

## Predictive Maxwell's demons

Nathaniel Rupprecht  and Dervis Can Vural   
 University of Notre Dame, South Bend, Indiana 46556, USA

 (Received 26 May 2020; accepted 7 December 2020; published 28 December 2020)

Here we study the operation efficiency of a finite-size finite-response-time Maxwell's demon, who can make future predictions. We compare the heat and mass transport rate of predictive demons to nonpredictive ones and find that predictive demons can achieve higher mass and heat transport rates over longer periods of time. We determine how the demon performance varies with response time, future sight, and the density of the gasses on which they operate.

DOI: [10.1103/PhysRevE.102.062145](https://doi.org/10.1103/PhysRevE.102.062145)

### I. INTRODUCTION

A Maxwell's demon is a device that uses information about the microstates of system to extract work from or reduce entropy in a system [1]. Maxwell's demons have served as a pivotal thought experiment to establish the close relationship between thermodynamics and information processing. The original thought experiment, which involved a demon sorting gas molecules according to their energy, has recently been expanded to include feedback control [2,3] and universal computation [4–6]. Some authors describe the computational process within the demon as manipulations of a tape bits [7–10] or qubits [11], while others consider demons that manipulate microstates by fully mechanical means [12–14]. Models of nonideal demons [7,8] account for the thermal equilibration of the demon with the system, the demon's finite mass [14], or finite size and response time [15]. Recently experimental constructions have also become feasible [16–27] and have been suggested as a practical means to produce energy [24,28–32] or sort atoms [33].

The close connection between information processing and entropy generation motivates us to ask if there might also be thermodynamic implications of predicting the future. Imagine two Maxwell's demons, each with finite response time, where one makes decisions based only on immediate molecular arrivals, whereas the other one predicts future arrivals and plans ahead. How much, if any, will a difference will this make in the rate at which they pump heat and mass?

Knowing the future would fundamentally change the actions of a demon. A predictive demon could make trade-offs such as rejecting a desirable particle if it knew that doing so would also keep out a large number of undesirable particles that would arrive afterwards, before the demon would be able to close the gate to keep them out. Previously, we studied how the finite size and response time constrains a demon's heat and mass transport rate [15]. However, our model demon only used local information, and did not take into consideration future arrivals.

In the present study, we establish heat and mass transport limitations of a predictive Maxwell's demon with finite size

and response time and compare these limitations to that of a nonpredictive demon. This way, we aim to begin exploring the thermodynamic consequences of prediction making.

### II. NONPREDICTIVE DEMONS

Here we start with a brief overview of the demon model with finite size  $A$  and response time  $\tau$ , which will be identical for both the predictive and non-predictive cases.

Our system consists of left and right subsystems of ideal gas with volumes  $V_l, V_r$ , energies  $E_l, E_r$ , numbers  $N_l, N_r$  of particles of mass  $m$ . Throughout, we will subscript variables that refer to the left and right subsystem by  $l$  and  $r$ , and variables that depend on a generic side with  $s$ , e.g.,  $E_s$  for the energy of the  $s$ -side, with  $s \in \{l, r\}$ . We theoretically consider  $d = 1, 2$ , or 3-dimensional systems, but only simulate  $d = 2$ . The subsystems are separated by a gate of area  $A$  ( $A \equiv 1$  for a one-dimensional system and to be the length of the gate for a two-dimensional system) controlled by the demon. We assume that the subsystems are large enough that each subsystem acts as a self-averaging canonical distribution. We take  $N_s, E_s \rightarrow \infty$  with fixed  $\rho_s \equiv N_s/V_s$  and  $\bar{E}_s \equiv E_s/V_s$ , where the temperature,  $1/\beta_s$ , and energy per particle,  $\bar{E}_s = E_s/N_s$ , are related by  $\beta_s E_s = N_s d/2 \equiv \beta_s N_s \bar{E}_s$ , as required by the equipartition theorem.

We analyze two types of demons. (1) An *energy* demon opens the gate whenever the net flow of energy from left to right would be positive. (2) A *number* demon opens the gate whenever the net particle transfer from left to right is positive.

