

Optimal control of a mesoscopic information engine

Emanuele Panizon *Area Science Park, Località Padriciano 99, 34149 Trieste, Italy*

(Received 10 April 2026; accepted 7 August 2026; published 28 September 2026)

We analytically solve the finite-time control problem of driving an overdamped particle via an optical trap under costly measurement. By formulating this mesoscopic information engine within the partially observable Markov decision process framework, we demonstrate that the underlying linear-quadratic-Gaussian dynamics decouples the optimal measurement and driving protocols. We derive the optimal feedback control law for the trap placement, which recovers the discontinuous Schmiedl-Seifert driving protocol in the open-loop limit and extends it to any measurement scheduling. For a costly binary (on-off) sensor, we evaluate the optimal measurement protocol and derive physical bounds on the maximum net gain that can be extracted from thermal fluctuations. We show the emergence of deadline-induced blindness, a phenomenon where all measurements cease as the deadline approaches regardless of their cost. Taking the infinite-horizon limit, we find the exact periodic measurement schedules for the steady state as a function of the measurement cost C and derive the macroscopic velocity envelopes beyond which viscous drag forces the engine into a net-dissipative regime. Finally, we generalize the results to a variable-precision sensor.

DOI: [10.1103/85qf-nwj5](https://doi.org/10.1103/85qf-nwj5)

I. INTRODUCTION

Thermal fluctuations dominate mesoscopic dynamics. Stochastic thermodynamics and fluctuation theorems [1–3] formalized the physical relations governing nonequilibrium trajectories, expanding the second law from macroscopic averages to individual microscopic realizations.

When extended to systems with feedback, this framework established how information acquisition can be used to extract work from fluctuations [4–7]. This concept, originating with Maxwell’s demon [8,9] and bounded by the thermodynamics of information [10,11], has recently been realized in experimental laboratories. Specifically, colloidal particles manipulated by harmonic optical traps serve as prototypical information engines, where measurement-driven feedback loops have been successfully used to convert spatial fluctuations into directed motion or stored energy [12–14]. Heuristic continuous-feedback ratchets, for example, perform well empirically in both thermal [15,16] and nonequilibrium baths [17].

However, how to construct optimal feedback loops is still an open problem, even when measurement costs are not considered explicitly. Exact analytical solutions are sparse, as they typically require solving complex continuous Hamilton-Jacobi-Bellman integrodifferential equations [18,19]. Recent advances in optimal transport theory for fully observable continuous Langevin systems [20,21] and generalized geometric frameworks [22] have established discontinuous end-point

jumps [23] as a generic physical mechanism to minimize driving dissipation.

For the specific case of finite-time operation of optical traps, analytical solutions focus on the optimization of the physical control law conditional on a fixed measurement architecture. Current models have described analytically the optimal trap protocol in the Schmiedl-Seifert regime (no feedback) [24], in the Abreu-Seifert limit (a single initial observation) [25], under periodic measurement loops [26,27], and (approximately) for continuous perfect measurement [28].

Since the literature optimizes the trap trajectory in isolation, rather than treating physical control and measurement scheduling as a single conjoint problem, the exact finite-time dynamics of costly measurements remains unsolved.

In this work we formulate the problem of an overdamped particle in a harmonic trap in discrete time in the framework of a partially observed Markov decision process [29,30]. As such, we exploit linear-quadratic-Gaussian (LQG) control theory [31–33] to bypass the complex integrodifferential equations of standard continuous-time partial observability. We show that the structure of thermodynamic work in a harmonic potential allows the measurement policy and the physical control law to decouple into simple algebraic equations. This decoupling yields a closed-form mapping of the demon’s complete thermodynamic phase space.

We derive the analytical control law λ_k^* for the optimal trap trajectory, demonstrating how discrete finite-time feedback recovers and extends the continuous limits $\Delta t \rightarrow 0$ previously found in [25]. By studying the extraction power limits of this information engine under two formulations of measurement costs, we map the thermodynamic boundaries of the system. Beyond general cost thresholds, beyond which observations can never be profitable, we show the existence of so-called

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

deadline blindness for finite-time operations. This condition is triggered when the threshold for profitable information, which diverges close to the terminal point, crosses the thermal equilibrium variance. We further establish the optimal periodic measurement schedule at the steady state for binary, i.e., on-off, sensors and derive the macroscopic velocity boundaries that bankrupt the information engine via viscous drag. Finally, we define the optimal precision acquisition rate for variable-precision sensors that transforms the discrete demon into a continuous so-called information thermostat.

II. MODEL AND POMDP FORMULATION

We consider an overdamped Brownian particle in a one-dimensional harmonic trap $U(x, \lambda) = \frac{1}{2}\kappa(x - \lambda)^2$ with stiffness κ , immersed in a heat bath at temperature T and with friction coefficient γ . The objective is to steer the trap center from an initial position $\lambda_0 = 0$ to a target λ_f over a finite-time horizon t_f , minimizing the expected thermodynamic work. Time t is discretized into discrete steps $k\Delta t$, $k \in \{0, 1, \dots, N\}$. We cast this as a partially observable Markov decision process (POMDP) [34,35] where the hidden system state is the particle's true position x_k . Because the agent cannot observe x_k for free, it tracks a belief state consisting of a Gaussian distribution, parametrized by its mean μ_k and variance Σ_k , augmented by the historical trap position λ_{k-1} .

At each time step, the temporal sequence of actions unfolds as follows. First, the agent decides how to actively measure its state, possibly paying a cost $C_k = C(\Sigma_k^+, \Sigma_k^-)$ to update its belief from a premeasurement state (μ_k^-, Σ_k^-) into (μ_k^+, Σ_k^+) . We consider this measurement cost to be purely parametric, without explicitly modeling the underlying energetics. Without relying on external entropy reservoirs [36,37] such as additional thermal baths [38], this parametric approach admits regimes that would violate the generalized second law of thermodynamics, which forbids true net work extraction from a single isothermal bath. The optimization methodology presented here, however, can accommodate for any physical measurement-cost function and is therefore general. Two different possibilities for the sensors will be discussed, both a binary, on-off, perfect sensor and a variable-sensitivity one. After the measurement, and exploiting the updated belief, the agent executes a physical control action by setting the new trap position λ_k . The step cost associated with this physical action is the expected thermodynamic work done on the system during the instantaneous trap shift:

$$\mathbb{E}[W_k] = -\kappa(\lambda_k - \lambda_{k-1})\left(\mu_k^+ - \frac{\lambda_k + \lambda_{k-1}}{2}\right). \quad (1)$$

Following standard thermodynamic convention, $\mathbb{E}[W_k] < 0$ denotes net work extracted from the thermal bath. The quadratic state terms cancel out, making the expected physical work independent of the posterior variance Σ_k^+ . Finally, the physical system evolves and the belief is modified following the Ornstein-Uhlenbeck updates [39,40]

$$\mu_{k+1} = \mu_k^+ \alpha + \lambda_k(1 - \alpha), \quad (2)$$

$$\Sigma_{k+1} = \Sigma_k^+ \alpha^2 + \Sigma_\Delta, \quad (3)$$

where $\alpha = \exp(-\kappa\Delta t/\gamma)$ and the equilibrium variance is $\Sigma_\Delta = (k_B T/\kappa)(1 - \alpha^2) = \Sigma_{\text{eq}}(1 - \alpha^2)$.

To study the continuous limit, we will consider the limit $\Delta t \rightarrow \infty$. We will approach the steady regime by having a finite macroscopic velocity $v = \lambda_f/t_f$ with the limits $t_f, \lambda_f \rightarrow \infty$.

III. BELLMAN OPTIMALITY EQUATION

The objective of the agent is to determine the feedback policy that minimizes the total expected cost over the finite horizon $\mathcal{J}$. The total cost accumulates the thermodynamic work done on the system (W_k) and the energy spent to measure (C_k) at all time steps k :

$$\mathcal{J} = \mathbb{E}\left[\sum_{k=0}^N (W_k + C_k)\right]. \quad (4)$$

We can formalize this minimization objective through the optimal cost function $V_n(\mu^-, \Sigma^-, \lambda)$, which evaluates the minimum (optimal) expected cost to go for the remaining steps $n = N - k$, conditioned on the current prior spatial estimate μ^- , the prior uncertainty Σ^- , and the trap position λ .

The optimal joint protocol of measurement and physical trap displacement satisfies the recursive Bellman optimality equation [41,42]

$$\begin{aligned} V_n(\mu^-, \Sigma^-, \lambda) = \min_{\Sigma^+} \{ & C(\Sigma^+, \Sigma^-) \\ & + \mathbb{E}_{\mu^+}[\min_{\lambda'} \{\mathbb{E}_x[W] + V_{n-1}(\mu'^-, \Sigma'^-, \lambda')\}] \}, \end{aligned} \quad (5)$$

where (μ'^-, Σ'^-) denote the prior belief state at the subsequent time step.

