





<https://doi.org/10.1038/s42003-022-04235-6>

OPEN

Sunk cost sensitivity during change-of-mind decisions is informed by both the spent and remaining costs

A. David Redish¹ [✉], Samantha V. Abram², Paul J. Cunningham¹, Anneke A. Duin^{1,9}, Romain Durand-de Cuttoli³ , Rebecca Kazinka⁴, Adrina Kocharian⁵, Angus W. MacDonald III⁶, Brandy Schmidt¹, Neil Schmitzer-Torbert⁷ , Mark J. Thomas¹  & Brian M. Sweis⁸

Sunk cost sensitivity describes escalating decision commitment with increased spent resources. On neuroeconomic foraging tasks, mice, rats, and humans show similar escalations from sunk costs while quitting an ongoing countdown to reward. In a new analysis taken across computationally parallel foraging tasks across species and laboratories, we find that these behaviors primarily occur on choices that are economically inconsistent with the subject's other choices, and that they reflect not only the time spent, but also the time remaining, suggesting that these are change-of-mind re-evaluation processes. Using a recently proposed change-of-mind drift-diffusion model, we find that the sunk cost sensitivity in this model arises from decision-processes that directly take into account the time spent (costs sunk). Applying these new insights to experimental data, we find that sensitivity to sunk costs during re-evaluation decisions depends on the information provided to the subject about the time spent and the time remaining.

¹ Department of Neuroscience, University of Minnesota, Minneapolis, MN 55455, USA. ² San Francisco Veterans Affairs Medical Center, San Francisco, CA 94121, USA. ³ Nash Family Department of Neuroscience, Icahn School of Medicine at Mount Sinai, New York, NY 10029, USA. ⁴ Department of Psychiatry and Behavioral Sciences, University of Minnesota, Minneapolis, MN 55454, USA. ⁵ Graduate Program in Neuroscience and Medical Scientist Training Program, University of Minnesota, Minneapolis, MN 55455, USA. ⁶ Department of Psychology, University of Minnesota, Minneapolis, MN 55455, USA. ⁷ Department of Psychology, Wabash College, Crawfordsville, IN 47933, USA. ⁸ Department of Psychiatry, Icahn School of Medicine at Mount Sinai, New York, NY 10029, USA. ⁹ Present address: Epic Systems, 1979 Milky Way, Verona, WI 53593, USA. ✉email: redish@umn.edu

Although, logically, decisions should be made based on future expectations, many foraging decisions show an escalation of continued pursuit of reward as a function of spent resources (sunk costs)^{1–6}. Although sunk cost sensitivity has been observed in non-human animals as well as in humans^{7–13}, the human and non-human experiments were done on computationally different tasks. Translating decision processes between species depends on the computational alignment between how each subject makes decisions on a given task¹⁴.

In a suite of previous studies, we developed a family of new computationally-aligned foraging tasks in which subjects made *accept/skip* decisions for consumable rewards. In these tasks, subjects face a decision of whether to wait out a fully-signaled delay before getting the reward or leaving that offer and trying their luck at a subsequent offer for a different reward. Subjects have a limited time to gather these rewards within a session, making time a limited resource, and this an economic task. Multiple variants of these tasks exist, including ones in which the consumable reward is food (*Restaurant Row* [mice, rats]^{15–22}, *Candy Row* [humans]²³) or videos (*Web-Surf* [humans]^{19,24,25}, *Movie Row* [humans]²³). In the *Restaurant Row* task, mice or rats physically run around an environment, encountering four different “restaurants” cyclically, each providing a different flavor of food reward. In the *Web-Surf* and *Movie Row* tasks, humans encounter four video galleries cyclically, each providing either still photographs²⁴ or short videos^{23–26}. The *Web-Surf* results have been replicated both in-person^{23–25} and online^{23,26}, under multiple variations. In the *Candy Row* task, humans are given the opportunity to receive candy/snacks from four dispensers cyclically²³. In all of these tasks, the delay to be waited out for the reward is signaled (by pitch of tone for mice and rats, by a number or a “download” bar for humans) and is random (uniformly distributed between 1 s and 30 s) and unknown until entry into the restaurant or gallery.

The computational alignment across species within these tasks offers a platform for translational research aimed at understanding how underlying decision processes change in psychiatric disorders. Recent studies have found that humans with various psychiatric disorders behave differently on these tasks^{18,27–29}. Understanding the decision-making processes that underlie the observed behavior on these tasks can provide insight as to how these computational processes change in psychiatric disorders.

Multiple variants of these tasks exist, including versions in which there are separate offer and wait stages^{16,18,19,23,26}. In these versions, rodents were given two zones at each restaurant, while humans were given two stages to each gallery (Fig. 1a, b). For simplicity, we will refer to these as the “offer zone” and “wait zone”, with the understanding that for humans the two components are separated logically rather than physically. On entering the offer zone, the subject was informed of the delay, but the countdown did not decrease while subjects remained in the offer zone. Subjects could skip as before, leaving the offer zone for the next restaurant/gallery or they could accept by entering the wait zone, at which point the delay began to countdown. Importantly, subjects in this variant could quit the wait zone even though they had accepted the offer by entering into the wait zone, at which point the countdown stopped and the offer was rescinded.

Analyzing behavior on this two-zone variant, we found that the decision to *quit* depended not only on the time remaining in the countdown but also on the time that had already been spent in the wait zone on a given trial—that the longer the subject had spent within the wait zone, the less willing they were to quit, even for identical future conditions in which the time remaining to receive reward was the same^{19,26}. Time already spent waiting is the definition of sunk costs^{1,2,4}.

Two escalations of commitment are observed on these tasks. First, there is an escalation of commitment as subjects approach

the reward³⁰ (Fig. 2a). The longer that a subject has waited within the wait zone, the more likely the subject is to wait out the delay (Fig. 2b). While this escalation can be partially explained by the subject getting closer to the reward with passing time, the time course of this escalation (Fig. 2c) is an interaction between the time spent and the time remaining. (Linear model fit, $df = 431$, F-vs constant model 182, $p = 10^{-76}$, effect of time remaining, $p = 10^{-13}$, effect of time spent, $p = 10^{-17}$, interaction, $p = 0.0016$, adjusted- R^2 relative to constant model = 0.55), implying that it is not simply that the subject is getting closer to the goal. If the escalation were simply due to getting closer to the goal, it should depend solely on the nearness of the goal; it should not depend on the time spent, nor on an interaction between time spent and time remaining.

Second, we can see a difference in commitment between two conditions in which the future path to the goal is identical, but the past is different. Imagine a subject who one time accepts a 15 s offer, waits 5 s, and thus has 10 s more to wait to get to a given goal, and another time accepts a 25 s offer, waits 15 s, and thus has 10 s left to wait for that same goal. Note that the future conditions at these two analysis points are the same—the subject has 10 s to wait for a given reward (the same reward in both conditions). However, when the subject has 10 s to go in the first condition, it has been waiting 5 s (has 5 s of “sunk costs”), while when the subject has 10 s to go in the second condition, it has been waiting 15 s (has 15 s of “sunk costs”). (See Fig. 2d). In the original study¹⁹, this was measured as the change in linearly fit slopes to the probability of waiting out the delay as the time already waited increased (Fig. 2e), but it is simpler to directly measure the effect through the changes in the probability as an interaction effect between time remaining and time already waited (Fig. 2f). (Linear model fit, $df = 461$, F vs constant model = 3880, $p < 10^{-100}$, effect of time remaining, $p < 10^{-100}$, effect of time spent $p = 10^{-42}$, interaction, $p < 10^{-100}$, adjusted- R^2 relative to constant model = 0.962).

This sensitivity to sunk costs is seen in both the *Restaurant Row* and *Web-Surf* tasks and their follow-up variants (such as variants with and without an offer zone^{15,19,21}, in the *Movie Row* task²³, in the *Candy Row* task²³, and in the *Known-Delay/Randomized-Delay* tasks, in which we manipulated the predictability of the upcoming delay in each restaurant³¹).

Results

Sunk cost sensitivity appears primarily in situations where agents are correcting error-made decisions. We do not observe sunk cost sensitivity in all conditions. Nor do we observe sunk cost sensitivity at all times within the wait zone. Nor do we observe sunk cost sensitivity uniformly within the waiting zone. Instead, we find that the sensitivity to sunk costs in the wait zone depends on how economically disadvantageous the decision to accept the offer was in the first place.

The *Restaurant Row* and *Web-Surf* tasks (and their variants) are economic foraging tasks, in which subjects are attempting to maximize their consumed rewards within a limited time. Thus, we can identify how worthwhile an offer is to an agent in terms of the delay to the goal (the cost from the limited time budget available) and the (subjective) preference for a given reward. In practice, subjects on these tasks generally reveal a *threshold* for each flavor/gallery, such that they accept delays below that threshold, but reject delays above that threshold. Thresholds can be measured by fitting a logistic regression to the accept vs skip choices, providing both the midpoint (where the accept/skip probability is 0.5) and a slope parameter (measuring the sharpness of the decision logistic). Agents typically showed an individual threshold for each flavor, but these were generally

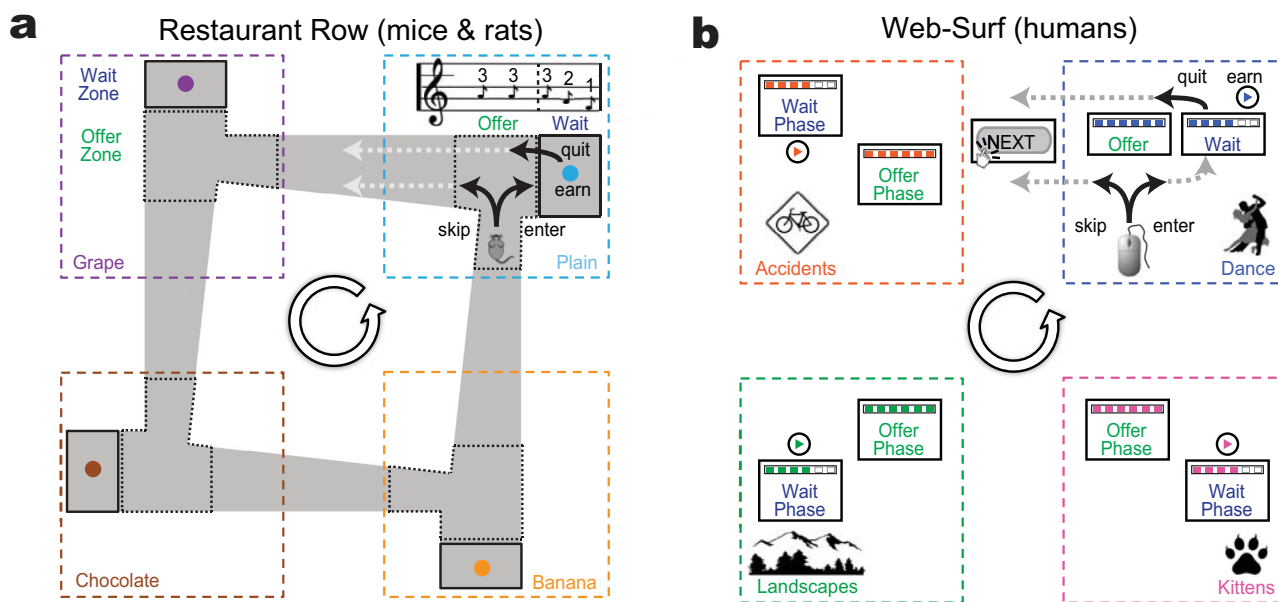


Fig. 1 The Restaurant Row and Web-Surf tasks. **a** The Restaurant Row task for mice and rats. As the animal enters the offer zone, the delay is indicated by a tone, which only starts counting down on entering the wait zone. Animals encounter restaurants serially by proceeding around the cycle counterclockwise. **b** The Web-Surf task for humans. An offer is provided to the human as a download bar with a set delay, but the download does not start to countdown until they select “enter” to enter the wait zone. Humans encounter the video galleries serially by clicking through a sequence of buttons. Note the topological analogies. Figure reprinted from¹⁹ with permission of the publisher.

consistent across sessions for rats and mice that ran multiple sessions^{15,16,19,22}, and consistent with stated preferences (rankings, ratings) in humans^{23,24,26}. Under the assumption that the thresholds measure the individual preferences for each subject, we can thus measure the *value* of an offer, corrected for that individual preference, by subtracting off the threshold. We thus define *value* as the difference between the threshold that the agent has shown for that flavor/gallery and a given offer: **Value = threshold – offered delay**. Thus, offers with delays longer than the threshold for that flavor/gallery have low value and offers with delays shorter than the threshold have high value.

In Fig. 3, we rearranged the sunk cost sensitivity measure (the increased likelihood of waiting out the delay, the change in $p(\text{Earn})$, measured as the difference between $p(\text{Earn})$ at a given delay and $p(\text{Earn})$ with 0 s invested) by the value that was left at the countdown at the time of quitting. The probability of waiting out the delay to earn is enhanced through sunk costs primarily in the bad deals (deals where the delay remaining was above the typical threshold for that reward site [restaurant or gallery]). (Fig. 3 statistics, Mann–Whitney test, comparing change in $p(\text{Earn})$ for values <0 and values >0: 3a (Mice¹⁹), $p < 10^{-100}$; 3b (Mice²⁰), $p = 10^{-100}$, 3c (Rats¹⁹), $p = 10^{-62}$; 3d (Humans in person¹⁹), $p = 0.27$; 3e (Humans online²³), $p = 10^{-19}$; 3f (Humans online, NCST dataset), $p = 10^{-94}$).

These data suggest that subjects quit out of the wait-zone due to a re-evaluation process, whereby the subject recognizes that they (incorrectly) accepted an offer that was larger than the threshold for that reward. This suggests that one should be able to model this decision with a change-of-mind decision-making model, in which a subject commits to a choice, but then can change its mind and make a different choice if a re-evaluation process occurs and completes before a time-limit^{32,33}.

These data suggest that a sensitivity to sunk costs comes through the conflicting decisions of whether to wait out an excessive delay that was accepted incorrectly or to quit. Once the delay during the countdown has crossed the typical threshold, the deal is now seen as reasonable, even if the already spent time is

lost, and the sensitivity to sunk costs disappears. We argue that the conflict of whether to rectify the incorrectly accepted deal is driving the sensitivity to sunk costs because the wasted costs lost by rectifying the incorrectly accepted deal are the time already spent (the sunk costs). These hesitation effects are commonly seen in situations of oncoming scarcity threat³⁴.

Sunk cost sensitivity does not solely arise from attrition biases.

The strongest evidence for a sensitivity to sunk costs on these tasks is the escalation of commitment seen when comparing the same future outcome, but different times already waited (sunk costs) (Fig. 2d). An important concern, however, is that this effect can only be measured when comparing two different offered delays. Assuming that a subjects’ motivation for each reward can vary over the course of a session, on every encounter with a decision, we can assume that the “willingness to wait” (the motivation for the goal) is drawn from a distribution with some noise. Because subjects are assumed to take deals only when their motivation (their willingness to wait) is above threshold, the distribution of motivations that will drive a subject to accept the longer delay will be different than the distribution of motivations that will drive a subject to accept a shorter delay. Assuming that the agent has variability in its motivation to accept different offers, acceptance of a 25 s deal will require a stronger motivation than acceptance of a 15 s deal.

