



Supplementary Materials for

Sensitivity to “Sunk Costs” in Mice, Rats, and Humans

Brian M. Sweis, Samantha V. Abram, Brandy J. Schmidt, Kelsey D. Seeland, Angus W. MacDonald III, Mark J. Thomas, A. David Redish

correspondence to: redish@umn.edu

This PDF file includes:

Materials and Methods
Supplementary Text S1 to S7
Figs. S1 to S14
Table S1
References (32-75)

Materials and Methods

Mice

32 C57BL/J6 male mice, 13 weeks of age, were trained in Restaurant Row. Mice were single-housed in a temperature- and humidity-controlled environment with a 12-hr-light/12-hr-dark cycle with water ad libitum. Mice were food restricted to a maximum of 85% free feeding body weight and trained to earn their entire day's food ration during their 1-hr Restaurant Row testing session. A replication cohort of an additional 32 mice were run (Fig.S12). These mice were intentionally food restricted to a lesser extreme (a maximum of 90% free feeding body weight). All experiments were approved by the University of Minnesota Institutional Animal Care and Use Committee.

Rodent flavored-pellet training (mice).

Mice underwent 1 week of pellet training prior to the start of being introduced to the Restaurant Row maze. During this period, mice were taken off of regular rodent chow and introduced to a single daily serving of BioServ full nutrition 20mg dustless precision pellets in excess (5g). This serving consisted of a mixture of chocolate-, banana-, grape-, and plain-flavored pellets. Next, mice (hungry, before being fed their daily ration) were introduced to the Restaurant Row maze 1 day prior to the start of training and were allowed to roam freely for 15min to explore, get comfortable with the maze, and familiarize themselves with the feeding sites. Restaurants were marked with unique spatial cues. Feeding bowls in each restaurant were filled with excess food on this introduction day.

Restaurant Row procedure (mice).

Task training was broken into 4 stages. Each daily session lasted for 1hr. At test start, one restaurant was randomly selected to be the starting restaurant where an offer was made if mice entered that restaurant's T-shaped offer zone from the appropriate direction in a counter-clockwise manner. During the first stage (day 1-7), mice were trained for 1 week being given only 1s offers. Brief low pitch tones (4000Hz, 500ms) sounded upon entry into the offer zone and repeated every second until mice skipped or until mice entered the wait zone after which a pellet was dispensed. To discourage mice from leaving earned pellets uneaten, motorized feeding bowls cleared any uneaten pellets upon restaurant exit. Left over pellets were counted after each session and mice quickly learned to not leave the reward site without consuming earned pellets. The next restaurant in the counter-clockwise sequence was always and only the next available restaurant where an offer could be made such that mice learned to run laps encountering offers across all four restaurants in a fixed order serially in a single lap. During the second stage (day 8-12), mice were given offers that ranged from 1s to 5s (4000Hz to 5548Hz, in 387Hz steps) for 5 days. Offers were pseudo-randomly selected such that all 5 offer lengths were encountered in 5 consecutive trials before being re-shuffled, selected independently between restaurants. Again, offer tones repeated every second in the offer zone indefinitely until either a skip or enter decision was made. In this stage and subsequent stages, in the wait zone, 500ms tones descended in pitch every second by 387Hz steps counting down to pellet delivery. If the wait zone was exited at any point during the countdown, the tone ceased and the trial ended, forcing mice to proceed to the next restaurant. Stage 3 (day 13-17) consisted of offers from 1s to 15s (4000Hz to

9418Hz) for another 5 days. Stage 3 is the timepoint used in Fig.S10. Stage 4 (day 18-70) offers ranged from 1s to 30s (4000Hz to 15223Hz) and lasted until mice showed stable economic behaviors. Our early 1-30s and well-trained 1-30s timepoints used in Fig.S9 include the first 5 and last 5 days of stage 4. We used 4 Audiotek tweeters positioned next to each restaurant powered by Lepy amplifiers to play local tones at 70dB in each restaurant. We recorded speaker quality to verify frequency playback fidelity. We used Med Associates 20mg feeder pellet dispensers and 3D-printed feeding bowl receptacles fashioned with mini-servos to control automated clearance of uneaten pellets. Animal tracking, task programming, and maze operation was powered by AnyMaze (Stoelting). Mice were tested at the same time every day in a dim-lit room, were weighed before and after every testing session, and were fed a small post-session ration in a separate waiting chamber on rare occasions as needed to prevent extremely low weights according to IACUC standards (not <85% free-feeding weights).

Rats.

10 Fisher Brown-Norway rats (4 male, 6 female), aged between 8-12 months, were trained to run the Restaurant Row task variant with an offer zone. 22 Brown-Norway rats (male), aged between 8-12 months, were trained to run the Restaurant Row task variant without an offer zone. Rats were single-housed in a temperature and humidity-controlled environment and kept on a 12hr light/dark cycle with water ad libitum. Rats were food restricted to a maximum of 80% free feeding body weight and earned their food each day during their 1-hr Restaurant Row session. All experiments were approved by the University of Minnesota Institutional Animal Care and Use Committee.

Rodent flavored-pellet training (rats).

All rats were given 8 days of handling and pellet training prior to being introduced to the Restaurant Row task. Handling consisted of roaming freely on the experimenter's lap for 15-20 minutes daily. For pellet training, rats were taken off ad libitum access to Teklad rodent chow and given 1hr of access to 15g of TestDiet full nutrition 45mg purified rodent tablets in four unique flavors (chocolate, banana, cherry, plain). Rats were allowed to eat freely, either while being handled, in a bedding-free cage, or a combination of both.

Restaurant Row procedure (rats).

Training in the task variant with an offer zone consisted of four training phases. Each session lasted 60min. Rats ran each phase for five days. Phase one consisted of 1s delays at each restaurant. Phase two consisted of randomly selected delays from 1-5s. Phase three included 1-15s offers, and phase four had the final delay range of 1-30s offers. Training in the task variant without an offer zone in a separate cohort of rats was slightly different. These rats were initially trained twice a day in 30min sessions. Training began with 5 days of 1s offers at all feeder sites. Then, the randomized list of delays presented to animals was expanded to 1-2s, 1-3s, 1-4s, and 1-5s delays over 4 consecutive days. Rats then received 10 days of 1-30s delays. Next, rats switched to once a day 60min testing sessions using 1-30s delays. All delays were randomly selected and varied between day and restaurant. In the task variant with an offer zone, maze contingences were similar to mice described above. In the task variant without an offer

zone, delay countdowns began immediately upon entry into the restaurant. Rats ran in a counter clockwise direction, where offers were only triggered if the rats passed through each restaurant in serial order, such that trials would not be triggered when running backwards. After triggering a trial, the next available restaurant where an offer could be made was always the restaurant immediately after the last restaurant triggered, regardless of if an offer was accepted or declined. Rats were run at the same time each day in a very dimly lit room. At the start of the task, rats were always placed on the maze in the same place. Rats were weighed before running the task. If a rat was at or near their 80% weight and did not receive enough food on the track during their running session, they were fed no sooner than a half an hour after completing their running session. Maze operation was done by Matlab. We used 45mg Med Associated feeders to deliver the pellets into in-house 3D printed feeding bowls.

Humans.

65 undergraduate students from the University of Minnesota completed the Web-Surf Task (24 male, 41 female, mean age = 20.23 years), and an additional 17 completed the task variant without an offer zone (4 male, 13 female, mean age = 19.63 years); of note, 14 of these 17 subjects represent a subset of the data presented in Abram et al. 2016. 24 Participants received compensation in the form of extra credit towards psychology courses. Ethnicity of subjects included 73% White, 16.5% Asian, 4.5% Black/African American, 2.5% Hispanic/Latino, 0.5 American Indian/Alaska Native, 0.5% Native Hawaiian/Pacific Islander, and 2.5% other. The University of Minnesota Institutional Review Board approved the human study procedures, and all undergraduate subjects provided written informed consent.

Web-Surf procedure.

In the task variant with an offer phase, subjects had 30 minutes to travel between four galleries that included video rewards; we used the same categories described in Abram et al. 2016 (15): kittens, dance landscapes, and bike accidents. Offers were presented in text and with a webpage-like progress bar. When subjects arrived at a gallery, they first had the option to stay or skip. If they chose to skip, they traveled on to the next gallery and encountered a new offer. If they stayed, they entered the wait phase, and the progress bar begins to count down. At any point before the delay finished, the subject could elect to quit, which again led the process of traveling to the next gallery. If the subject stayed through the entire delay, they were shown a video reward for 4 seconds, after which they rated the video from 1-4 (4 = most enjoyed) according to how much they liked that video. When traveling between galleries, subjects pressed a series “next” buttons as they randomly appeared around a gray screen; numbers were presented in a slightly darker shade of gray to increase task difficulty. For training, subjects completed 3 forced-choice trials to illustrate what happens in the enter, skip, and quit conditions. The forced-choice trials are followed by 8 additional practice trials where subjects make decisions.

In the task variant without an offer phase, subjects similarly had 30 minutes to travel between the same four categories, with offers presented in the same manner. In this variant, the delay began to countdown immediately upon gallery arrival. Subjects then had the option to quit and move on to the next gallery or continue to wait for the delay to

finish. Videos were again shown for 4 seconds, and subjects rated each viewed video on a 1-4 scale (4 = most enjoyed). When passing between galleries, subjects clicked a series of “next” buttons that randomly appeared around the gray screen; the buttons were shown in dark gray to blend into the background to increase task difficulty. For training, subjects completed 2 forced-choice trials to illustrate quit decisions, followed again by 8 practice trials where they could decide whether to quit or earn. Regardless of the task variant, subjects ranked the categories by preference (again from 1-4, 4 = most preferred) after the testing session in a post-testing debriefing survey.

Sunk Cost Data Analysis.

All data were processed in Matlab and statistical analyses were carried out using JMP Pro 13 Statistical Discovery software package from SAS. The main analysis carried out in the manuscript critical to our findings included computing linear regression models of earning probability in the wait zone as a function of either time remaining in the wait zone (Fig.2,S8,S10,S12,S13,S14) or time spent in the offer zone (Fig.3,S12,S13). Fig.2,S8,S10,S12,S13,S14 compare the slope coefficients of these regressions interacting across various sunk cost conditions using an ANOVA with p(earn) as the dependent variable and time-remaining x sunk cost condition as factors. Importantly, control datasets were re-calculated to control for subtle skews in different dataset availability distributions between the various sunk cost scenarios (described in Fig.S5). Comparisons of interest included testing regression coefficients of the zero sunk cost condition against each sunk cost condition (Fig.2,S8,S10,S12,S13,S14 – data originating from black dataset against color datasets after taking into account proper adjusted control datasets illustrated in Fig.S5A-F) as well as testing regression coefficients of each sunk cost condition against other sunk cost conditions (Fig.2,S8,S10,S12,S13,S14 – color data against color data, again, after taking into account proper adjusted control datasets, illustrated in Fig.S5G).

