Supplemental Text: Optimal Control in Soft and Active Matter

José Alvarado, John Bechhoefer, David Sivak, Erin Teich

In this supplement, we illustrate some of the basic techniques for optimal control, in hopes that a simple example solved several ways can illuminate and make concrete the broader discussion given in the main text. Although our choice is just the "physicists' favorite toy model," the simple harmonic oscillator, we will find that it is similar to and sheds light on many examples of current interest in soft and active-matter physics. For more examples, see [1, 2].

I. A BASIC EXAMPLE

Consider the transport of a particle in a harmonic potential from one position to another in a specified time interval (protocol duration) τ. The equation of motion and initial and final conditions are

$$\underbrace{\ddot{x} + x = u}_{\text{dynamics}} \qquad \underbrace{x(0) = x_0\,,\ \dot{x}(0) = 0}_{\text{initial}} \qquad \underbrace{x(\tau) = \dot{x}(\tau) = 0}_{\text{final}} \qquad , \tag{S1}$$

where x is the one-dimensional position of the particle and $u(t)$ is the applied force at time t. All units are scaled to eliminate explicit parameters such as mass, stiffness, and angular frequency. The initial condition at $t = 0$ corresponds to a stationary particle at position x_0, and the final condition at the end of the protocol is at $t = \tau$ and corresponds to a stationary particle at position $x(\tau) = 0$. For now, we neglect damping and stochastic (thermal) forces.

We seek the control u that moves the particle from x_0 to the origin, in time τ, while minimizing the total cost

$$J = \frac{1}{2} \int_0^\tau \mathrm{d}t\, u^2(t)\,. \tag{S2}$$

Here, J is the integral of the *running cost* $\frac{1}{2}u^2$. Here, the cost depends only on the input u and implies moving from x_0 to 0 in time τ while using a minimum of "control effort." Absent any control, the particle will oscillate forever about the desired position, $x = 0$ (Fig. S1c.) In Sec. VII, we discuss the relation between effort and physical quantities such as energy.

The control can be given either as a *feedforward command* $u(t)$ or as a *feedback command* $u(\boldsymbol{x}(t), t)$ that also depends on the instantaneous particle state. Here, $\boldsymbol{x}$ denotes the two-dimensional state vector formed from $x_1 = x$ and $x_2 = \dot{x}$. In vector-matrix notation, the equations of motion are $\dot{\boldsymbol{x}} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{B}u$. More explicitly,

$$\frac{\mathrm{d}\boldsymbol{x}}{\mathrm{d}t} = \frac{\mathrm{d}}{\mathrm{d}t}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \underbrace{\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}}_{\boldsymbol{A}} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} + \underbrace{\begin{pmatrix} 0 \\ 1 \end{pmatrix}}_{\boldsymbol{B}} u\,, \tag{S3}$$

with $\boldsymbol{A}$ the dynamics matrix and $\boldsymbol{B}$ the input coupling. The input u (*control parameter*, in the physics literature) here is a scalar because there is only one input function, but more generally would be a vector.