Ott et al.³⁵ argued that this different set of initial motivations would make the set of subjects encountering these two similar analysis points different—the set of subjects encountering the 10 s to go after waiting 15 s would have had to start with a higher subset of the distribution of motivations than the set of subjects encountering the 10 s to go after waiting 5 s. Subjects willing to wait 15 s, but not 16 s or 20 s or 24 s would still have taken the 15 s deal but not the 25 s deal. They point out that this would mean that subjects reaching the 10 s analysis point after waiting 15 s were more likely to have more “willingness to wait” (more motivation) than subjects reaching the 10 s analysis point after

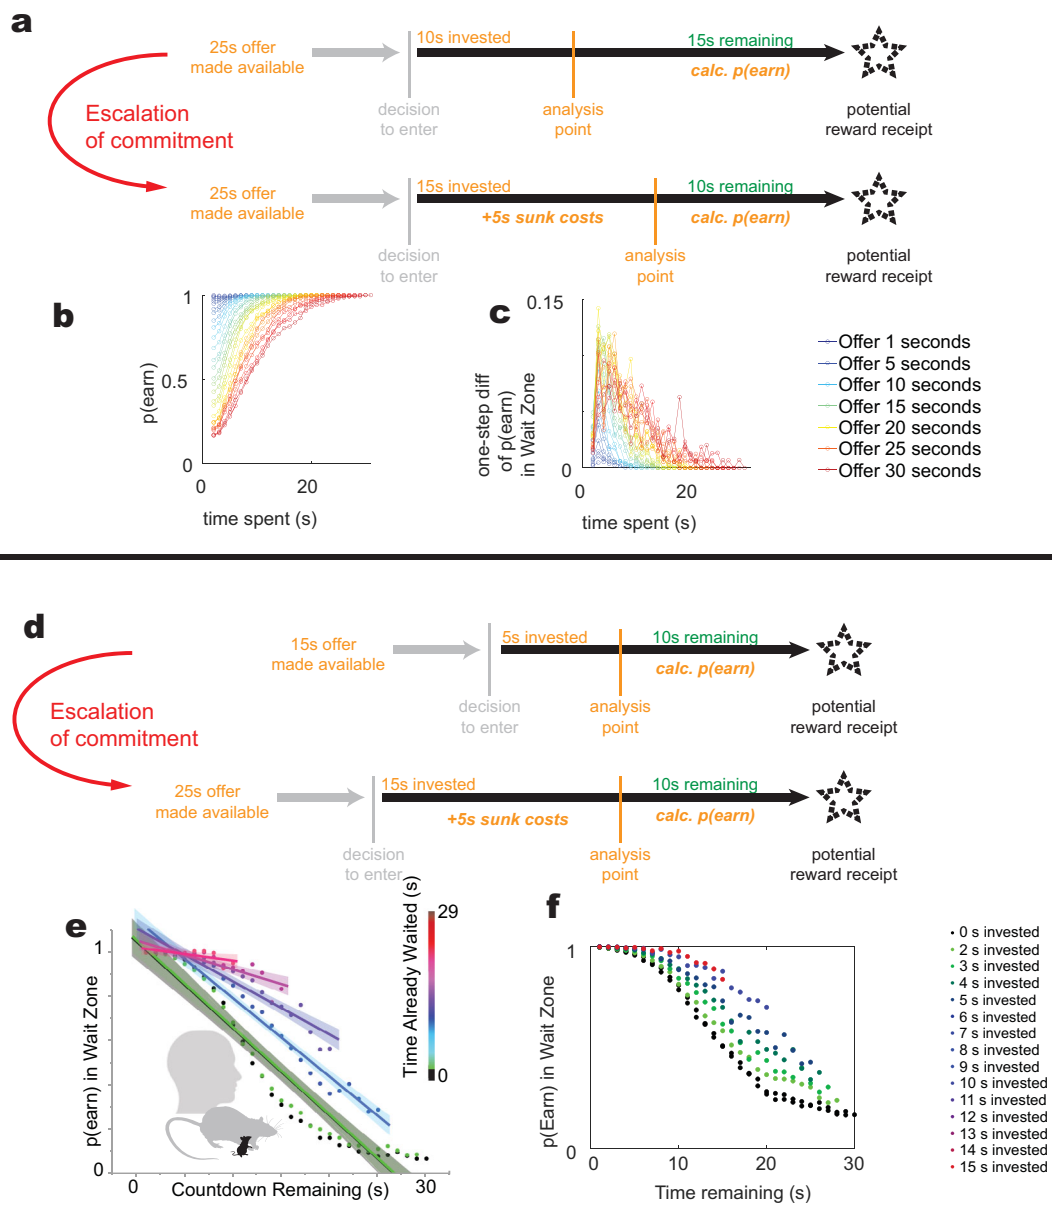


Fig. 2 Escalations of commitment on foraging tasks. Once a subject is waiting out a delay, two escalations of commitment can be measured on these tasks. At a given point within the countdown in the wait zone, we ask what is the likelihood that the subject will wait out the delay to receive the reward. This is $p(\text{Earn})$. **a** As an agent approaches the reward, it becomes less likely to quit ($p(\text{Earn})$ increases). **b** The probability of waiting out the delay ($p(\text{Earn})$) increases with the time spent. **c** This increase is not a simple exponential decrease of quitting, but increases non-linearly with time-spent. **d** Comparing the same future (10 s remaining before receiving the same reward) with different times already invested can also reveal a sensitivity to sunk costs. **e** The evidence for increased $p(\text{Earn})$ with time spent (time already waited) originally reported in ref. ¹⁹. The graph shows the probability of earning, $p(\text{Earn})$, as a function of the time remaining in the wait zone. Each colored dot and each line aligned to the colored dots corresponds to the $p(\text{Earn})$ given that the subject has waited a certain amount of time in the wait zone already. Notice that the slopes of the lines decrease, indicating that having waited longer, subjects are more likely to wait out the delay, even for a given countdown remaining. Thus, following the logic in **d**, we conclude that the subjects are behaviorally sensitive to sunk costs. **f** The full extent of the data from **e**, showing the continuous “bubble” of increased willingness to wait with time already waited. Models and future comparisons will be plotted in this form, without the linear slope assumption used in **b**, **c**. Replotted from ref. ³⁰. **d** Reprinted from ref. ¹⁹ with permission from the publisher.

only 5 s. Thus, this increase in staying after waiting longer may actually reflect a difference in the willingness to wait for reward as a function of the offered delay to the reward rather than an escalation based on the delay waited. They proposed that the behavioral sensitivity to sunk costs seen in these tasks could be an epiphenomenon consequence of these attrition statistics³⁵. They argue that the escalation of commitment shown in Fig. 2e could

arise from the distribution of motivations (the willingness to take the deal in the offer zone) in the first place.

Ott et al.³⁵ test this theory with a model in which they assume (1) that the subject has a *willingness to wait* W which initially distributes normally around the threshold ($W0 = W_{\text{threshold}} + N(\sigma_W)$), (2) that the decision to enter an offer occurs when $W0$ is greater than the threshold for that offer T_{OZ} , (3) that W then

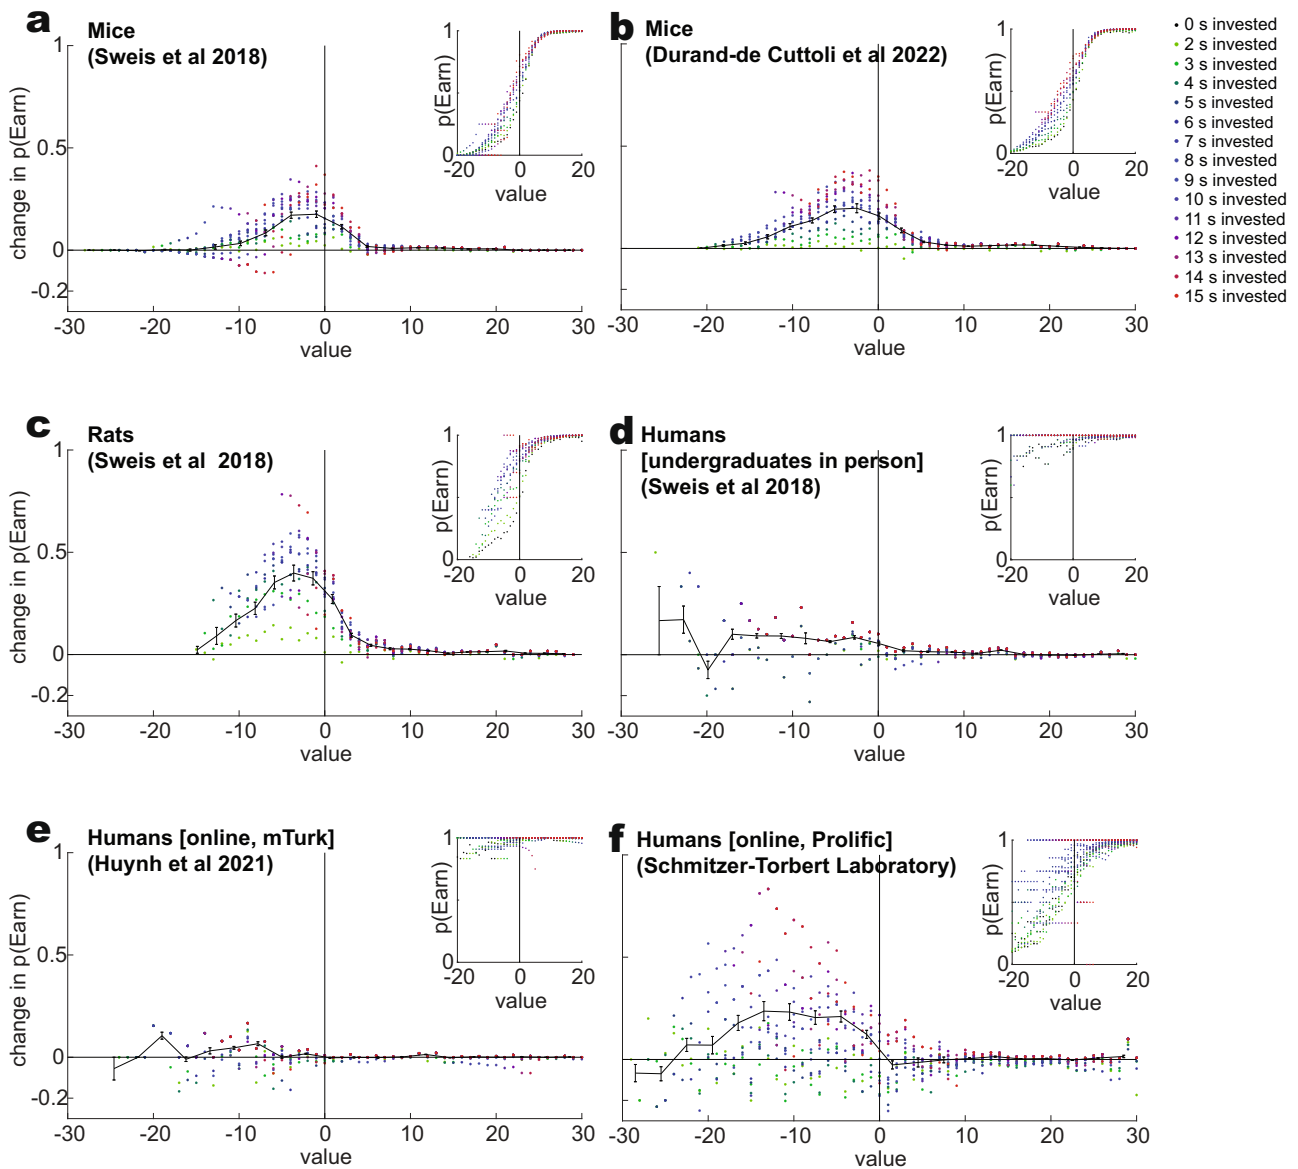


Fig. 3 Sensitivity to sunk costs depends on the value remaining. Value is defined as (threshold - offer). Each primary panel shows the increased probability of remaining to earn given the value of a current moment in the countdown. The x axis shows the value left in the countdown at the time of quit. The y axis shows the change in $p(\text{Earn})$ between the time invested and the baseline condition (0 s invested), with time invested marked in color. Insets show the raw probability of remaining to earn at a given value within a countdown. Probability is measured over all subjects. The vertical line marks the transition to where the value crosses zero, where the delay crosses the threshold. Thus, to the right of the line, the remaining time is below the threshold for that restaurant/gallery. Mice, rats, and humans were only sensitive to sunk costs until the value crossed that line. Statistics measured if the change in $p(\text{Earn})$ was greater before than after value crossed from negative to positive (Mann-Whitney sign rank test). **a** Mouse data from^{16,19}, $p < 10^{-100}$ **b** Mouse dataset from²⁰, $p < 10^{-100}$. **c** Rat dataset from¹⁹, $p = 10^{-62}$. **d** Human dataset from¹⁹, $p = 0.27$. **e** Human dataset from²³, $p = 10^{-19}$. **f** Online human data from Schmitzer-Torbert laboratory, $p = 10^{-94}$.

wanders with a given variance (σ_N), and the subject quits if W ever crosses a quit-threshold in the wait zone T_{WZ} . Although they do not discuss the importance of the shape of the quit threshold T_{WZ} in their model, they assume that the T_{WZ} decreases with time so that it starts at the entry threshold $T_{WZ}(t = 0) = T_{OZ}$ and reaches 0 at the time of reward $T_{WZ}(t = T_{OZ}) = 0$. Without loss of generality, we can assume a given threshold and examine the sunk costs seen at a single restaurant/gallery. (Ott et al.³⁵ assume $W_{\text{threshold}} = 18$ s, which provides a threshold of 18 s for that modeled restaurant).

We concur that their model does show a sensitivity to sunk costs (see Fig. 4). However, two important questions remain: First, *what are the factors within their model that create this*

sensitivity to sunk costs? Second, how well does their model describe the data we have observed in mice, rats, and humans?

We start by noting that the Ott et al.³⁵ model is a standard drift-diffusion change-of-mind model with some additional parameters (Fig. 4a). Change-of-mind models assume that decision variables continue to drift after a decision is made and that change-of-mind occurs when the decision variable crosses a threshold^{32,33}. Their model starts after the decision has been made, and thus assumes that the motivation must have been larger than the threshold. It then drifts until it crosses the quit threshold or the agent receives reward.