Modeling Sub-Optimality

In the Restaurant Row and Web-Surf Tasks, decisions to abandon an on-going investment appear highly sub-optimal at face value, particularly if subjects were cued of the offer cost at trial onset before accepting. A potential optimal strategy could include making smart offer zone decisions to skip vs. enter informed by cued offer cost and never quitting once in the wait zone. Taking advantage of the economic nature of these tasks, we characterized the efficiency of observed behaviors by generating a computer model that predicted optimal number of rewards that could be earned if subjects were behaving as efficiently as possible. Optimal behavior was based on each individual’s behavioral performance (e.g., reaction time and subjective threshold of willingness to wait in each restaurant) if they were behaving as efficiently as possible (using the fastest quartile of reaction and consumption times, by following their revealed preferences strictly, by never quitting). Proportion of actual observed earnings relative to model-predicted maximally optimal earnings was calculated.

Supplementary Text S1: Cross-validating economic decisions and revealed preferences across species.

Mice, rats, and humans learned to forage economically on these tasks (Fig.S1). We characterized multiple valuation metrics that revealed subjective flavor preferences similarly across dimensions in all species (Fig.S2). We ranked flavors from least- to most-preferred by summing the number of rewards earned in each flavor at the end of every session (Fig.S2A-C). We characterized how subjects budgeted their time within the limited session by calculating economic “thresholds” of willingness to earn rewards as a function of offer length in seconds (Fig.S2D-F). Thresholds for different flavors varied across individuals in all three species (Fig.S1D-F). Thresholds in rodents remained stable across many days for a given flavor for an individual (Fig.S1D-E). Additionally, subjects rarely violated these thresholds (Fig.S1G-I). We also characterized how subjects evaluated rewards post-consumption (Fig.S2G-I, see additional supplementary text below). Lastly, in humans, we obtained stated preference rankings in a post-testing debrief survey (Fig.S2J). All valuation parameters significantly correlated with each other in each species (Pearson and Spearman correlation coefficients all $P < 0.001$, controlling for multiple comparisons).

Supplementary Text S2: Controlling for time differences in the offer zone vs. wait zone across task variants.

In the offer zone, the reward delay does not start counting down toward reward delivery. Therefore, two critical differences from the wait zone may explain why sunk costs do not accrue in the offer zone: (1) time spent in the wait zone is valued differently because it is actually counted toward reward-earning progress; (2) countdown tones that descend in the wait zone (i.e., melodic contours, 32) or a diminishing video download bar in the wait phase may carry added value learned through incentive salience (33). Both factors could weigh continued commitment in the wait zone disproportionately higher than quitting (i.e., enhance a loss-aversion bias) and could explain why sunk costs do not accrue in the offer zone.

To control for these two factors, we tested rodents and humans in variants of the Restaurant Row and Web-Surf Tasks in which the offer zone was not present (Fig.S7). In these alternative task variants, time started counting down as soon as the subject entered a restaurant or video gallery. In rats, we found that the sunk cost effect (ANOVA collapsing across all sunk cost conditions, $F=33.93$, $P < 0.0001$) did not begin to accrue until after an initial window of time had elapsed (Fig.S7B, post-hoc tests: +1s: $F=0.08$, $P=0.78$; +2s: $F=0.22$, $P=0.64$; +3s: $F=0.01$, $P=0.93$; +4s: $F=2.15$, $P=0.14$; +5s: $F=12.69$, $P < 0.001$). Despite detracting from the session’s total time budget and despite counting as “down payments” progressing toward reward delivery, an initial portion of time spent was not counted toward sunk costs, mimicking offer zone findings. Thus, in a task variant without an offer zone, the sunk-cost effect in the wait zone had a delayed onset. This suggests a serial process of (1) choosing-between followed by (2) opting-out was still engaged, despite the lack of an explicit offer zone. In support of this, the window when sunk costs did not accrue matched the average reaction time it took rats to turn down offers (Fig. S7E, ~5s). Humans on this variant of the task also displayed the sunk-cost effect (ANOVA collapsing across all sunk-cost conditions, $F=61.42$,

$P < 0.0001$), although a delayed onset was not detectable in these data (Fig.S7C, post-hoc test: +1s: $F = 12.29$, $P < 0.001$, unlike rats on this variant). However, humans were significantly faster than rodents at turning down offers, and, consistent with rodents, the onset of sunk-cost effects matched this average reaction time (Fig.S7F, ~1s). It is possible that humans similarly progressed through a serial process of choosing-between followed by opting-out but at a faster pace than rodents.

Supplementary Text S3: Assessing vicarious trial and error behavior in rodents (and across learning in mice).

In the offer zone, rodents displayed “pause and look” behaviors, known as “vicarious trial and error” (Fig.S8, 34-36). Vicarious trial and error behaviors reveal on-going deliberation and planning during moments of indecision (34-36). Numerous in vivo electrophysiology recording studies have demonstrated that during these behaviors, hippocampal representations sweep forward along the path of the animal, alternating between potential goals until the animal knows where to go (37). Such goal representations are synchronized to reward-value representations in the ventral striatum, suggesting outcome predictions are being evaluated serially (38-39). This work on deliberation and serial planning has been replicated in rats in other labs (70). Other labs have also demonstrated these neurophysiological representations of future outcome expectations in certain deliberative situations in mice (71) and in humans (72-74). Compare the computational distinctions made in Ref. 75. A recent review on the neurophysiological processes underlying deliberation can be found in (Ref. 35).

Restaurant Row training in mice (offer zone present) provided further insight into the development of a serial process decision stream: initially deliberating followed by re-evaluations to opt-out. Mice were tested daily up to 115 days, far longer than rats (typically run for 40 days) and humans (tested in a single session). We observed pronounced changes in decision strategies between early and late training as mice were still learning the structure of the task. Early in training, mice rapidly accepted all offers indiscriminately without factoring in the randomly selected 1-30s offer cost, and thus relied on quits to turn down expensive offers (Fig.S9). In other words, offer zone enter decisions early in training were automated. This was likely a learned behavior in even earlier stages of training when offers were relatively inexpensive (see Methods, progressive stages: 1s offers only, offers ranging between 1-5s, and then between 1-15s), when cost information could be discarded, and when mice could afford to accept all offers. Interestingly, on these early 1-30s training sessions, we found that the sunk cost effect (Fig.S7A, ANOVA collapsing across all sunk cost conditions, $F = 4.46$, $P < 0.0001$) did not begin to accrue in the wait zone until an initial window had elapsed (post-hoc tests: +1s: $F = 0.14$, $P = 0.71$; +2s: $F = 0.28$, $P = 0.60$; +3s: $F = 0.65$, $P = 0.43$; +4s: $F = 1.67$, $P = 0.20$; +5s: $F = 4.10$, $P < 0.05$) that matched the average reaction time it took mice to quit offers (Fig.S7D, ~5s). This is reminiscent of the delayed onset of sunk costs observed in rats on the task without an offer zone when vicarious trial and error behaviors typically occurred (compare to Fig.S7B,E). That is, despite separating each restaurant into an offer zone and wait zone, mice early in 1-30s training essentially ignored the offer zone. Importantly, mice did not display vicarious trial and error behaviors in the offer zone early in 1-30s training (Fig.S9E-F). This suggests that an initial deliberation

decision was not actually made in the offer zone and that a true initial choose-between decision was not made until after mice entered the wait zone. Therefore, this would delay the onset of a secondary opt-out process in the wait zone. Taken together inexperienced mice transitioned between deliberation and foraging decision modalities once in the wait zone while experienced mice separated this transition between the offer- and wait zones (35). Importantly, the sunk cost effect tracked this transition across training.

Well-trained mice were fully capable of discriminating randomly presented cued offer costs in the offer zone differently in the different restaurants, ensuring information uncertainty was not a confounding factor in rodents (Fig.S10). Furthermore, sunk cost effects occurred in all restaurants regardless of subjective flavor preference ranking, again only in the wait zone and not offer zone (Fig.S10). Environmental factors including different experimenters or severity of food restriction did not impact main sunk cost findings in a replication cohort of mice, despite influencing overall devalued economic behaviors likely due to satiety effects (Fig.S11).

Supplementary Text S4: Assessing optimality.

To determine whether behavior on these tasks was sub-optimal, we used a computer-simulation model to predict the maximal rewards subjects could earn if they behaved optimally by following their revealed preferences strictly by not abandoning accepted offers in the wait zone (see Methods). We found mice, rats, and humans indeed performed sub-optimally (Fig.S12A-C). This sub optimality arose from disadvantageous wait-zone decisions (Fig. S12D-F). We can use an individual's restaurant's threshold to calculate the relative value of each offer encountered. The value of the offer ($VO = \text{Threshold} - \text{Offer}$) plotted against the value of the time left in the countdown for a given trial at the moment of quitting ($VL = \text{Threshold} - \text{Time Left To Go}$) reveals the economic efficiency of wait zone decisions. Economically disadvantageous wait-zone quit decisions drive sub optimality through a susceptibility to sunk costs.

Supplementary Text S5: Assessing post-consumption hedonic valuations.

The value of a reward can be assessed in multiple ways. Often, this is measured in an instrumental manner, where how willing a subject is to take a reward (e.g., measured in amount of resources spent, effort expended, or behavioral invigoration) can reflect reward valuation. This is sometimes referred to as reward-seeking, reward-taking, or "wanting" valuations (33,40). In the Restaurant Row and Web-Surf Tasks, these are measurable in offer- and wait zone choice behaviors. These are distinct from post-consummatory valuations after an individual has earned a reward. Such post-consumption valuations are sometimes referred to as hedonic or "liking" valuations (40). In the Restaurant Row and Web-Surf Tasks, these can be measured after subjects earned rewards. In humans, this is more overtly assessed, since subjects were asked to rate each video on a scale from 1-4 (4=most enjoyed) immediately after viewing (or "consuming") the short (4s) video reward. These ratings matched other metrics of subjective valuation on this task (Fig.S2C-J). Importantly, naturally enjoyable videos were used to mimic the same immediate consummatory nature of inter-trial pellet eating used in many appetitive-driven rodent laboratory studies (which is unlike many human studies that use

hypothetical ratings, money, or other token-systems redeemed at the end of testing). Interestingly, we found that rodents, after earning and consuming food pellets on the Restaurant Row task, often lingered at the reward site before advancing to the next trial at the next restaurant (Fig.S2G,H). Surprisingly, rodents typically spent ~50% of the entire 1hr testing session engaging in this lingering behavior. This decision to linger rather than leave, where no overt reward is being sought out, may represent a conditioned-place-preference-like effect associated with each restaurant's unique context (33). Interestingly, rodents lingered longer in more-preferred restaurants (Fig.S2G,H). Therefore, we can take this behavioral metric as one that is (a) distinct from an instrumental reward-taking or "wanting" valuation, (b) appears to mimic post-consumption hedonic valuations measured in humans on the parallel Web-Surf Task, and (c) captures subjective flavor preferences that match other valuation metrics for unique flavors on this task, similar to humans (Fig.S2).

