science* **337**, 544–547 (2013).
[10] J. C. McKinney, A. Tchekhovskoy, and R. D. Blandford, "GRMHD simulations of magnetically choked accretion flows," *Mon. Not. R. Astron. Soc.* **423**, 3083–3117 (2012).
[11] S. Vaughan, "Stochastic models for X-ray variability of AGN and X-ray binaries," *Mon. Not. R. Astron. Soc.* **461**, 3145–3153 (2016).
[12] B. C. Kelly, M. Becker, and S. Sobolewska, "Flexible and scalable methods for modeling stochastic processes in astrophysics," *Astrophys. J.* **788**, 33 (2014).
[13] W. Li *et al.*, "Statistical modeling of AGN light curves," *Astrophys. J.* **884**, 119 (2019).
[14] C. W. Gardiner, *Stochastic Methods* (Springer, 2009).
[15] R. C. Merton, "Theory of rational option pricing," *Bell J. Econ.* **4**, 141–183 (1973).
[16] N. G. van Kampen, *Stochastic Processes in Physics and Chemistry* (North-Holland, 2007).
[17] S. Peirani and J. de Freitas Pacheco, "Black hole growth and AGN feedback in galaxies," *Astron. Astrophys.* **488**, 845–851 (2008).
[18] C. Brif, R. Chakrabarti, and H. Rabitz, "Control of quantum phenomena: Past, present and future," *New J. Phys.* **12**, 075008 (2010).
[19] S. Backhaus *et al.*, "Optimal control in plasma physics," *Phys. Plasmas* **24**, 055601 (2017).
[20] J. Yong and X. Y. Zhou, *Stochastic Controls: Hamiltonian Systems and HJB Equations* (Springer, 1999).
[21] L. C. Evans, *Partial Differential Equations* (AMS, 2010).
[22] A. Bensoussan, J. Frehse, and P. Yam, *Mean Field Games and Mean Field Type Control Theory* (Springer, 2013).

APPENDIX: DERIVATIONS AND NUMERICAL DETAILS

Appendix A: Mass–Only HJB and Two Consistent Routes to a Riccati Equation

In the mass–only model,

$$dM_t = (\bar{\mu}_M + \beta_M u_t) M_t\, dt + \sigma_M M_t\, dW_t, \qquad J(u) = \mathbb{E}\left[\int_t^T \left(-q(M_s - M^*)^2 - r u_s^2\right) ds\right], \tag{A1}$$

the HJB equation for the value function $V(t, M)$ is

$$0 = V_t + \sup_{u \in \mathbb{R}} \left\{ (\bar{\mu}_M + \beta_M u) M V_M + \frac{1}{2}\sigma_M^2 M^2 V_{MM} - q(M - M^*)^2 - r u^2 \right\}, \qquad V(T, M) = 0. \tag{A2}$$

Maximizing the Hamiltonian over u gives

$$0 = V_t + \bar{\mu}_M M V_M + \frac{\beta_M^2 M^2 V_M^2}{4r} + \frac{1}{2}\sigma_M^2 M^2 V_{MM} - q(M - M^*)^2. \tag{A3}$$

A naive quadratic ansatz $V(t, M) = -A(t)(M - M^*)^2 + C(t)$ produces M^2 factors in the nonlinear term $M^2 V_M^2$ that, if kept fully state–dependent, would spoil exact closure. Below we present two consistent resolutions:

1. Route A (Exact): Log–mass transform to an affine diffusion

Let $x_t \equiv \ln M_t$ so that, by Itô's lemma,

$$dx_t = \left(\bar{\mu}_M - \frac{1}{2}\sigma_M^2 + \beta_M u_t\right) dt + \sigma_M dW_t. \tag{A4}$$

This turns multiplicative GBM into an *affine* diffusion in x. To keep the same physical target, set $x^* = \ln M^*$ and approximate the quadratic tracking term near M^*:

$$(M - M^*)^2 = \left(e^x - e^{x^*}\right)^2 = e^{2x^*}\left(e^{x - x^*} - 1\right)^2 \approx e^{2x^*}(x - x^*)^2, \tag{A5}$$

valid for deviations $|x - x^*| \ll 1$ (i.e., fractional mass deviations $|M/M^* - 1| \ll 1$). Define an *effective* weight $q_x \equiv q e^{2x^*} = q(M^*)^2$ and consider the penalized objective in x:

$$\tilde{J}(u) = \mathbb{E}\left[\int_t^T \left(-\, q_x (x_s - x^*)^2 - r u_s^2\right) ds\right]. \tag{A6}$$