Their model has two interesting features: (1) it has an expanding boundary (threshold), such that the decision variable

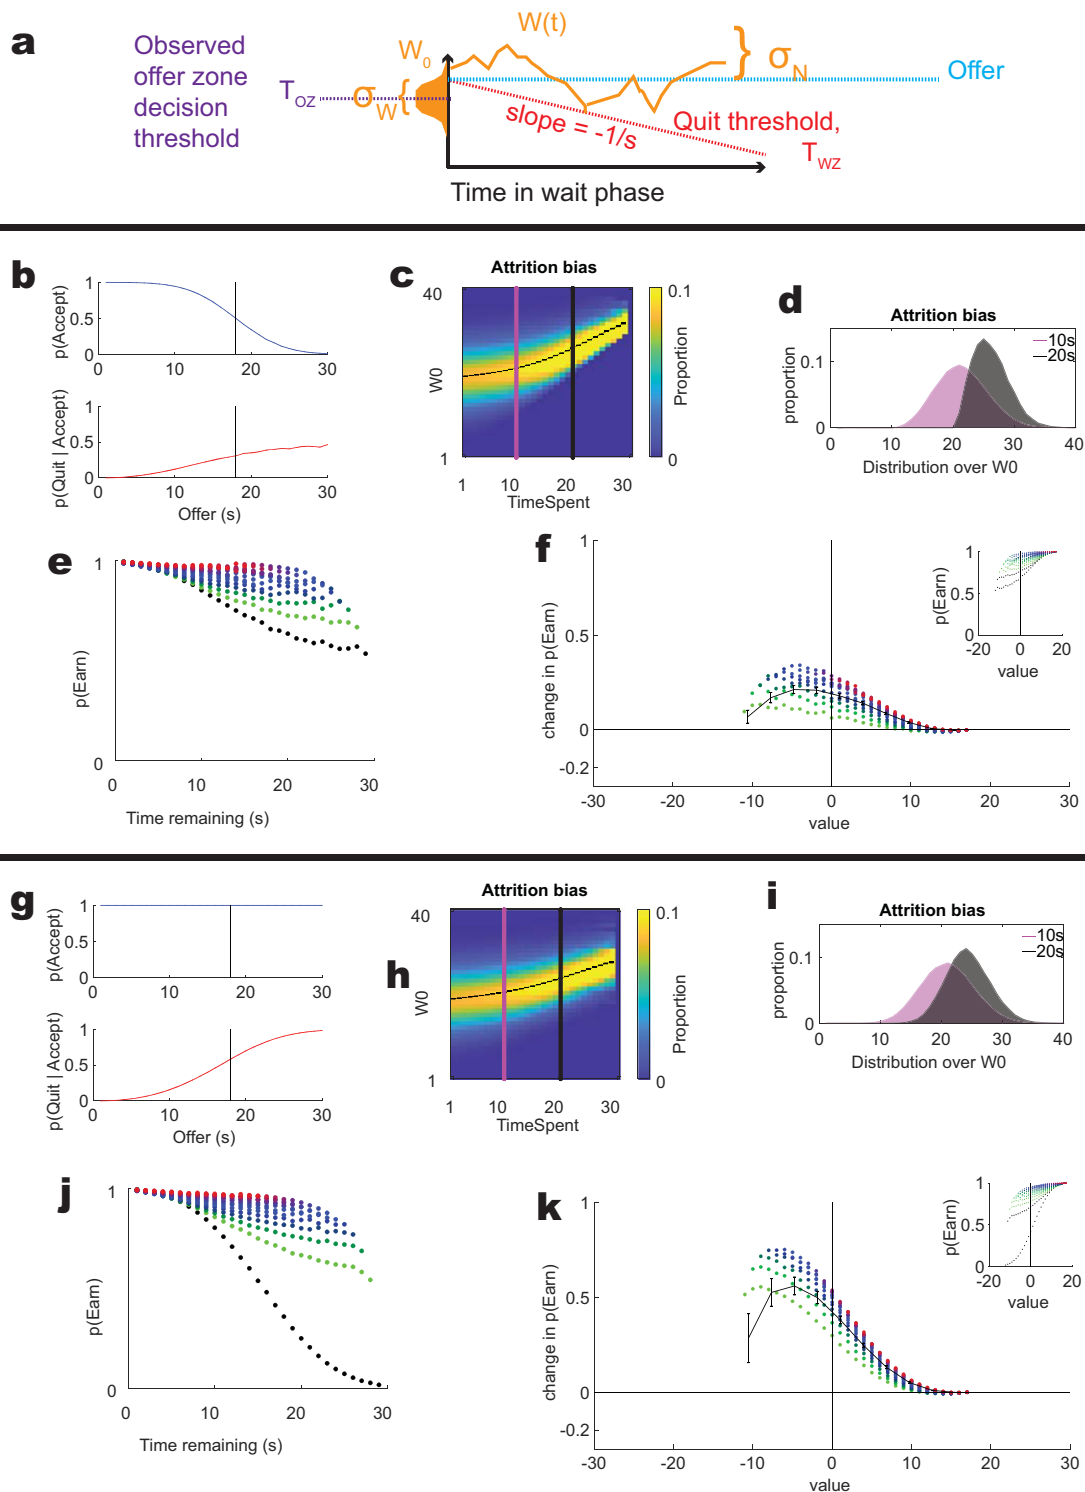


Fig. 4 The basic structure of the Ott et al.³⁵ model. **a** The basic layout of their model, separating the original enter decision and the leaving decision as dependent on willingness to wait parameters. **b** In the original model, agents probabilistically accept offers based on the initial $W_0 > offer$ and quit when the wandering W crosses the quit threshold T_{wz} . **c** This produces an attrition bias in their model, in that with longer times-spent in the wait zone, lower initial motivations are no longer part of the distribution. In the panel, each column sums to 1 and is thus a probability distribution over the remaining initial motivations (proportion of the distribution indicated by color as specified in the colorbar). **d** Two example distributions drawn from the set in **c**. Note that the distribution remaining at 20 s time-spent is shifted relative to that available after 10 s time-spent. **e** The model also shows a sensitivity to sunk costs, qualitatively comparable to those seen in Fig. 2f. **f** The model also shows that sunk costs are increased before the value crosses from negative to positive. **g–k** This effect is not due to selection bias, but rather to other factors in the model. In a version where the agent accepts all deals (**g**), the attrition bias is reduced (**h, i**), but the sensitivity to sunk costs are still as strong (**j**), and the model still shows an increased change in $p(\text{Earn})$ for negative over positive values (**k**). Results shown are from our simulations.

has to move farther away from its starting point in order to reach the boundary the longer the agent has been in the wait zone, and (2) there is only a lower, but not an upper boundary, allowing the decision variable to drift farther away from the boundary the longer the agent has been in the wait zone. Both of these components make the decision variable explicitly and directly sensitive to sunk costs.

These two components (the quit threshold T_{WZ} expanding away from the decision variable and the wandering W component of the drift diffusion model) contain information about the time spent within the wait zone. Because decisions are made based on these processes, the model is less likely to quit the longer the agent has waited (Fig. 4e, f). Thus, the model proposed by Ott et al.³⁵ shows a sensitivity to sunk costs because it makes decisions based on parameters that accrue with sunk costs. Interestingly, in the model, these two components do not start to accrue until the agent enters the wait zone, consistent with the observations that sunk costs do not start to accrue until the subject commits to the decision by entering the wait zone^{19,20,23,26}.

Ott et al.³⁵ claim that the sensitivity to sunk costs arises from the attrition bias that occurs in their model. As can be seen in Fig. 4b–d, there is a selection bias inherent in their model—accepting an offer with a high threshold implies that W_0 was selected from the higher part of the distribution $W_{\text{threshold}} + N(\sigma_W)$. As noted above, this implies an attrition in the distribution of W_0 with time spent because the shorter times-spent (which can be accepted with lower W_0) will fall away as the subject crosses that delay. That is, a subject can only reach a 15 s delay for offers of longer than 15 s—shorter delay offers would have provided reward already. To determine whether these effects arose directly from the selection process in taking deals (that the agent only enters the wait zone if W_0 was greater than the offer), we tested a version of the model in which the agent takes every choice (Fig. 4g–k). Note that these agents are still quitting if $W < T_{WZ}$, but that includes them in the measurement of sunk costs (in comparison to the agents who do not enter the wait zone at all, and thus are not included in that measurement). This change reduces, but does not eliminate the attrition bias (Fig. 4i); however, this change, if anything, increases the observed sunk costs (Fig. 4j, k), suggesting that there must be other factors in their model producing sunk cost sensitivity beyond the attrition bias.

The Ott et al. model contains two important parameters: σ_W , the variability in the initial willingness-to-wait that controls the spread of the W_0 Gaussian, and σ_N , the variability that controls the drift in W as the agent waits in the wait zone. The simulations reported in³⁵ used one pair of parameters ($\sigma_W = 5$, $\sigma_N = 3$), which produce results that are reminiscent of the data shown in Fig. 2e. A thorough exploration of the parameter space finds that the sensitivity to sunk costs are dependent on these two parameters (Fig. 5a), as is the attrition bias (Fig. 5b). However, the attrition bias and the sunk cost sensitivity in this model are unrelated.

We can measure three key parameters of the decisions made in one of these simulations: the “sunk cost bubble”, the “baseline slope”, and the “attrition bias”. We measure the attrition bias as the linear fit of the available initial willingness to wait W_0 as a function of time waited (Fig. 5c). We measure the baseline slope as the linear slope of the 0 s condition (Fig. 5d). And we measure the sunk cost bubble as the summed area comparing the increased likelihood of waiting out the delay relative to the 0s-waited condition (Fig. 5e).

As shown in Fig. 5, attrition bias is most strongly related to σ_W (Fig. 5f, linear model: $df = 33$, related to σ_W : $F = 92$, $p = 10^{-11}$, related to σ_N : $F = 0.2$, $p = 0.65$), while the baseline slope depends on both σ_N and σ_W (Fig. 5g, linear mode, $df = 33$: related to σ_W : $F = 34$, $p = 10^{-6}$, related to σ_N : $F = 65$, $p = 10^{-9}$). Sunk cost

sensitivity, however, is more related to σ_N and the relationship to σ_W is non-linear, because it is largest at large σ_N and mid-ranges of σ_W and smaller at high and low σ_W (Fig. 5h, linear model, $df = 33$: related to σ_W : $F = 0.01$, $p = 0.9$, related to σ_N : $F = 56$, $p = 10^{-8}$). Thus the highest sunk cost sensitivity appears at $\langle \sigma_W = 3, \sigma_N = 5 \rangle$. Most importantly, however, the sunk cost sensitivity is not linearly related to the attrition bias (linear model: $df = 34$, F vs constant model = 1.01, $p = 0.32$, adjusted- R^2 relative to constant model = 0.0004). Thus, attrition bias is neither necessary nor sufficient to create a sensitivity to sunk costs within their model.

So if the attrition bias of the willingness to wait parameter is not causing the sensitivity to sunk costs in this model, what is? As noted above, in the model, the **quit threshold** T_{WZ} decreases from being equal to the offer on entry into the wait zone to 0 at the time of reward, providing an expanding boundary to the drift-diffusion change-of-mind model. The quit threshold decreases at a constant rate of 1/s. Note that agents enter the wait zone with a willingness to wait, W_0 , that is going to be above the offer (because of the process by which the model accepts an offer). As the agent spends time in the wait zone, T_{WZ} decreases so that it will be 0 when the countdown stops. This means that T_{WZ} moves away from W_0 at a rate of 1/s as the agent waits in the wait zone, making the agent less likely to quit the longer it spends in the wait zone. (This is the literal definition of a “sensitivity to sunk costs”).

Importantly, this means that the changing threshold contains information about both the time remaining and the time spent (Fig. 6a–c). When we remove this rate of change (removing the information about time spent as we remove the information about time remaining), we find that this decreases the sensitivity to sunk costs. If we set the quit threshold to the offer, we decrease the sensitivity to sunk costs (Fig. 6d–f). If we set the quit threshold to 0, the agent never quits (Fig. 6g–i). A large part of the reason that their model is sensitive to sunk costs is that such a sensitivity is directly built into the model with the changing quit threshold. In order to produce an agent that quits but does not show sunk costs, we need to include an anti-sunk cost formulation, wherein the quit threshold increases with time spent rather than decreases, making it easier to quit the longer the agent is in the wait zone (Fig. 6j–l), suggesting that there is an additional factor critical to the presence of sunk costs in the model.

The other factor that is critical to the presence of sensitivity to sunk costs in this model is that the model has only a lower boundary but no corresponding upper boundary. This means that as time passes, the drifting willingness to wait is more likely to drift away from the threshold (Fig. 7). Thus, the variance of the term W contains information about the time spent. Because the decision process in their model is made sensitive to that information through the use of a lower but not an upper bound, their agent is less likely to quit the longer it spends in the wait zone. (Again, this is the literal definition of a “sensitivity to sunk costs”). To test the extent to which this variability drives the sensitivity to sunk costs in this model, we examined limiting the maximum drift of the willingness to wait W over the time waiting, that is, removing that sunk cost information from the decision process. Figure 7 shows that the sensitivity to sunk costs arises from the distribution of the decision variable W over time, which has been made directly sensitive to sunk costs in the model due to the expanding lower bound (the decreasing slope of T_{WZ}) and the increasing variance of W with time due to the lack of an upper bound. (Original model, Fig. 7a–c, limiting only the quit threshold, Fig. 7d–f, limiting the upper bound, Fig. 7g–i, limiting both, Fig. 7j–l).