Interestingly, we found that the degree of post-consumption hedonic valuations correlated with offer cost in mice, rats, and humans (Fig.S13). That is, after subjects earned rewards that required more expensive investments, they hedonically valued those consumed rewards higher compared to earned rewards that were inexpensive. This is related to but distinct from the sunk-cost phenomenon, which is rooted in enhanced reward valuations. The sunk-cost phenomenon is classically described as an escalation of continued reward investment or commitment as a function of irrecoverable past investments made toward reward receipt. Thus, the sunk-cost effect is linked to enhancements in instrumental reward-taking valuations. The sunk-cost phenomenon, classically, makes no prediction with regard to the hedonic value of already owned or consumed rewards. While the two types of valuations are certainly related, they are separable on the Restaurant Row and Web-Surf Tasks. At the surface, because findings in Fig.S13 are related to both post-earn hedonic valuations as well as instrumental reward-taking costs required to earn the reward itself, these data appear related to two other cognitive heuristics often compared to the sunk-cost phenomenon: the endowment effect and the post-purchase rationalization effect. The endowment effect posits that already owned objects are more highly valued than those not possessed (41). The post-purchase rationalization effect (sometimes described as "Buyer's Stockholm Syndrome" or "Marketing Placebo Effect") is a choice-supportive bias that posits more expensive rewards are more highly valued than less expensive rewards after purchasing them (42-44). Data in Fig.S13 is consistent with the post-purchase rationalization effect shared across species. While there have been reports of the endowment effect in non-human animals, to our knowledge this is the first report of evidence for the post-purchase rationalization effect in non-human animals (45-46).

Supplementary Text S6: Assessing the role of environmental scarcity.

The sunk-cost effect may also relate to the amount of reward available in the environment. In a reward-scarce environment, the importance of decision optimality is higher than in a reward-rich environment. In other foraging tasks, rodents have been shown to adopt sub-optimal foraging strategies exacerbated in reward-scarce environments (16). These reports hint at sunk-cost-like effects that depend on the subjective over-valuation of time invested as a function of the state of the environment.

Interestingly, mice training in relatively reward-rich environments (where offers ranged from 1-15s) did not demonstrate sunk-cost effects (Fig.S14, ANOVA collapsing across all sunk-cost conditions, $F=1.13$, $P=0.35$), consistent with these reports. This indicates another important determinant of the sunk-cost phenomenon – environmental reward scarcity – that may have been overlooked in past studies presenting conflicting findings across species. This data is consistent with notions in other work that demonstrate State-Dependent Valuation Learning (SDVL) depends on the subject being in a leaner state and not necessarily on lack of training (19-20). This helps explain why studies of Within-Trial Contrast (WTC) may have had unreliable outcomes.

If sunk-cost effects existed all the time in a leaner environment in our tasks (in the offer zone and wait zone, or during all moments in the wait zone in the task variants without an offer zone), sunk-cost effects on our task could be fully explainable by SDVL. The fact that decision processes seem to be interrupted or segregated into distinct stages within the same trial only in leaner environments and that these separate stages are differentially susceptible to sunk costs reveals dissociable valuation algorithms are at play within the same trial. Furthermore, aligning with the rationale posited in other work in response to reasoning for lack of WTC effects due to lack of sufficient training in other studies, effects reported in Fig.S14 (lack of sunk cost effects even in the wait zone after any duration of waiting) occurred after 17 days of consecutive training (20). This included exposure to roughly 4,500 trials overall on average per mouse by day 17, or 1,500 trials of 1-15s offers in that training block (days 13-17), with 100 trials per newly experienced offer stimulus in that block (tones representing 6-15s offers). The previous block of training (days 8-12) included 1-5s offers and thus added additional exposure (roughly 400 trials) to 5s tones (capable of eliciting sunk cost effects in leaner 1-30s environment), by the 1-15s offer block. Despite any of this, tones only provide extra information that SDVL and WTC, in theory, could function without in order to drive sunk cost effects in this task design.

Supplementary Text S7: Helping resolve discrepancies in previous work.

There is a large body of literature sparked by an early publication from Arkes & Ayton that not only suggested non-human species are incapable of displaying the sunk-cost fallacy, but that human children are also incapable of displaying the sunk cost fallacy (2). That is, children “outperform” adults by not factoring in added value due to sunk costs. This study postulated that children have not yet learned a socionormative “do not waste” rule in these early stages of development. The majority of studies building off of this work have primarily focused on phenomenology of the sunk cost fallacy in the context of aging and have largely focused on developmental processes from a psychology perspective, departing from comparative studies across species as well as advancements in neuroscience. As a result, the original notion put forth by Arkes & Ayton has remained largely unchallenged, with caveats, exceptions, and inconsistencies found in follow-up human studies informing newer theories of developmental psychology that further depart from cross-species and neurobiological mechanistic work on decision-making information processing (2).

For instance, work on impulsivity in children is mixed and at odds with Arkes & Ayton’s claim that children are more likely to resist sunk costs (2). Several studies in

children report sometimes positive, sometimes negative, or sometimes no correlation between decision-making competency or impulsivity inventory scales and susceptibility to honor sunk costs (47-49). Various explanations attempting to explain such discrepancies have been proposed, rooted in self-regulation, perseveration, learned rules to not waste, and ability to recognize when to activate appropriate heuristics (49-58). Such decision-making competency and impulsivity inventories have been well-validated, control for general mental ability and intelligence, and are capable of predicting, for example real-world negative socioeconomic and health outcomes attributed to lifestyle choices, including risk of substance abuse and delinquency later in life (59-61). Other studies show, at the other end of the spectrum, that older adults too are less sensitive to sunk costs than younger adults, adding an inverted-U shape trajectory to the variable prevalence of this cognitive bias (62). Such reports suggest that older adults are less likely to ruminate on past expenditures and prior negative events, are better at coping with failed plans, and are less likely to mention sunk costs when making decisions about the future (62-69). None of these studies include any mechanistic data nor interpretation of their work in the context of changing neural circuits of different, parallel decision-making valuation systems. We propose that future studies should integrate biologically plausible models of the neural mechanisms underlying nonlinear developmental changes in separable decision systems (22-24). Our data can offer a fresh lens through which we and others can revisit past literature and inform future experiments.