Since the continuous dynamics are governed by linear Langevin equations and the thermodynamic work is quadratic in the spatial coordinates, the system operates in the LQG regime [33,43,44]. The separation principle therefore guarantees that the informational and spatial controls can be separated. As shown in the Appendix A, the value function V_n decouples into two independent components concerning either the spatial or the informational degrees of freedom:

$$V_n(\mu^-, \Sigma^-, \lambda) = J_n(\mu^-, \lambda) + g_n(\Sigma^-). \quad (6)$$

Here J_n is the deterministic physical cost depending only on the spatial coordinates and g_n is the informational cost depending only on the estimation variance. This allows us to separately tackle the optimal trap control and the optimal measurement scheduling.

IV. RICCATI RECURRENCE

In standard discrete-time LQG control, finding the optimal physical policy reduces to propagating a quadratic value function backward in time (hence the use of the index n). This requires solving a discrete-time recurrence [32,42] for a cost matrix $\mathbf{P}_n$, called the Riccati equation. Given the augmented state space (μ, λ) , this is a 2×2 matrix. However, the thermodynamic work in a one-dimensional harmonic potential separates into terms dependent on either the current action λ_k or the previous trap position λ_{k-1} . The absence of cross terms

means that the optimal feedback law λ_k^* is independent of the previous trap placement.

This decoupling, as shown in the Appendix A, reduces the matrix Riccati recurrence to a single scalar sequence for the spatial cost coefficient P_n , where $n = N - k$ is the number of remaining steps. This scalar dictates the leading energetic cost $\frac{1}{2}P_n(\mu^+ - \lambda_f)^2$ for deviations from the final target. The optimal feedback law is computed by evaluating the stationary point of the value function with respect to the active control, $\partial/\partial\lambda(\cdot) = 0$. Exploiting the quadratic nature of the value function, we write the optimal control law of LQG control as a linear interpolation between the present (expected) position μ_k^+ and the final target λ_f : $\lambda_k^* = K_n\mu_k^+ + (1 - K_n)\lambda_f$. Here K_n is the optimal feedback gain that is computed as $K_n = (\kappa - \alpha P_{n-1})/[P_{n-1}(1 - \alpha) + 2\kappa]$. Substituting the closed-form solution for P_n (derived in the Appendix B),

$$P_n = -\kappa \left(\frac{n(1 - \alpha)}{1 + \alpha + n(1 - \alpha)} \right), \quad (7)$$

returns the explicit optimal trap placement

$$\lambda_k^* = \mu_k^+ + \frac{1}{n(1 - \alpha) + 1 + \alpha}(\lambda_f - \mu_k^+). \quad (8)$$

Physically, as mentioned above, the optimal control linearly interpolates between the current belief of the particle position μ_k^+ and the final target λ_f . Far from the deadline ($n\Delta t \gg \gamma/\kappa$), the time-dependent weight approaches zero, and the trap closely follows the inferred position ($\lambda_k^* \approx \mu_k^+$) to optimally exploit thermal fluctuations. As the deadline approaches ($n \rightarrow 1$), the weighting shifts to drive the system toward the target λ_f .

V. CONTINUOUS-TIME LIMIT

To demonstrate that our finite-time discrete control reproduces and extends established continuous-time frameworks, we take the continuous limit $\Delta t \rightarrow 0$. Expanding the relaxation parameter to first order, $1 - \alpha \approx \Delta t/\tau$ (where $\tau = \gamma/\kappa$), we obtain the continuous-time limit for Eq. (8):

$$\lambda^*(t) = \mu(t) + \frac{\tau}{(t_f - t) + 2\tau}[\lambda_f - \mu(t)]. \quad (9)$$

It is important to recall that by the separation principle of LQG control, the optimal trap policy $\lambda^*(t)$ is independent of the measurement architecture; it is defined entirely by the instantaneous state estimate $\mu(t)$. As a consequence, this optimal law is always optimal: Different costs and definitions of the sensors will result in different optimal measuring schedules and different distributions of trajectories $\mu(t)$, but for each of them this same control will apply.

As an example, we can see that the control law at the start of the protocol ($t = 0$) recovers the known [24] initial trap displacement of $\Delta\lambda = [\tau/(t_f + 2\tau)]\Delta\mu$, where $\Delta\mu = \lambda_f - \mu(0)$. Assuming no further measurements are taken, the trajectory $\bar{\mu}(t)$ is given by the deterministic evolution $\tau\dot{\bar{\mu}}(t) = \lambda^*(t) - \bar{\mu}(t)$:

$$\bar{\mu}(t) = \bar{\mu}(0) + \left(\frac{\lambda_f - \bar{\mu}(0)}{t_f + 2\tau} \right)t. \quad (10)$$




FIG. 1. Optimal finite-time control dynamics of the information engine with binary measurements. (a) Simulated spatial trajectory illustrating the true particle position x (thin gray line), the inferred belief mean μ (blue dashed line) with 1σ uncertainty bounds (blue shaded area), and the optimal deterministic trap placement λ (red solid line). Black markers denote discrete measurement events. The red dot shows the final target λ_f . (b) Corresponding temporal evolution of the belief variance Σ (purple solid line). Thermal diffusion drives the uncertainty upward until it intersects the dynamically computed optimal threshold $\Sigma_{th}(n)$ (orange dashed line), triggering a measurement that collapses the variance. As the deadline approaches $t \rightarrow t_f$, the required threshold for profitable measurement rigorously diverges, demonstrating the onset of finite-time deadline blindness, here shown with grey shading when $\Sigma_{th}(n) > \Sigma_{eq}$. The simulation parameters are $k_B T = 1$, $\kappa = 1$, $\gamma = 1$, $\Delta t = 0.01$, $t_f = 20$, and $C = 0.6 C_{max}$.

Thus, as $\Delta t \rightarrow 0$, the discrete-time POMDP framework recovers the constant-velocity trajectory and the discontinuous jump sizes characteristic of the Schmiedl-Seifert open-loop driving protocol [24]. In the case of initial measurement $\bar{\mu}(0) = x_0$, the Abreu-Seifert protocol for fixed stiffness [25] is also recovered. For the case of periodic perfect measurements, the $\bar{\mu}(t)$ trajectory is a piecewise linear function with discontinuous jumps located at the measurement times, and the optimal control is a periodic reset of an open-loop control, consistent with [26,27] (see also Fig. 1). The general control law, as derived within the POMDP framework, shows that the optimal agent chooses at each step the same microscopic control it would select if no further measurement were expected. This is also qualitatively consistent with the numerical approximations found in [28].

VI. MEASUREMENT PROTOCOL AND FINITE-TIME TRIGGER

We can now tackle the problem of optimal measurement scheduling. The decision to measure reduces to the

independent optimization of the informational cost $g_n(\Sigma^-)$ shown in Eq. (6).

First, we consider an instantaneous binary measurement $a_{\text{obs}} \in \{0, 1\}$, incurring a macroscopic cost C . Here $a_{\text{obs}} = 1$ denotes a full observation, while for $a_{\text{obs}} = 0$ the sensor is shut. While a perfect measurement of a continuous variable in principle requires infinite thermodynamic work [11,45], C here represents the macroscopic energy cost of activating a sensor of high enough resolution that any error is negligible relative to the physical length scales of the system dynamics.

The optimal policy can be reduced (see Appendixes A and C) to a discrete minimization over the two branches of the observation space (on or off):

$$g_n(\Sigma^-) = \min_{a_{\text{obs}} \in \{0, 1\}} \{ (1 - a_{\text{obs}})g_{n-1}(\alpha^2 \Sigma^- + \Sigma_\Delta) + a_{\text{obs}}[C - A_n \Sigma^- + g_{n-1}(\Sigma_\Delta)] \}. \quad (11)$$

For clarity we have introduced here the term $A_n = -P_n/2$, which measures the thermodynamic gain that can be extracted from the particle variance. Rearranging the minimization isolates the decision boundary for measurements:

$$g_{n-1}(\alpha^2 \Sigma^- + \Sigma_\Delta) - g_{n-1}(\Sigma_\Delta) \geq C - A_n \Sigma^-. \quad (12)$$

The right-hand side represents the net immediate cost of the measurement step: the observation cost C reduced by the immediate profit $-A_n \Sigma^-$. The left-hand side represents the difference in future informational costs for the different measurement choices.

Since work can be extracted from reduction of uncertainty, the informational cost $g_n(\Sigma)$ is a decreasing function of Σ : A larger variance yields a better (more negative) cost (see the Appendix C). The choice of measurement results from a compromise: Should the demon pay a cost C to collapse the variance now and immediately harvest a gain $A_n \Sigma^-$ or allow the variance to grow for a larger potential gain in the future?