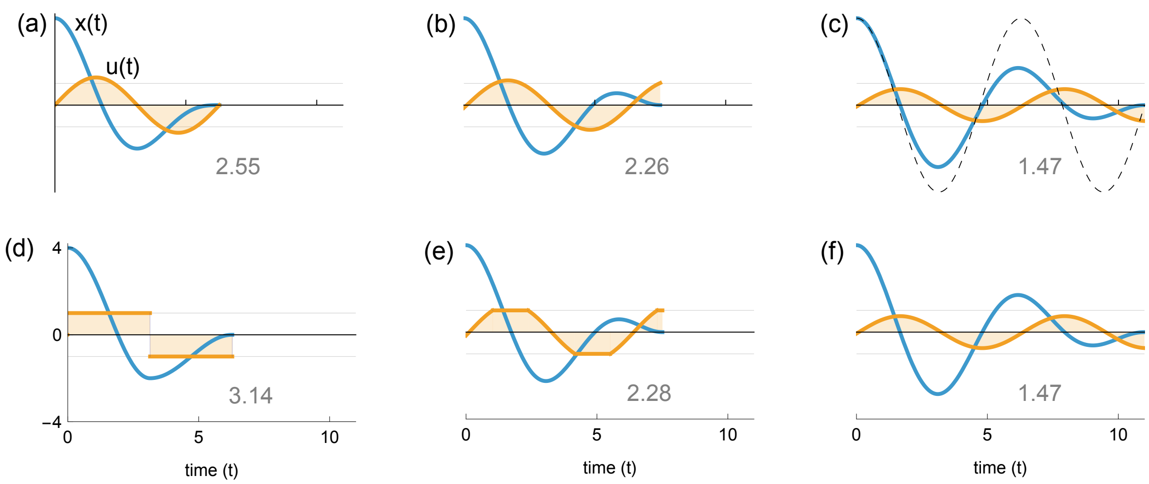

FIG. S1. Unconstrained and constrained feedforward control. Protocol duration τ is (a,d) 2π, (b,e) 7.5, or (c,f) 11. Blue curve is the particle position $x_1^*(t)$. Orange-brown filled curve is the control input $u^*(t)$. Initial position is $x_0 = 4$. Control costs J for each case are in gray at bottom center-right. (a–c): Unconstrained input. (d–f): Input constrained to $|u(t)| \leq 1$. The black dashed curve in (c) shows the uncontrolled sinusoidal oscillations.

Note that in traditional problems of classical mechanics, we specify the initial conditions and input $u(t)$ and then determine the trajectory $\boldsymbol{x}(t)$ and, with it, the final condition $\boldsymbol{x}(\tau)$. Here, instead of an *initial value problem*, we have a *boundary value problem*: we specify $\boldsymbol{x}$ at two different times but seek a $u(t)$ that connects them. The goal for optimal control is to choose the $u(t) = u^*(t)$ from the infinity of possible solutions so that $u^*(t)$ minimizes the cost functional J of Eq. (S2).

A helpful check and aid in solving problems is that when H has no explicit time dependence (i.e., when both the running cost and the equations of motion are time invariant), the Hamiltonian is constant when evaluated on optimal trajectories $(\boldsymbol{x}^*(t), \boldsymbol{\lambda}^*(t), u^*(t))$. For smooth solutions, the argument is the same as used in classical mechanics to identify a Hamiltonian with the total energy: $\dot{H} = \partial_x H \dot{x} + \partial_\lambda H \dot{\lambda} + \partial_u H \dot{u} = 0$, after substituting Hamilton's equations.

II. UNCONSTRAINED CONTROL: FEEDFORWARD SOLUTION

If there are no constraints on the magnitude of the allowed control $u(t)$, then we can solve the optimal-control problem by the calculus of variations to find the optimal feedforward control $u(t)$. As in the main text, the Hamiltonian form is convenient. The control Hamiltonian H is

$$H(\boldsymbol{x}, \boldsymbol{\lambda}, u) = \frac{1}{2}u^2 + \lambda_1 x_2 + \lambda_2(-x_1 + u)\,. \tag{S4}$$

The equation $\dot{\boldsymbol{x}} = \partial_{\boldsymbol{\lambda}} H$ reproduces Eq. (S3). The adjoint equation $\dot{\boldsymbol{\lambda}} = -\partial_{\boldsymbol{x}} H$ gives $\dot{\lambda}_1 = \lambda_2$ and $\dot{\lambda}_2 = -\lambda_1$. The compatibility equation $\partial_u H = 0$ leads to $u = -\lambda_2$. Substituting for u then gives coupled linear differential equations for the enlarged state $(x_1, x_2, \lambda_1, \lambda_2)^\mathsf{T}$ with the four boundary conditions $x_1(0) = x_0$ and $x_2(0) = x_1(\tau) = x_2(\tau) = 0$.

Figure S1 shows three examples of an unconstrained protocol and the resulting particle position, for $x_0 = 4$ and $\tau = 2\pi$, 7.5, and 11. Notice that in all three cases, the derivative is flat (the particle is stationary) at $t = 0$ and τ. It is easy to confirm that the control cost (shown in gray at bottom center of each plot) decreases asymptotically as $1/\tau$ [3, 4]. Notice that the control parameter (force) is not applied monotonically; rather, it mostly brakes the oscillatory motion, providing a controlled feedback damping for the otherwise-dissipationless system.