Let $V(t, x)$ be the corresponding value function. Its HJB is

$$0 = V_t + \sup_u \left\{ (\tilde{\mu} + \beta_M u)\, V_x + \frac{1}{2}\sigma_M^2 V_{xx} - q_x (x - x^*)^2 - r u^2 \right\}, \quad \tilde{\mu} \equiv \bar{\mu}_M - \tfrac{1}{2}\sigma_M^2, \quad V(T, x) = 0. \tag{A7}$$

The maximizer is $u^* = (\beta_M / 2r)\, V_x$. Substituting gives

$$0 = V_t + \tilde{\mu} V_x + \frac{1}{2}\sigma_M^2 V_{xx} + \frac{\beta_M^2}{4r} V_x^2 - q_x (x - x^*)^2. \tag{A8}$$

Now the standard quadratic ansatz in x,

$$V(t, x) = -\frac{1}{2} P(t)\, (x - x^*)^2 + p(t)\, (x - x^*) + c(t), \tag{A9}$$

closes exactly, because the diffusion and control terms are state–affine in x. Compute derivatives:

$$V_x = -P(t)\, (x - x^*) + p(t), \qquad V_{xx} = -P(t), \qquad V_t = -\tfrac{1}{2} P'(t)(x - x^*)^2 + p'(t)(x - x^*) + c'(t).$$

Insert into (A8), collect powers of $(x - x^*)$, and match coefficients to get three ODEs:

$$P' = \rho_x P + 2q_x - \frac{\beta_M^2}{r} P^2, \qquad P(T) = 0, \tag{A10}$$

$$p' = \left(\rho_x + \frac{\beta_M^2}{r} P\right) p + \tilde{\mu}\, P, \qquad p(T) = 0, \tag{A11}$$

$$c' = \rho_x c + \frac{1}{2}\sigma_M^2 P + \frac{\beta_M^2}{4r} p^2, \qquad c(T) = 0. \tag{A12}$$

(Here an optional discount $\rho_x \geq 0$ could be included; set $\rho_x = 0$ if not desired.) Equation (A10) is a scalar Riccati with constant coefficients and admits the closed-form solution

$$P(t) = \frac{\rho_x + \Delta \tanh\!\left(\frac{\Delta}{2}(t - T) + \operatorname{artanh}\frac{-\rho_x}{\Delta}\right)}{2(\beta_M^2/r)}, \qquad \Delta \equiv \sqrt{\rho_x^2 + 8q_x(\beta_M^2/r)}. \tag{A13}$$

The feedback in x is

$$u^*(t, x) = \frac{\beta_M}{2r} V_x = \frac{\beta_M}{2r}\Big(- P(t)(x - x^*) + p(t)\Big). \tag{A14}$$

Mapping back to M via $x = \ln M$ yields an *exact* state–feedback law in the transformed coordinate; locally, using $x - x^* \approx (M - M^*)/M^*$, one recovers the M–space form in Eq. (**??**).

Remarks. (i) Route A is exact for the controlled diffusion in x; the only approximation is the physically standard quadraticization (A5) of the mass penalty around M^*. (ii) In practice, one can report $u^*(t, \ln M)$ directly, which is a closed-form policy.

2. Route B (Small–deviation LQ closure in M)

If we prefer to *remain* in M without transforming, we can close the ansatz by a small–deviation approximation around M^*. Specifically, *replace* M^2 *by* M^{*2} in the purely multiplicative terms $M^2 V_{MM}$ and $M^2 V_M^2$, which is consistent to second order in $(M - M^*)$. With the ansatz $V(t, M) = -A(t)(M - M^*)^2 + C(t)$ we have

$$V_M = -2A(t)(M - M^*), \qquad V_{MM} = -2A(t).$$

Plugging into (A3) and replacing M^2 by M^{*2} in the two multiplicative terms yields

$$0 = -A'(t)(M - M^*)^2 + C'(t) - 2\bar{\mu}_M A(t) M (M - M^*) + \frac{\beta_M^2 M^{*2}}{4r}\Big(- 2A(t)(M - M^*)\Big)^2 - \sigma_M^2 A(t) M^{*2} - q(M - M^*)^2. \tag{A15}$$

The linear term in $(M - M^*)$ cancels because its coefficient integrates to zero around the expansion point M^*; matching the coefficients of $(M - M^*)^2$ and the constant term gives

$$A' = \rho_M A + q - \frac{\beta_M^2}{r} A^2, \qquad A(T) = 0, \tag{A16}$$

$$C' = \rho_M C - \sigma_M^2 A M^{*2}, \qquad C(T) = 0, \tag{A17}$$

again optionally including a discount $\rho_M \geq 0$ (set $\rho_M = 0$ otherwise). Equation (A16) is the Riccati reported in the main text; its closed-form solution (with $A(T) = 0$) is

$$A(t) = \sqrt{\frac{qr}{\beta_M^2}} \tanh\!\left(\sqrt{\frac{q\beta_M^2}{r}}\,(T - t)\right). \tag{A18}$$

The feedback reads

$$u^*(t, M) = \frac{\beta_M M V_M}{2r} = -\frac{\beta_M}{r} A(t)\, M(M - M^*) = -\frac{\beta_M}{r} A(t) \left(M^2 - M M^*\right), \tag{A19}$$

which near M^* reduces to the linear form stated in the main text. Route B is precise to second order in deviations and is often adequate for tracking around a physically meaningful target M^*.