The probability that  $n$  particles are incident on the gate during a length of time  $\tau$  follows a Poisson distribution with a Poisson parameter,

$$\kappa_s = \frac{\rho_s \tau A}{\sqrt{2\pi \beta_s m}} = \rho_s \tau A \sqrt{\frac{\bar{E}_s}{d \pi m}} \equiv v_s \tau. \quad (1)$$

The rates  $v_l, v_r$  are crucial in characterizing the demons' performance. For example, the average energy and number currents for nonpredictive demons with small response time

$\tau$  are [15]

$$\dot{N}_\tau^{(d)} = \frac{d+1}{2\beta_l} v_l e^{-v_l \tau} + \mathcal{O}[\tau e^{-(v_l+v_r)\tau}]$$

$$\dot{N}_\tau^{(d)} = v_l e^{-v_l \tau} + \mathcal{O}[\tau e^{-(v_l+v_r)\tau}].$$

As mentioned, a nonpredictive demon simply “seizes the day” by making decisions only according to the mass or energy flux during a present time interval  $\tau$ . We will see that a predictive demon, on the other hand, will be able to forgo short-term success for higher average transmission rates in the long run. It is known that having to sacrifice short-term gains can be beneficial in the long run in similar scenarios, for example, see Refs. [2,34].

### III. PREDICTIVE DEMONS

We assume that a predictive demon can know when particles will be incident on the gate area from the left or right subsystem during the next (possibly infinite) amount of time,  $t_c$ . We call  $t_c$  the clairvoyance time of the demon and consider demons that run for a finite amount of time,  $T$ . We divide the total number of particles transported by the demon by  $T$  to obtain the time averaged current. In summary, the three timescales in the system are the clairvoyance time of the demon,  $t_c$ , its reaction time,  $\tau$ , and the total time of operation,  $T$ .

Given a sequence of times at which individual particles hit the gate area from the left or right, the predictive demon must determine the optimal sequence of gate openings and closings, subject to the response time constraint, i.e., the gate cannot change its open or close state faster than  $\tau$ .

We assume that particle arrivals to the gate area are independent events, and we assign a “score” to each arrival according to a suitable characteristic of the particle, which the demon aims to maximize. For a number demon, the score will be 1 and  $-1$  for arrivals from the left and right sides of the gate. For an energy demon, the score associated with a particle arrival is the energy of the particle, again with a sign that depends on its direction.

To quantify how well the predictive demon operates, we will compare it to the nonpredictive demon. To do so, we first allow the nonpredictive demon to decide what offset  $0 \leq t_0 < \tau$  its time division should have. When the total time of the simulation is long, the offset does not effect the average score. The demon will then bin events into the bins  $[k\tau - t_0, (k+1)\tau - t_0]$ ,  $k \geq 0$ , and decide for each bin whether it is better for the gate to be open or closed. Figure 1 illustrates how the two types of demons will treat the same set of events and why the predictive demon is able to achieve better performance than the nonpredictive demon.

### IV. SIMULATING A PREDICTIVE DEMON

We created a program that determines the schedule of gate openings and closings for predictive and nonpredictive demons for a given set of random left and right particle arrivals with Boltzmann-distributed energies [35]. To implement the scheduling, the predictive demon divides time into very small *microbins* and then converts scheduling into a discrete problem by placing particle arrival events into these

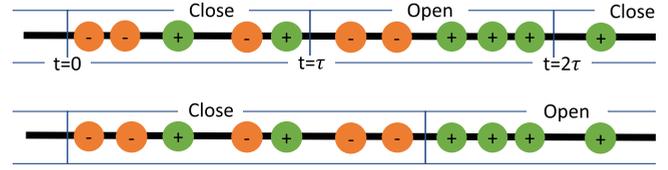


FIG. 1. A schematic example illustrating the difference between a nonpredictive demon (top row) and a predictive (bottom row) demons, both with response time  $\tau$ . The nonpredictive demon decides on whether to open or close the gate based on tallying the desirable (green plus) and undesirable (red minus) particles within the present  $\tau$  interval. In contrast, the predictive demon optimizes its open or close intervals globally, while still satisfying the constraint that the state of its gate cannot change faster than  $\tau$ . In this example, the nonpredictive demon generates a net current of one particle (three particles in the desired direction and two particles in the undesired direction while the gate is open), whereas the predictive demon transports 4 particles within the same time.

microbins. We call the number of microbins per  $\tau$  the *resolution*,  $g$ , of the simulation. We take the large- $g$  limit to approximate the continuous time problem. More details on the simulation can be found in Appendix B. We have found that demon performance does not change much beyond  $g = 50$  (see Appendix C).

Using our program, we numerically study how the efficiency of predictive demons change and compares to the nonpredictive demons as we change subsystem temperatures, number densities, and demon response time. We do so by generating a realization of particle arrival times and scores and pass this to the scheduler program. The energy of a particle given that it hits the gate area follows a Gamma distribution (cf. Appendix A and [15]),

$$P(E) = \frac{\beta_s^{(d+1)/2} E^{(d-1)/2} e^{-\beta_s E}}{\Gamma[(d+1)/2]}, \quad (2)$$

which we use to assign energy scores.

The *performance* of the number and energy demons is quantified by the average mass and heat transfer (average score per unit time) they can achieve.