Figures 6 and 7 show that the sensitivity to sunk costs in the model is due to the decision processes that were made sensitive to the time spent waiting (the sunk costs). The decision process is

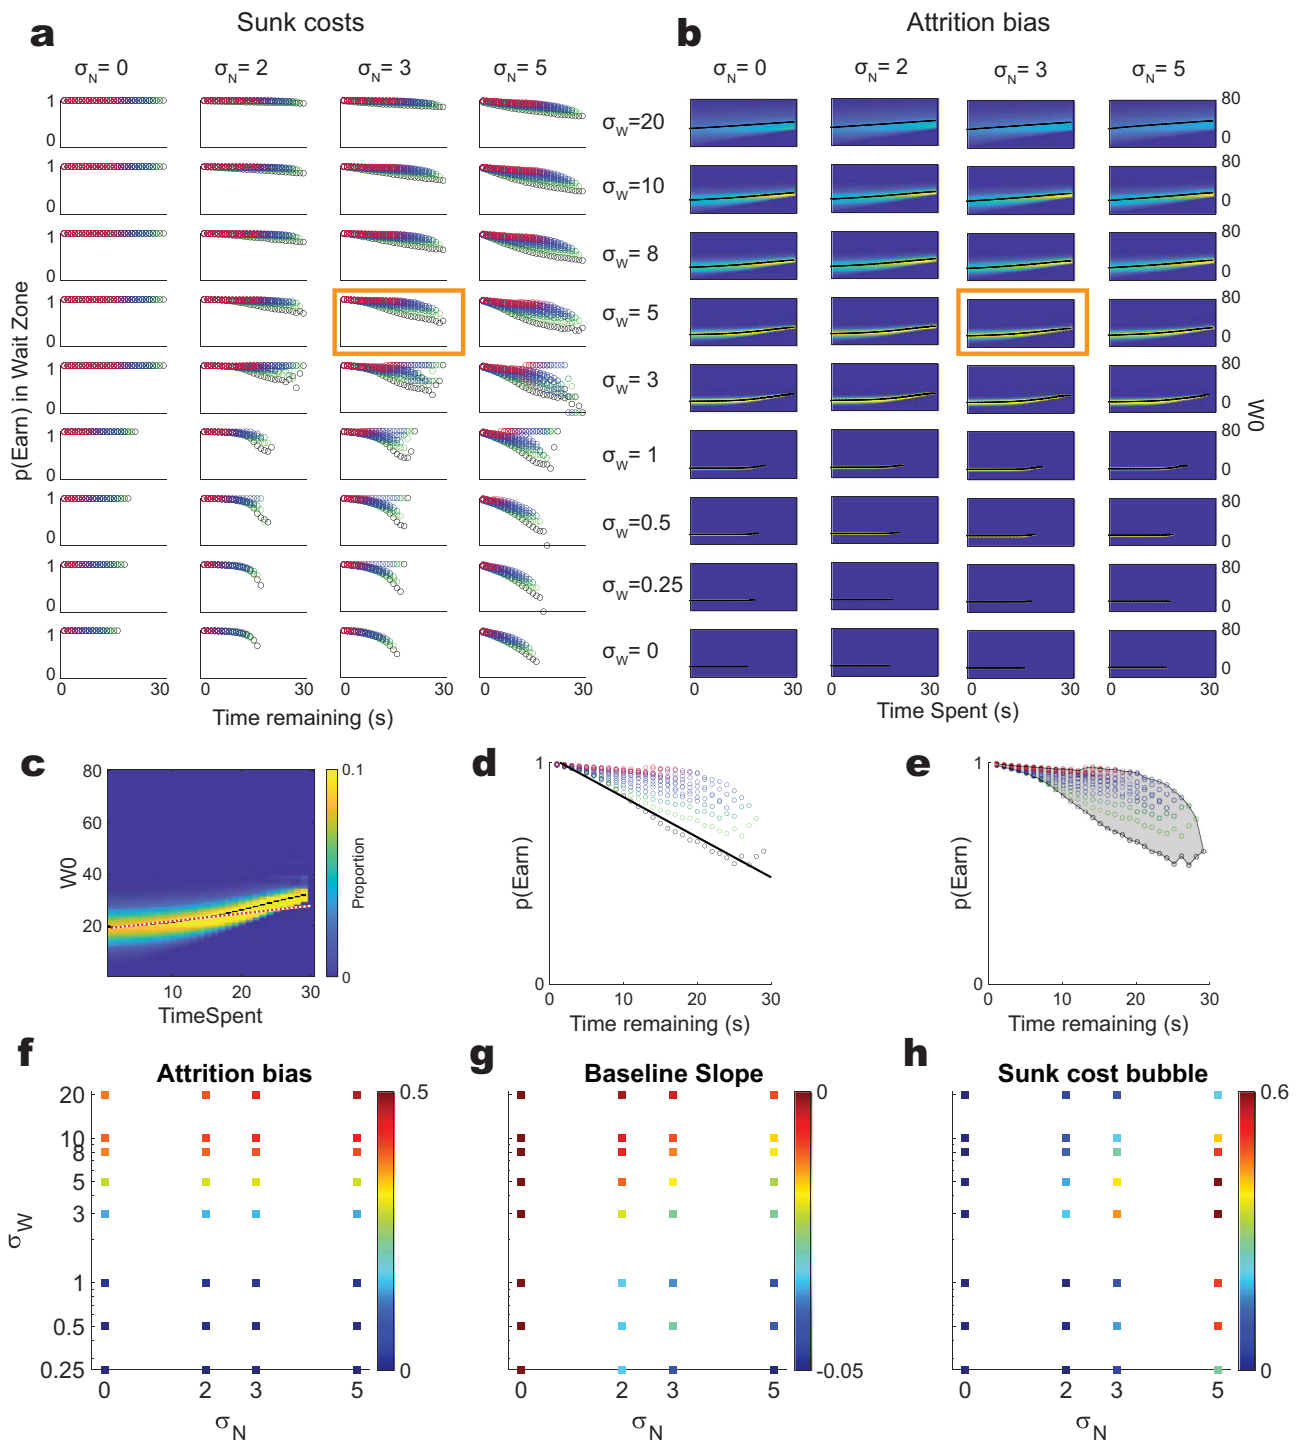


Fig. 5 A parameter exploration of the Ott et al.³⁵ model. The two key parameters of the model are σ_W , the variability in the initial motivation W_0 , and σ_N , the variability in how that motivation changes with time waited in the wait zone. We tested these parameters over a large parameter range ($\sigma_W \in \{0, 0.25, 0.5, 1, 3, 5, 8, 10, 20\} \times \sigma_N \in \{0, 2, 3, 5\}$). **a** The sunk cost “bubbles” for each parameter pair. Note that some parameter pairs show more sunk cost sensitivity than others. Colors as per Fig. 2f. **b** The attrition bias histograms for each parameter pair. Note that some parameter pairs show more attrition bias than others. **c-e** Measurements used to quantify these simulations. The orange box in a and b indicate $\langle \sigma_W = 5, \sigma_N = 3 \rangle$, which are the parameters used in³⁵ and shown in Fig. 4. **c** The attrition bias measured as a linear function of all of the samples providing data to the histogram. Attrition bias slope was measured from the actual points included in the histogram rather than the post-calculated histogram. **d** The “baseline slope” measures how much more likely the agent will wait out the delay as a function of time remaining with 0 s invested. **e** The sunk cost “bubble”, measures sunk cost sensitivity as the summed increased likelihood of waiting as a function of time waited. **f-h** How these measures change across the $\sigma_W \times \sigma_N$ parameter space. **f** Attrition bias is primarily related to σ_W . **g** The baseline slope is related to both σ_W and σ_N . **h** The sunk cost bubble is largest at high σ_N and a mid-range of σ_W . Results shown are from our simulations.

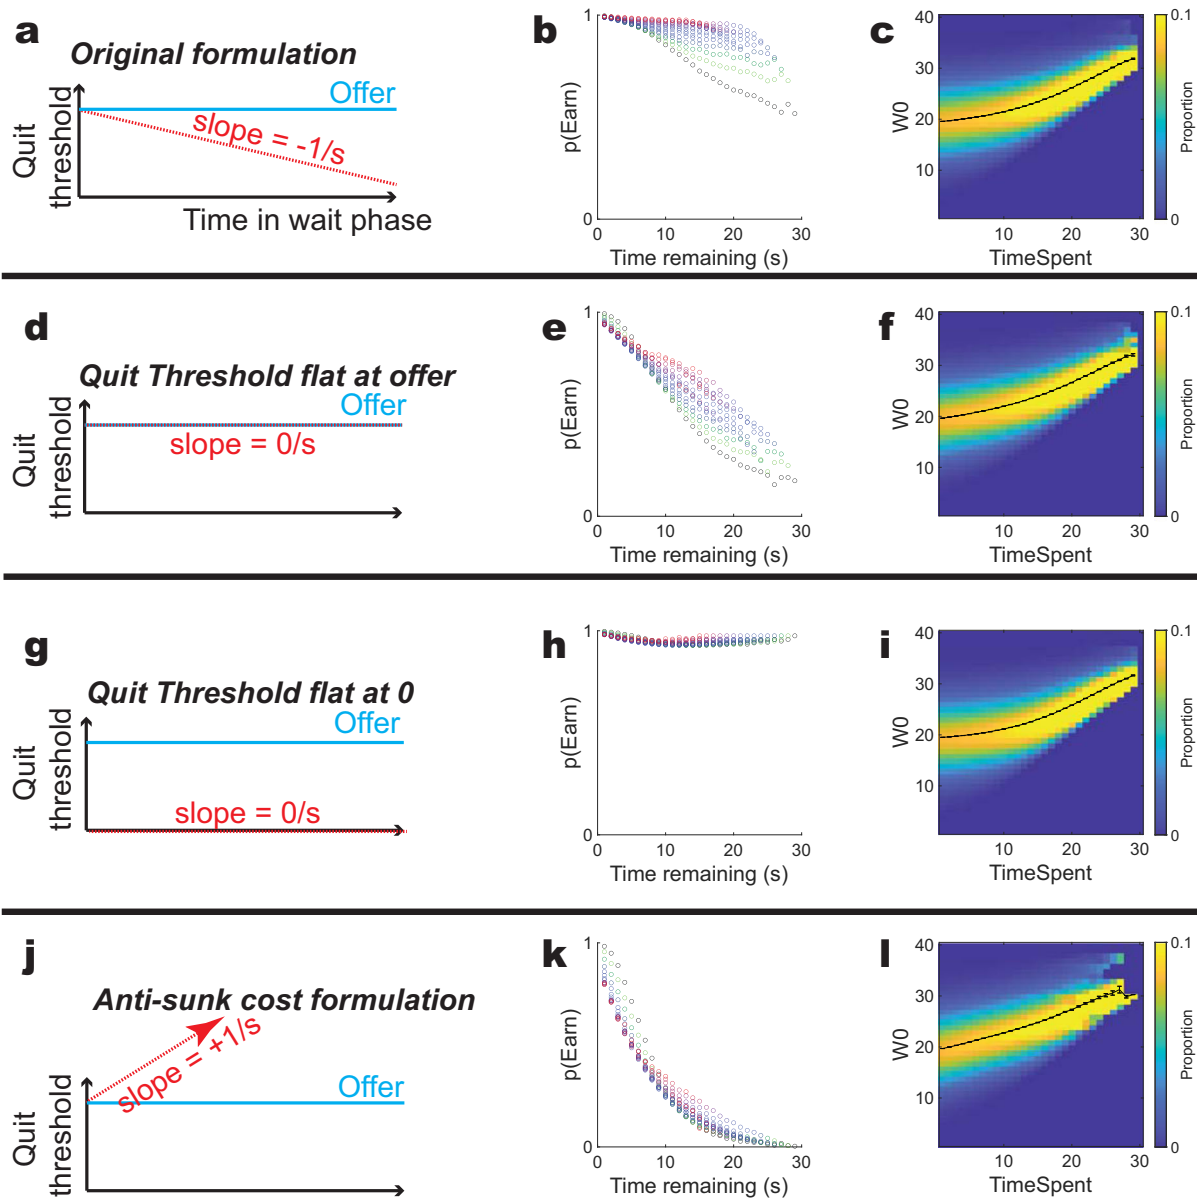


Fig. 6 Effect of changing the slope of the quit threshold T_{wz} in the Ott et al. model. **a** The original formulation. **b** The sunk cost sensitivity seen in the original formulation. **c** The attrition bias seen in the original formulation. **d-f** Effect of removing the slope of T_{wz} , while starting T_{wz} at the offer, thus making the quit threshold equal the offer for the entire time in the wait zone. **g-i** Effect of removing the slope of T_{wz} , while starting T_{wz} at 0, thus setting the quit threshold to 0 for the entire time in the wait zone. Note that because W can wander to a number less than 0, it is still possible for the agent to quit. **j-l** An anti-sunk-cost formulation in which the T_{wz} bound increases at a rate of $1/s$, providing a collapsing bound to the drift-diffusion model. The sunk cost bubble is inverted with lower $p(\text{Earn})$ the longer the agent waits in the zone. Note that even the anti-sunk-cost formulation shows attrition biases. Time-invested colors in **b, e, h, k** as in Fig. 2f. Results shown are from our simulations.

sensitive to the expanding bounds, which make it harder to quit the longer the agent is in the wait zone—a lower bound that expands away from the offer with a rate of $-1/s$ controlled by T_{wz} , and an upper bound that is driven by the ability of the W to walk away from the initial willingness-to-wait $W0$ that is controlled by σ_N .

Comparing the model to data. The preceding sections illustrate that sunk cost sensitivity within the model proposed by Ott et al. is not related to attrition bias but instead depends on decision variables/parameters that represent the time waited (the costs that have been sunk). Nevertheless, their model can serve as a starting approach to examine decision processes underlying sunk cost in mice, rats, and humans. To address this possibility, we assessed

how well the parameterizations in their model can fit the data from individual subjects.

Individual variability in the sensitivity to sunk costs. Importantly, we also find that there is a lot of individual variability in each subject’s sensitivity to sunk costs, with some mice, rats, and humans being particularly sensitive, and others not^{16,19,20}. These differences likely arise from variations in cognitive and behavioral factors, such as attention, predictability, baseline motivation (such as hunger), and other factors. Kazinka et al.²⁶ tested humans online under two conditions, one in which they were asked to attend to a button that could change color subtly, and another without such an attention-check. When humans were distracted from the delay countdown, we found that the sensitivity to sunk