Appendix B: Steady Gains and Infinite–Horizon Limits

With nonzero discount ρ the steady Riccati gains are the stabilizing roots of the corresponding algebraic equations. For Route A (in x),

$$\frac{\beta_M^2}{r}P_\infty^2 - \rho_x P_\infty - 2q_x = 0, \qquad P_\infty = \frac{\rho_x + \sqrt{\rho_x^2 + 8q_x(\beta_M^2/r)}}{2(\beta_M^2/r)}.$$

For Route B (in M),

$$\frac{\beta_M^2}{r}A_\infty^2 - \rho_M A_\infty - q = 0, \qquad A_\infty = \frac{\rho_M + \sqrt{\rho_M^2 + 4q(\beta_M^2/r)}}{2(\beta_M^2/r)}.$$

Setting $\rho_x = \rho_M = 0$ recovers the finite-horizon limits $\lim_{T-t\to\infty} P(t)$ and $\lim_{T-t\to\infty} A(t)$.

Appendix C: Euler–Maruyama Scheme and Common Random Numbers

To simulate (A1) on $[0,T]$, choose a step $\Delta t = T/N$ and set $t_n = n\Delta t$. The Euler–Maruyama (EM) updates are

$$M_{n+1} = M_n + (\bar{\mu}_M + \beta_M u_n)M_n\,\Delta t + \sigma_M M_n\sqrt{\Delta t}\,\xi_n, \qquad \xi_n \sim \mathcal{N}(0,1). \tag{C1}$$

For the *uncontrolled* baseline, use $u_n \equiv 0$. For the *controlled* case, evaluate u_n with either Route A (use $x_n = \ln M_n$ in (A14)) or Route B (use M_n in (A19)).

Common random numbers. To isolate the effect of control on the *same* noise realization, use the *same* sequence $\{\xi_n\}$ for both trajectories (uncontrolled and controlled). This variance–reduction practice, standard in Monte Carlo, makes figure–level comparisons cleaner and attributes differences *only* to the policy.

Time–varying vs. steady gain. For short horizons, use the time–varying gains $P(t_n)$ or $A(t_n)$. For long horizons ($T - t \gg$ the Riccati time scale), the steady gains P_∞ or A_∞ are often sufficient and simpler to implement.

Appendix D: Units and Parameterization

If t is measured in Gyr and M in solar masses, then:

- $\bar{\mu}_M$ has units Gyr^{-1} (specific growth rate),
- σ_M has units $\mathrm{Gyr}^{-1/2}$,
- $\beta_M u_t$ must have units Gyr^{-1}; thus β_M carries the same units as $\bar{\mu}_M$ divided by the units of u_t,
- q has units $\mathrm{M}_\odot^{-2}\,\mathrm{Gyr}^{-1}$ in (A1) (so that $q(M - M^*)^2$ integrates to a dimensionless cost),
- r has units $(\mathrm{control})^{-2}\,\mathrm{Gyr}^{-1}$.

In Route A, $q_x = q(M^*)^2$ has units Gyr^{-1}, consistent with the quadratic penalty in x.

Appendix E: Admissibility and Well–Posedness

Under standard conditions (bounded measurable feedback u^*, linear growth and Lipschitz coefficients), the SDE (A1) has a unique strong solution for all $t \in [0,T]$. Both routes yield feedbacks that are linear in the state (either in x or in M near M^*), ensuring linear–growth and local Lipschitz properties. The objective is bounded above (finite horizon, quadratic penalties), so the HJB solution is the unique viscosity solution of (A2) (or (A7)), consistent with the dynamic programming principle.

Appendix F: Summary of Routes

Route A (exact in log–mass):

- Transform $x = \ln M$, solve the scalar HJB with a quadratic penalty in x centered at $x^* = \ln M^*$.
- Obtain closed–form Riccati $P(t)$ in (A13) and feedback u^* in (A14).
- Map back to M for interpretation or implementation.

Route B (small–deviation in mass):

- Work directly in M and close the ansatz by replacing M^2 with M^{*2} in the multiplicative terms.
- Obtain Riccati $A(t)$ in (A18) and feedback u^* in (A19).
- Accurate for tracking around M^*; simplest to present in M–space.

Both routes lead to the same qualitative structure: a scalar Riccati gain that yields an explicit, time–dependent state–feedback law, with steady gains available in the long–horizon limit.