### V. THERMODYNAMIC CONSEQUENCES OF PREDICTION MAKING

We start by studying the basic properties of a who can see infinitely far into the future ( $t_c \rightarrow \infty$ ) and determine how it compares to a nonpredictive demon.

First, we show how the demons’ performance varies with the system parameters (Fig. 2). The number density of the left subsystem is fixed at  $\rho_l = 50$ . We find that for both number and energy demons, the predictive demon has the maximal advantage over the nonpredictive demon when the right subsystem is denser, but not too much denser, than the left subsystem.

Second, we study how the demons’ performance changes as the total time they must schedule over,  $T$ , increases (again,  $t_c \rightarrow \infty$ ). Intuitively, as  $T$  increases, both demons should do worse. To see why, suppose that the average score on an interval of time of length  $t$  is  $S_t$ . Then the average performance

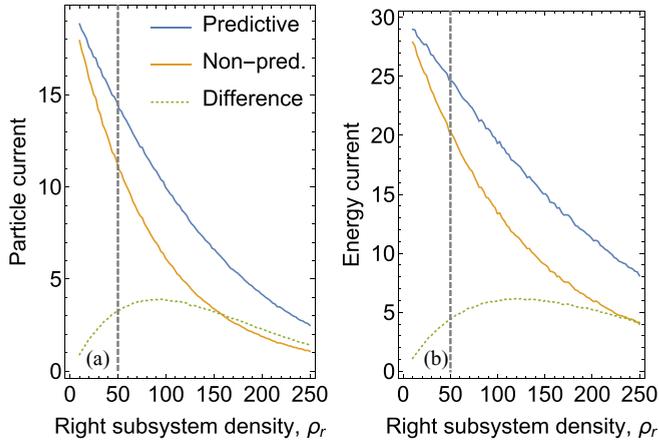


FIG. 2. Performance versus density. The performance of number demons (a) and energy demons (b) for predictive and nonpredictive (orange) cases are plotted as a function of right subsystem density, while the left subsystem density is kept constant (vertical dashed line), averaged over 500 particle arrival sequences.

on two separate intervals of length  $t$  is just  $2S_r$ . However, if these time intervals are not separated, but contiguous, then the “boundary conditions” of the best schedules on each time interval will in general not align.

The demon performance as a function of total time  $T$  is shown for several different right subsystem parameters in Fig. 3. The heat and mass currents for the predictive demons quickly asymptote (blue curves), whereas that for the nonpredictive demons asymptote far later (red curves). A predictive demon’s advantage over the nonpredictive demon continues to grow long after the change in its performance has stabilized. We also see that while the performance of both demons drop when the density of the right subsystem increases, the difference in performance increases for both the number and energy demons.

If we had restricted the nonpredictive demon so that its binning offset,  $t_0$ , was zero, then changing the length of time it must operate over would not change the performance—it

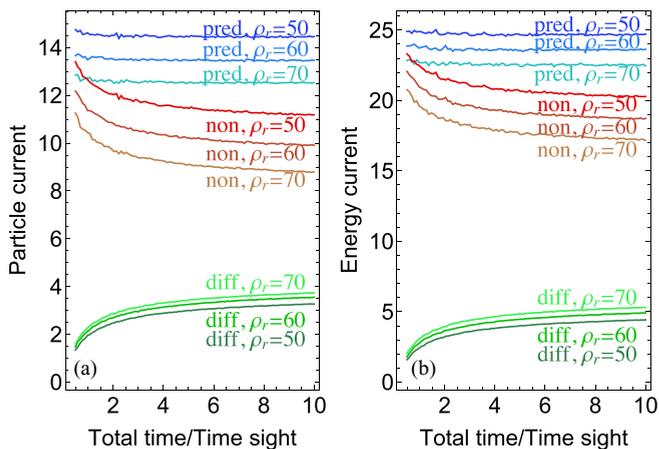


FIG. 3. Performance vs total time. Performance was average over 5000 random sets of events. (a) Predictive (blue) and nonpredictive (red) demons and their difference (green) is plotted for number (a) and energy demons (b).