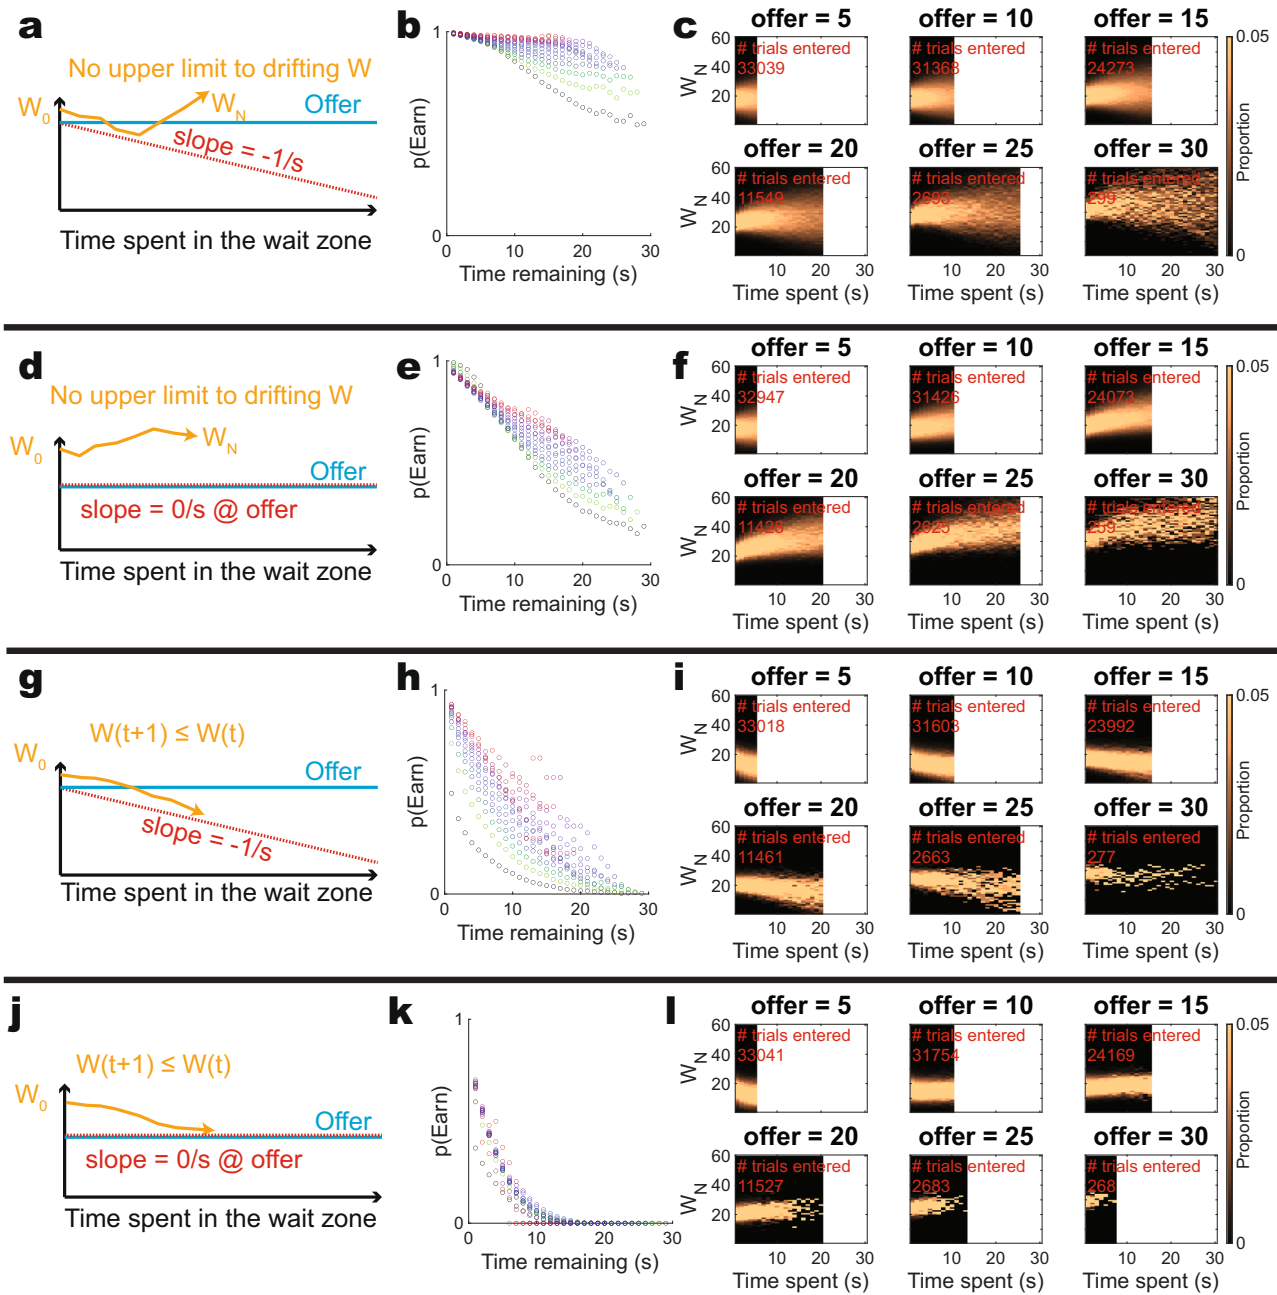


Fig. 7 Limiting the drift of the willingness to wait. **a** The original formulation limits the drift of the willingness to wait at the lower bound by T_{wz} , decreasing by $1/s$. **b** This produces sunk cost sensitivity. **c** The distribution of \mathbf{W} over time as a function of offer. Note the increasing variance with time, containing information about the time spent. **d-f** The formulation used in Fig. 6d. The distribution of \mathbf{W} over time as a function of offer continues to show the increasing variance with time, containing information about the time spent. Samples drift away from the lower bound, making it more difficult to quit with time (i.e., making the decision sensitive to sunk costs). **g-i** A formulation with the original lower bound, but a new upper bound, restricting the change in \mathbf{W} to always decrease. If $\mathbf{W}(t)$ was slated to rise above $\mathbf{W}(t-1)$, it was set to $\mathbf{W}(t-1)$. Note that there is still information about sunk costs due to the time information in the descending lower bound slope. **j-l** If we include both flat upper and lower bounds, the sensitivity to sunk costs disappears. Time-invested colors in **b, e, h, k** as in Fig. 2f. Results shown are from our simulations.

costs disappeared, but that it remained intact in the version without the attention-check²⁶. In rodents, we found differences in the sensitivity to sunk costs as a function of environmental richness (where the distribution of delays was shifted lower so more food could be gathered more quickly or shifted higher requiring more time to gather food)^{16,36}. Rodents also exhibit differences in sunk cost sensitivity as a function of the predictability of the upcoming delay—delays that could be predicted

produced less sensitivity to sunk costs than delays that were only revealed on entry into the restaurant³¹.

As noted above, the Ott et al.³⁵ model provides two key parameters (σ_N and σ_W) that produce different sunk cost sensitivities (Fig. 5). While it is clear that the parameters chosen by Ott et al.³⁵ [$\sigma_N=3$ and $\sigma_W=5$] would not produce this variation, it is possible that different parameter configurations could capture the variability in sunk cost sensitivity seen across subjects.

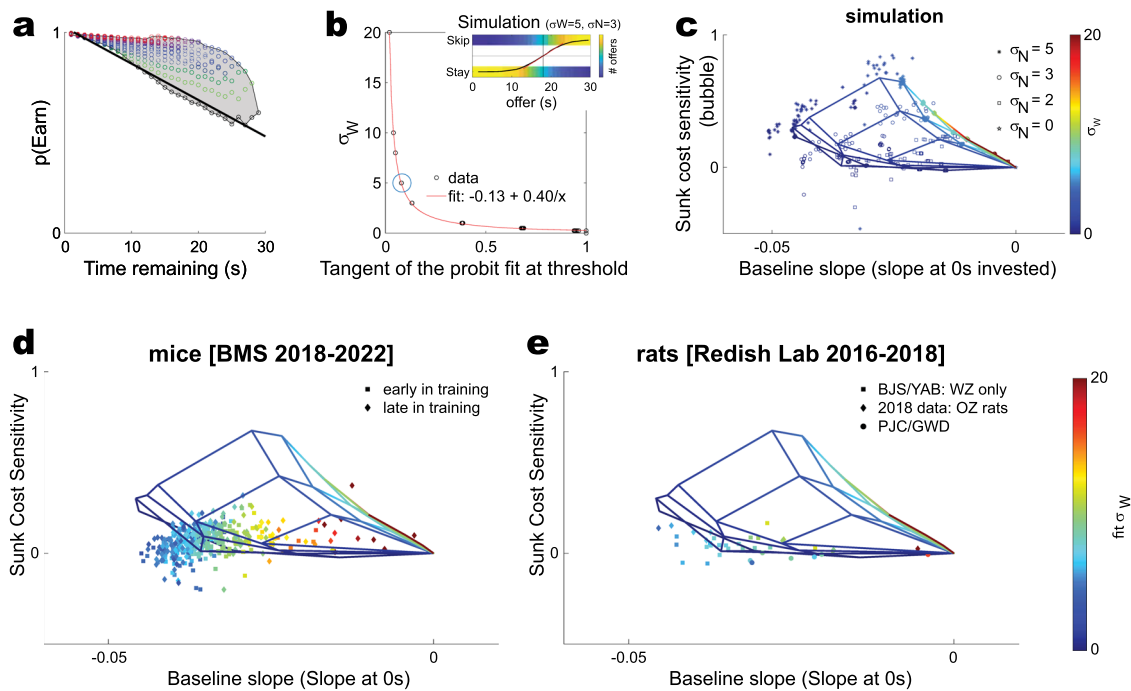


Fig. 8 Individual variability. The basic model proposed in³⁵ has two parameters σ_W (which controls the offer zone accept/skip decision) and σ_N (which controls the drift rate within the wait zone). **a** As shown in Fig. 5, each of these parameters produces a sunk cost sensitivity, measurable as the area of the increased likelihood of earning as a function of the time invested (the “sunk cost bubble”), and a specific slope of the likelihood of quitting at 0 (“baseline slope”). **b** Although we know of no way to measure σ_N from the behavioral data directly, σ_W is directly related to the slope of the psychophysics curve of accepting an offer. Inset panel shows the original model ($\langle \sigma_W = 5, \sigma_N = 3 \rangle$), the histogram of offers for which the agent stayed, offers for which it skipped, the fit probit model (in black), and the slope at threshold (in red). **c** Thus, for each simulation we can measure the “sunk cost bubble”, the “baseline slope”, and σ_W . We ran 20 simulations at each part of the parameter space from Fig. 5. Symbols for each simulation are shown. The mesh grid indicates the average for each parameter pair. Color on the mesh is interpolated from σ_W . **d** Mice from ref.^{19,20}. Each mouse provides one symbol from early in training and one symbol from late in training, placed at the measured baseline slope and sunk cost sensitivity. Color is the fit σ_W , given the slope at threshold. $n(\text{mice}) = 208$. **e** Rats from^{19,21} ($n = 21$ rats with wait zone only, 10 rats with both wait and offer zones) and new data from the Redish lab ($n = 15$ rats).

Although we know of no behavioral way to directly measure the σ_N parameter, the σ_W parameter is related to the willingness to take economically inconsistent offers above and below the observed threshold, which appears as the slope of the psychophysics decision curve of accept vs skip decisions made in the offer zone. That is, we can consider a subject’s decisions to accept or reject an offer as a standard decision process which shows a typical psychophysics curve with a slope β . The flatter the slope, the more likely the subject will pick a choice inconsistent with its personal threshold. Because the decision in the Ott et al. model is based on whether $W_0 > \text{threshold}$ or not and W_0 is drawn from a distribution with variance σ_W , the slope will be directly related to that variance σ_W . To confirm this, we ran the simulation with 10 samples each of varying σ_W parameters and measured the tangent at the 50% threshold of a fit probit model. We found a remarkably good relationship (Fig. 8b), such that $\sigma_W \approx -0.13 + 0.40/x$, where x is the tangent of the probit model at the threshold (i.e., the point where the probability of skipping versus accepting an offer as a function of the delay offered is 50%). Importantly, this tangent can be directly measured from data to derive an estimate of σ_W for each subject.

Humans only experience one session each, and thus we do not have enough data to fit sunk cost sensitivity for each individual subject. (The human data sets measure sunk cost sensitivity over the whole population^{19,23,26}. However, the rats and mice each run the Restaurant Row task for many days^{16–18,20,21,30}, and thus, we can measure the three key parameters for each individual subject. We can measure (1) the sunk cost bubble and (2) the baseline

slope (Fig. 8a) from the choices made by the subject. (3) We can estimate σ_W from the tangent at the probit fit of the subject’s threshold (Fig. 8b). Thus, each subject can be placed on a three-dimensional plot of baseline slope, the sensitivity to sunk costs, and the fit σ_W .

To compare the model’s ability to capture the individual variability, we first ran 20 simulations of each parameter instantiation of the model (σ_W ranging over $\{0, 0.25, 0.5, 1, 3, 5, 8, 10, 20\}$, and σ_N ranging over $\{0, 2, 3, 5\}$, as shown in Fig. 5). We then measured the average baseline slope and sunk cost sensitivity for each simulation. We constructed a mesh grid of the mean at each $\langle \sigma_W, \sigma_N \rangle$ pair, color coded by σ_W (Fig. 8c). Repeating this grid for the mice (Fig. 8d) and the rats (Fig. 8e) shows that the actual observed data does not track the grid very well.

To compare these statistically, we found the closest mean point on the simulation matrix mesh for the $\langle \sigma_W, \sigma_N \rangle$ with the fit σ_W for the individual, and then measured the distance between that point and the individual’s baseline slope and the distance between that point and the individual’s sunk cost sensitivity. If these differences are due to noise, they should center at 0 with a Mann-Whitney sign rank test. Every dataset was significantly different in either the baseline slope or the sunk cost sensitivity or both. (Fig. 8d: $n = 208$ mice early in training¹⁹ $p(\text{sunk cost sensitivity}) = 0.12$, $p(\text{baseline slope}) = 10^{-35}$; $n = 208$ mice late in training¹⁹ $p(\text{sunk cost sensitivity}) = 0.17$, $p(\text{baseline slope}) = 10^{-35}$; Fig. 8e: $n = 21$ rats without an offer zone¹⁹ $p(\text{sunk cost sensitivity}) = 0.5$, $p(\text{baseline slope}) = 10^{-5}$; $n = 10$ rats with both zones¹⁹ $p(\text{sunk cost$

sensitivity) = 0.04, $p(\text{baseline slope}) = 0.002$; $n = 15$ rats with both zones (new RedishLab rat data) $p(\text{sunk cost sensitivity}) = 0.0001$, $p(\text{baseline slope}) = 0.00001$). As can be seen in these statistics, the model can only match the sunk cost sensitivity by showing completely incompatible baseline slopes.

Sunk cost sensitivity and the offer zone. The most important result that came from the introduction of the offer zone variant¹⁹ is that we did not see a sensitivity to sunk costs in the offer zone. Subjects spent time deciding in the offer zone even though time spent in the offer zone counted against the total economic time budget that the subjects could use to forage for food. However, time spent in the offer zone did not show a sensitivity to sunk costs; the calculation of time spent (the cost that had been sunk) did not start to accrue until the agent entered the wait zone. This result has been replicated in online samples^{23,26} and subsequent experiments²⁰. In general, time spent in other portions of the task did not influence the probability of quitting after accepting an offer, implying that time spent in the wait zone carries a unique sensitivity to sunk costs³⁰.

As we have argued elsewhere^{16,19,26}, the most parsimonious explanation is that behavior in the offer and wait zones arise from different decision processes, consistent with both multiple-decision theories^{37,38} and commitment hypotheses^{39,40}. Importantly, it is not rational to spend any time in the offer zone, as any time spent making decisions in the offer zone could be more productively spent being made in the wait zone, while the delay is counting down. However, all three species spend more time than necessary in the offer zone, and mice and rats, at least, increase their time spent in the offer zone with experience^{16,19}. One possible explanation for this is that it is difficult to quit out of the wait zone due to the accumulation of sunk costs^{16,17,19,36}.

As shown in Figs. 4–7, the reason that the model laid out in ref. 35 shows an escalation of commitment with time spent is not due to the theoretical arguments made by Ott et al. therein (that sunk cost sensitivity seen on the Restaurant Row and WebSurf tasks is due to a statistical attrition of initial motivation) but rather to decision processes built into the model that are explicitly dependent on time spent in the wait zone (i.e., sunk cost). Nevertheless, even though the model proposed in³⁵ is incompatible with the theory proposed in³⁵, the model itself can provide an interesting explanation for the difference between the offer and wait zones because it includes change-of-mind decision processes^{32,33} that are explicitly sensitive to time spent in the wait zone, but not the offer zone. In that model, the offer zone decision is made based on the initial willingness-to-wait W_0 , which starts to drift once the agent enters the wait zone. Similarly, the quit threshold T_{WZ} only starts decreasing in the wait zone. Thus, the decision process is sensitive to the time spent in the wait zone, but not the offer zone.