The optimal demon delays measurement, waiting until the expected thermal fluctuations are large enough to compensate for both the absolute observation cost C and the forfeited future thermodynamic yield (see Fig. 1). Equation (12) allows one to calculate backward, from the trivial final step $n = 0$, the exact threshold variance $\Sigma_{\text{th}}(n)$ above which an optimal agent will observe. Since the trajectory of the observed and unobserved thermal expansion is deterministic, the discrete sequence of optimal measurement times can be computed offline, proceeding forward from the initial state $\Sigma_0 = \Sigma_{\text{eq}}$, before the transport protocol begins.

VII. DEADLINE DEVALUATION

The value of information [5] consists in extracting thermodynamic work from measured thermal fluctuations. Measurements can occur only if this gain compensates for the immediate cost of observation. Interestingly, we find that regardless of the measurement cost or type, observations become sparser and halt entirely before the deadline is reached. Because the trap must arrive at the target position at the terminal time, the system gradually loses the spatial leeway required to exploit favorable thermal fluctuations. The potential gain decays to zero, meaning that there exists a threshold where the cost of measurement, no matter how small, can

never be repaid. To find this value mathematically, we recall that the immediate available profit $-A_n \Sigma$ is governed by the time-dependent coefficient $A_n = -P_n/2$ [see Eq. (7)]. As the deadline approaches ($n \rightarrow 1$), the thermodynamic yield decreases from $A_\infty = \kappa/2$ toward its terminal value $A_0 = 0$. The break-even variance threshold required to compensate for the measurement cost C/A_n diverges. The term “deadline-induced blindness” refers to the temporal region close to t_f , shaded in gray in Fig. 1, where $\Sigma_{\text{th}}(n)$ exceeds the maximum physical variance of the thermal bath ($\Sigma_{\text{eq}} = k_B T/\kappa$). Since the bath cannot physically generate an expected fluctuation large enough to compensate for C , observations here are never profitable.

Since A_n monotonically decays with n , this blind region stretches backward from t_f across the time horizon. As C increases from 0, the boundary of the blindness region expands. The limit $A_\infty = \kappa/2$ sets a clear phase transition. The maximum cost $C_{\text{max}} = \frac{1}{2}k_B T$ is a general threshold, beyond which the observation cost can never be repaid: The demon’s optimal policy is to remain permanently blind. Conversely, costs $C < \frac{1}{2}k_B T$ can sustain an active steady-state limit cycle.

VIII. OPTIMAL MEASUREMENT SCHEDULING AND THE LAMBERT W TRANSITION

We consider now the steady-state regime, where both $\lambda_f \rightarrow \infty$ and $t_f \rightarrow \infty$ but at a finite macroscopic velocity $v = \lambda_f/t_f$. In practice, we can consider the limit $n \rightarrow \infty$ in Eq. (8). In this regime the optimal trap placement reduces as expected to

$$\lambda_k^* = \mu_k^+ + \frac{v \Delta t}{1 - \alpha}. \quad (13)$$

The classical hydrodynamic lag [24] is recovered, taking the limit $\Delta t \rightarrow 0$:

$$\lambda^*(t) = \mu(t) + \frac{\gamma v}{\kappa} = \mu(t) + \tau v. \quad (14)$$

Again, this holds for any trajectory $\mu(t)$, irrespective to the history of measurements.

For the binary sensor the optimal measurement scheduling is periodic, since the variance threshold is fixed and the evolution of the variance is deterministic. To calculate the exact measurement period $\tau_{\text{meas}} = N^* \Delta t$ in the active steady-state regime, we can maximize the average thermodynamic profit rate of the demon. We obtain the optimal continuous period (see the Appendix D)

$$N^* = \frac{1 + \mathcal{W}_{-1}\left(-\frac{1}{e}\left(1 - \frac{C}{C_{\text{max}}}\right)\right)}{2 \ln \alpha}, \quad (15)$$

where $\mathcal{W}_{-1}$ denotes the lower branch of the Lambert $\mathcal{W}$ function [46] and $\alpha = \exp(-\kappa \Delta t/\gamma)$. The mathematical singularity of the Lambert $\mathcal{W}$ function ($N^* \rightarrow \infty$) corresponds to the physical threshold at $C_{\text{max}} = \frac{1}{2}k_B T$, while as one expects the period monotonically decrease as $C \rightarrow 0$.

As a technical aside, the continuous-time Lambert $\mathcal{W}$ solution provides the analytical envelope for the optimal period, but the original discrete-time optimal policy cannot implement fractional intervals. Simply, the global dynamic programming algorithm must phase lock to integer cycle lengths $N_{\text{cycle}} \approx N^*$.

IX. MACROSCOPIC PROFITABILITY AND THE ENGINE BOUNDARY

We have proven that, in the steady-state regime, there is an optimal measurement frequency $1/N^*$ which depends exclusively on the observation cost C . This means that the optimal frequency is independent of the trap's macroscopic velocity $v = \lambda_f/t_f$. On the other hand, the thermodynamics of the demon cannot be independent of v .

Indeed, for the system to operate as a net-positive information engine, the average microscopic power extracted from thermal fluctuations, $\mathcal{P}_{\text{fluct}} = [C_{\text{max}}(1 - \alpha^{2N^*}) - C]/N^* \Delta t$, must exceed the deterministic macroscopic power dissipated by viscous drag, $\mathcal{P}_{\text{drag}} = \gamma v^2$. We define the thermodynamic viability ratio as $\mathcal{P}_{\text{fluct}}/\mathcal{P}_{\text{drag}}$ and use it to find the bounded active sensing regime in the (v, C) phase space. In the continuous-time limit $\Delta t \rightarrow 0$, we find the profitability boundary (see the Appendix D)

$$C_{\text{env}}(v) = \frac{1}{2} k_B T [1 - (v/v_{\text{max}})^2 + (v/v_{\text{max}})^2 \ln(v/v_{\text{max}})^2],$$

where $v_{\text{max}} = \sqrt{\kappa k_B T / \gamma^2}$. Beyond this envelope, the demon bankrupts the engine, as it cannot operate at a net-positive power with any measurement policy, as shown in Fig. 2.

X. VARIABLE-PRECISION SENSOR

Generalizing the binary sensor to a continuous variable-precision mechanism, we assume the thermodynamic observation cost grows linearly with the precision of a Bayesian measurement: $C(\Sigma^+, \Sigma^-) = c/\Sigma_{\text{meas}} = c(1/\Sigma^+ - 1/\Sigma^-)$. The reduction of variance can be quantified by the Kalman gain [47] $L_k = \frac{\Sigma^- - \Sigma^+}{\Sigma^-}$. Again, the optimal trap placement is the same as before, as it is fully independent of the measurement protocol and modality. To obtain the optimal measuring policy, we have to minimize the total informational cost:

$$g_n(\Sigma^-) = \min_{\Sigma^+ \leq \Sigma^-} [c(1/\Sigma^+ - 1/\Sigma^-) - A_n(\Sigma^- - \Sigma^+)] + g_{n-1}(\alpha^2 \Sigma^+ + \Sigma_\Delta). \quad (16)$$

We derive in the Appendix E that the optimal precision $\Sigma_{\text{opt}}^+(n)$ is independent of the prior variance Σ^- at all times n . This means that the optimal demon either pays the cost to reduce Σ^- to $\Sigma_{\text{opt}}^+(n)$ or if $\Sigma^- < \Sigma_{\text{opt}}^+(n)$ it chooses not to use the sensor. In the finite-time regime, the spatial extraction coefficient A_n monotonically decays as the deadline approaches ($n \rightarrow 0$). Again this deadline devaluation forces the required target $\Sigma_{\text{opt}}^+(n)$ to sweep upward.

Deadline blindness is triggered kinetically: The sensor fully deactivates the moment the growth of the target variance outpaces the forward thermal diffusion of the unmeasured system, i.e., when $\Sigma_{\text{opt}}^+(n-1) \geq \alpha^2 \Sigma_{\text{opt}}^+(n) + \Sigma_\Delta$. This can be seen in Fig. 3, where the Kalman gain L_k remains approximately stable for the majority of the trajectory but suddenly collapses close to the terminal point.

Conversely, in the steady-state $n \rightarrow \infty$ and continuous-time $\Delta t \rightarrow \infty$ limits, the posterior variance locks into a constant that can be calculated in closed form (see the




FIG. 2. Thermodynamic phase space of the information engine with binary measurements at steady state. Heatmaps display (a) the optimal measurement frequency and (b) the thermodynamic viability ratio $\mathcal{P}_{\text{fluct}}/\mathcal{P}_{\text{drag}}$ as a function of macroscopic driving velocity v and measurement cost C . The velocity envelope $C_{\text{env}}(v)$ (white dashed line) separates the active, profitable engine regime from the net-dissipative regime. This operational space is bounded by two fundamental physical limits: the steady-state general threshold ($C = k_B T/2$) at $v = 0$ and the speed limit ($v = v_{\text{max}}$) at $C = 0$, beyond which macroscopic viscous drag overwhelms the maximum extractable microscopic fluctuation power. The simulation parameters are $k_B T = 1$, $\kappa = 1$, and $\gamma = 1$.