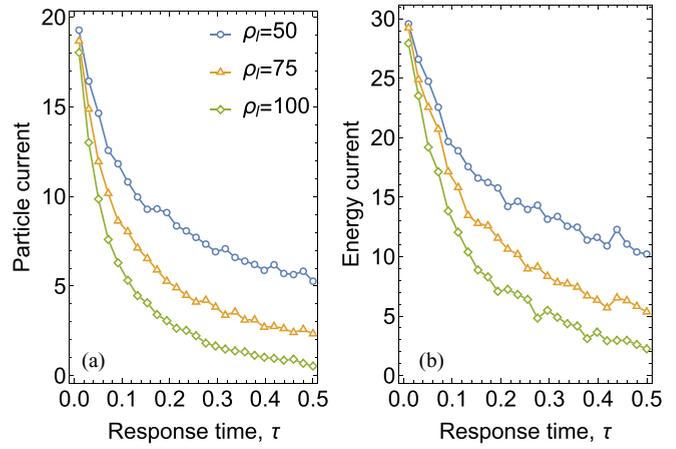


FIG. 4. Number and energy currents generated by predictive demons, as a function of response time. For all runs, the total time was  $T = 10$ , and the time sight was  $t_c = 1$ . (a) The performance of a number demon as response time changes. (b) The performance of the energy demon as response time changes.

performance would always be at a minimum. Clearly, since the performance of the nonpredictive demon decreases by a large amount as  $t_{\text{sched}}$  (see Fig. 3), it is a big advantage for the nonpredictive demon to be able to determine its time offset, at least for short times. Indeed, in the limiting case where  $t_{\text{sched}} \leq \tau$ , both types of demons operate in exactly the same way.

Third, we move on to demons who have finite future sight  $t_c$ , operating for a long time  $T \gg t_c$ . Simulating this demon is more complicated than simply scheduling gate openings and closings over a fixed amount of time that contains all the events. Since the demon only has knowledge of the events for a small part of the total time, it must revise its belief of what the best course of action is whenever it becomes aware of new information. The demon also must contend with the fact that the gate will most likely have to remain open or closed for some period of time beyond  $t_c$ . If the subsystems parameters are such that it is very likely that the net flow of particles will be negative during this time, then the demon should consider choosing a schedule that results in a lower score during the interval  $[0, t_c]$  but does not require the gate to be open as long during the period after  $t_c$ .

Since part of the schedule extends into a time the demon cannot see, the score is now a random variable that can be decomposed into a sum of the (deterministic) score  $X_d$ , of the schedule in the interval  $[0, t_c]$ , and of the (probabilistic) score  $\hat{X}_r$  of the part of the schedule beyond  $t_c$ .

When calculating the projected score for the schedule, we take the expectation value,  $X_d + \langle \hat{X}_r \rangle$ , to determine which schedule is best. We have observed that not taking the random part of the scheduling score into account results in very low or even negative demon performance as  $\tau$  increases.

The performance of demons with fixed time sight ( $t_c = 1.0$ ) and varying response times is shown in Fig. 4 and is similar to that in Ref. [15], where for  $\tau \approx 0$ , the rates decrease exponentially, but as  $\tau$  increases, additional terms of higher orders slow down the decrease. We also check what happens when the time sight is kept fixed and the response time varies

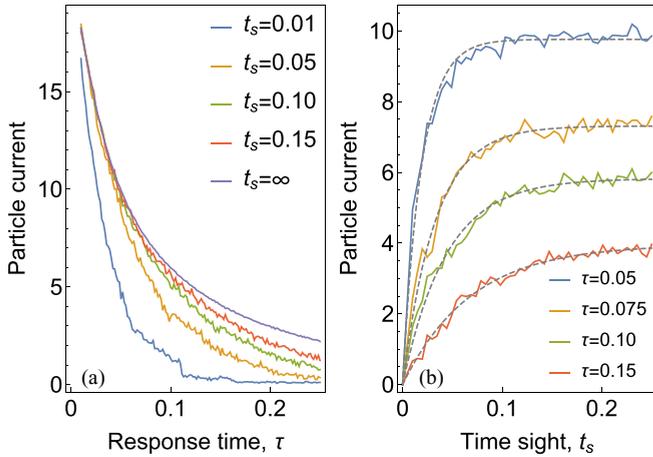


FIG. 5. Number demon performance, varying response time and time sight. The systems had parameters  $\rho_l = 50$ ,  $\rho_r = 100$ . (a) Number rate vs response time for several fixed time sights. An averaging of 500 runs with  $t = 1$  second was used. (b) Number rate vs time sight for several fixed response times. The best fitting functions of the form  $\dot{N}_\tau(t_c) = a_\tau(1 - e^{-\lambda_\tau t_c})$  for each  $\tau$  are shown in gray.

[Fig. 5(a)] and when the response time is fixed and the time sight varies [Fig. 5(b)]. The qualitative behavior of predictive energy demons is similar (see Fig. 4).

In Fig. 5(a), we see that having larger time sight is strictly better. Figure 5(b) shows demons with fixed response time as  $t_c$  is varied. Of course, faster response time is strictly better. While most of the curves that we found in Ref. [15] and here, including  $\dot{N}_\tau(\infty)$ , are “smoothly decreasing,” they cannot be fit to simple exponential or rational functions. Interestingly, while the infinite time sight rates,  $\dot{N}_\tau(\infty)$ , are nontrivial functions, the correction for finite time sight appears to be quite simple,

$$\dot{N}_\tau(t_c) = (1 - e^{-\lambda_\tau t_c})\dot{N}_\tau(\infty). \quad (3)$$

Despite our best efforts, we were not able to derive this elementary relationship.