Sensitivity to sunk costs is delayed in variants without an offer zone. In a recent variant without an offer zone, we directly manipulated the upcoming uncertainty of the future outcomes³¹. In two parallel tasks, rats approached a series of reward sites which either had a known delay (different across the four reward sites, but constant within each reward site for a given day, changing from day to day, thus predictable after experiencing one lap around the maze) or had a random delay to reward (1s–30s random on each entry, thus unpredictable). In the first task (*Known-Delay [KD]*), rats knew what the upcoming delay would be and showed behaviors indicating that they had already made their decision before entering the reward zone—they maintained a fast speed on approach to offers they skipped. In the second task (*Random-Delay [RD]*), rats slowed down on every entry into an

offer and then only sped up to leave after a few seconds of consideration. Furthermore, decisions in the KD task were more self-consistent, that is, rats were more likely to either choose to stay or skip for a given delay on the KD task. We found that sunk cost sensitivity was decreased in the KD task, but were still present. Importantly, however, we also found that sunk costs did not start accruing in the RD task until after the decision-time (~5 s), while sunk costs started accruing in the KD task immediately.

That subjects make a separate decision before committing to a choice and then show a sensitivity to sunk costs only after commitment^{16,19,26,39,40} suggests that one should see this difference, even without an explicit offer zone. And, in fact, that is what has been found. The early versions of the Restaurant Row and Web-Surf tasks did not have a separate offer zone, and only included a wait zone^{15,21,24}. A reanalysis of these data found that sunk costs did not start accruing for several seconds¹⁹. We hypothesized that animals were spending the first few seconds making an *accept/skip* decision as if they were in an offer zone rather than a *quit/remain* decision. We hypothesized that they then covertly transitioned to a *quit/remain* decision and began building sunk costs. We found robust sunk cost sensitivity in these animals, but only after a delay, consistent with typical decision times made in variants with an explicit [separate] offer zone.

Delayed accrual provides an opportunity to test the reset hypothesis. This delayed latency provides an opportunity to address a question raised by Ott et al. in their paper — when the subject finally makes the decision to stay in these single-zone variants, they have already been in the zone for several seconds. So we can ask whether the subsequent quit or remain decisions take into account the initial decision time. (Remember that we do not see sunk cost accrual during similar decisions made in the separate offer zone). We can test this by comparing two variants of the Ott et al. model: one in which T_{WZ} starts descending only after completion of the self-imposed (hidden) decision-time (Fig. 9a) and one in which T_{WZ} has already been descending for those several seconds (Fig. 9e). In other words, *does the increased difficulty to quit with time (the escalation of commitment) include that decision time or not?*

We modeled the single-zone construction by having the agent accept all offers (thus ignoring the offer zone, effectively modeling a task variant without an offer zone) and the decision time by simply preventing the agent from quitting for the first few seconds of waiting in the wait zone. Fascinatingly, the two simulations produced fundamentally distinct shapes of sunk cost outcomes, more specifically, the relationship between differences in the latency to accrue sunk costs and maximal sensitivity to sunk costs (Fig. 9d, h). These two models produced a subtle difference in the maximum sunk cost sensitivity levels seen as a function of the time spent deciding. The first model (T_{WZ} is at offer at the end of the decision time and starts decreasing only after the first-stage decision (to stay) has completed) showed a convex curve as a function of the time spent in the pre-wait-zone/hidden decision time, where particularly short and long decision times were less sensitive to sunk costs than mid-range decision times. (Fig. 9b–d, linear fit adjusted- $R^2 = 0.98$, quadratic fit adjusted- $R^2 = 0.99$). In contrast, the second model (T_{WZ} starts decreasing immediately) showed a concave curve, in which mid-range decision times were less sensitive to sunk costs. (Fig. 9f–h, linear fit adjusted- $R^2 = 0.98$, quadratic fit adjusted- $R^2 = 1.00$).

Comparing these simulations to new analyses of data from rats on versions of the task without an offer zone^{19,21,22}, and the subsets of mice who took all offers and effectively ignored the offer zone choice stage¹⁶, we found a fascinating difference between the rats and mice.

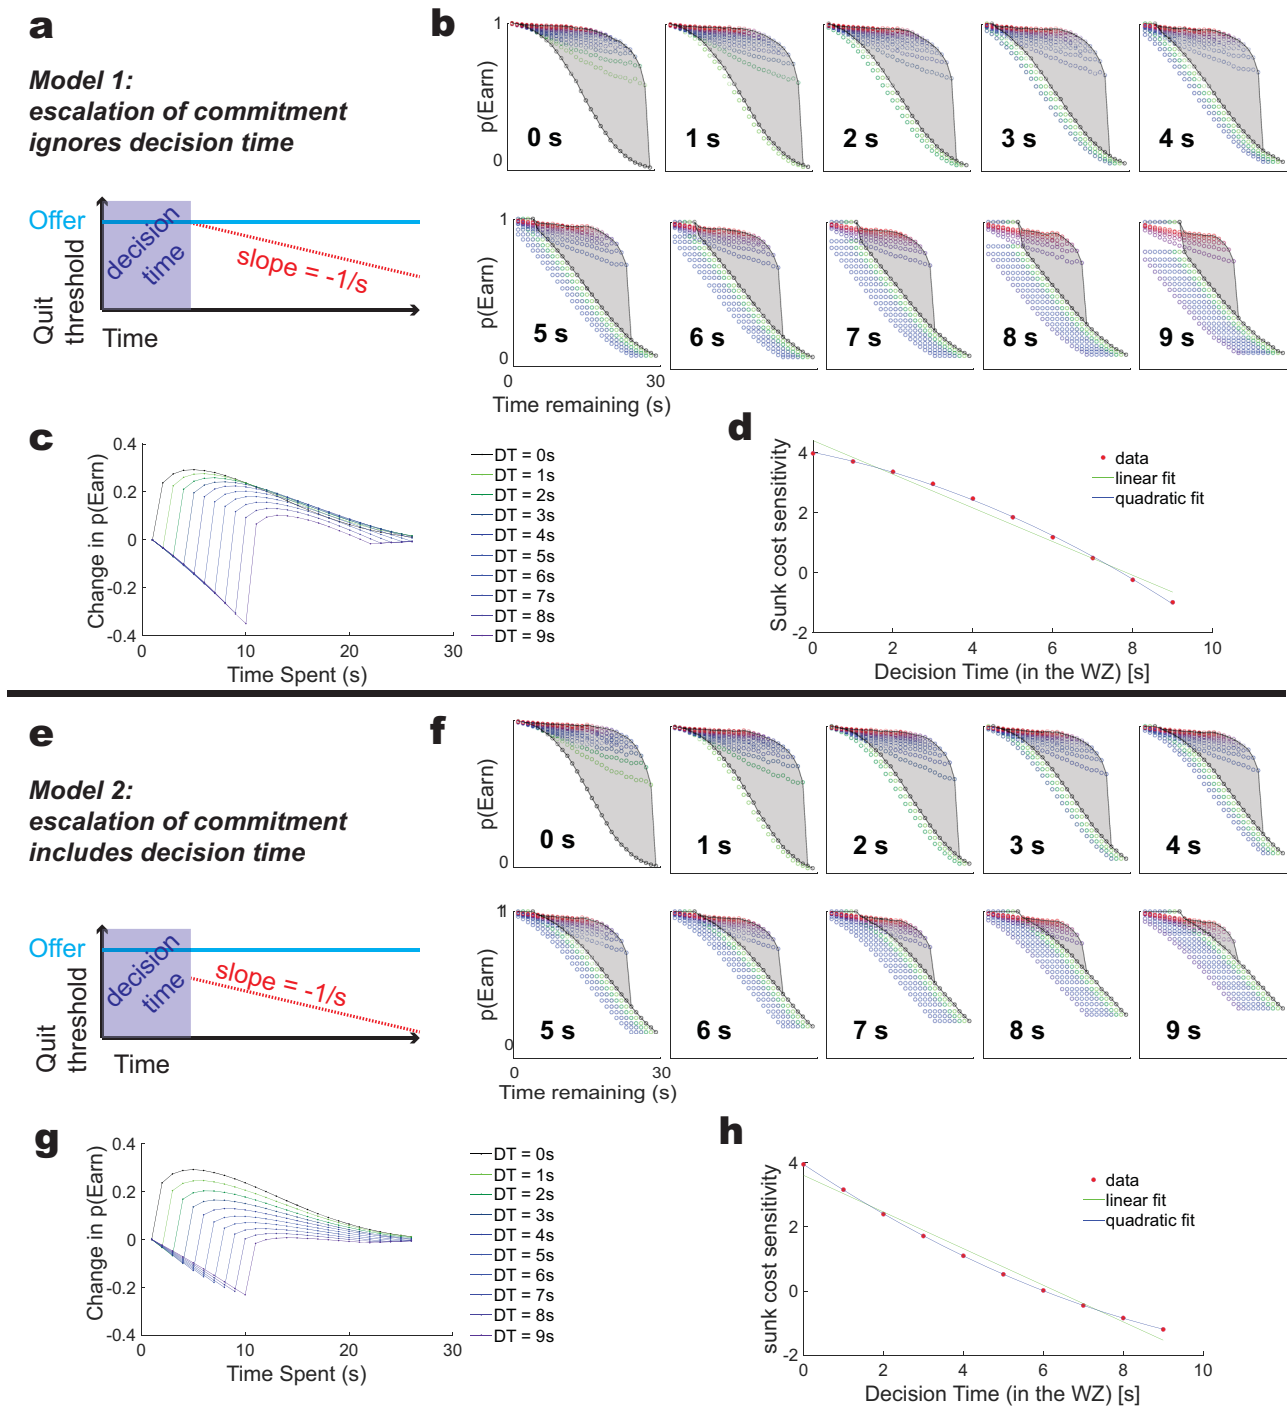


Fig. 9 A delay to accrual allows us to test the reset hypothesis. In variants without an offer zone, and in some subjects with limited experience with a separate offer zone, subjects do not quit for a short period of time in the wait zone (a few seconds). This short non-quit time in these subjects can be theorized to be a separate decision process^{15,16,19,36}. Modeling this as an inability to quit for a few seconds opens up two models: *Does the escalation of commitment take that decision time into account or not?* **a** Model in which sunk cost accrual starts only when the decision time (DT) is over. **b** The “sunk cost bubble” changes as a function of the delay before the agent can start quitting. Each panel shows a decision time of x seconds. **c** Comparing the total change in p(Earn) [the size of the sensitivity to sunk costs] as a function of time spent and a given decision time shows that these sensitivities are dependent on both time spent (time invested, costs sunk) and the decision time. **d** For each given decision time, we summed the total sunk cost sensitivity. This shows a significantly convex curved fit (linear fit adjusted- $R^2 = 0.98$, quadratic fit adjusted- $R^2 = 0.99$). **e** Model in which sunk cost accrual includes the decision time. **f-h** This produces a significantly concave curved fit (linear fit adjusted- $R^2 = 0.98$, quadratic fit adjusted- $R^2 = 1.00$).

The rats showed evidence that they did not start to accrue sunk costs until the decision had been made (Fig. 10a-c). Figure 10b shows that rats without an offer zone (i.e., with only a wait zone) delayed their sunk cost accrual by a few seconds^{15,19,31}. We thus

measured whether the sunk cost sensitivity was better explained by model 1 (no decay in the quit threshold through the decision time, Fig. 9a) or model 2 (quit threshold decays through decision time, Fig. 9e). There was one data point well apart from the other

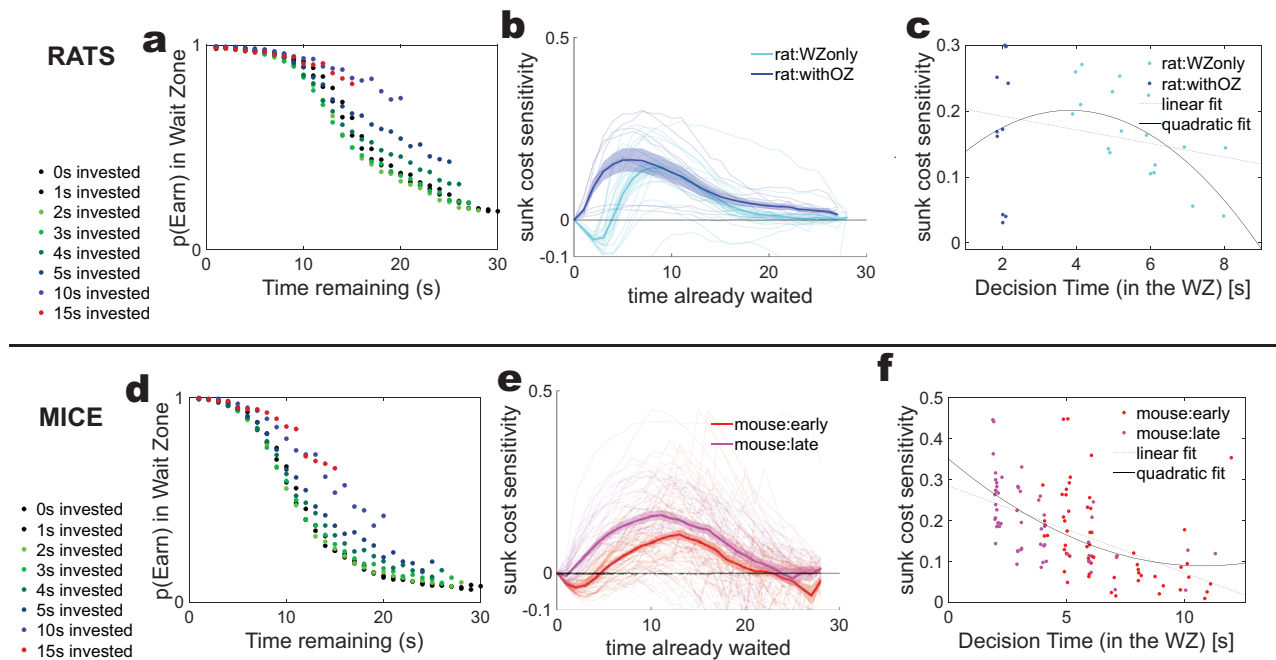


Fig. 10 Fitting sunk cost sensitivity by decision time for rats and mice. **a** The sunk cost bubble for one example session from rats without an offer zone. Note that the first few seconds invested are unchanged relative to the 0 s invested condition. Compare Fig. 2f. **b** The sunk cost sensitivity as a function of time spent (time already waited). Note that the rats with the offer zone start accruing sunk costs immediately, but rats without an offer zone do not start accruing sunk costs for a few seconds. Compare Fig. 9c. **c** Fitting this curve by decision time finds a significantly convex curve ($n = 21$ rats without an offer zone and 10 rats with, quadratic adjusted- $R^2 = 0.13 >$ linear adjusted- $R^2 = 0.04$), suggesting Model 1 (Fig. 9a) that rats without an offer zone are not including the decision time in the expanding T_{wz} boundary. **d** One example mouse session showing that the probability of waiting out the delay given that the mouse has already invested a few seconds is not different from the 0 s invested data. **e** Early in training mice do not accrue sunk costs immediately, but take a few seconds before accruing them. Late in training, mice start accruing more immediately. **f** Fitting this curve by decision time finds a significantly concave curve ($n = 32$ mice, each mouse provides one data point from early training and one from late training, quadratic adjusted- $R^2 = 0.31 >$ linear adjusted- $R^2 = 0.29$), suggesting Model 2 (Fig. 9d) that mice early in training are including the decision time in the expanding T_{wz} boundary. See text for discussion of outlier removal.

samples (Latency to Accrual = 29 s, all other data in the range of 4 s to 8 s). With this outlier included, the quadratic fit was not better than the linear fit (linear fit adjusted- $R^2 = 0.14$, quadratic fit adjusted- $R^2 = 0.11$); however, with this outlier excluded, the quadratic fit was better and was convex (linear fit adjusted- $R^2 = 0.04$, quadratic fit adjusted- $R^2 = 0.13$). Given the single point that was far away from the main distribution, we treat it as an outlier. The rest of the data suggest that rats were treating this “wait-zone-only” version of the task as having a conceptual “offer zone”, and thus that the decision time is not incorporated into the perception of sunk-costs. In other words, the rats did not start to accrue sunk costs until after the decision was made and they had left their self-imposed “offer zone” (consistent with model 1, Fig. 9a). This observation is consistent with previous experiments on rats without an offer zone^{15,19,21,31}.