Appendix E):

$$\Sigma_\infty^* = \sqrt[3]{\frac{2ck_B T}{\kappa^2} + \sqrt{\frac{4c^2(k_B T)^2}{\kappa^4} + \frac{8c^3}{27\kappa^3}}} + \sqrt[3]{\frac{2ck_B T}{\kappa^2} - \sqrt{\frac{4c^2(k_B T)^2}{\kappa^4} + \frac{8c^3}{27\kappa^3}}}. \quad (17)$$

In this regime, unlike the cyclic regime of the binary sensor, the optimal policy is to operate as an information thermostat. Utilizing a variable-precision Kalman filter [47], the demon constantly measures with a precision rate $\dot{\Sigma}_{\text{meas}}/\Sigma_\infty^{*2}$ (in the continuous-time limit) which exactly offset natural thermal diffusion, locking the posterior variance to the optimal target Σ_∞^* . This balance requires $\dot{\Sigma}_{\text{meas}} = \frac{2\kappa}{\gamma}(\Sigma_\Delta - \Sigma_\infty^*)$.




FIG. 3. Optimal finite-time control dynamics of the information engine with variable-precision measurements. (a) Simulated spatial trajectory illustrating the true particle position x (thin gray line) and the optimal deterministic trap placement λ (red solid line). The agent's belief mean is color coded by the instantaneous Kalman gain L_k , showing the active measurement intensity. The shaded blue region denotes the 1σ posterior uncertainty. The red dot shows the final target λ_f . (b) Corresponding temporal evolution of the belief variance Σ (purple solid line). Throughout the trajectory, the demon continuously measures with just enough precision to keep the the variance locked to the dynamic programming (DP) target Σ_{opt}^+ (orange dashed line). As the deadline approaches, the target diverges triggering the terminal deadline blindness phase (gray shading): The sensor is completely deactivated ($L_t \rightarrow 0$, visible as the dark shift in the belief mean colormap). Afterward the variance converges to the thermal equilibrium variance Σ_{eq} (gray dotted line). The simulation parameters are $k_B T = 1$, $\kappa = 1$, $\gamma = 1$, $\Delta t = 0.01$, $t_f = 20$, and $c = 0.6 c_{\text{max}}$.

XI. MACROSCOPIC PROFITABILITY AT THE STEADY STATE

To derive the state space of the (v, c) plane where the variable-precision sensor can act at a positive-net power, we must again balance the optimal power extractable from fluctuations against the macroscopic viscous drag γv^2 . In the Appendix E we derive this new envelope $c_{\text{env}}(v)$ for the steady state, in the continuous-time limit $\Delta t \rightarrow 0$:

$$c_{\text{env}}(v) = c_{\text{max}} \frac{[s^*(v/v_{\text{max}})]^3}{2 - s^*(v/v_{\text{max}})}, \quad (18)$$

where $v_{\text{max}} = \sqrt{\kappa k_B T / \gamma^2}$ and the maximum cost threshold is $c_{\text{max}} = \frac{(k_B T)^2}{2\kappa} = \frac{k_B T}{2} \Sigma_{\text{eq}}$. The envelop is defined by $s^*(v/v_{\text{max}})$, which is the physically viable ($0 < s \leq 1$) root of the equation

$$s^*(v/v_{\text{max}}) = \frac{4 - (v/v_{\text{max}})^2 - \sqrt{(v/v_{\text{max}})^4 + 8v}}{4}. \quad (19)$$




FIG. 4. Thermodynamic phase space of the information engine with variable-precision measurements at steady state. Heatmaps display (a) the continuous precision rate $\dot{\Sigma}_{\text{meas}}/\Sigma_{\infty}^2$ and (b) the thermodynamic viability ratio $\mathcal{P}_{\text{fluct}}/\mathcal{P}_{\text{drag}}$ as a function of macroscopic driving velocity v and the variable precision cost scalar c . The analytical envelope $c_{\text{env}}(v)$ (white dashed line) separates the active, profitable engine regime from the net-dissipative regime. This operational space is bounded by two fundamental physical limits: the steady-state general threshold [$c = (k_B T)^2 / 2\kappa$] at $v = 0$ and the speed limit ($v = v_{\text{max}}$) at $c = 0$, beyond which macroscopic viscous drag overwhelms the maximum extractable continuous fluctuation power. The simulation parameters are $k_B T = 1$, $\kappa = 1$, and $\gamma = 1$.

The velocity envelop shape can be seen as the white dashed line in Fig. 4. While both $C_{\text{env}}(v)$ and $c_{\text{env}}(v)$ share the same limiting velocity v_{max} , they differ in their asymptotic scaling near the stationary limit ($v \rightarrow 0$). The former approaches the maximum cost with a flat plateau [$-(v/v_{\text{max}})^2$], while for the continuous thermostat boundary the envelop drops linearly. This can be understood by considering that at the boundary, the two measurement models differ in how they reduce variance: A full observation yields a thermodynamic gain that is always proportional to the full equilibrium variance. On the other hand, the continuous observation harvests near equilibrium, where variance is saturated, thus limiting the gain.

XII. CONCLUSION

In this work we applied control theory methods to information engines, tackling the control and measurement

problem for an overdamped harmonic trap with costly sensors. Framing the problem as a discrete-time POMDP, we were able to solve it analytically in closed form. We showed that the optimal feedback law recovers the continuous-time Schmiedl-Seifert limit and generalizes it to any measurement scheduling. We computed the optimal measurement protocols for two different realizations of sensors. In the case of a binary sensor, the optimal scheduling is a sequence of intermittent measurements, while the variable-precision sensor locks the variance onto an optimal time-dependent target.

Our analytical baselines can be used to gain insights into the finite-time boundaries of mesoscopic work extraction and the fundamental thermodynamic limits of an active information engine. Beyond the general cost thresholds and deadline-induced blindness, we calculated the macroscopic velocity envelopes $C_{\text{env}}(v)$ and $c_{\text{env}}(v)$ for the two sensors, which bound the regions of net-positive power output of the optimal agent.

The remarkable simplicity of our results derives from the one-dimensional nature of the Riccati recursive equation and the decoupling of the control and measurements tasks. These are direct consequences of the definition of the physical problem, i.e., the harmonicity of the trap and the separation principle, guaranteed in the linear-quadratic-Gaussian framework within which we operate. While seemingly strict constraints, we note that these conditions should hold also for other, more general, physical setups such as the underdamped particle and the active Ornstein-Uhlenbeck particles [26,48–51], though with the increased complexity of having higher-dimensional state representations.

In contrast, severe deviations from this setup, such as anharmonic potentials or intrinsically non-Markovian fluctuations (e.g., a viscoelastic medium), or treating the trap stiffness $\kappa(t)$ as an active control [24,25], will couple the measurement policy to the physical state and break our analytical solution. By setting the analytical boundaries of the uncoupled regime, our work helps define the baseline against which approximate numerical solvers and reinforcement learning algorithms must be compared to further advance the problem of optimal transport with costly measurements.

ACKNOWLEDGMENTS

The author also acknowledges Antonio Celani and Alessandro H. Ingrosso for fruitful discussions on control theory and Luis Reinalter and Clemens Bechinger for interesting discussions on the measurement scheduling problem. The author acknowledges financial support from Regione FVG, MUR/MIUR, MISE under the project “E-ARGO,” Grant No. J95F21002190001.

The author acknowledges the use of Gemini 3.1 Pro to assist with algebraic expansions, Python simulations, and LaTeX formatting. The author conceptualized the physical framework, verified all mathematical derivations and generated code, and assumes full responsibility for the originality, accuracy, and scientific integrity of this manuscript.

DATA AVAILABILITY

The physical simulation code and plotting scripts required to reproduce the figures and thermodynamic boundaries presented in this paper are openly available [52].

APPENDIX A: DECOUPLING AND DIMENSIONAL REDUCTION OF THE COST FUNCTION

To distinguish the state transitions within a single step, we denote the prior state before measurement by (μ^-, Σ^-) , the posterior state after measurement by (μ^+, Σ^+) , and the prior state at the next step by (μ'^-, Σ'^-) .

In the following we demonstrate by induction that the POMDP cost function V_n separates into an informational cost function $g_n(\Sigma)$ and a deterministic, quadratic physical cost $J_n(\mu, \lambda)$. We further show that the LQG matrix for this spatial cost reduces to a one-dimensional scalar solution of a recurrence equation.