III. CONSTRAINED CONTROL

In a physical experiment, there will always be limits on the control input $u(t)$. The Pontryagin Minimum (Maximum) Principle (PMP) gives a systematic way to derive minimum-cost solutions in the presence of such constraints. As explained in the main text, if the value of $u(t)$ minimizing H lies within the allowed set, it can be determined from the compatibility condition using $\partial_u H(\boldsymbol{x}, \boldsymbol{\lambda}, u) = 0$. If not, one should check the boundary values of u and see which minimizes H. Here, we let the set of allowable (scaled) forces be $|u| \leq 1$ and check the cases $u = \pm 1$. Since $u^2 = 1$ in this case, the only relevant term in H is $\lambda_2 u$, implying $u = -\text{sign}(\lambda_2)$. Putting together the "interior" and boundary solutions, $u^*(t) = -\text{sat}(\lambda_2(t))$ is valid for all values of λ_2, where $\text{sat}(\cdot)$ clips its argument whenever its magnitude exceeds 1. Thus, in this case, the only change in the coupled state-adjoint system of equations is that the equation $\dot{x}_2 = -x_1 - \lambda_2$ becomes $\dot{x}_2 = -x_1 - \text{sat}(\lambda_2)$.

In Fig. S1(d–f), we plot the constrained solutions corresponding to the protocol durations used in (a–c), where there are three cases. For large τ (Case f), the solution is continuous and identical to the uncontrolled solution. Indeed, in this problem, as the protocol duration is extended, solutions with decreasing amplitude for $u(t)$ become possible. The amplitude scales as $u \sim 1/\tau$, meaning the cost scales as

$$J \sim \int_0^\tau \mathrm{d}t\, u^2 \sim (\tau)(1/\tau^2) = 1/\tau \,, \tag{S5}$$

as mentioned at the end of Sec. II. The costs J for the constrained solutions are equal or greater than the costs of the unconstrained solutions.

Figure S1d is an example of *bang-bang control*, since $u(t)$ alternates exclusively between its limits ± 1. The protocol duration $\tau = 2\pi$ was chosen to be the minimum at which there is a solution meeting the boundary conditions at $t = 0$ and τ. This could be found numerically by computing the solutions for different values of τ; alternatively, in Sec. IV we use the PMP to analytically find the minimum duration.

The last case, Fig. S1e, shows a complicated mixed solution, with repeated alternations between continuously varying solutions from the interior problem and boundary solutions.

A feature common to all the solutions is that there are discontinuities in either $u(t)$ or its temporal derivative $\dot{u}(t)$. These are typical in optimal-control problems, and it is worth understanding their origin and implications. We begin by noting that even the unconstrained solutions in Fig. S1(a–c) can show discontinuities at the initial or final times of the protocol. These are induced by the particular choice of boundary conditions. For example, at the end, we require that $x(\tau) = \dot{x}(\tau) = 0$. In a slightly modified problem, we might be satisfied if at $t = \tau$ the particle is *near* the origin and not moving too quickly, achievable via explicit end-time penalties. Alternatively, if we add a term $\sim x_1^2(t)$ to the running cost L, the final conditions become $\lambda_1(\tau) = \lambda_2(\tau) = 0$. This follows easily from an analysis using the calculus of variations for an augmented cost that adds the equations of motion as a constraint imposed by a Lagrange multiplier $\boldsymbol{\lambda}(t)$. More intuitively, for finite cost of ending up near (but not exactly at) the desired final state, it is better to shut off the control at the end, setting $u(\tau) = 0$: In a system with inertia, the benefit of a control applied at time t is only accrued at a later time, whereas the cost is applied at the current time. The exception is the final condition, which implies that the cost for violating a final condition is infinite. The control then seeks to satisfy the condition, whatever the future costs.

The other discontinuities present in $u(t)$ result from switching among interior and exterior solutions. Bang-bang control results from switching suddenly from one extreme to another. The slope discontinuities arise when passing from exterior to interior solutions for u, or vice versa.

Note that there is nothing physically implausible about a control that changes discontinuously. We expect that physical system states should change continuously, but the control we desire can jump "instantaneously." Of course, the control is implemented by a physical system that cannot change instantaneously, but often the times involved are sufficiently short to be negligible. For example, light intensity, which can affect optimal trapping strength or the rate of chemical reactions, can easily be varied at nanosecond to microsecond time scales. Those scales are often much faster than those of the physical system under control, and their effects can be neglected. If not, one can extend the physical model to include the actuator dynamics. Still, the input to the augmented system could change instantaneously.