In contrast, the mice showed evidence that sunk costs accrued through the whole time in the wait zone (consistent with the second model that T_{wz} starts decreasing immediately) Fig. 10d–f. Figure 10e shows that early in their experience, the mice did not start accruing sunk costs for a few seconds, while late in their training, sunk costs started accruing on entry into the wait zone. This is likely due to the mice learning to use the offer zone to avoid sunk costs in the wait zone¹⁶. The data shown has outliers removed as per the rat data, below, however, for the mice, both analyses with and without outliers removed showed concave quadratic fits that were better than the linear fits. (Data shown, linear adjusted- $R^2 = 0.29$, quadratic adjusted- $R^2 = 0.31$). These results suggest that the mice started accruing their sunk costs immediately and that the escalation of commitment (i.e., increased difficulty to quit) includes

the complete time spent in the wait zone (consistent with model 2, Fig. 9e). This has important implications for how the internal value state of the agent has changed during the passage of time while waiting and how this state compares to the perceived value of time left in the countdown registered at the transition point when sunk costs begin to accrue. In other words, when a mouse looked up from the decision time and found itself already in the wait zone for several seconds, it took that total time already waited in its sunk cost calculation, in effect resetting its sunk costs based on the actual time spent.

Importantly, the subset of rats analyzed in Fig. 10 came from an experiment in which they never had an offer zone;²¹ they only had a wait zone. This makes their experience very different from the other subjects analyzed in this paper, including the mice shown in Fig. 10, who had an offer zone available, but had not learned to use it yet, although they did eventually learn to make their decisions in the offer zone¹⁶.

Discussion

A sensitivity to sunk costs can be operationally defined as an escalation of commitment with continued time spent in an operation^{1,2,4}. An escalation of commitment with time spent waiting for reward (a sensitivity to sunk cost) has been observed in some variants of the Restaurant Row and Web-Surf tasks, particularly in the wait zone but not the offer zone of these tasks^{19,23,26}.

Ott et al.³⁵ proposed that this observed escalation of commitment may be due to a statistical attrition bias in the initial motivation, such that longer times-spent waiting were correlated

with increased initial motivations, thereby resulting in a higher probability of earning reward with more time spent. Our thorough simulations of their model find that while this attrition bias does exist, it is unrelated to the presence or absence of a sensitivity to sunk costs. Although theory in ref. ³⁵ proposed that the escalation of commitment seen in these tasks is due to statistical attrition, we show that this is not a viable explanation for the observed phenomenon. Ott et al. also proposed a specific simulation model that shows a sensitivity to sunk costs. Our rigorous analysis of their model finds that sunk cost sensitivity appears in their model due to two phenomena in which decision processes were explicitly made sensitive to time spent (sunk costs): the escalation of commitment arises from the expanding lower boundary provided by the decreasing quit threshold T_{WZ} and from the increasing range available to the drifting motivation (willingness to wait, W) over time.

This is not the first model to show how a simple decision process can produce an escalation of commitment and a sensitivity to sunk costs. In fact, several previous animal learning theory models have been put forward that produce similar sunk cost sensitivity effects, including the state-dependent-valuation learning [SDVL]^{7,9,41} and within-trial-contrast [WTC]⁴² models. The SDVL model notes that ongoing energy expenditures imply that food rewards will become more valuable to an agent as time progresses. Thus, reward valuation should depend on the time since the last reward was received. Of course, the SDVL model cannot explain why humans foraging for videos show similar effects, but we can hypothesize a more general “value process” that depends on recency of experience. The WTC model notes that distant rewards are discounted so that approaching rewards are seen as increasing in value, which can provide an increasing contrast between the current state of the animal (changing through SDVL) and the approaching reward⁴².

As noted in ref. ¹⁹, neither the SDVL nor the WTC models can explain why sunk costs only seem to begin accruing after an initial commitment decision has been made (entering the wait zone of the tasks) and why a parallel sensitivity to time spent is not also present in the offer zone. Moreover, neither model can explain why a lack of attention to the decreasing delay reduces sensitivity to sunk costs²⁶, why predictability of upcoming costs reduces the sensitivity to sunk costs³¹, or why the richness of the environment affects the sensitivity to sunk costs^{16,20,36}.

The new model proposed in ref. ³⁵ may provide an interesting hypothesis. Although it is unclear why the quit threshold expands away from the offer at a rate of $1/s$, if the offer zone and wait zone are decided upon by different decision processes^{16,18,19,23,26,29,31,43}, this quit threshold that only starts decreasing on entry into the wait zone might provide an explanation for the accrual of sunk costs after commitment. In the versions of the task with separate offer and wait zones, the delay does not begin counting down until entry into the wait zone^{16,19,23,26}. Similarly, if a new drift diffusion process begins only with commitment, change of mind behaviors^{32,33} might be sensitive to sunk costs only after that commitment^{39,40}, which would be well modeled by the increasing variance of the drift-diffusion process providing more opportunities to have increased motivation with time spent (an escalation of commitment). Applying their model to versions with a single zone (wait zone only, no offer zone) suggests that rats who never experienced an offer zone did not start to accrue sunk costs until making the commitment to stay in the wait zone, effectively treating the initial decision time as a separate offer zone. However, mice that had not yet learned to use the available offer zone showed an accrual of sunk costs that took into account all of the time within that wait zone. Whether this is a species difference (rats vs mice), a difference in the tasks (no offer zone available vs have not yet learned to use the offer zone), or both, remains unknown.

Also present within the data are observations that are incompatible with these simple drift-diffusion models, including the observation that sunk cost sensitivity is stronger when subjects have taken a bad deal²⁰ and that some subjects show this escalation of commitment, while others do not^{18,26}. We suspect that one potential explanation of the data is the increased motivation provided by Pavlovian associations between the countdown in the wait zone and the value of the outcome. This would predict a decreased sensitivity to sunk costs with reduced attention (consistent with the online human data²⁶), as well as a heightened sensitivity to factors that increase that relationship to the reward zone (such as conditioned place preference)^{44–47}. This is a form of the endowment effect^{48–50} and is similar to the deliberative vs implementational mindset constructs in the human literature^{39,40}. Essentially, this hypothesis suggests that quitting out of the wait zone is a recognition of a mistake and staying is due to an unwillingness to leave that mistake. Because the extent of the mistake depends on the effort spent (more effort spent was a larger mistake), the decision becomes related to the effort already spent and shows a sensitivity to the sunk costs.

It is worth noting that there are situations where an increased motivation from effort already spent can actually be economically useful (see ref. ^{45,51,52} for discussion), which could further provide evolutionary drive to make decisions based on past costs. There are many cases where sunk costs can provide the additional incentive to push through difficult choices. For example, when pushing through the last part of a marathon run, a common coaching line is “you’ve come so far already, you can finish!” Similar incentives can be seen in the last segment of any long task that requires pushing through burnout and exhaustion, such as a PhD thesis. In truth, the correct comparison is between quitting (with no reward but no additional costs) and continuing (with the reward of finishing but with the additional costs of pushing through the burnout). It is possible that the costs of pushing through may appear over-burdensome because of economic myopia. The additional motivation provided by the refusal to quit (due to sunk costs) can allow one to achieve goals that may be difficult to achieve otherwise and may be linked to resilience. This motivational effect can be interpreted as the following logic: in a long marathon, there is uncertainty in that future valued goal. Thus it becomes easy to conclude that the goal is not worth the additional effort, particularly as that additional effort increases with exhaustion. But because the sunk costs provide additional evidence that the goal will be worth it, the past effort spent provides increased motivation.

Methods

Measures used. **Threshold** was found by a simple maximum-likelihood estimation over all possible thresholds from 0 s to 30 s. Value is defined as the difference between that threshold and the offered delay. (Thus an offer of 10 s with a threshold of 18 s provides a value of +8 s, while an offer of 20 s with a threshold of 18 s provides a value of –2s).

p(Earn) was measured as the probability that an agent (mouse, rat, human, simulation) in a given (time-remaining, time-spent) pair remained to eventually earn the reward. Thus, if an agent has taken a 5 s offer, waited 2 s, it is in a [time-remaining = 3 s, time-spent = 2 s] pair. If the agent waited out the remaining 3 s to earn the reward on that offer, we counted that as an “earn”, while if the agent quit on the next second (after 3 s, with 1 s remaining), we counted that as a “not-earn”.

The “sunk cost bubble”. From the pEarn matrix, for each time-invested (time-spent), we summed the difference in probability of waiting out the delay (earning) between each time-remaining point and the probability of earning at that time-remaining point with 0 s invested. The measure was then the sum (over time-spent conditions) of those summed measures.

Baseline slope. We measured a linear fit to the probability of earning with 0 s invested. The measure was the slope of this fit line.

Attrition bias. The simulations provide **W0** values for each second that the agent is “in the wait zone” (i.e., took the deal, has not earned, and has not yet quit). We measured a linear fit to this scatter plot of these **W0** samples as a function of the time spent. The measure was the slope of this line.

Simulations. Simulations were based on those of ref. 35. A simulated agent was designed that made two decisions: first an enter/skip decision (modeling the offer zone in Restaurant Row and WebSurf) and then a repeated quit/remain decision (modeling the wait zone) made once per second.

Offer zone decision process: In a given trial, the agent started with a draw of a willingness to wait, $W_0 = W_T + N(\sigma_W)$, drawn from a normal distribution around the threshold ($W_T = T_{OZ}$ defined as 18 s) with standard deviation σ_W . If W_0 was greater than the threshold ($W_0 \geq T_{OZ}$), the agent accepted the offer and “entered the wait zone”.

Wait zone decision process: The willingness to wait $W(t)$ drifted each second by adding in normally distributed random noise with standard deviation σ_N . $W(t+1) = W(t) + N(\sigma_N)$ once per second. If $W(t)$ decreased below the quit-threshold ($W(t) < T_{WZ}$), the agent quit at time t . If the agent did not quit through the full delay, the agent “earned reward”.

The quit-threshold was defined as the linear function such that $T_{WZ} = W_0$ at the start of the waiting time and decreasing by 1 s each second until it reached $T_{WZ} = 0$ at the end of the waiting countdown. That is, $\Delta T_{WZ} = -1/s$. If the cumulative willingness to wait (integrating the drift over time) ever fell below the quit-threshold TWZ, the agent “quit” out of the wait zone.

The **attrition bias** was measured as the distribution of initial W_0 willingness-to-wait that remained after a given time spent.

Trials were independent. The simulation did not do any learning and was assumed to follow the above decision process. Offers were uniformly distributed over 1s–30s inclusive. 1 million trials were run for each experiment.

Figure 4b–f: The initial parameter set used by Ott et al.³⁵ was tested:

$W_T = T_{OZ} = 18$ s. $\sigma_W = 5$ s. $\sigma_N = 3$ s.

Figure 4g–k: The offer zone decision process was changed to always accept the offer.

Figure 5: We did a parameter sweep running a full experiment (1 M samples evenly divided over 1s–30s offers inclusive) for each combination of ($\sigma_W \in \{0, 0.25, 0.5, 1, 3, 5, 8, 10, 20\}$ x $\sigma_N \in \{0, 2, 3, 5\}$).

Figure 6a–c: The initial parameter set and the initial model as described above.

Figure 6d–f: $T_{WZ} = \text{offer}$ and $\Delta T_{WZ} = 0/s$.

Figure 6g–i: $T_{WZ} = 0$ and $\Delta T_{WZ} = 0/s$.

Figure 6j–l: T_{WZ} was initialized to the offer and $\Delta T_{WZ} = +1/s$.

Figure 7a–c: The initial parameter set and the initial model as described above.

Figure 7d–f: The modification used in Fig. 6d–f: $T_{WZ} = \text{offer}$ and $\Delta T_{WZ} = 0/s$.

Figure 7g–i: The initial parameter set as described above, with the modification that we forced $W(t) \leq W(t-1)$. If the random draw pulled $\Delta W > 0$, $W(t)$ was set to be $W(t-1)$.

Figure 7j–l: Combining both modifications used in 7d, g.

Figure 8b, c: 20 simulations of 1 M offers each were made for each combination of ($\sigma_W \in \{0, 0.25, 0.5, 1, 3, 5, 8, 10, 20\}$ x $\sigma_N \in \{0, 2, 3, 5\}$). The total sunk cost bubble and slope of the line at 0 s was measured as above. From each simulated session, we calculated the tangent of the probit fit at threshold and plotted the σ_W that generated it. The function was fit with MATLAB 2021a’s fit function (“a + b/x”), adjusted-R2 as reported from the goodness-of-fit from that function. A hyperbolic fit was found (adjusted-R2 = 0.9997).

Figure 9: Simulations used the initial parameter set as described above, but agents accepted every offer and were prevented from leaving the wait zone for DZ seconds, after which $W(t)$ was set to W_0 and drifted normally with standard deviation σ_N . In the first model (9a), T_{WZ} was set to **offer** at the conclusion of DZ seconds and decreased at a rate of 1/s. In the second model (9e), T_{WZ} was set to **offer** at entry into the wait zone and decreased at a rate of 1/s as per the original model. Linear and quadratic fits were made with Matlab 2021a’s fit function, adjusted-R2 as reported from the goodness-of-fit from that function.

Figure 10: For rat and mouse data, latency to accrual was identified as the delay in seconds to when the sunk cost bubble began to grow. Outliers were defined as points with a latency to accrual larger than 15 s. For the rats, this included 97% of the data (all but one sample at 29 s, all other samples were in the range of 4s–8s). For the mice, this included 96% of the data (all but 5/115 samples). Linear and quadratic fits were made with Matlab 2021a’s fit function, adjusted-R2 as reported from the goodness-of-fit from that function.

Subjects. Mice (Sweis et al.¹⁹) 32 M C57B6J mice, age 13 weeks

Mice (Durand-de Cuttoli et al.²⁰) 32 M C57BL6J mice, age 10 weeks

Rats (Sweis et al.¹⁹) 4M 6F FBNF-1 rats.

Humans (Sweis et al.¹⁹) 65 undergraduate students were recruited from the University of Minnesota (24 M, 41 F, mean age = 20 yrs). Ethnicity as reported by the students: 73% white, 16.5% asian, 4.5% Black/African American, 2.5% Hispanic/Latinx, 3.5% other or declined to answer.