1. Base case $n = 0$

At the deadline, the terminal spatial cost is defined by the expected work of the constrained final jump to the target λ_f . Evaluating this expectation cancels the spatial variance terms:

$$V_0(\mu^-, \lambda) = \kappa \mu^- \lambda - \frac{1}{2} \kappa \lambda^2 - \kappa \lambda_f \mu^- + \frac{1}{2} \kappa \lambda_f^2. \quad (\text{A1})$$

Because this terminal cost is independent of the posterior variance Σ^+ , the value function decouples into $g_0(\Sigma^-) = 0$ and $J_0(\mu^-, \lambda) = V_0(\mu^-, \lambda)$.

2. Inductive step and decoupling

Assume at step $n-1$ the cost function is decoupled: $V_{n-1}(\mu^-, \Sigma^-, \lambda) = J_{n-1}(\mu^-, \lambda) + g_{n-1}(\Sigma^-)$. At step n , the Bellman optimality equation evaluates the measurement and control actions. Using λ for the historical trap position and λ' for the current action, we have

$$\begin{aligned} V_n(\mu^-, \Sigma^-, \lambda) = & \min_{\Sigma^+} \{C(\Sigma^+, \Sigma^-) \\ & + \mathbb{E}_{\mu^+}[\min_{\lambda'} \{\mathbb{E}_x[W] + V_{n-1}(\mu'^-, \Sigma'^-, \lambda')\}]\}. \end{aligned} \quad (\text{A2})$$

Substituting the decoupled cost-to-go function gives

$$\begin{aligned} V_n(\mu^-, \Sigma^-, \lambda) = & \min_{\Sigma^+} \{C(\Sigma^+, \Sigma^-) \\ & + \mathbb{E}_{\mu^+}[\min_{\lambda'} \{\mathbb{E}_x[W] + J_{n-1}(\mu'^-, \lambda') \\ & + g_{n-1}(\Sigma'^-)\}]\}. \end{aligned} \quad (\text{A3})$$

To separate the inner minimization over the physical action λ' from the informational cost, we rely on three physical independences.

(i) The expected thermodynamic work of instantaneously translating a harmonic trap is independent of the spatial variance. Since the trap stiffness κ is constant, the $\mathbb{E}[x^2]$ terms cancel exactly: $\mathbb{E}_x[W] = -\kappa \mu^+ \lambda' + \frac{1}{2} \kappa (\lambda')^2 + \kappa \mu^+ \lambda - \frac{1}{2} \kappa \lambda^2$.

(ii) The transition of the prior mean is deterministic and invariant to the measurement uncertainty Σ^+ : $\mu'^- = \alpha \mu^+ + (1 - \alpha) \lambda'$.

(iii) The information dynamics $\Sigma'^- = \alpha^2 \Sigma^+ + \Sigma_\Delta$ is deterministic and independent of the physical trap placement λ' .

Since Σ'^- is independent of λ' , the future informational cost $g_{n-1}(\Sigma'^-)$ acts as a constant with respect to the inner minimization over λ' and factors out. We define the remaining inner minimization as the intermediate physical cost $\tilde{J}_n(\mu^+, \lambda)$:

$$\tilde{J}_n(\mu^+, \lambda) = \min_{\lambda'} \{ \mathbb{E}_x[W] + J_{n-1}(\alpha\mu^+ + (1-\alpha)\lambda', \lambda') \}. \quad (\text{A4})$$

The system is linear and quadratic, so $\tilde{J}_n$ is quadratic in the posterior mean μ^+ , possessing a spatial curvature defined as $P_{xx,n} \equiv \partial^2 \tilde{J}_n / \partial (\mu^+)^2$.

We can evaluate the outer expectation over the stochastic measurement outcome μ^+ , which is distributed as $\mathcal{N}(\mu^-, \Sigma^- - \Sigma^+)$. Applying the identity $\mathbb{E}[X^2] = \mathbb{E}[X]^2 + \text{Var}(X)$ to the quadratic intermediate cost $\tilde{J}_n$ gives

$$\mathbb{E}_{\mu^+}[\tilde{J}_n(\mu^+, \lambda)] = \tilde{J}_n(\mu^-, \lambda) + \frac{1}{2} P_{xx,n}(\Sigma^- - \Sigma^+). \quad (\text{A5})$$

We define the physical cost evaluated at the prior mean as $J_n(\mu^-, \lambda) \equiv \tilde{J}_n(\mu^-, \lambda)$. Substituting this expansion back into the full Bellman equation separates the state space into independent spatial and informational recurrences:

$$\begin{aligned} V_n(\mu^-, \Sigma^-, \lambda) &= J_n(\mu^-, \lambda) \\ &+ \min_{\Sigma^+} \{ C(\Sigma^+, \Sigma^-) + \frac{1}{2} P_{xx,n}(\Sigma^- - \Sigma^+) \\ &+ g_{n-1}(\alpha^2 \Sigma^+ + \Sigma_\Delta) \}. \end{aligned} \quad (\text{A6})$$

The minimization over Σ^+ defines the current informational cost $g_n(\Sigma^-)$, confirming the inductive hypothesis.

3. Riccati dimensional reduction

Standard discrete-time LQG control requires tracking an augmented joint state space $\mathbf{z} = [\mu^+, \lambda]^T$ and solving a 2×2 Riccati cost matrix $\mathbf{P}_n$.

Evaluating the stationary point $\partial/\partial \lambda'(\cdot) = 0$ inside the intermediate cost $\tilde{J}_n$ [Eq. (A4)] removes the previous trap position λ , as the expected work contains no cross-coupling term $\lambda'\lambda$. The optimal control law $\lambda'^*(\mu^+)$ is therefore independent of the previous trap position.

Substituting λ'^* back into the objective function allows λ to exit the minimization operator. Mapped to the standard 2×2 augmented LQG formulation, the off-diagonal coupling term is constant at $P_{x\lambda,n} = \kappa$, and similarly $P_{\lambda\lambda,n} = -\kappa$ for all finite n . Thus, the multivariable Riccati matrix equation is unnecessary; the spatial dynamics is captured entirely by the scalar recurrence of the spatial curvature $P_{xx,n}$, which for simplicity we can rename, $P_{xx,n} \equiv P_n$.

APPENDIX B: ANALYTICAL SOLUTION OF THE DISCRETE RICCATI RECURRENCE

The dimensional reduction of the POMDP reduces the spatial cost to a single scalar recurrence [33]

$$P_n = P_{n-1}\alpha^2 - \frac{\Gamma^2}{\Omega}, \quad (\text{B1})$$

where $\Gamma = (1-\alpha)(P_{n-1}\alpha - \kappa)$ and $\Omega = (1-\alpha)[P_{n-1}(1-\alpha) + 2\kappa]$.

Substituting these explicit coefficients into the recurrence, expanding the algebraic fraction, and grouping terms by P_{n-1} reduces the nonlinear equation to a Möbius transformation:

$$P_n = \frac{2\alpha\kappa P_{n-1} - (1-\alpha)\kappa^2}{(1-\alpha)P_{n-1} + 2\kappa}. \quad (\text{B2})$$

This recurrence be linearized via the substitution

$$y_n = \frac{1}{P_n + \kappa}. \quad (\text{B3})$$

Inserting this substitution into Eq. (B2) transforms it into a simpler arithmetic progression

$$y_n = y_{n-1} + \frac{1-\alpha}{\kappa(1+\alpha)}. \quad (\text{B4})$$

For a finite-time transport protocol, the cancellation of the quadratic terms in the terminal cost dictates the boundary condition $P_0 = 0$, which forces the terminal sequence value to $y_0 = 1/\kappa$. Evaluating the arithmetic progression backward from this deadline yields

$$y_n = \frac{1}{\kappa} + n \frac{1-\alpha}{\kappa(1+\alpha)}. \quad (\text{B5})$$

Inverting the initial substitution via $P_n = \frac{1}{y_n} - \kappa$ provides the closed-form analytical solution provided in the main text for the dynamic coefficient at any n :

$$P_n = -\kappa \left(\frac{n(1-\alpha)}{1+\alpha+n(1-\alpha)} \right). \quad (\text{B6})$$

APPENDIX C: MEASUREMENT POLICY

1. Informational cost function

As derived in Appendix A, the mathematical decoupling of the Bellman equation isolates the informational cost into the standalone minimization shown in Eq. (A6). The average over the posterior mean, $\mathbb{E}_{\mu^+}[\cdot]$, shown in Eq. (A5) brings out the additional term $\frac{1}{2}P_n(\Sigma^- - \Sigma^+)$.