## VI. ANALYTICAL CONSIDERATIONS

Assuming that a demon has  $t_c = \infty$  and a total scheduling time  $T$ , for any response time  $\tau$ , microbin time  $dt = \tau/g$ , and a random variable  $\hat{v}$  that encodes the distribution and score of events, we can define the discrete Scheduling process,  $\hat{\mathbf{S}}_t = \hat{\mathbf{S}}_t[\tau, dt, \hat{v}, T] \equiv \{\hat{S}_t, \hat{S}_t^*\}$ , which is the best score possible during the remaining time in the scheduling interval given that the door was open for at least the previous  $\tau$  or closed for at least the previous  $\tau$  ( $\hat{S}_t$  and  $\hat{S}_t^*$ , respectively). The scheduling process is defined by the set of equations

$$\hat{S}_t = \max(\hat{S}_{t+dt} + \hat{v}_t(dt), \hat{S}_{t+\tau}^*), \quad (4)$$

$$\hat{S}_t^* = \max(\hat{S}_{t+\tau} + \hat{v}_t(\tau), \hat{S}_{t+dt}^*), \quad (5)$$

where  $\hat{v}_t(\Delta t)$  is the event score between  $t$  and  $t + \Delta t$ .

If we could solve for the average growth rate,  $\langle \hat{S}_T/T \rangle$ , of these equations as  $T \rightarrow \infty$ , then we would obtain the performance of a predictive demon with infinite  $t_c$ ; and since we

know from (3) that the finite sight demon is related to the infinite sight demon in a simple way, we would essentially have a full solution to the behavior of cyclically operating predictive demons (though we would have to find a way to obtain  $\lambda_\tau$ ). We could also take  $dt \rightarrow 0$  to obtain the continuous time limit.

Equations like (4) and (5) are called *Bellman equations*, used extensively in mathematical optimization and dynamic programming [36,37]. These equations break the problem down into recursively computable pieces and are solved backwards in time. They have also been used in statistical physics to, e.g., find globally optimal control protocols for Brownian particles in a flashing ratchet [38], and where it was also found that using Bellman’s principle beats a local greedy algorithm, or to study the cost of dynamics in nonequilibrium density currents [39]. Here we set  $\hat{S}_t = 0$  and  $\hat{S}_t^* = 0$  for  $t > T$  and solve (4) and (5) backwards in time. The two component processes represent the best possible score from time  $t$  onward *given* that the door was previously open ( $\hat{S}_t$ ) or given that it was previously closed ( $\hat{S}_t^*$ ). At each time point, we have already computed the values of  $\hat{S}_{t'}$ ,  $\hat{S}_{t'}^*$  for  $t' > t$ , and know the random variable  $\hat{v}$ . We simply have to decide whether it would be better to have an open door or closed door, and the maximum score is just the larger of these potential scores.

Equation (4) says that if the door was open, then the demon can keep it open for the next  $dt$  and reevaluate if the door should be opened or closed (resulting in a score of  $\hat{S}_{t+dt} + \hat{v}_t(dt)$ ), or the demon can shut the door, which must remain shut until  $t + \tau$ , and then decide what the best course of action is at that point in time, given that the door had been shut (resulting in a score of  $\hat{S}_{t+\tau}^*$ ). The best score at this point in time is simply the max of these two scores. The reasoning behind (5) is similar. Our simulation simply automatically solves these equations for given realizations of the events. To ensure that our algorithm operates correctly, we have run brute-force searches over all possible sequences of door openings or closings for some feasible times and microbin sizes. If the number of total microbins becomes too large, then the brute-force method quickly becomes intractable. For the realizations that we checked, we see that the Bellman algorithm does correctly compute the best possible schedule and score for the demon.

While there is literature on the asymptotic behavior of stochastic algorithms [40], on differential equations containing max or min terms [41], and on stochastic Bellman equations [42,43], it seems like an analytical solution to (4) and (5) would be very hard to come by, especially since the problem involves coupled stochastic processes, and we have not yet been able to solve for the average behavior of  $\hat{S}_t$  or  $\hat{S}_t^*$ .

## VII. DISCUSSION

We have developed an optimal protocol for *predictive* Maxwell’s demons, determined their heat and mass transfer rates, and compared these to the performance of their non-predictive counterparts. Knowing the future greatly enhances heat and mass transport performance.