Humans (Huynh et al.²³) 178 humans, including 31 undergraduates recruited from Wabash College, and the rest recruited online. Online recruited from Amazon mTurk, restricted to workers from the US. 97 M 80 F 1 non-binary as identified by the subject. Mean age = 39 yrs. Ethnicity as reported by the subject: 3% Black/African American, 62% white, 8.5% asian, 5.5% Hispanic/Latinx, 21% other or declined to answer.

Humans (Kazinka et al.²⁶) 259 humans in the original version; 280 humans in the distractor version. Online recruited through Amazon mTurk. Inclusion criteria detailed in the original study²⁶. Original version: 47% M, 53% F, <1% non-binary as reported by subject, mean age = 37 yrs; Distractor version: 50% M, 49% F, <1%

non-binary as reported by subject, mean age = 38 yrs. Ethnicity as reported by the subject: Original version, 73% White, 10% Black/African American, 10% mixed/other, 7% Asian, 1% refused. Distractor version: 77% White, 12% Black/African American, 5% Asian, 4% mixed/other, 1% refused.

Humans (NCST dataset) 232 humans. Online recruited from Prolific, restricted to participants from the US. 110 M 117 F 5 non-binary as identified by the subject. Mean age 37. Ethnicity as defined by the subject: 9% Black/African American, 68.5% white, 7% asian, 8% Hispanic/Latinx, 7.5% other or declined to answer.

Rats (PJC/GWD dataset): 7 M 8 F FBNF-1 rats, age 10–17 mos.

Rats (BJC/YAB dataset): 21 M BN rats, age 8–12 mos.

Data reporting and ethics statements. Experimental data analyzed here were drawn from the following experiments:

Figure 3: **a** Mice, Sweis et al. 2018;¹⁹ **b** Mice, Durand-de Cuttoli et al. 2022;²⁰ **c** Rats, Sweis et al. 2018;¹⁹ **d** Humans, (UG in person) Sweis et al. 2018;¹⁹ **e** Humans, (online) Huynh et al. 2021;²³ **f** Humans, (online) [NCST dataset].

Figure 8: **d** Mice, Sweis et al. 2018¹⁹, Durand-de Cuttoli et al. 2022,²⁰ **e** Rats, including Sweis et al. 2018¹⁹ and data from the Redish lab [PJC/GWD dataset].

Figure 10: **a** Rats without an offer zone from Sweis et al.¹⁹ **b, c** Rats, Sweis et al.¹⁹ and data from the Redish lab [BJS/YAB dataset]; **d** mice with limited experience, including mice C57BL/6 from Sweis et al. 2018a,b^{17,19}. **e, f** Each mouse provides two lines, one early and one late.

All experiments were approved by the appropriate boards and comply with the ethical regulations of their respective universities as well as US National Institute of Health guidelines. All human subjects provided informed consent as in the original studies and as follows:

- Approved by the University of Minnesota IACUC: Sweis et al. 2018¹⁹, PJC/GWD dataset, BJS/YAB dataset.
- Approved by the Mount Sinai IACUC: Durand-de Cuttoli et al. 2022²⁰.
- Approved by the University of Minnesota IRB: Sweis et al. 2018¹⁹, Kazinka et al. 2021²⁶.
- Approved by the Wabash College IRB: Huynh et al. 2021;²³ NCST dataset.

All human studies were designed to include as diverse a population as possible. Inclusion criteria were described in the original studies and are provided above.

Reporting summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

Simulation code, data, and figure generation code are all available at <https://github.com/adredish/SunkCostModelsAndData2022>.

Received: 25 July 2022; Accepted: 8 November 2022;

Published online: 07 December 2022

References

1. Staw, B. Knee-deep in the big muddy: a study of escalating commitment to a chosen course of action. *Organ. Behav. Hum. Perform.* **16**, 27–44 (1976).
2. Staw, B. M. & Fox, F. V. Escalation: the determinants of commitment to a chosen course of action. *Hum. Relat.* **30**, 431–450 (1977).
3. Arkes, H. & Blumer, C. The psychology of sunk cost. *Organ. Behav. Hum. Decis. Process.* **35**, 124–140 (1985).
4. Staw, B. M. & Ross, J. Understanding behavior in escalation situations. *Science* **246**, 216–220 (1989).
5. McAfee, R. P., Mialon, H. M. & Mialon, S. H. Do sunk costs matter? *Econ. Inq.* **48**, 323–336 (2010).
6. Thaler, R. *Misbehaving*. (Editora Intrinseca, 2019).
7. Pompilio, L., Kacelnik, A. & Behmer, S. T. State-dependent learned valuation drives choice in an invertebrate. *Science* **311**, 1613–1615 (2006).
8. Magalhães, P., White, K. G., Stewart, T., Beeby, E. & van der Vliet, W. Suboptimal choice in nonhuman animals: rats commit the sunk cost error. *Learn. Behav.* **40**, 195–206 (2012).
9. Kacelnik, A. & Marsh, B. Cost can increase preference in starlings. *Anim. Behav.* **63**, 245–250 (2002).
10. Navarro, A. D. & Fantino, E. The sunk cost effect in pigeons and humans. *J. Exp. Anal. Behav.* **83**, 1–13 (2005).
11. Pattison, K. F., Zentall, T. R. & Watanabe, S. Sunk cost: pigeons (*Columba livia*), too, show bias to complete a task rather than shift to another. *J. Comp. Psychol.* **126**, 1–9 (2012).
12. Magalhães, P. & White, K. G. The sunk cost effect across species: a review of persistence in a course of action due to prior investment. *J. Exp. Anal. Behav.* **105**, 339–361 (2016).

13. Magalhães, P. & White, K. G. Sunk cost and work ethic effects reflect suboptimal choice between different work requirements. *Behav. Processes* **94**, 55–59 (2013).
14. Redish, A. D. et al. Computational Validity: Using computation to translate behaviors across species. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **377**, 20200525 (2022).
15. Steiner, A. P. & Redish, A. D. Behavioral and neurophysiological correlates of regret in rat decision-making on a neuroeconomic task. *Nat. Neurosci.* **17**, 995–1002 (2014).
16. Sweis, B. M., Thomas, M. J. & Redish, A. D. Mice learn to avoid regret. *PLoS Biol.* **16**, e2005853 (2018).
17. Sweis, B. M., Larson, E. B., Redish, A. D. & Thomas, M. J. Altering gain of the infralimbic-to-accumbens shell circuit alters economically dissociable decision-making algorithms. *Proc. Natl. Acad. Sci. USA* **115**, E6347–E6355 (2018).
18. Sweis, B. M., Redish, A. D. & Thomas, M. J. Prolonged abstinence from cocaine or morphine disrupts separable valuations during decision conflict. *Nat. Commun.* **9**, 2521 (2018).
19. Sweis, B. M. et al. Sensitivity to ‘sunk costs’ in mice, rats, and humans. *Science* **361**, 178–181 (2018).
20. Durand-de Cuttoli, R. et al. Distinct forms of regret linked to resilience versus susceptibility to stress are regulated by region-specific CREB function in mice. *Sci. Adv.* **8**, eadd5579 (2022).
21. Schmidt, B., Duin, A. A. & Redish, A. D. Disrupting the medial prefrontal cortex alters hippocampal sequences during deliberative decision making. *J. Neurophysiol.* **121**, 1981–2000 (2019).
22. Schmidt, B. & Redish, A. D. Disrupting the medial prefrontal cortex with designer receptors exclusively activated by designer drug alters hippocampal sharp-wave ripples and their associated cognitive processes. *Hippocampus* **31**, 1051–1067 (2021).
23. Huynh, T., Alstatt, K., Abram, S. V. & Schmitzer-Torbert, N. Vicarious trial-and-error is enhanced during deliberation in human virtual navigation in a translational foraging task. *Front. Behav. Neurosci.* **15**, 586159 (2021).
24. Abram, S. V., Breton, Y.-A., Schmidt, B., Redish, A. D. & MacDonald, A. W. 3rd The Web-Surf Task: a translational model of human decision-making. *Cogn. Affect. Behav. Neurosci.* **16**, 37–50 (2016).
25. Abram, S. V., Hanke, M., Redish, A. D. & MacDonald, A. W. 3rd Neural signatures underlying deliberation in human foraging decisions. *Cogn. Affect. Behav. Neurosci.* **19**, 1492–1508 (2019).
26. Kazinka, R., MacDonald, A. W. 3rd & Redish, A. D. Sensitivity to sunk costs depends on attention to the delay. *Front. Psychol.* **12**, 604843 (2021).
27. Sweis, B. M. et al. Human cocaine users, like cocaine-treated mice, display unique disruptions on a neuroeconomic task. *Society for Neuroscience* 335.05 (2019).
28. Haynos, A. F. et al. Measuring and modeling excess cognitive control: an example in anorexia nervosa. *Psychiatric Research Society* (2021).
29. Haynos, A. F., Abram, S. V., MacDonald, A. W., Redish, A. D. & Crow, S. J. Parsing decision-making systems in anorexia nervosa using a computationally-informed translational neuroeconomic paradigm. *Annual meeting for the Association for Behavioral and Cognitive Therapies* (2020).
30. Sweis, B. M. Beyond simple tests of value: a neuroeconomic, translational, disease-relevant, and circuit-based approach to resolve the computational complexity of decision-making. (PhD Thesis, University of Minnesota, 2018).
31. Duin, A. A., Aman, L., Schmidt, B. & Redish, A. D. Certainty and uncertainty of the future changes planning and sunk costs. *Behav. Neurosci.* **135**, 469–486 (2021).
32. Resulaj, A., Kiani, R., Wolpert, D. M. & Shadlen, M. N. Changes of mind in decision-making. *Nature* **461**, 263–266 (2009).
33. Stone, C., Mattingley, J. B. & Rangelov, D. On second thoughts: changes of mind in decision-making. *Trends Cogn. Sci.* **26**, 419–431 (2022).
34. Mullainathan, S. & Shafir, E. *Scarcity: Why Having Too Little Means So Much*. (Macmillan, 2013).
35. Ott, T., Masset, P., Gouvêa, T. S. & Kepecs, A. Apparent sunk cost effect in rational agents. *Sci. Adv.* **8**, eabi7004 (2022).
36. Wikenheiser, A., Stephens, D. W. & Redish, A. D. Subjective costs drive overly-patient foraging strategies in rats on an intertemporal foraging task. *Proc. Natl. Acad. Sci. USA* **110**, 8308–8313 (2013).
37. Redish, A. D. *The Mind within the Brain: How we make decisions and how those decisions go wrong*. (Oxford, 2013).
38. Redish, A. D. Vicarious trial and error. *Nat. Rev. Neurosci.* **17**, 147–159 (2016).
39. Gollwitzer, P. M., Heckhausen, H. & Steller, B. Deliberative and implemental mind-sets: cognitive tuning toward congruous thoughts and information. *J. Pers. Soc. Psychol.* **59**, 1119–1127 (1990).
40. Achtziger, A. & Gollwitzer, P. M. Rubicon model of action phases. *Encyclopedia of Social Psychology* **2**, 769–770 (2007).
41. Aw, J. M., Vasconcelos, M. & Kacelnik, A. How costs affect preferences: experiments on state dependence, hedonic state and within-trial contrast in starlings. *Anim. Behav.* **81**, 1117–1128 (2011).
42. Singer, R. A. & Zentall, T. R. Preference for the outcome that follows a relatively aversive event: contrast or delay reduction? *Learn. Motiv.* **42**, 255–271 (2011).
43. Abram, S. V., Redish, A. D. & MacDonald, A. W. 3rd Learning from loss after risk: dissociating reward pursuit and reward valuation in a naturalistic foraging task. *Front. Psychiatry* **10**, 359 (2019).
44. Robinson, T. E. & Berridge, K. C. The neural basis of drug craving: an incentive-sensitization theory of addiction. *Brain Res. Rev.* **18**, 247–336 (1993).
45. Berridge, K. Food reward: Brain substrates of wanting and liking. *Neurosci. Biobehav. Rev.* **20**, 1–25 (1996).
46. Clark, J. J., Hollon, N. G. & Phillips, P. E. M. Pavlovian valuation systems in learning and decision-making. *Curr. Opin. Neurobiol.* **22**, 1–8 (2012).
47. Fanselow, M. S. & Wassum, K. M. The origins and organization of vertebrate pavlovian conditioning. *Cold Spring. Harb. Perspect. Biol.* **8**, a021717 (2015).
48. Kahneman, D., Knetsch, J. L. & Thaler, R. H. Experimental tests of the endowment effect and the Coase theorem. *J. Polit. Econ.* **98**, 1325–1348 (1990).
49. Kahneman, D., Knetsch, J. L. & Thaler, R. H. The endowment effect, loss aversion, and status quo bias. *J. Econ. Perspect.* **5**, 193–206 (1991).
50. Redish, A. D., Schultheiss, N. W. & Carter, E. C. The computational complexity of valuation and motivational forces in decision-making processes. *Curr. Top. Behav. Neurosci.* **27**, 313–333 (2016).
51. Inzlicht, M., Shenhav, A. & Olivola, C. The effort paradox: effort is both costly and valued. *Trends Cogn. Sci.* **22**, 337–349 (2017).
52. Kanodia, C., Bushman, R. & Dickhaut, J. Escalation errors and the sunk cost effect: an explanation based on reputation and information asymmetries. *J. Account. Res.* **27**, 59 (1989).

Acknowledgements

We thank Geoff Diehl for the use of the additional data he collected. We thank Torben Ott and Adam Kepecs for continued discussions on these issues. We thank Chris Heinz for discussion related to these issues. Funding: R01 MH080318, R01 MH12688 (ADR), University of Minnesota Johnston Chair in Neuroscience (ADR), Department of Veterans Affairs CX002355 (SVA), F32 DA038392, Diversity supplement NIDA DA030672S1, funding from the Society for Neuroscience Scholars Program (BJS), T32 DA007234 (PJC), F30 MH124404 (AK), L40MH127601, R01MH051399-31S1 (BMS).

Author contributions

Conceptualization: A.D.R., B.M.S. Code and simulations: A.D.R. Data collection: S.V.A., P.J.C., A.A.D., R.D.C., R.K., A.K., B.S., B.M.S. Data analysis: A.D.R., S.V.A., P.J.C., B.M.S. Data discussion and review: A.D.R., S.V.A., P.J.C., A.A.D., R.D.C., R.K., A.K., A.W.M., B.S., N.S.T., M.J.T., B.M.S. Writing—original draft: A.D.R. Writing—review & editing: A.D.R., S.V.A., P.J.C., A.A.D., R.D.C., R.K., A.K., A.W.M., B.S., N.S.T., M.J.T., B.M.S.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42003-022-04235-6>.

Correspondence and requests for materials should be addressed to A. David Redish.

Peer review information *Communications Biology* thanks the anonymous reviewers for their contribution to the peer review of this work. Primary Handling Editor: Karli Montague-Cardoso.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons

Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022