By defining the thermodynamic energy bonus as $A_n = -P_n/2$, this term becomes

$$-A_n(\Sigma^- - \Sigma^+). \quad (\text{C1})$$

One can think of this negative cost as the immediate spatial energy bonus extracted by localizing the particle prior to executing the trap displacement. The Bellman equation for the informational cost reduces to minimizing the sum of the observation cost, this immediate spatial energy bonus, and the expected future informational cost:

$$\begin{aligned} g_n(\Sigma^-) &= \min_{\Sigma^+} [C(\Sigma^+, \Sigma^-) \\ &- A_n(\Sigma^- - \Sigma^+) + g_{n-1}(\alpha^2 \Sigma^+ + \Sigma_\Delta)]. \end{aligned} \quad (\text{C2})$$

2. Optimality for the binary sensor

For the binary measurement, the minimization separates into the two branches of Eq. (6):

$$\begin{aligned} g_n(\Sigma^-) &= \min_{a_{\text{obs}} \in \{0,1\}} \{ (1-a_{\text{obs}})g_{n-1}(\alpha^2 \Sigma^- + \Sigma_\Delta) \\ &+ a_{\text{obs}}[C - A_n \Sigma^- + g_{n-1}(\Sigma_\Delta)] \}. \end{aligned} \quad (\text{C3})$$

In the measured branch ($a_{\text{obs}} = 1$), the residual variance collapses, leaving only the thermal diffusion over a single step, Σ_{Δ} . The trigger condition dictates that the controller measures if and only if

$$g_{n-1}(\alpha^2 \Sigma^- + \Sigma_{\Delta}) - g_{n-1}(\Sigma_{\Delta}) \geq C - A_n \Sigma^-. \quad (\text{C4})$$

In a thermodynamic harmonic engine, variance represents extractable potential energy. The informational cost function $g_n(\Sigma)$ can be proven to be monotonically decreasing in Σ : A larger variance yields a better (more negative) cost.

Proof of monotonicity. We prove by induction that $\partial g_n / \partial \Sigma \leq 0$.

(i) *Base case* ($n = 0$). At the terminal deadline, there is no future extraction, so $g_0(\Sigma) = 0$. The derivative is therefore zero.

(ii) *Step 1* ($n = 1$). Evaluating the Bellman equation one step before the deadline yields $g_1(\Sigma) = \min\{0, C - A_1 \Sigma\}$. Because the physical extraction coefficient $A_1 > 0$, the measured branch decreases with Σ . Thus, $\partial g_1 / \partial \Sigma \leq 0$.

(iii) *Inductive step* (n). Assume $g_{n-1}(\Sigma)$ is monotonically decreasing ($\partial g_{n-1} / \partial \Sigma \leq 0$). We evaluate the derivatives of both branches of $g_n(\Sigma)$. For the unmeasured branch, $\frac{\partial}{\partial \Sigma} g_{n-1}(\alpha^2 \Sigma + \Sigma_{\Delta}) = \alpha^2 \frac{\partial g_{n-1}}{\partial \Sigma} \leq 0$. For the measured branch, $\frac{\partial}{\partial \Sigma} [C - A_n \Sigma + g_{n-1}(\Sigma_{\Delta})] = -A_n < 0$. Because both branches are nonincreasing functions of Σ , $g_n(\Sigma)$ must be monotonically decreasing.

APPENDIX D: MACROSCOPIC PROFITABILITY BOUNDARY FOR PERFECT MEASUREMENT

We derive the profitability boundary $C_{\text{env}}(v)$ for an information engine with a perfect binary sensor. We evaluate the continuous-time limits where the maximum achievable fluctuation power is limited by the macroscopic viscous drag.

1. Continuous-time extracted power

In the discrete-time formulation, the average microscopic power extracted from thermal fluctuations is $\mathcal{P}_{\text{fluct}} = [C_{\text{max}}(1 - \alpha^{2N^*}) - C] / N^* \Delta t$, where $C_{\text{max}} = \frac{1}{2} k_B T$. To evaluate the continuous-time measurement period $t_p = N^* \Delta t$, we can simply write explicitly the relaxation parameter $\alpha = \exp(-\kappa \Delta t / \gamma)$. This yields the power equation

$$\mathcal{P}_{\text{fluct}}(t_p) = \frac{C_{\text{max}}(1 - e^{-2\kappa t_p / \gamma}) - C}{t_p}. \quad (\text{D1})$$

2. Optimality condition

For the system to operate at the information engine boundary, it must maximize its average thermodynamic profit rate with respect to the measurement period. We enforce the stationarity condition $\frac{\partial \mathcal{P}_{\text{fluct}}}{\partial t_p} = 0$:

$$\frac{\partial}{\partial t_p} \left(\frac{C_{\text{max}}(1 - e^{-2\kappa t_p / \gamma}) - C}{t_p} \right) = 0. \quad (\text{D2})$$

Defining the dimensionless parameter $x = 2\kappa t_p / \gamma$, which represents the ratio of the measurement period to the trap relaxation time, the derivative requires

$$C_{\text{max}} e^{-x} (x + 1) = C_{\text{max}} - C. \quad (\text{D3})$$

Isolating the required thermodynamic measurement cost C yields

$$C = C_{\text{max}} [1 - e^{-x} (x + 1)]. \quad (\text{D4})$$

Solving for the optimal measurement period t_p^* , we use the known inversion of the equation $Y e^Y = Z$. Rearranging Eq. (D4), we have

$$-(x + 1) e^{-(x+1)} = -\frac{1}{e} \left(1 - \frac{C}{C_{\text{max}}} \right). \quad (\text{D5})$$

This maps to the defining equation of the Lambert W function, $W(z) e^{W(z)} = z$ [46]. Applying the function to both sides gives

$$-(x + 1) = W \left(-\frac{1}{e} \left(1 - \frac{C}{C_{\text{max}}} \right) \right). \quad (\text{D6})$$

Since the physical time ratio $x = 2\kappa t_p / \gamma$ must be positive ($x > 0$), the argument $-(x + 1)$ is less than -1 . The Lambert W function is double valued on the interval $[-1/e, 0)$ [46] for which $C_{\text{max}} > C > 0$, but the principal branch $W_0(z) \geq -1$ yields nonphysical negative times $x \leq 0$. The solution is found by the lower real branch, $W_{-1}(z) \leq -1$. Solving for x ,

$$-(x + 1) = W_{-1} \left(-\frac{1}{e} \left(1 - \frac{C}{C_{\text{max}}} \right) \right), \quad (\text{D7})$$

and substituting the physical constants back into the parameter yields Eq. (15).

3. Maximized extracted power

Substituting the optimality constraint calculated in Eq. (D4) back into the power equation defines the maximum theoretically extractable power for a given dimensionless period x :

$$\mathcal{P}_{\text{max}}(x) = \frac{2\kappa}{\gamma x} [C_{\text{max}}(1 - e^{-x}) - C_{\text{max}}(1 - e^{-x} - x e^{-x})], \quad (\text{D8})$$

$$\mathcal{P}_{\text{max}}(x) = \frac{2\kappa C_{\text{max}}}{\gamma} e^{-x}. \quad (\text{D9})$$

4. Speed limit

The physical speed limit of the engine occurs when the observation cost is zero. Here the optimal policy dictates measuring each step, driving the period $t_p \rightarrow 0$, and therefore $x \rightarrow 0$. In this limit, the limiting factor is the macroscopic drag

$$\gamma v_{\text{max}}^2 = \frac{2\kappa C_{\text{max}}}{\gamma}. \quad (\text{D10})$$

Solving for the limiting velocity yields

$$v_{\text{max}} = \sqrt{\frac{2\kappa C_{\text{max}}}{\gamma^2}} = \sqrt{\frac{\kappa k_B T}{\gamma^2}}. \quad (\text{D11})$$

5. The $C_{\text{env}}(v)$ envelope

At the profitability boundary, the optimized fluctuation power equals the power dissipated by viscous drag

$$\mathcal{P}_{\text{max}}(x) = \gamma v_{\text{max}}^2 e^{-x} = \gamma v^2. \quad (\text{D12})$$

Solving this equality for the dimensionless period x yields identities for the boundary constraints

$$e^{-x} = \left(\frac{v}{v_{\max}}\right)^2 \Rightarrow x = -\ln\left(\frac{v}{v_{\max}}\right)^2. \quad (\text{D13})$$

Substituting these identities back into the stationary cost constraint [Eq. (D4)] allows us to obtain the velocity envelope

$$C_{\text{env}}(v) = C_{\max} \left\{ 1 - \left(\frac{v}{v_{\max}}\right)^2 \left[-\ln\left(\frac{v}{v_{\max}}\right)^2 + 1 \right] \right\}. \quad (\text{D14})$$

APPENDIX E: VARIABLE-PRECISION SENSOR

In this appendix we derive the optimal measurement policy for the variable-precision sensor, for both the finite-time horizon and the steady state. We evaluate Σ_{∞}^* and the profitability boundary $c_{\text{env}}(v)$ in the continuous-time limit $\Delta t \rightarrow 0$.