In closing, we should emphasize that (1) the limitations on heat or mass currents reported here stem from the finite response time of the demons (which may be due to the inertia of the gate or time required to measure and process

information). A demon who could measure and haul particles at infinite velocity could of course achieve infinitely large entropy reduction rates, whether predictive or nonpredictive. The fact that fundamental physics prohibits infinitely fast measurement and gate motion suggests to us that information driven heat and mass transfer (and thus entropy reduction rate) is bounded by fundamental physics. (2) Being able to predict the future does not, of course, provide additional free negative entropy. The total entropy that can be pumped out of the system is set by the number of erasures the demon must carry out during measurement and information processing, as set by Landauer's principle. Rather, prediction making improves the *rate* of entropy reduction and heat or mass transport.

#### APPENDIX A: ENERGY DISTRIBUTION OF PARTICLES

Here we derive Eq. (2), which states that distribution of a particle's energy, conditioned on the event that it hits the gate, is

$$P(E) = \frac{\beta^{(d+1)/2} E^{(d-1)/2} e^{-\beta E}}{\Gamma[(d+1)/2]}.$$

We start with the Maxwell-Boltzmann distribution in  $d$  dimensions,

$$p_d(v) = \left(\frac{m\beta}{2\pi}\right)^{d/2} \Omega_d v^{d-1} e^{-\beta m v^2/2},$$

where  $\Omega_d$  is the surface area of a  $(d-1)$  sphere and is  $2$ ,  $2\pi$ , and  $4\pi$  in dimensions 1, 2, and 3.

Let  $\mathcal{A}$  be the event that a particle hits the gate. In Ref. [15], we showed that  $p_d(\mathcal{A}|v) = c_d v \tau A/V$  and  $p(\mathcal{A}) = \kappa/N$ , where  $c_d = 1/2, 1/\pi, 1/4$  in dimensions 1, 2, and 3. Recall that  $\kappa = \frac{\rho \tau A}{\sqrt{2\pi\beta m}}$ .

By Bayes's law,

$$P(E) = P(E|\mathcal{A}) = \frac{P(\mathcal{A}|E)P(E)}{P(\mathcal{A})}.$$

Changing variables in the Maxwell-Boltzmann distribution and evaluating  $P(\mathcal{A}|E)/P(\mathcal{A})$  gives us that

$$p_d(E) = \frac{1}{2}(\beta/\pi)^{d/2} \Omega_d E^{(d-2)/2} e^{-\beta E}$$

$$P(\mathcal{A}|E)/P(\mathcal{A}) = 2 c_d \sqrt{\pi \beta E}.$$

Combining this with the fact that  $c_d \Omega_d = \pi^{(d-1)/2} / \Gamma((d+1)/2)$  yields Eq. (2).

#### APPENDIX B: SIMULATION DETAILS

In this section, we give more details on how we simulate the predictive demon. As we noted above, the general idea is that we discretize time into "microbins" that are some fraction,  $1/g$  of the demon response time. We call  $g$  the resolution of the simulation, with the intuitive effect that the demon operating continuously is better and better approximated as  $g \rightarrow \infty$ .

Suppose that  $T = t_c$ , so that the demon can accurately predict when all particles will hit the gate area. The particle hit times and energies are generated according to the distributions (1) and (2) and can be represented as two lists of time or energy pairs  $(t_0^{(l)}, E_0^{(l)})$ ,  $\dots$  and  $(t_0^{(r)}, E_0^{(r)})$ ,  $\dots$ . To simulate the

demon, we divide the time interval  $[0, T]$  into the microbins  $[0, \tau/g)$ ,  $[\tau/g, 2\tau/g)$ ,  $\dots$ ,  $[T - \tau/g, T)$ , and bin the particle hit lists by time. For each microbin, we compute a "score" for the microbin, which will either be the net energy or the net particle flow during the microbin (depending on whether we are simulating an energy or number demon).

To carry out the Bellman calculation, we use exactly Eqs. (4) and (5), modified to be nonstochastic. Given our discretization,  $t$  is restricted to be an integer multiple of  $\tau$ ,  $dt = \tau/g$ , and  $\hat{v}_t(\Delta t)$  is the total score of the microbins between times  $t$  and  $t + \Delta t$ . We simply initialize  $S_T = S_T^* = 0$  and calculate  $S_{T-\tau/g}, S_{T-\tau/g}^*, S_{T-2\tau/g}, S_{T-2\tau/g}^*$ , and so on, until we get to  $t = 0$ . At that point, the maximum performance of the demon is simply  $\max(S_0, S_0^*)$ .