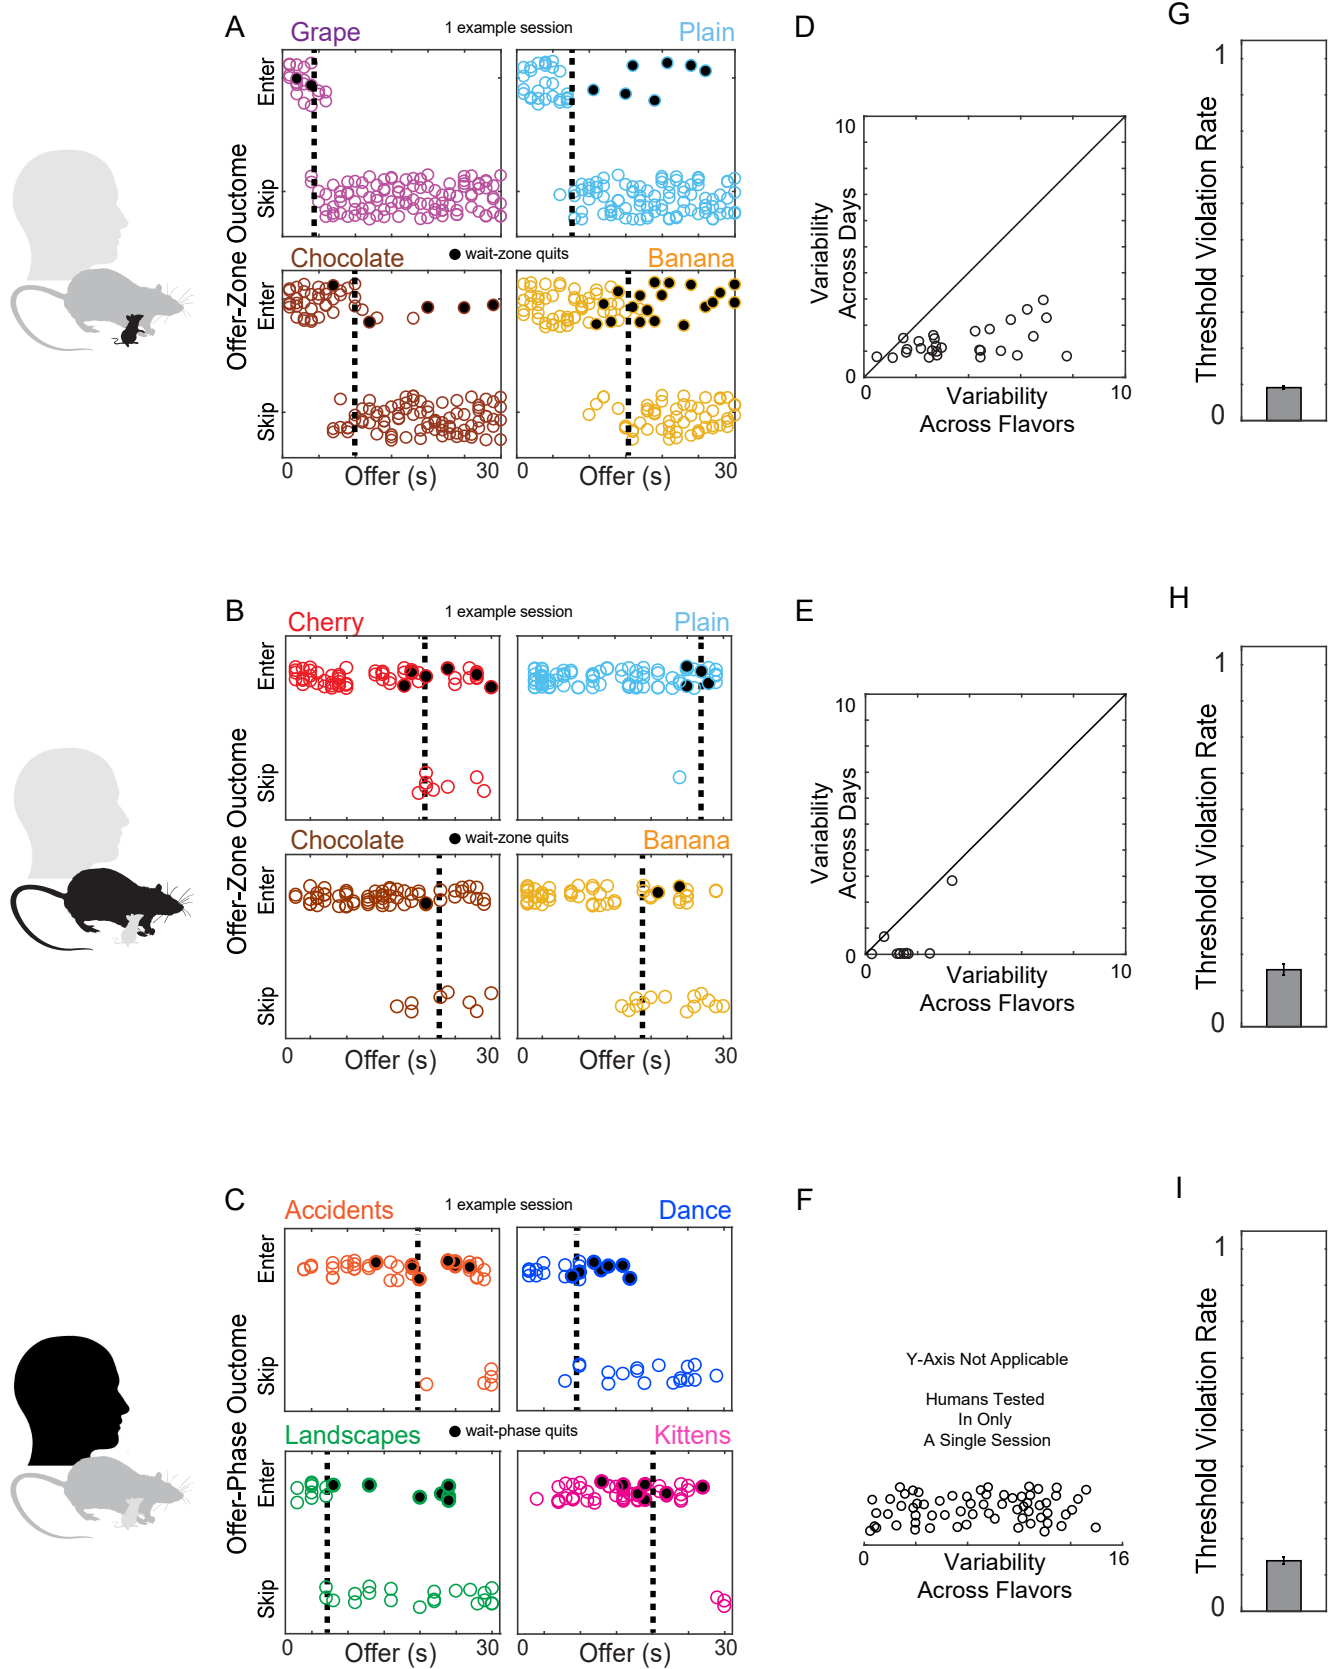


Fig. S1. Economic thresholds and budgets in Restaurant-Row and Web-Surf Tasks.

(A-C) Mice (A), rats (B), and humans (C) entered low delays and skipped high delays in the offer zone, while infrequently quitting once in the wait zone (black dots). Dashed vertical black lines represent calculated offer zone and wait zone “thresholds” of willingness to budget time. Thresholds were measured from the inflection point of fitting a sigmoid curve to earns vs. not-earns as a function of offer cost. (D-F) Threshold variability between flavors along the x-axis in mice (D), rats (E), and humans (F) as well as within flavor across days along the y-axis in mice (D) and rats (E). Threshold variability across days in rodents was relatively stable. (G-I) Probability of trials that were either earned when the offer was greater than threshold or not earned when the offer was less than threshold in mice (G), rats (H) and humans (I).

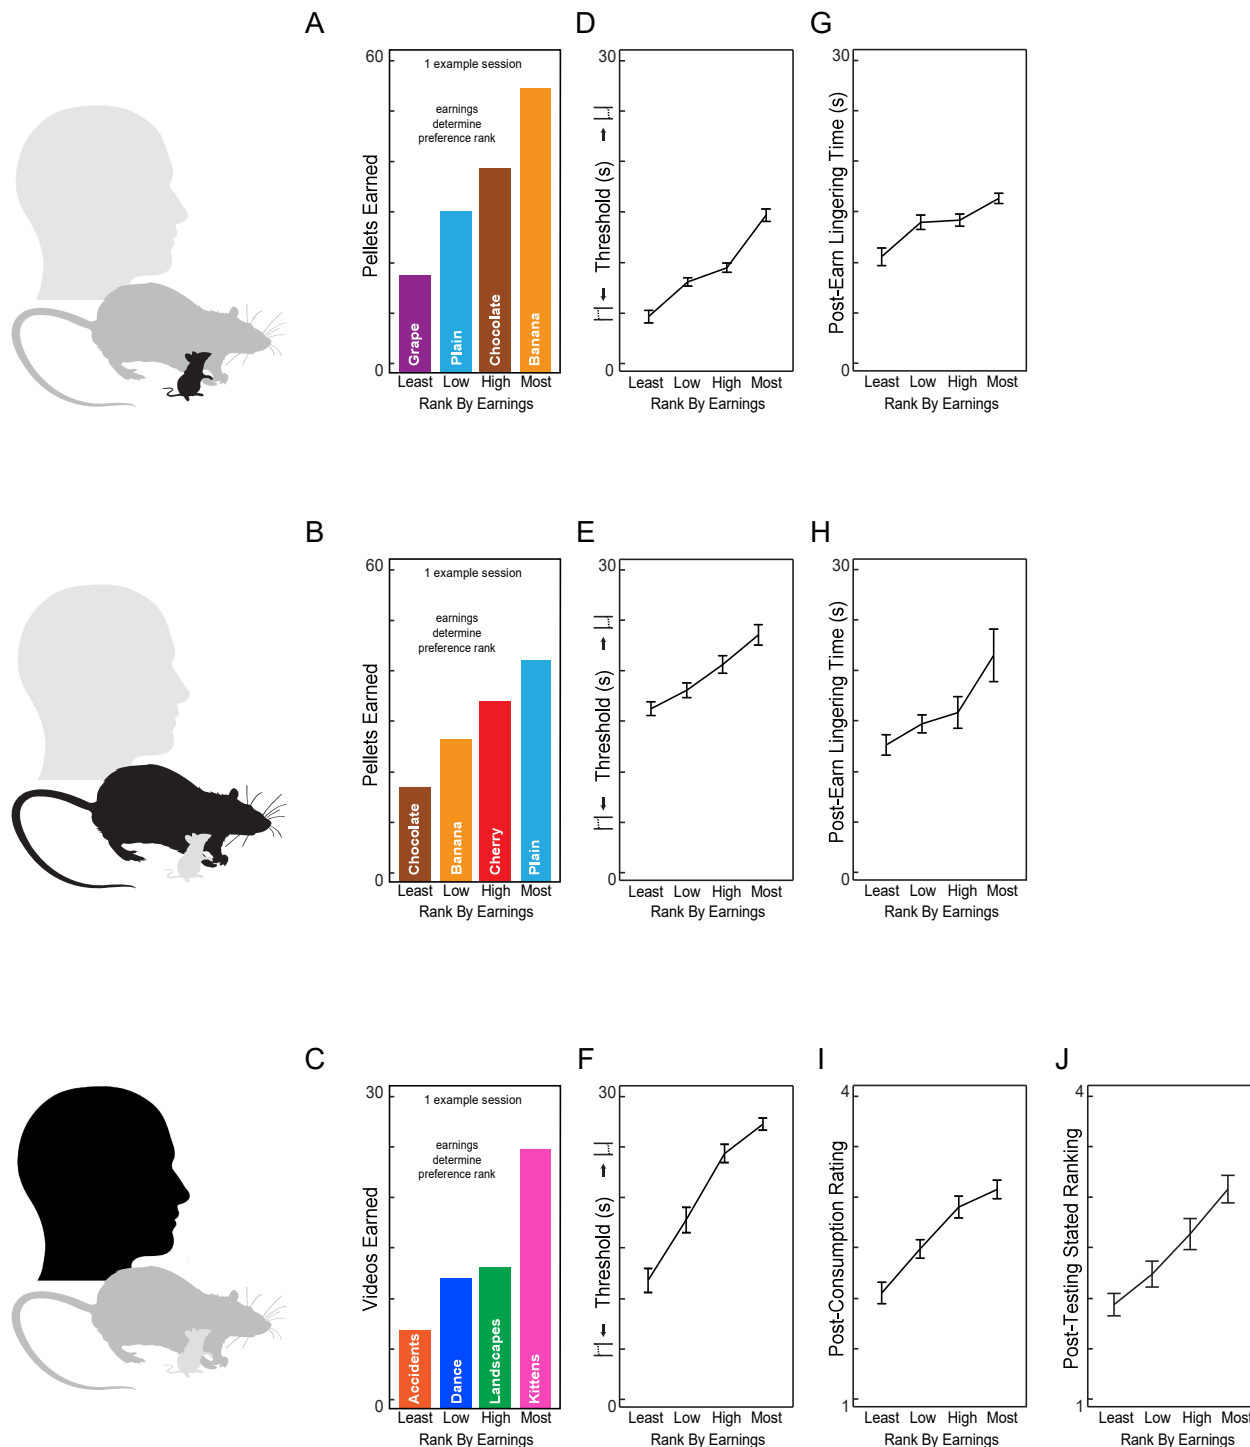


Fig. S2. Multiple valuation metrics of subjective preferences.

(A-C) Flavors were ranked from least- to most-preferred by summing the number of rewards earned in each restaurant or gallery in a single session. Panels show one example session in mice (A), rats (B), and humans (C). (D-F) Average thresholds sorted by rankings as defined in (A-C) were higher for more-preferred flavors in mice (D), rats (E), and humans (F). (G-I) Post-consumatory hedonic valuations sorted by rankings as defined in (A-C) were higher for more-preferred flavors in mice (G), rats (H), and humans (I). See Supplemental Discussion. (J) After the testing session was completed, human subjects stated rankings in a survey debrief session. Stated preferences were higher for more-preferred genres sorted by rankings as defined in (C). Multivariate Pearson and Spearman correlation analysis controlling for multiple comparisons found that all valuation metrics (earn rankings, calculated thresholds, post-consumption behaviors, and stated preferences [humans only]) significantly correlated with each other within each species (all correlations, $P < 0.001$).