1. Finite-time optimal measurement policy

For a variable-precision sensor, the measurement cost is inversely proportional to the resolved spatial uncertainty: $C(\Sigma^+, \Sigma^-) = c(1/\Sigma^+ - 1/\Sigma^-)$. To determine the optimal measurement policy at step n , the agent minimizes the total informational cost:

$$g_n(\Sigma^-) = \min_{\Sigma^+ \leq \Sigma^-} \left[c \left(\frac{1}{\Sigma^+} - \frac{1}{\Sigma^-} \right) - A_n(\Sigma^- - \Sigma^+) + g_{n-1}(\alpha^2 \Sigma^+ + \Sigma_{\Delta}) \right]. \quad (\text{E1})$$

To find the minimum, we differentiate the bracketed objective function with respect to the measurement outcome Σ^+ and set it to zero:

$$-\frac{c}{(\Sigma^+)^2} + A_n + \alpha^2 g'_{n-1}(\alpha^2 \Sigma^+ + \Sigma_{\Delta}) = 0. \quad (\text{E2})$$

Note that the terms containing the prior variance Σ^- , being constants during the differentiation over Σ^+ , have vanished from the stationarity condition: The optimal measurement depends exclusively on the system parameters c , α , and Σ_{Δ} ; the physical extraction coefficient A_n ; and the future value function gradient g'_{n-1} .

We define this optimal root as the dynamic target precision $\Sigma_{\text{opt}}^+(n)$. Since it is independent of the prior state Σ^- , the optimal control policy reduces to a simple binary choice.

(i) *Measurement regime* $\Sigma^- > \Sigma_{\text{opt}}^+(n)$. If thermal diffusion has pushed the prior spatial uncertainty above the optimal set point, the agent injects the measurement energy required to reduce the posterior variance down to the target: $\Sigma^+ = \Sigma_{\text{opt}}^+(n)$.

(ii) *Idle regime* $\Sigma^- \leq \Sigma_{\text{opt}}^+(n)$. If the prior uncertainty is already at or below the optimal set point, measuring incurs a cost without yielding a sufficient energetic return. The agent leaves the state unaltered: $\Sigma^+ = \Sigma^-$.

2. Steady-state net power and optimal precision

For the variable-precision sensor, the measurement cost scales as $C(\Sigma^+, \Sigma^-) = c(1/\Sigma^+ - 1/\Sigma^-)$. To derive the

continuous-time thermodynamic profit rate from the discrete-time control framework, we evaluate the discrete profit generated during a single interval Δt and take the asymptotic limit $\Delta t \rightarrow 0$.

Let us consider a steady state, where we have locked into a given precision Σ_{ss} . Each step, the variance evolves via thermal diffusion for Δt to $\Sigma' = \alpha^2 \Sigma_{\text{ss}} + (1 - \alpha^2) \Sigma_{\Delta}$, gaining a (fixed) amount $\Delta \Sigma = (1 - \alpha^2)(\Sigma_{\Delta} - \Sigma_{\text{ss}})$ and paying a cost of $c(1/\Sigma_{\text{ss}} - 1/\Sigma')$.

3. Thermodynamic profit for a discrete step

The net thermodynamic profit generated during a single discrete step is the spatial energy bonus minus the observation cost:

$$\Delta \mathcal{P} = A_{\infty} \Delta \Sigma - c \left(\frac{1}{\Sigma_{\text{ss}}} - \frac{1}{\Sigma'} \right). \quad (\text{E3})$$

Factoring out $\Delta \Sigma$ from the cost term gives the discrete profit per step

$$\Delta \mathcal{P} = \Delta \Sigma \left(A_{\infty} - \frac{c}{\Sigma_{\text{ss}} \Sigma'} \right). \quad (\text{E4})$$

4. Continuous limit $\Delta t \rightarrow 0$

The continuous profit rate is defined as the limit of the discrete profit over the time step: $\mathcal{P}_{\text{fluct}} = \lim_{\Delta t \rightarrow 0} \frac{\Delta \mathcal{P}}{\Delta t}$.

To evaluate this, we apply the small- Δt limits to the discrete components. The relaxation parameter gives $1 - \alpha^2 \approx \frac{2\kappa}{\gamma} \Delta t$ and the variance rate $\lim_{\Delta t \rightarrow 0} \frac{\Delta \Sigma}{\Delta t} = \frac{2\kappa}{\gamma} (\Sigma_{\Delta} - \Sigma)$. Furthermore, we have that $A_{\infty} = \frac{\kappa}{2}$ and, in the denominator, $\Sigma_{\text{ss}} \Sigma' = \Sigma_{\text{ss}}^2$.

Substituting these into the profit rate equation gives

$$\begin{aligned} \mathcal{P}_{\text{fluct}}(\Sigma_{\text{ss}}) &= \lim_{\Delta t \rightarrow 0} \frac{\Delta \Sigma}{\Delta t} \left(A_{\infty} - \frac{c}{\Sigma_{\text{ss}} \Sigma'} \right) \\ &= \frac{2\kappa}{\gamma} (\Sigma_{\Delta} - \Sigma_{\text{ss}}) \left(\frac{\kappa}{2} - \frac{c}{\Sigma_{\text{ss}}^2} \right). \end{aligned} \quad (\text{E5})$$

The optimal Σ_{∞}^* satisfies the stationarity condition $\partial \mathcal{P}_{\text{fluct}} / \partial \Sigma_{\text{ss}}|_{\Sigma_{\infty}^*} = 0$:

$$\frac{c}{\Sigma_{\infty}^{*2}} \left(\frac{2\kappa_B T}{\kappa \Sigma_{\infty}^*} - 1 \right) = \frac{\kappa}{2}. \quad (\text{E6})$$

Multiplying by $2\kappa \Sigma_{\infty}^{*3}$ reorganizes this into a cubic polynomial

$$\kappa^2 \Sigma_{\infty}^{*3} + 2c\kappa \Sigma_{\infty}^* - 4c\kappa_B T = 0. \quad (\text{E7})$$

Since the physical parameters are positive, the discriminant is positive, guaranteeing a single real root. Applying Cardano's formula returns the expression for the optimal continuous precision Σ_{∞}^* shown in the main text, resulting in

$$\begin{aligned} \Sigma_{\infty}^* &= \sqrt[3]{\frac{2c\kappa_B T}{\kappa^2} + \sqrt{\frac{4c^2(\kappa_B T)^2}{\kappa^4} + \frac{8c^3}{27\kappa^3}}} \\ &\quad + \sqrt[3]{\frac{2c\kappa_B T}{\kappa^2} - \sqrt{\frac{4c^2(\kappa_B T)^2}{\kappa^4} + \frac{8c^3}{27\kappa^3}}}. \end{aligned} \quad (\text{E8})$$

5. Dimensionless formulation

To evaluate the macroscopic transport limits, we define the dimensionless variance $s = \kappa \Sigma_{\infty}^*/k_B T$, constrained to $s \in (0, 1]$. We can thus rewrite the optimality condition of Eq. (E6) to isolate the measurement cost c :

$$c = \frac{(k_B T)^2}{2\kappa} \frac{s^3}{2-s}. \quad (\text{E9})$$

The absolute limit occurs at the boundary of net-zero power ($s \rightarrow 1$), defining the critical sensor cost:

$$c_{\max} = \frac{(k_B T)^2}{2\kappa}. \quad (\text{E10})$$

The required operating cost for any optimal dimensionless variance s therefore reduces to $c(s) = c_{\max} \frac{s^3}{2-s}$.

6. Envelope $c_{\text{env}}(v)$

Substituting $c(s)$ back into the continuous power equation (E5), we get the optimal net power as a function of the dimensionless variance:

$$\mathcal{P}_{\text{fluct}}(s) = \frac{2\kappa k_B T}{\gamma} \frac{(1-s)^2}{2-s}. \quad (\text{E11})$$