Finally, we can confirm numerically that both the constrained *and* unconstrained solutions shown in Fig. S1 have Hamiltonians that are constant when evaluated on the solutions. That is, even for bang-bang solutions where $u(t)$ has jump discontinuities and the simple calculus-of-variations argument given above breaks down, H nonetheless remains constant.[1]

IV. MINIMUM-TIME CONTROL

A natural application of optimal control is to find how to accomplish a task as fast as possible. For linear dynamics with additive control, such problems have bang-bang solutions set by the actuator limits. In the example discussed here, we find the minimum-time solution by changing the running cost from $\frac{1}{2}u^2$ to 1, meaning that the cost $J = \tau$. The control Hamiltonian becomes

$$H(\boldsymbol{x}, \boldsymbol{\lambda}, u) = \tau + \lambda_1 x_2 + \lambda_2(-x_1 + u)\,. \tag{S6}$$

Since H is now linear in u, the solution always has the bang-bang form. Proceeding as with the original problem (but reverting to second-order notation for simplicity), we solve the system of equations

$$\ddot{x} + x = -\mathrm{sign}(\dot{\lambda}), \qquad \ddot{\lambda} + \lambda = 0\,, \tag{S7}$$

with boundary conditions $x(0) = x_0$, $\dot{x}(0) = x(\tau) = \dot{x}(\tau) = 0$. Here, τ is unknown, and we seek the smallest τ for which a solution exists. It is easy to verify that, for $x_0 = 4$, the solution is $\tau = 2\pi$. The control $u^*(t) = -1$ up to $t = \pi$ and $+1$ between $t = \pi$ and 2π. The particle trajectory

$$x^*(t) = \begin{cases} 1 + 3\cos t\,, & 0 \le t < \pi\,, \\ -1 + \cos t\,, & \pi \le t \le 2\pi\,, \end{cases} \tag{S8}$$

matches the one found numerically in Fig. S1d. Notice that the value of the optimal cost found numerically, ≈ 3.14, is consistent with $J^* = \int_0^{2\pi} \mathrm{d}t\, \frac{1}{2}(\mp 1)^2 = \pi$. The form $x^*(t)$ of the bang-bang state trajectory is pleasing when plotted in phase space: the motion consists of circular arcs, with the first half centered on $+1$ and the second on -1, joining at $t = \pi$ (Fig. S2).

[1] More precisely, if $u^*(t)$ has a jump discontinuity at $t = t_{\text{jump}}$, then $H(t \to t^-_{\text{jump}}) = H(t \to t^+_{\text{jump}})$.

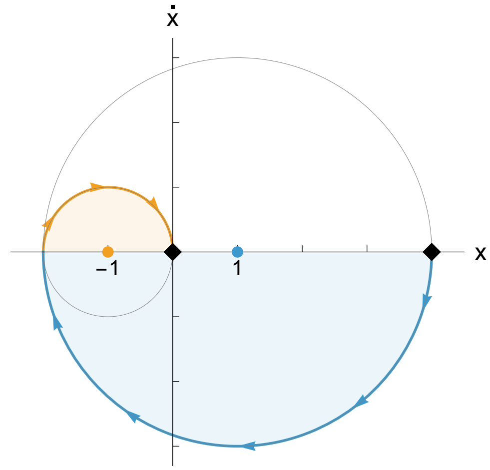

FIG. S2. Phase-space plot of the bang-bang solution plotted in Fig. S1d. Black diamonds denote the start and end states. The first stage (blue), from $t = 0$ to π is a circle of radius 3, centered on $x = +1$. The second stage (orange), from $t = \pi$ to 2π, is a circle of radius 1, centered on $x = -1$.

V. OPTIMAL CONTROL: FEEDBACK FORM

The above discussion illustrated how one can find feedforward solutions $u^*(t)$ to optimal-control problems. As discussed in the main text, the Hamilton-Jacobi-Bellman equation in principle leads to feedback solutions of the form $u^*(x(t), t)$. For linear dynamics, the solution can be carried through (semi-) analytically and reproduces the feedforward solution. For simplicity, we assume a continuous, interior solution, but the generalization to constrained control follows a path similar to that taken for the feedforward case.