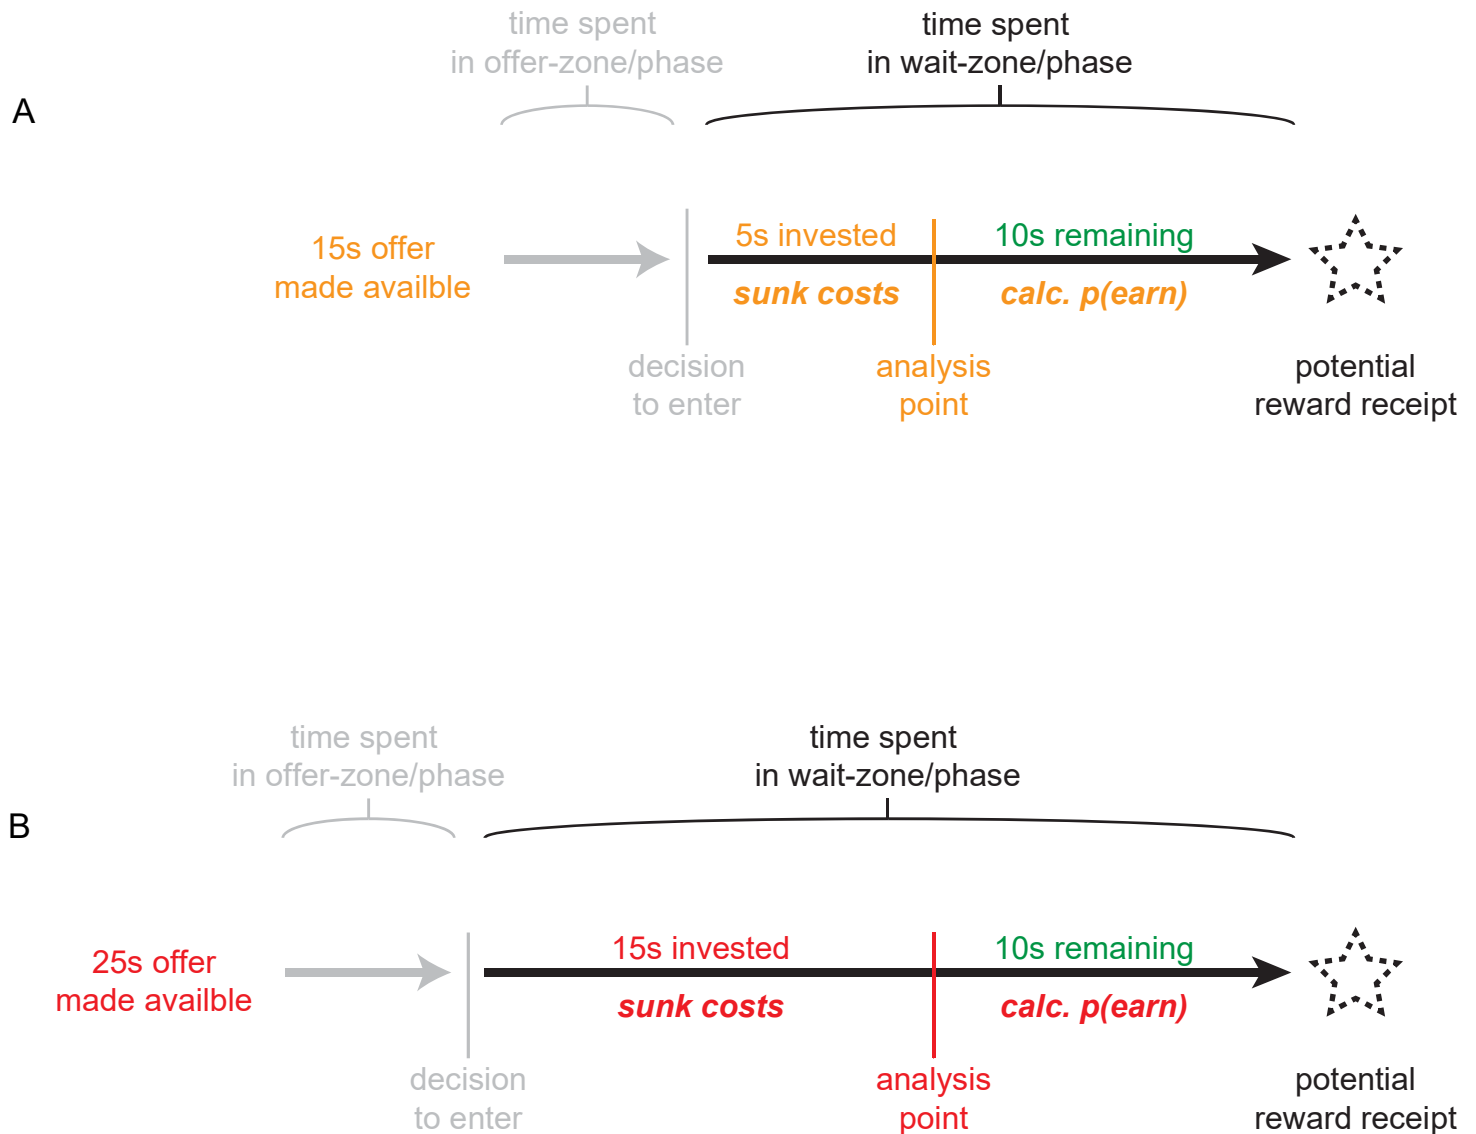


Fig. S3. Example economic scenarios in the wait zone sunk-cost analysis.

(A) Example in which a subject is offered a 15s-reward and chooses to enter the wait zone. [e.g., 100 cases of 15s offers entered]. Once in the wait zone, subjects could choose to quit at any moment. For example, taking trials where subjects had invested 5s in the wait zone and had not yet quit from that point onward [e.g., 25 trials already quit with <5s spent waiting, leaving 75 / 100 trials “surviving” so far], we calculated the probability of earning a reward in the 10s window remaining in the countdown [e.g., 25 / 75 trials survived from that point forward until countdown completion = 0.33]. (B) Contrast with an example with a different initial starting offer: 25s. In this example, trials where subjects had invested 15s in the wait zone and had not yet quit have the same 10s window remaining in the countdown. We compared the $p(\text{earn})$ probabilities calculated in scenario A with the $p(\text{earn})$ probabilities calculated in scenario B. The sunk cost effect would predict an observed increase in $p(\text{earn})$ after investing 15s (B) compared to after investing 5s (A). The Restaurant Row and Web-Surf Tasks provide multiple initial starting delays of the offer as well as shifts in the analysis point, highly parameterizing the sunk cost effect along a continuum, along past-(sunk)- and future-cost dimensions.

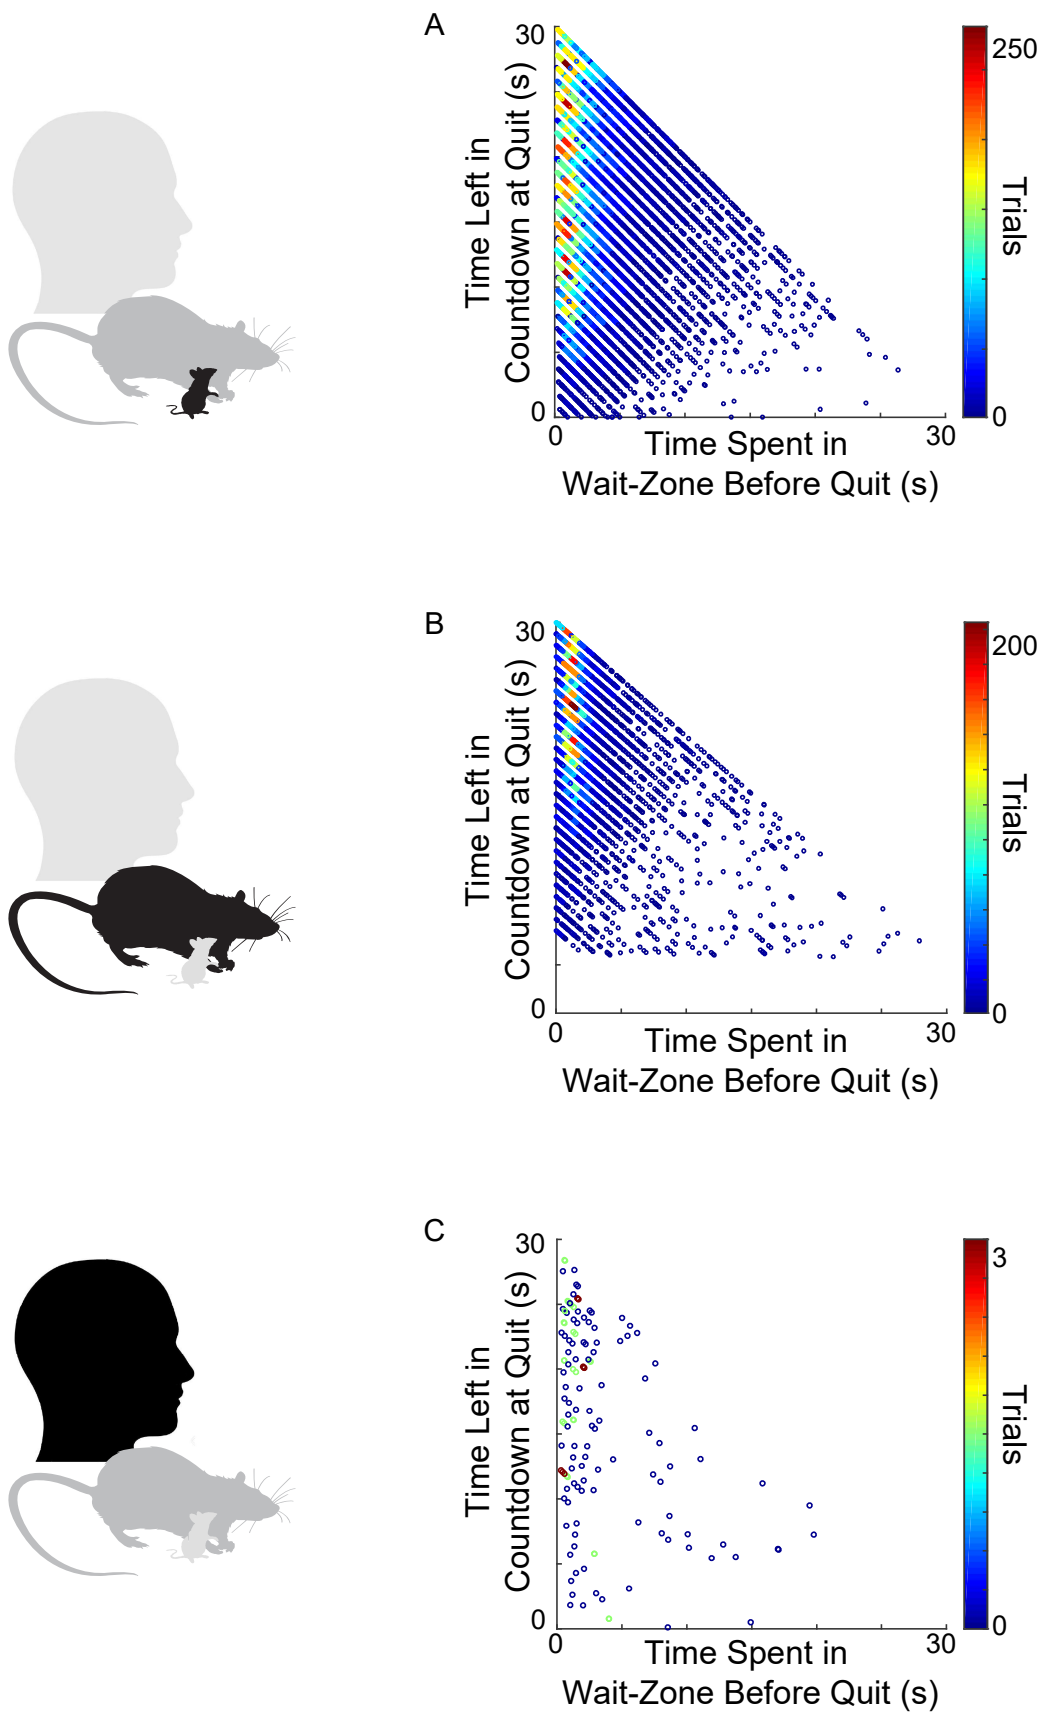


Fig. S4. Distribution of trials with varying economic scenarios for use in the wait zone sunk cost analysis. Each panel shows enter-then-quit trials plotting where time spent in the wait zone before quitting is plotted on the x-axis and time remaining in the countdown at the moment of quitting is plotted on the y-axis. Color axis indicates number of trials binned in each unique economic scenario for mice (A), rats (B), and humans (C).