To find the operational envelope for any arbitrary macroscopic velocity v , we equate the optimal net power to the required drag dissipation:

$$\gamma v^2 = \gamma v_{\max}^2 \frac{2(1-s)^2}{2-s}. \quad (\text{E12})$$

Defining the squared velocity ratio $v = (v/v_{\max})^2$, this expands into a quadratic condition for the required steady-state variance:

$$2s^2 + (v-4)s + 2(1-v) = 0. \quad (\text{E13})$$

Taking the physical root ($s \leq 1$) yields the explicit spatial threshold required to sustain velocity v :

$$s^*(v) = \frac{4-v-\sqrt{v^2+8v}}{4}. \quad (\text{E14})$$

Substituting this dynamic root back into the cost function closes the system, providing the profitability boundary for the variable-precision engine:

$$c_{\text{env}}(v) = c_{\max} \frac{[s^*(v)]^3}{2-s^*(v)}. \quad (\text{E15})$$

-
- [1] C. Jarzynski, Nonequilibrium equality for free energy differences, *Phys. Rev. Lett.* **78**, 2690 (1997).
 - [2] G. E. Crooks, Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences, *Phys. Rev. E* **60**, 2721 (1999).
 - [3] U. Seifert, Stochastic thermodynamics, fluctuation theorems and molecular machines, *Rep. Prog. Phys.* **75**, 126001 (2012).
 - [4] T. Sagawa and M. Ueda, Generalized Jarzynski equality under nonequilibrium feedback control, *Phys. Rev. Lett.* **104**, 090602 (2010).
 - [5] M. Bauer, D. Abreu, and U. Seifert, Efficiency of a Brownian information machine, *J. Phys. A: Math. Theor.* **45**, 162001 (2012).
 - [6] T. Sagawa and M. Ueda, Nonequilibrium thermodynamics of feedback control, *Phys. Rev. E* **85**, 021104 (2012).
 - [7] J. Ehrich and D. A. Sivak, Energy and information flows in autonomous systems, *Front. Phys.* **11**, 1108357 (2023).
 - [8] J. C. Maxwell, *Theory of Heat* (Longmans, Green, and Co., London, 1871).
 - [9] L. Szilard, über die entropieverminderung in einem thermodynamischen system bei eingriffen intelligenter wesen, *Z. Phys.* **53**, 840 (1929).
 - [10] R. Landauer, Irreversibility and heat generation in the computing process, *IBM J. Res. Dev.* **5**, 183 (1961).
 - [11] J. M. Parrondo, J. M. Horowitz, and T. Sagawa, Thermodynamics of information, *Nat. Phys.* **11**, 131 (2015).
 - [12] S. Toyabe, T. Sagawa, M. Ueda, E. Muneyuki, and M. Sano, Experimental demonstration of information-to-energy conversion and validation of the generalized Jarzynski equality, *Nat. Phys.* **6**, 988 (2010).
 - [13] A. Bérut, A. Arakelyan, A. Petrosyan, S. Ciliberto, R. Dillenschneider, and E. Lutz, Experimental verification of Landauer's principle linking information and thermodynamics, *Nature (London)* **483**, 187 (2012).
 - [14] I. A. Martínez, É. Roldán, L. Dinis, and R. A. Rica, Colloidal heat engines: A review, *Soft Matter* **13**, 22 (2017).
 - [15] G. Paneru and H. Kyu Pak, Colloidal engines for innovative tests of information thermodynamics, *Adv. Phys. X* **5**, 1823880 (2020).
 - [16] T. K. Saha, J. N. Lucero, J. Ehrich, D. A. Sivak, and J. Bechhoefer, Maximizing power and velocity of an information engine, *Proc. Natl. Acad. Sci. USA* **118**, e2023356118 (2021).
 - [17] T. K. Saha, J. Ehrich, M. Gavrilov, S. Still, D. A. Sivak, and J. Bechhoefer, Information engine in a nonequilibrium bath, *Phys. Rev. Lett.* **131**, 057101 (2023).
 - [18] A. Bensoussan, *Stochastic Control of Partially Observable Systems* (Cambridge University Press, Cambridge, 1992).
 - [19] J. Alvarado, E. G. Teich, D. A. Sivak, and J. Bechhoefer, Optimal control in soft and active matter, *Annu. Rev. Condens. Matter Phys.* **17**, 327 (2026).
 - [20] E. Aurell, C. Mejía-Monasterio, and P. Muratore-Ginanneschi, Optimal protocols and optimal transport in stochastic thermodynamics, *Phys. Rev. Lett.* **106**, 250601 (2011).
 - [21] T. Kamijima, A. Takatsu, K. Funo, and T. Sagawa, Optimal finite-time Maxwell's demons in Langevin systems, *Phys. Rev. Res.* **7**, 023159 (2025).
 - [22] A. T. Mohite and H. Rieger, Generalized finite-time optimal control framework in stochastic thermodynamics, *arXiv:2511.00974*.
 - [23] S. Blaber, M. D. Louwerse, and D. A. Sivak, Steps minimize dissipation in rapidly driven stochastic systems, *Phys. Rev. E* **104**, L022101 (2021).
 - [24] T. Schmiedl and U. Seifert, Optimal finite-time processes in stochastic thermodynamics, *Phys. Rev. Lett.* **98**, 108301 (2007).
 - [25] D. Abreu and U. Seifert, Extracting work from a single heat bath through feedback, *Europhys. Lett.* **94**, 10001 (2011).

- [26] R. Garcia-Millan, J. Schüttler, M. E. Cates, and S. A. M. Loos, Optimal closed-loop control of active particles and a minimal information engine, *Phys. Rev. Lett.* **135**, 088301 (2025).
- [27] J. Schüttler, R. Garcia-Millan, M. E. Cates, and S. A. M. Loos, Active particles in moving traps: Minimum work protocols and information efficiency of work extraction, *Phys. Rev. E* **112**, 024119 (2025).
- [28] S. Whitlam, Demon in the machine: Learning to extract work and absorb entropy from fluctuating nanosystems, *Phys. Rev. X* **13**, 021005 (2023).
- [29] J. Bechhoefer, Hidden Markov models for stochastic thermodynamics, *New J. Phys.* **17**, 075003 (2015).
- [30] M. Biehl and N. Virgo, Communications in Computer and Information Science, in *Active Inference*, edited by C. L. Buckley, D. Cialfi, P. Lanillos, M. Ramstead, N. Sajid, H. Shimazaki, and T. Verbelen (Springer, Cham, 2022), Vol. 1721, pp. 16–31.
- [31] L. Meier, J. Peschon, and R. Dressler, Optimal control of measurement subsystems, *IEEE Trans. Autom. Control* **12**, 528 (1967).
- [32] B. D. Anderson and J. B. Moore, *Optimal Control: Linear Quadratic Methods* (Courier, North Chelmsford, 2007).
- [33] J. Bechhoefer, *Control Theory for Physicists* (Cambridge University Press, Cambridge, 2021).
- [34] K. J. Åström, Optimal control of Markov processes with incomplete state information, *J. Math. Anal. Appl.* **10**, 174 (1965).
- [35] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, Planning and acting in partially observable stochastic domains, *Artif. Intell.* **101**, 99 (1998).
- [36] J. M. Horowitz and M. Esposito, Thermodynamics with continuous information flow, *Phys. Rev. X* **4**, 031015 (2014).
- [37] G. Falasco and M. Esposito, Macroscopic stochastic thermodynamics, *Rev. Mod. Phys.* **97**, 015002 (2025).
- [38] J. V. Koski, A. Kutvonen, I. M. Khaymovich, T. Ala-Nissila, and J. P. Pekola, On-chip Maxwell’s demon as an information-powered refrigerator, *Phys. Rev. Lett.* **115**, 260602 (2015).
- [39] G. E. Uhlenbeck and L. S. Ornstein, On the theory of the Brownian motion, *Phys. Rev.* **36**, 823 (1930).
- [40] D. T. Gillespie, Exact numerical simulation of the Ornstein-Uhlenbeck process and its integral, *Phys. Rev. E* **54**, 2084 (1996).
- [41] R. Bellman, *Dynamic Programming* (Princeton University Press, Princeton, 1957).
- [42] D. Bertsekas, *Dynamic Programming and Optimal Control*, 4th ed. (Athena Scientific, Nashua, 2012), Vol. 1.
- [43] H. A. Simon, Dynamic programming under uncertainty with a quadratic criterion function, *Econometrica* **24**, 74 (1956).
- [44] Y. Bar-Shalom and E. Tse, Dual effect, certainty equivalence, and separation in stochastic control, *IEEE Trans. Autom. Control* **19**, 494 (1974).
- [45] T. Sagawa and M. Ueda, Minimal energy cost for thermodynamic information processing: Measurement and information erasure, *Phys. Rev. Lett.* **102**, 250602 (2009).
- [46] R. M. Corless, G. H. Gonnet, D. E. Hare, D. J. Jeffrey, and D. E. Knuth, On the Lambert W function, *Adv. Comput. Math.* **5**, 329 (1996).
- [47] R. Kalman, A new approach to linear filtering and prediction problems, *J. Basic Eng.* **82**, 35 (1960).
- [48] L. L. Bonilla, Active Ornstein-Uhlenbeck particles, *Phys. Rev. E* **100**, 022601 (2019).
- [49] L. K. Davis, Optimal multi-parameter control of trapped active matter, [arXiv:2603.16778](https://arxiv.org/abs/2603.16778).
- [50] L. Cocconi, J. Knight, and C. Roberts, Optimal power extraction from active particles with hidden states, *Phys. Rev. Lett.* **131**, 188301 (2023).
- [51] C. Casert and S. Whitlam, Learning protocols for the fast and efficient control of active matter, *Nat. Commun.* **15**, 9128 (2024).
- [52] E. Panizon, RitAreaSciencePark/information-engine: v1.2.0: Rebuttal manuscript submission, Zenodo, 2026, <https://doi.org/10.5281/zenodo.21774633>.