The simulation of the demons with  $T \gg t_c$  is not much more difficult. The demon repeatedly calculates what course of action will lead to the best possible score, performs that action, and moves  $\tau/g$  forward in time to repeat this procedure. The best course of action for the demon can be calculated as described above, but instead of the maximum performance being computed by  $\max(S_0, S_0^*)$ , we record when the demon opens and closes the gate as time goes by and do not allow the gate to have its state change until the time delay  $\tau$  expires. For example, if the gate was opened by the demon  $\tau/2$  ago, then the maximum performance achievable by the demon is  $v_0(\tau/2) + \max(S_{\tau/2}, S_{\tau/2}^*)$  since the demon must keep the gate open until at least  $+\tau/2$  (recall that  $v_0(\tau/2)$  is the event score from time 0 to time  $\tau/2$ ). The one additional component to simulating what is the best possible action for the demon is that for times closer than  $\tau$  to the end of the demon's time sight, we just add the expected number of particles to

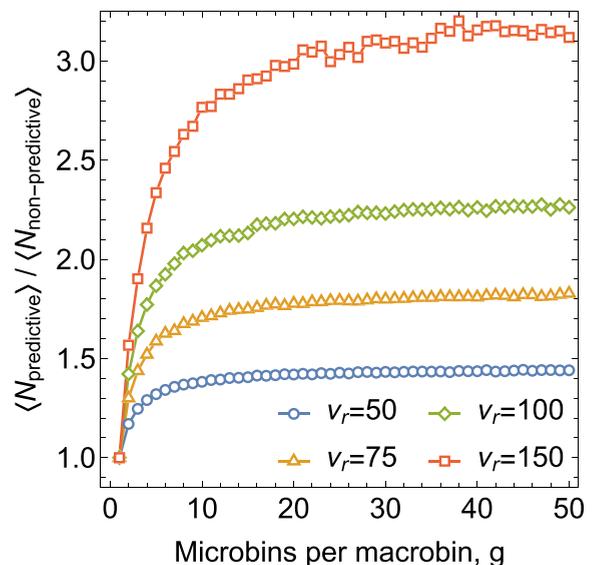


FIG. 6. Convergence of a number demon's performance as the number of microbins increases, normalized by the performance of a nonpredictive number demon. As  $g \rightarrow \infty$ , the discrete algorithm for the demon better and better approximates the continuous demon model. Clearly, the discrete algorithm with  $g = 50$  is a good approximation to the continuous case, as the score ratio has stabilized by that point.

pass through the door in the hidden part of time to  $S_t$ , where  $T < t' < t + \tau$ .

### APPENDIX C: TIME RESOLUTION AND DEMON PERFORMANCE

In this section, we show how the resolution of the demon's scheduling procedure affects its performance. Recall that we divide time into "microbins" with a certain resolution  $g$ , which is the number of microbins in a response time,  $\tau$  (so the length of a microbin is  $\tau_{\text{mb}} = \tau/g$ ). In principle, if a real demon was trying to schedule its sequence of gate openings or closing with an algorithm similar to ours, then there would be a

trade-off between the amount of time and computational resources necessary to do the scheduling and the resolution the demon would use, but in the paper, we simply use a "large" resolution to approximate the continuous limit.

As visible in Fig. 6, demon score increases rapidly as the resolution increases, especially for the case where the subsystems have similar parameters.

When there is only one microbin, predictive and nonpredictive demons act exactly the same, hence the ratio of scores being 1. As the number of microbins increases, the predictive demon has more freedom to adjust exactly when it opens and closes its gate. Because of this, the performance of the demon rapidly increases with granularity, before asymptotically approaching the score of an ideal, continuous predictive demon.