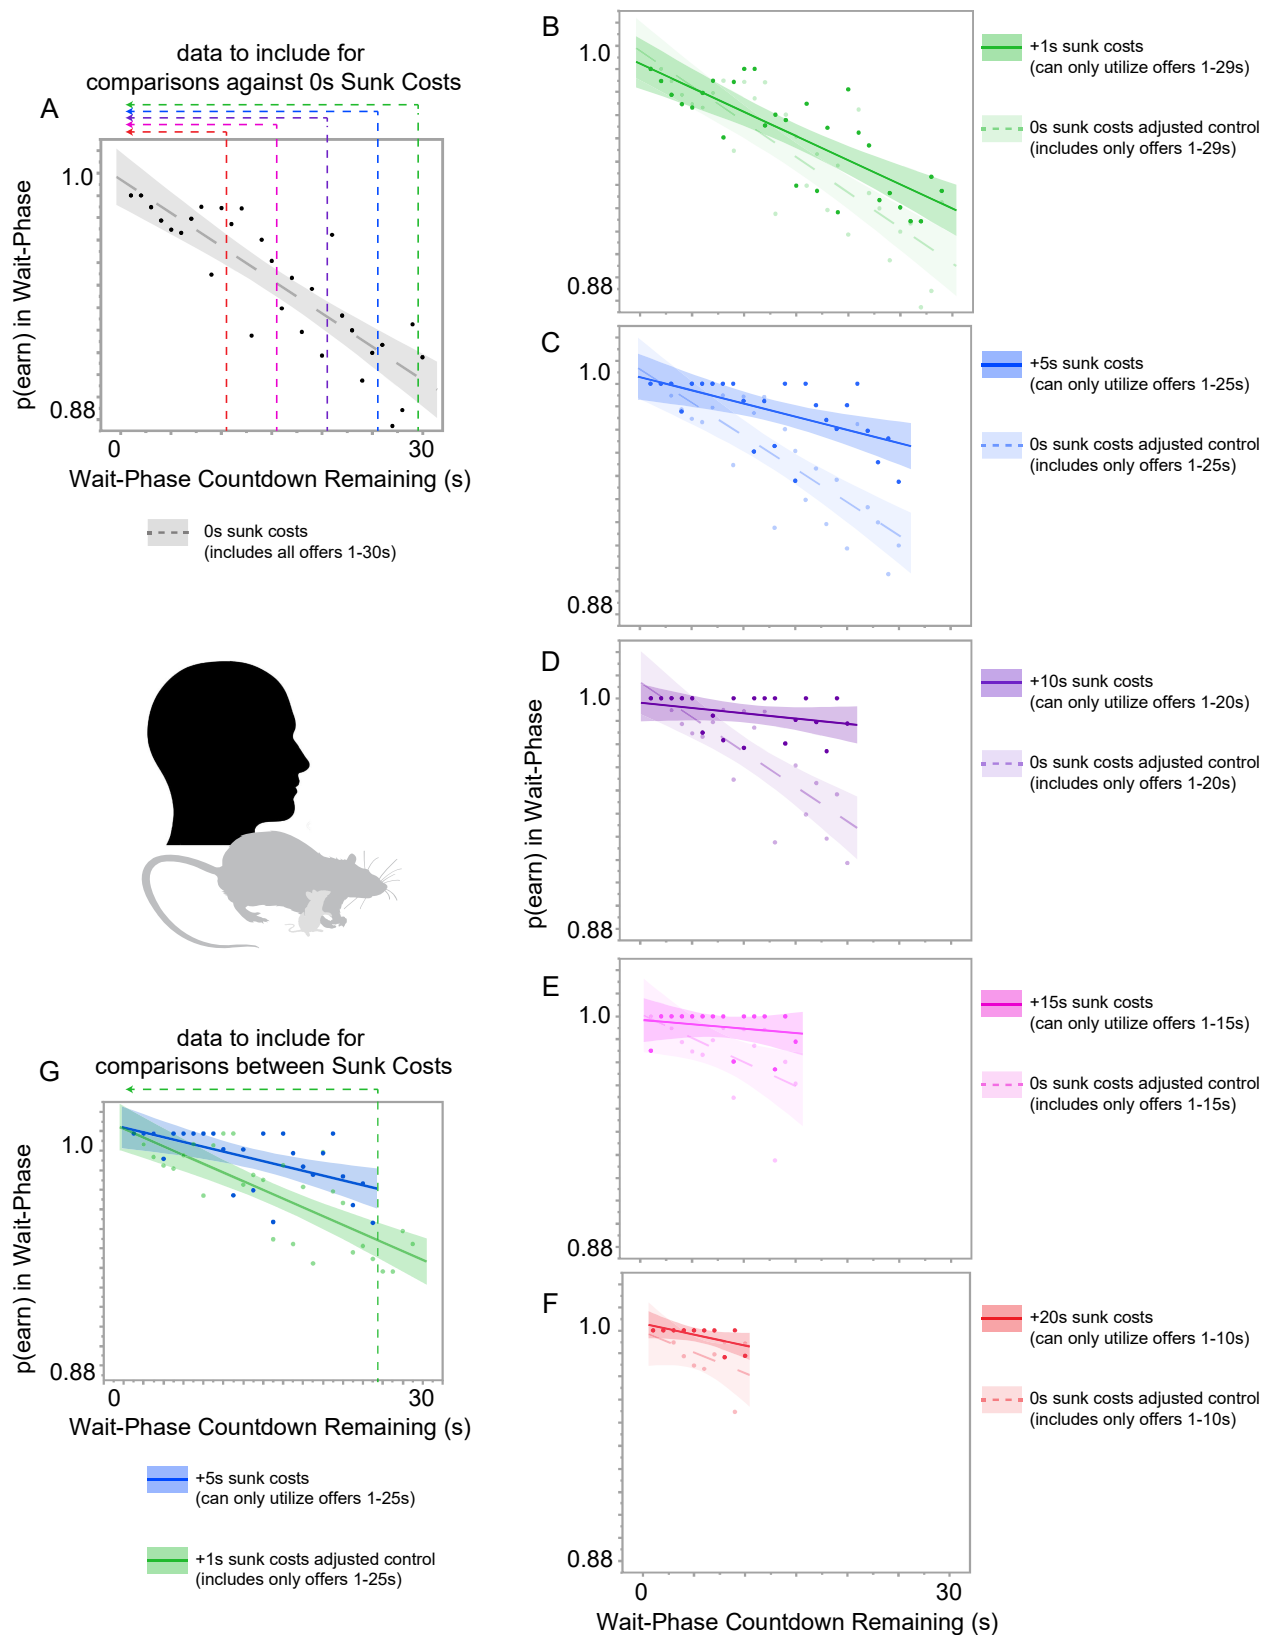


Fig. S5. Visualization of method to control comparisons between economic conditions in wait zone sunk cost analysis.

Using human data as an example, we ensured that our interpretation of slopes from linear regressions was not skewed by the different distributions of available data in the different sunk cost conditions. For instance, data from “30s remaining in the countdown” does not exist for the +1s sunk cost condition while it does exist for the zero sunk cost condition. To correct for this potential confound, we used regressions on the zero sunk cost condition iteratively leaving out the right most data points successively as the underlying control (A). Colored arrows indicate data that was included for each adjusted control regression: +1s-sunk costs (green), +5s-sunk costs (blue), +10s-sunk costs (purple), +15s-sunk costs (magenta), and +20s-sunk costs (red). (B-F) Examples of comparing regressions between each sunk cost condition and the zero sunk cost condition adjusted for the progressively smaller dataset illustrated in (A). (G) Similar concept of using regressions on progressively shortened datasets instead here to compare two sunk cost conditions against each other. In this example, the +5 sunk cost condition can only include offers from 1-25s. Thus, comparison against the +1 sunk cost condition, when adjusted, includes that same limited range of offers.