It will be just as easy to do the derivation for general linear dynamics $\dot{\boldsymbol{x}} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{B}\boldsymbol{u}$ and general quadratic running cost $L(\boldsymbol{x}, \boldsymbol{u}) = \frac{1}{2}(\boldsymbol{x}^\mathsf{T}\boldsymbol{Q}\boldsymbol{x} + \boldsymbol{u}^\mathsf{T}\boldsymbol{R}\boldsymbol{u})$. From the main text, the HJB equation for linear dynamics is

$$\inf_{\boldsymbol{u}} \left[L(\boldsymbol{x}, \boldsymbol{u}) + (\partial_{\boldsymbol{x}} J^*)\ \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{u})\right] = 0\,, \tag{S9}$$

where $J^*(\boldsymbol{x}, t)$ is the optimal cost-to-go function starting from $\boldsymbol{x}$ at time t. Then,

$$\inf_{\boldsymbol{u}} \left[\tfrac{1}{2}(\boldsymbol{x}^\mathsf{T}\boldsymbol{Q}\boldsymbol{x} + \boldsymbol{u}^\mathsf{T}\boldsymbol{R}\boldsymbol{u}) + (\partial_{\boldsymbol{x}} J^*)\ (\boldsymbol{A}\boldsymbol{x} + \boldsymbol{B}\boldsymbol{u})\right] = -\partial_t J^*\,. \tag{S10}$$

The suggested *ansatz* $J^* = \frac{1}{2}\boldsymbol{x}^\mathsf{T}\boldsymbol{S}(t)\boldsymbol{x}$ implies that $\boldsymbol{S}$ is symmetric, since any antisymmetric component will contribute 0 to J^*. Thus,

$$(\partial_{\boldsymbol{x}} J^*) = \boldsymbol{x}^\mathsf{T}\boldsymbol{S}\,. \tag{S11}$$

Hence,

$$\inf_{\boldsymbol{u}} \left[\tfrac{1}{2}(\boldsymbol{x}^\mathsf{T}\boldsymbol{Q}\boldsymbol{x} + \boldsymbol{u}^\mathsf{T}\boldsymbol{R}\boldsymbol{u}) + \left(\boldsymbol{x}^\mathsf{T}\boldsymbol{S}\right)\ (\boldsymbol{A}\boldsymbol{x} + \boldsymbol{B}\boldsymbol{u})\right] = -\tfrac{1}{2}\boldsymbol{x}^\mathsf{T}\dot{\boldsymbol{S}}\boldsymbol{x}\,. \tag{S12}$$

Because $\boldsymbol{u}$ is unbounded, the infimum is found by taking $\partial_{\boldsymbol{u}}$ and setting to zero:

$$\boldsymbol{u}^{\mathsf{T}}\boldsymbol{R} + \boldsymbol{x}^{\mathsf{T}}\boldsymbol{S}\,\boldsymbol{B} = \boldsymbol{0}^{\mathsf{T}} \,. \tag{S13}$$

Taking a transpose and remembering that $\boldsymbol{R}$ and $\boldsymbol{S}$ are symmetric gives

$$\boldsymbol{u} = -(\boldsymbol{R}^{-1}\boldsymbol{B}^{\mathsf{T}}\boldsymbol{S})\boldsymbol{x} \equiv -\boldsymbol{K}^{\mathsf{T}}\boldsymbol{x} \,. \tag{S14}$$

Substituting $\boldsymbol{u}$ back into the HJB, Eq. (S9), gives

$$\tfrac{1}{2}\left[\boldsymbol{x}^{\mathsf{T}}\boldsymbol{Q}\boldsymbol{x} + (\boldsymbol{x}^{\mathsf{T}}\boldsymbol{S}\boldsymbol{B}\boldsymbol{R}^{-1})\boldsymbol{R}(\boldsymbol{R}^{-1}\boldsymbol{B}^{\mathsf{T}}\boldsymbol{S}\boldsymbol{x})\right] + \left(\boldsymbol{x}^{\mathsf{T}}\boldsymbol{S}\right)\left[\boldsymbol{A}\boldsymbol{x} - \boldsymbol{B}(\boldsymbol{R}^{-1}\boldsymbol{B}^{\mathsf{T}}\boldsymbol{S})\boldsymbol{x}\right] = -\tfrac{1}{2}\boldsymbol{x}^{\mathsf{T}}\dot{\boldsymbol{S}}\boldsymbol{x} \tag{S15a}$$