- 
- [1] A. Rex, *Entropy* **19**, 240 (2017).
- [2] F. J. Cao, L. Dinis, and J. M. R. Parrondo, *Phys. Rev. Lett.* **93**, 040603 (2004).
- [3] U. Seifert, *Rep. Prog. Phys.* **75**, 126001 (2012).
- [4] W. Zurek, *Algorithmic Randomness, Physical Entropy, Measurements, and the Demon of Choice* (Perseus Books, Reading, MA, 1999).
- [5] C. M. Caves, *Phys. Rev. Lett.* **64**, 2111 (1990).
- [6] C. M. Caves, W. G. Unruh, and W. H. Zurek, *Phys. Rev. Lett.* **65**, 1387 (1990).
- [7] D. Mandal and C. Jarzynski, *Proc. Natl. Acad. Sci. USA* **109**, 11641 (2012).
- [8] D. Mandal, H. T. Quan, and C. Jarzynski, *Phys. Rev. Lett.* **111**, 030602 (2013).
- [9] A. C. Barato and U. Seifert, *Europhys. Lett.* **101**, 60001 (2013).
- [10] A. Hosoya, K. Maruyama, and Y. Shikano, *Sci. Rep.* **5**, 17011 (2015).
- [11] S. Deffner, *Phys. Rev. E* **88**, 062128 (2013).
- [12] L. Gordon, *Found. Phys.* **13**, 989 (1983).
- [13] P. A. Skordos and W. H. Zurek, *Am. J. Phys.* **60**, 876 (1992).
- [14] Z. Tu, *J. Phys. A: Math. Theor.* **41**, 312003 (2008).
- [15] N. Rupprecht and D. C. Vural, *Phys. Rev. Lett.* **123**, 080603 (2019).
- [16] P. Strasberg, G. Schaller, T. Brandes, and M. Esposito, *Phys. Rev. Lett.* **110**, 040601 (2013).
- [17] É. Roldán, I. A. Martínez, J. M. Parrondo, and D. Petrov, *Nat. Phys.* **10**, 457 (2014).
- [18] N. Cottet, S. Jezouin, L. Bretheau, P. Campagne-Ibarcq, Q. Ficheux, J. Anders, A. Auffèves, R. Azouit, P. Rouchon, and B. Huard, *Proc. Natl. Acad. Sci. USA* **114**, 7561 (2017).
- [19] J. J. Thorn, E. A. Schoene, T. Li, and D. A. Steck, *Phys. Rev. Lett.* **100**, 240407 (2008).
- [20] S. T. Bannerman, G. N. Price, K. Viering, and M. G. Raizen, *New J. Phys.* **11**, 063044 (2009).
- [21] J. V. Koski, A. Kutvonen, I. M. Khaymovich, T. Ala-Nissila, and J. P. Pekola, *Phys. Rev. Lett.* **115**, 260602 (2015).
- [22] M. D. Vidrighin, O. Dahlsten, M. Barbieri, M. S. Kim, V. Vedral, and I. A. Walmsley, *Phys. Rev. Lett.* **116**, 050401 (2016).
- [23] P. A. Camati, J. P. S. Peterson, T. B. Batalhao, K. Micadei, A. M. Souza, R. S. Sarthour, I. S. Oliveira, and R. M. Serra, *Phys. Rev. Lett.* **117**, 240502 (2016).
- [24] K. Chida, S. Desai, K. Nishiguchi, and A. Fujiwara, *Nat. Commun.* **8**, 15310 (2017).
- [25] G. Schaller, C. Emary, G. Kiesslich, and T. Brandes, *Phys. Rev. B* **84**, 085418 (2011).
- [26] M. Esposito and G. Schaller, *Europhys. Lett.* **99**, 30003 (2012).
- [27] G. Schaller, J. Cerrillo, G. Engelhardt, and P. Strasberg, *Phys. Rev. B* **97**, 195104 (2018).
- [28] L. Gammaitoni, *Contemp. Phys.* **53**, 119 (2012).
- [29] S. Bo, M. Del Giudice, and A. Celani, *J. Stat. Mech.: Theory Exp.* (2015) P01014.
- [30] B. Sothmann, R. Sánchez, and A. N. Jordan, *Nanotechnology* **26**, 032001 (2014).
- [31] B. Roche, P. Roulleau, T. Jullien, Y. Jompol, I. Farrer, D. A. Ritchie, and D. Glattli, *Nat. Commun.* **6**, 6738 (2015).
- [32] F. Hartmann, P. Pfeffer, S. Höfling, M. Kamp, and L. Worschech, *Phys. Rev. Lett.* **114**, 146805 (2015).
- [33] M. G. Raizen, *Sci. Am.* **304**, 54 (2011).
- [34] L. Dinis and J. M. Parrondo, *Europhys. Lett.* **63**, 319 (2003).
- [35] See <https://github.com/nrupprecht/Scheduling-Process> for our implementation.
- [36] E. Barron and H. Ishii, *Nonlinear Anal.: Theory Methods Appl.* **13**, 1067 (1989).
- [37] D. E. Kirk, *Optimal Control Theory: An Introduction* (Courier Corporation, New York, 2004).
- [38] F. Roca, J. Villaluenga, and L. Dinis, *Europhys. Lett.* **107**, 10006 (2014).
- [39] V. Y. Chernyak, M. Chertkov, J. Bierkens, and H. J. Kappen, *J. Phys. A: Math. Theor.* **47**, 022001 (2013).
- [40] M. Pelletier, *Stoch. Process. Appl.* **78**, 217 (1998).
- [41] Y. Liu and J. Wu, *Adv. Differ. Eq.* **2015**, 379 (2015).
- [42] L. C. Evans and A. Friedman, *Trans. Am. Math. Soc.* **253**, 365 (1979).
- [43] N. Turhan, Deterministic and stochastic Bellman's optimality principles on isolated time domains and their applications in finance, Masters theses and Specialist Projects, paper 1045, 2011, <https://digitalcommons.wku.edu/theses/1045>.