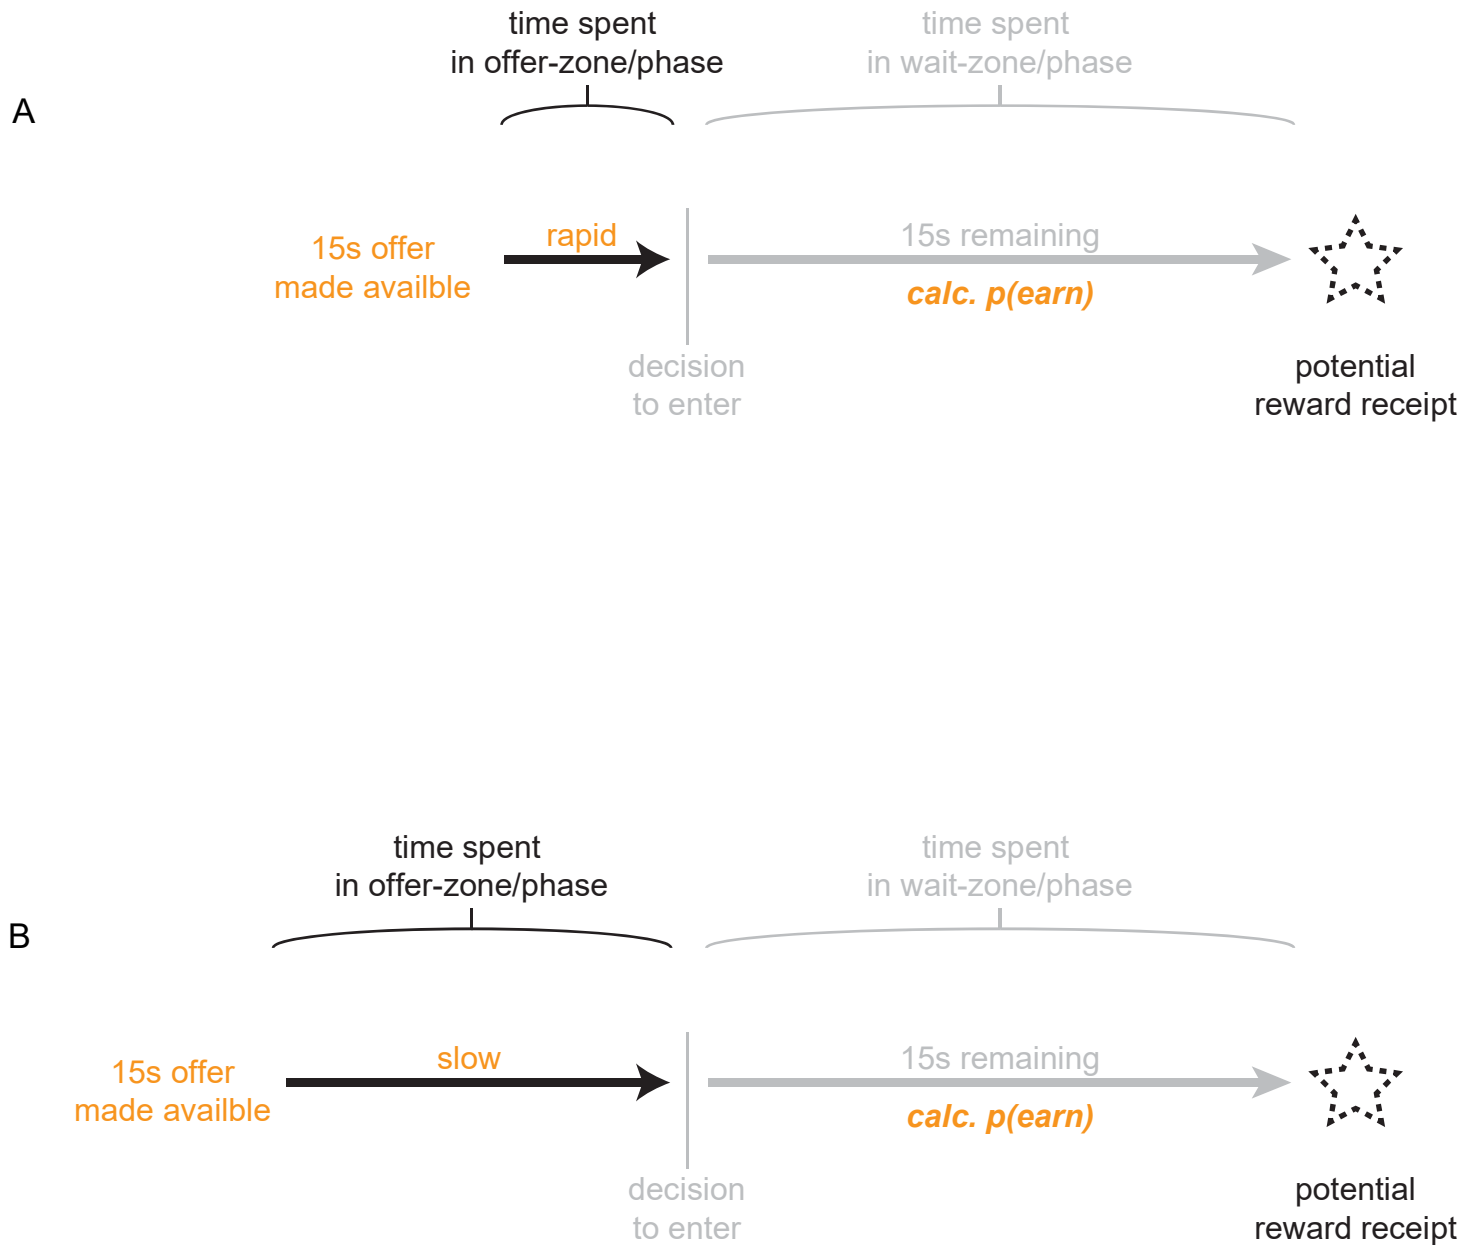


Fig. S6. Example economic scenarios in the offer zone sunk-cost analysis.

In these two scenarios, subjects were offered a 15s-reward and chose to enter the wait zone. Scenario A illustrates a rapid decision to enter whereas scenario B illustrates a slow decision to enter where subjects invested more time in this initial offer zone decision. In both examples, the delay is the same (both 15s). We calculated the probability of earning once in the wait zone. The sunk cost effect, based on a resource-depletion/wastefulness-avoidance would predict an observed increase in $p(\text{earn})$ in (B) compared to (A). However, we found no such changes in mice, rats, or humans.

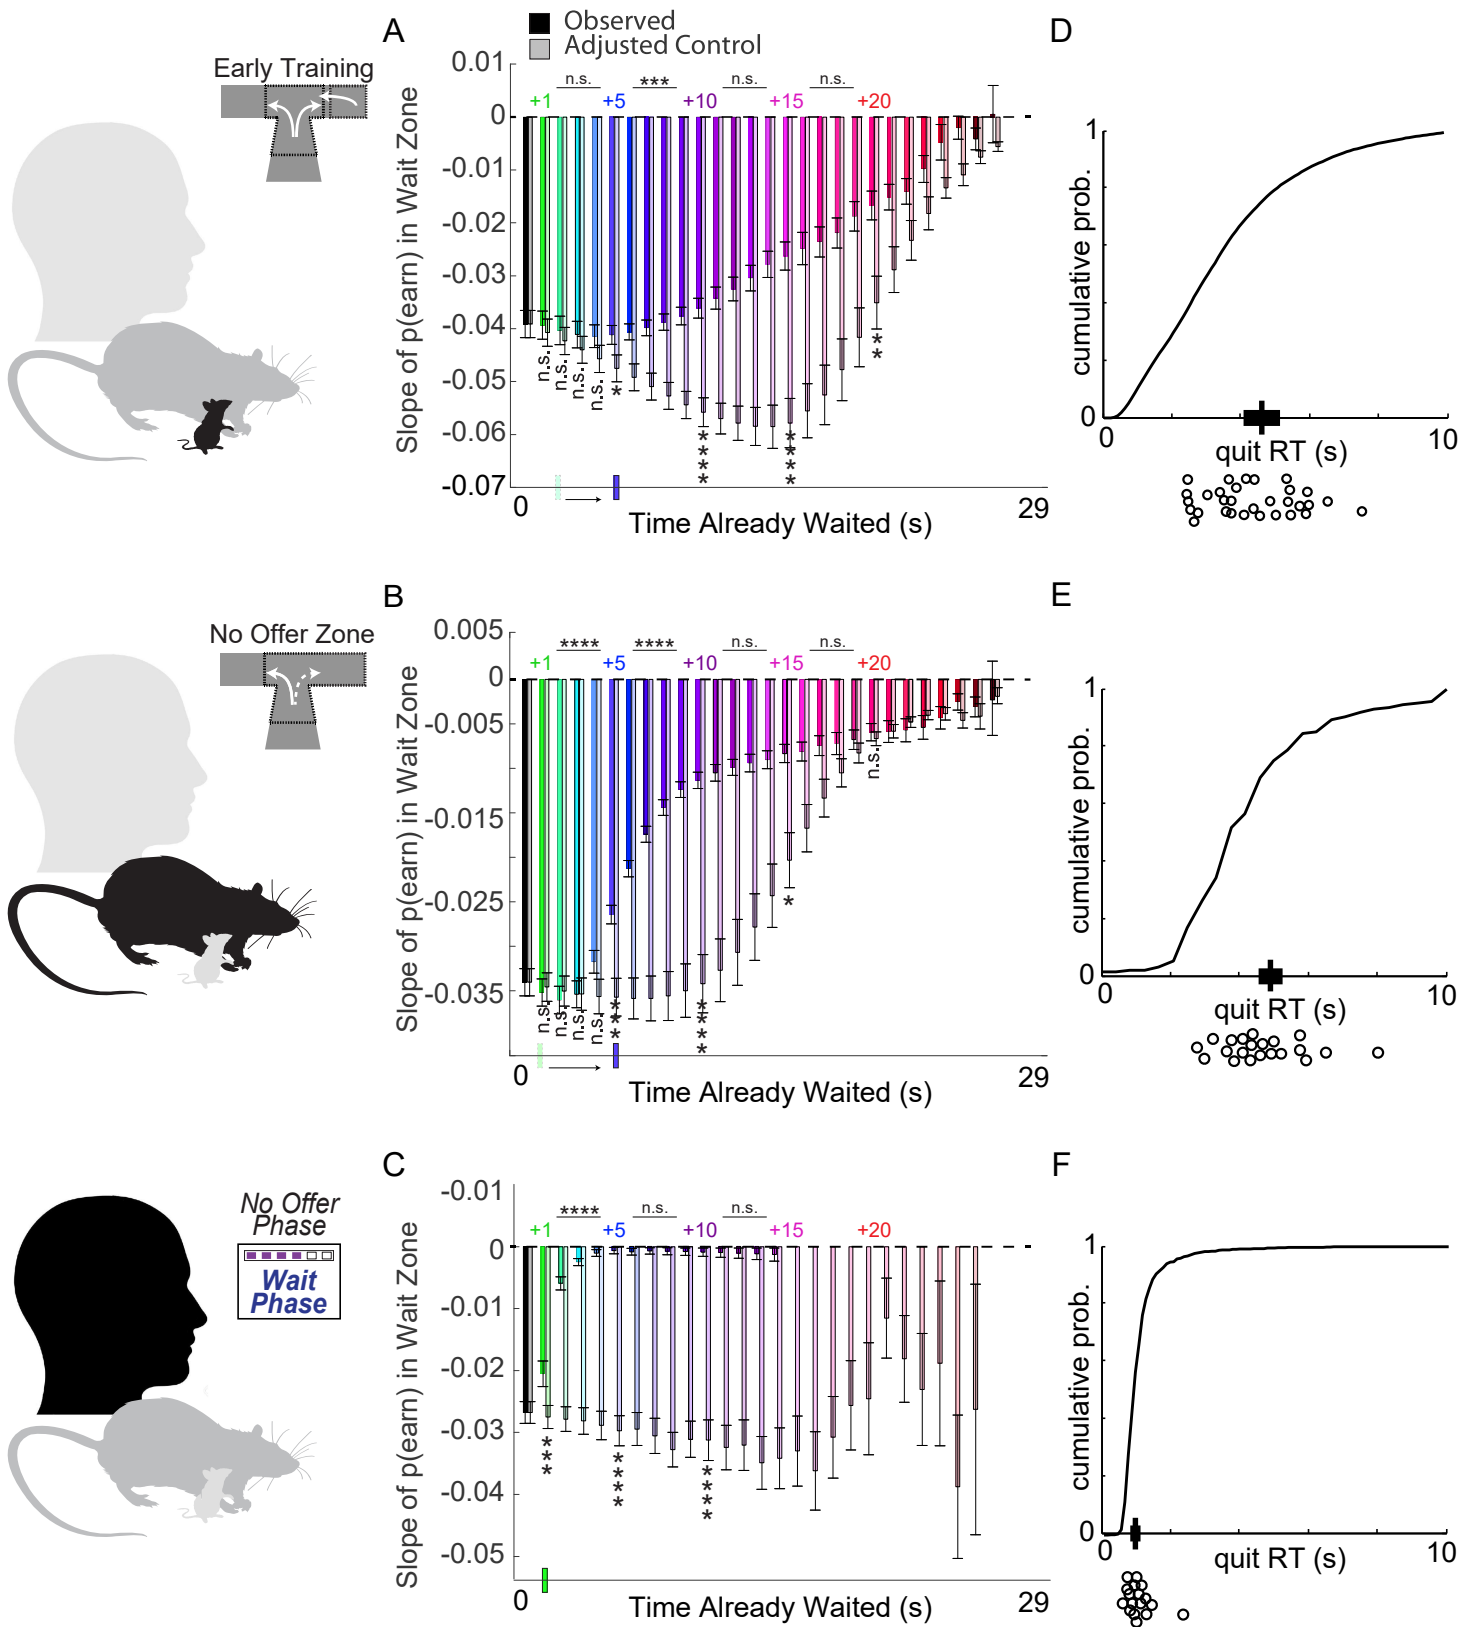


Fig. S7. Additional sunk cost analyses across training and in other task variants.