$$\boldsymbol{x}^{\mathsf{T}}\left[\tfrac{1}{2}\left(\boldsymbol{Q} + \boldsymbol{S}\boldsymbol{B}\boldsymbol{R}^{-1}\boldsymbol{B}^{\mathsf{T}}\boldsymbol{S}\right) + \tfrac{1}{2}\left(\boldsymbol{S}\boldsymbol{A} + \boldsymbol{A}^{\mathsf{T}}\boldsymbol{S}\right) - \boldsymbol{S}\boldsymbol{B}\boldsymbol{R}^{-1}\boldsymbol{B}^{\mathsf{T}}\boldsymbol{S}\right]\boldsymbol{x} = -\tfrac{1}{2}\boldsymbol{x}^{\mathsf{T}}\dot{\boldsymbol{S}}\boldsymbol{x} \,, \tag{S15b}$$

which implies that $\boldsymbol{S}$ obeys the (matrix) Riccati equation,

$$\boldsymbol{Q} + \boldsymbol{S}\boldsymbol{A} + \boldsymbol{A}^{\mathsf{T}}\boldsymbol{S} - \boldsymbol{S}\boldsymbol{B}\boldsymbol{R}^{-1}\boldsymbol{B}^{\mathsf{T}}\boldsymbol{S} = -\dot{\boldsymbol{S}} \,. \tag{S16}$$

Note the decomposition $\boldsymbol{S}\boldsymbol{A} \to \frac{1}{2}(\boldsymbol{S}\boldsymbol{A} + \boldsymbol{A}^{\mathsf{T}}\boldsymbol{S})$, which follows because the condition $\boldsymbol{x}^{\mathsf{T}}[\cdots]\boldsymbol{x} = 0$ implies that the *symmetric* part of the bracketed terms $[\cdots]$ equals zero. The linear combination isolates the symmetric part of $\boldsymbol{S}\boldsymbol{A}$. There is no constraint placed on antisymmetric terms.

In the problem here, the dynamics is $\boldsymbol{A} = \left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)$, the input coupling is $\boldsymbol{B} = \left(\begin{smallmatrix} 0 \\ 1 \end{smallmatrix}\right)$, the cost factors are $\boldsymbol{Q} = \boldsymbol{0}$, and $\boldsymbol{R} = 1$. The Riccati equation (S16) then reduces to three equations for the components of the symmetric matrix $\boldsymbol{S}$:

$$\dot{s}_{11} = 2s_{12} + s_{12}^2 \tag{S17a}$$

$$\dot{s}_{22} = -2s_{12} + s_{22}^2 \tag{S17b}$$

$$\dot{s}_{12} = -s_{11} + s_{22} + s_{12}s_{22} \,, \tag{S17c}$$

with final conditions $s_{11}(\tau) = s_{22}(\tau) = \infty$ and $s_{12}(\tau) = 0$. The infinite costs arise because of the final conditions placed on $\boldsymbol{x}(\tau)$; if deviations from the desired state instead have finite costs, the final condition would be $\boldsymbol{S}(\tau) = \boldsymbol{0}$. Solving these coupled nonlinear equations (the *Matrix Riccati equation*) numerically and setting $u^* = -\boldsymbol{K}^{\mathsf{T}}\boldsymbol{x}$ with $\boldsymbol{K}^{\mathsf{T}} = \boldsymbol{B}^{\mathsf{T}}\boldsymbol{S}$, we find precisely the different unconstrained u^* shown in Fig. S1a–c,f.

VI. HEURISTIC (PD) FEEDBACK

The form of the feedback solution found above is that of a time-dependent linear feedback for the two-component vector $\boldsymbol{K}^{\mathsf{T}} \equiv (K_p\ K_d)$. Writing out the feedback law in components and reverting to $x_1 \to x$ and $x_2 \to \dot{x}$ gives

$$u(x, \dot{x}, t) = -K_p(t)x - K_d(t)\dot{x} \,. \tag{S18}$$

Equation (S18) has the form of proportional-derivative (PD) feedback, one of the heuristic algorithms discussed in the main text. The important difference is that the heuristic PD control has constant gains K_p and K_d; by contrast, the optimal control has time-dependent gains.

Figure S3 compares the optimal and heuristic feedback control algorithms on our example oscillator, for $x_0 = 4$ and $\tau = 11$. The solid curves in Fig. S3a reproduce the results from Fig. S1c,f, but they were calculated in a completely different way, from the Riccati equation (S17c). Figure S3b

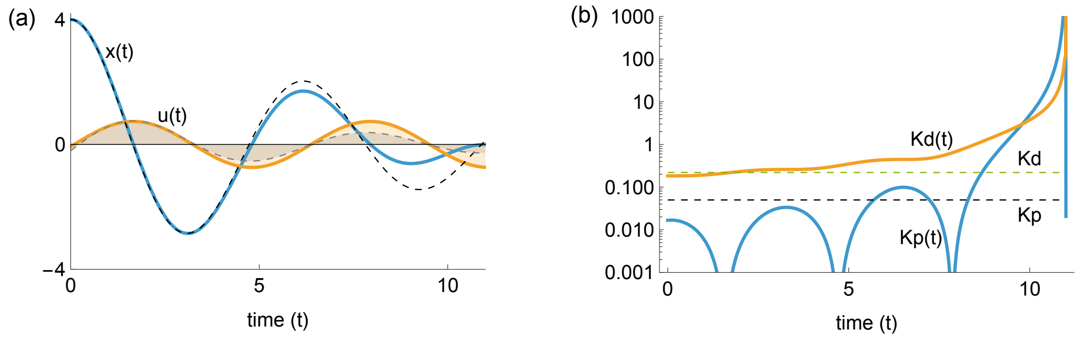

FIG. S3. Comparison of optimal and heuristic feedback control. (a) Solid blue and orange curves replicate using feedback the feedforward curves of Fig. S1c, with $x_0 = 4$ and $\tau = 11$. The dashed curves use proportional-derivative (PD) control. (b) Feedback gains for optimal and heuristic control. The time-dependent solid curves result from the optimal-control calculation, Eqs. (S14) and (S16). The horizontal dashed lines are empirically tuned, constant PD gains K_p and K_d.

shows the time-dependent gains $K_p(t)$ and $K_d(t)$. Note how they diverge at the end of the protocol, when $t \to \tau$.

To understand the divergence, Fig. S3b also shows the result of the heuristic PD control, using constant gains K_p and K_d tuned so that the initial decay of $x(t)$ follows the optimal-control solution. The agreement is reasonable until $t \approx 4$ and then diverges increasingly. The action of the PD control is easy to understand, if we substitute the feedback $u = -K_p x - K_d \dot{x}$ into the equations of motion, $\ddot{x} + x = u$. The closed-loop equations are

$$\ddot{x} + K_d \dot{x} + (1 + K_p) x = 0 \,. \tag{S19}$$

Thus, proportional gain speeds up the natural frequency by a factor $\sqrt{1 + K_p}$, while the derivative gain introduces damping. The result is a damped harmonic oscillator, whose amplitude oscillates and decays, within an envelope $\mathrm{e}^{-(K_d/2)t}$. However, to stop at finite time (reach $x = \dot{x} = 0$ at time τ), the gain must diverge: a constant-gain control cannot stop a particle in finite time. If the protocol time is long, the exponential decay means that there is little difference between the protocols, but at shorter times there is a large difference.

Returning to optimal control, one might wonder, If the feedback and feedforward solutions are identical, why seek feedback solutions? Given perfect knowledge of dynamics and absent unexpected disturbances, there would be no difference; however, feedback can deal much more robustly with these uncertainties. Indeed, the best control strategies typically combine feedforward and feedback: feedforward uses the model to implement the optimal control; the feedback around the feedforward solution then deals with the unknowns of modeling errors in the dynamics and external disturbances.

VII. RELATIONS TO STOCHASTIC, SOFT, AND ACTIVE-MATTER PROBLEMS

Above, we chose the problem of moving a harmonic oscillator by an applied force because it was a simple physical setting that can be intuitively understood. This setting is very close to the experiment of Le Cunuder, et al., who studied experimentally an atomic force microscope cantilever whose end position is controlled by forces generated by an external electric potential [5]. Compared to our toy problem here, the experimental system is subject both to a linear damping term $\gamma \dot{x}$ and to thermal fluctuations, modeled as white noise $\eta(t)$ that has zero mean and variance $\langle \eta(t)\eta(t') \rangle = 2\gamma k_\mathrm{B} T \delta(t-t')$. Here γ is the linear damping coefficient, T is the environment temperature, and k_B is the Boltzmann constant. The authors did not attempt to minimize the (average) work required to move the cantilever but instead used inverse engineering to find a protocol $u(t)$ that moved the system from an initial equilibrium state at x_0 to a final equilibrium state at x_τ. Because the equations are linear and the noise Gaussian, the distributions of all quantities remain Gaussian and are thus specified entirely by their mean and variance. Thus, relative to the problem considered above in this Supplement, one needs to match not only the mean positions and velocities at start and end but also the initial and final variances of these two quantities. The authors chose a polynomial whose order matched the number of boundary-condition constraints.

Second-order linear systems have also been used to model several other experimental situations. For example, Loos, et al. examined, theoretically and in experiment, the transport of a colloidal particle by a moving harmonic potential in a viscoelastic medium modeled as a Maxwell fluid with a single time constant [6]. The system was then described by

$$\tau_p \dot{x}_p = -\frac{\kappa}{\kappa_b}(x_p + u) - (x_p - x_b) + \xi_p \tag{S20a}$$

$$\tau_b \dot{x}_b = -(x_b - x_p) + \xi_b \,, \tag{S20b}$$

where quantities subscripted by p denote the colloidal particle and quantities denoted by b denote a fictitious "bath particle" that models the viscoelastic response of the fluid medium. The dynamics for x_p are that of an overdamped particle, where inertial effects relax so quickly that they may be neglected. The bath is modeled as another first-order system. Together, the two form a second-order system of equations with two independent noise sources. Similar equations also describe active-matter models such as Active Ornstein-Uhlenbeck Particles (AOUP), where a colloidal particle moves in a simple fluid that is nonetheless subject to active fluctuations in the bath [7–9].

One other generalization that arises when we consider soft-matter problems is that the cost function J may have a more complicated form than considered above. For example, let us consider a colloidal particle moving in a harmonic potential in one dimension. For the deterministic problem, we used $\frac{1}{2}u^2(t)$ as an instantaneous control "cost" that was to be minimized. For the stochastic problem, the mean work required to move a harmonic oscillator in a time τ is given by [10]

$$W = \int_0^\tau \mathrm{d}t \, \dot{u}(u - \langle x \rangle) \,, \tag{S21}$$

where $\langle x \rangle$ is the time-dependent mean position of the particle, averaged over an ensemble of systems in a fluctuating thermal bath. The expression for the mean work W can be manipulated to be closer to our $\int \mathrm{d}t \, \frac{1}{2}u^2$ using integration by parts. The result includes boundary terms ("terminal costs") and quadratic combinations of u and $\langle x \rangle$. Thus, the cost function falls into the same class as those considered here.

---

[1] J. Bechhoefer, *Control Theory for Physicists* (Cambridge Univ. Press, 2021).
[2] D. E. Kirk, *Optimal Control Theory: An Introduction* (Dover Publications, 1998).
[3] K. Sekimoto and S.-I. Sasa, Complementarity relation for irreversible process derived from stochastic energetics, J. Phys. Soc. Jpn. **66**, 3326 (1997).

[4] T. Schmiedl and U. Seifert, Optimal finite-time processes in stochastic thermodynamics, Physical Review Letters **98**, 108301 (2007).
[5] A. Le Cunuder, I. A. Martínez, A. Petrosyan, D. Guéry-Odelin, E. Trizac, and S. Ciliberto, Fast equilibrium switch of a micro mechanical oscillator, Appl. Phys. Lett. **109**, 113502 (2016).
[6] S. A. M. Loos, S. Monter, F. Ginot, and C. Bechinger, enUniversal Symmetry of Optimal Control at the Microscale, Physical Review X **14**, 021032 (2024).
[7] G. Szamel, Self-propelled particle in an external potential: Existence of an effective temperature, Phys. Rev. E **90**, 012111 (2014).
[8] D. Martin, J. O'Byrne, M. E. Cates, É. Fodor, C. Nardini, J. Tailleur, and F. van Wijland, Statistical mechanics of active Ornstein-Uhlenbeck particles, Phys. Rev. E **103**, 032607 (2021).
[9] T. K. Saha, J. Ehrich, M. Gavrilov, S. Still, D. A. Sivak, and J. Bechhoefer, Information engine in a nonequilibrium bath, Physical Review Letters **131**, 057101 (2023).
[10] A. Gomez-Marin, T. Schmiedl, and U. Seifert, Optimal protocols for minimal work processes in underdamped stochastic thermodynamics, J. Chem. Phys. **129**, 024114 (2008).