(A) Mice early on in training still displayed sunk cost effects, however did so with a delayed onset compared to late in training (see Fig.2D). Blue colored tick on x-axis indicates time in wait zone until first significant sunk-cost effect was observed. Transparent green tick replotted from Fig.2D for reference in mice with extensive training. (B-C) A separate cohort of rats (B) and humans (C) were trained on variants of the Restaurant Row and Web-Surf Tasks without an offer zone where countdown began immediately and subjects only made quit decisions. Colored tick on x-axis indicates time in wait zone until first significant sunk-cost effect was observed. (B) Transparent green tick replotted from Fig.2E for reference in rats on the task variant with an offer zone. (D-F) Cumulative probability distribution of quit reaction time in the wait zone in mice (D), rats (E), and humans (F). Black tick on x-axis indicates cohort average reaction time ± 1 SEM. Dot scatter below axis represent individual subjects' average reaction time. ANOVAs were used to compare slopes of linear regression models, testing for interactions with sunk cost conditions and controls, correcting for multiple comparisons. Not-significant (n.s.) $P > 0.05$, * $P < 0.05$, ** $P < 0.01$, *** $P < 0.001$, **** $P < 0.0001$.

Offer-Zone Path Trajectory




Fig. S8. Measurement of Vicarious Trial and Error in rodents.

Illustrated here are X-Y-locations of a rodent's path-trajectory in the offer zone over time during two example trials. Vicarious trial and error measures the absolute integrated angular velocity over time and is measured in IdPhi units. (A) Example of a low vicarious trial and error event without any re-orientations. (B) Example of a high vicarious trial and error event with re-orientations at the choice-point.

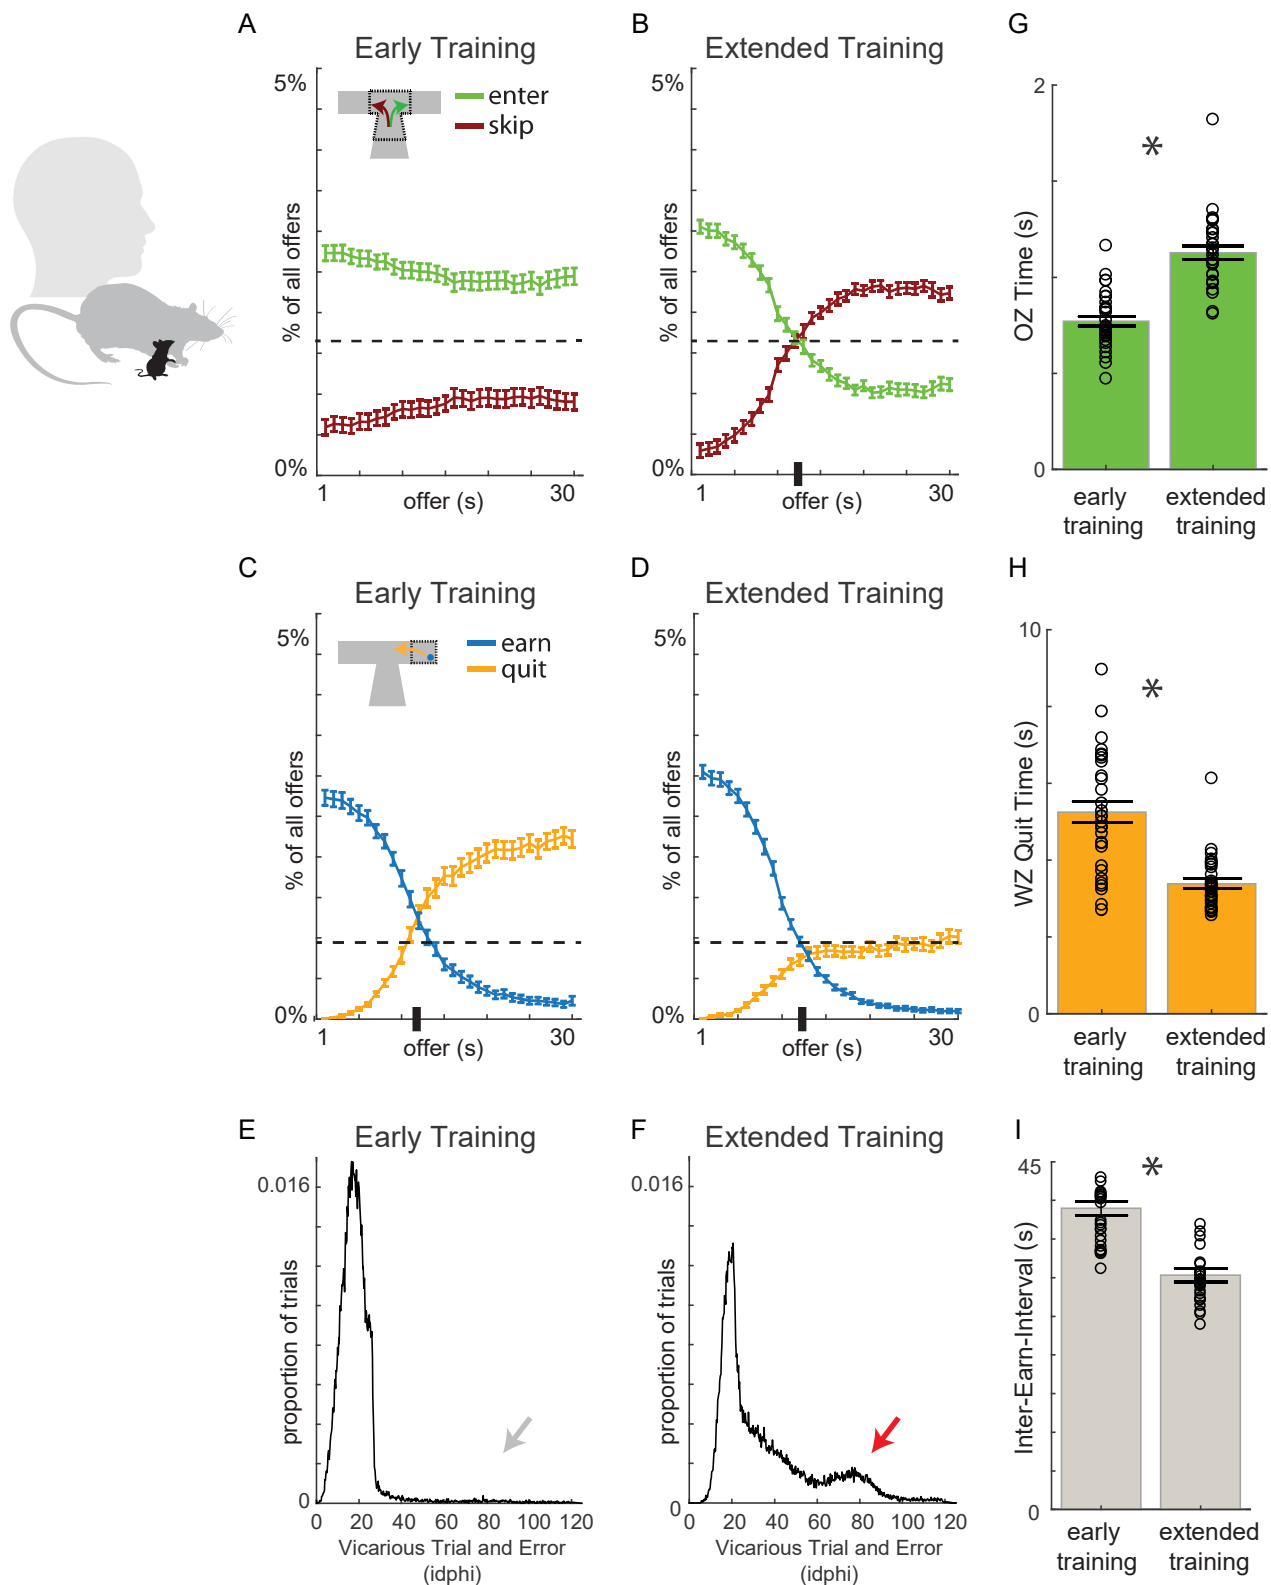


Fig. S9. Change in decision-making behavior over training in mice on the Restaurant Row task.

(A-B) Offer zone choice probability between skipping vs. entering as a function of offer length early in training (see Supplementary Discussion). (A) and after extended training (B). Mice entered nearly all offers indiscriminately early in training (A) while later learning to skip expensive offers (B). (C-D) Wait zone choice probability between quitting vs. earning as a function of offer length early in training (C) and after extended training (D). Mice learned to quit less with extended training, as they were more likely to accept offers in the offer zone they would be willing to earn. Black tick on the x-axis in B-D indicate cross over point of skip/enter or quit/earn decisions. Horizontal dashed black line indicates choice probabilities if decisions were random. (E-F) Probability density plots of vicarious trial and error behavior in the offer zone across early (E) vs. extended training (F). Red arrow in F emphasizes the presence of high vicarious trial and error events as mice learned to deliberate in the offer zone between cheap and expensive offers, illustrated in B. Gray arrow in (E) emphasis the absence of these high vicarious trial and error events as mice made snap judgments to accept all offers indiscriminately, illustrated in A. (G-I) As a consequence of learned changes in decision-making strategies over training, offer zone reaction time increased (G, $F=67.38$, $*p<0.0001$), wait zone quit time decreased (H, $F=36.48$, $*p<0.0001$), and reinforcement rate increased (I, decrease in inter-earn-interval time between pellet consumption, $F=45.26$, $*p<0.0001$). (open circles in (G-I) represent individual animals, error bars ± 1 SEM)



Restaurant Preference Ranking

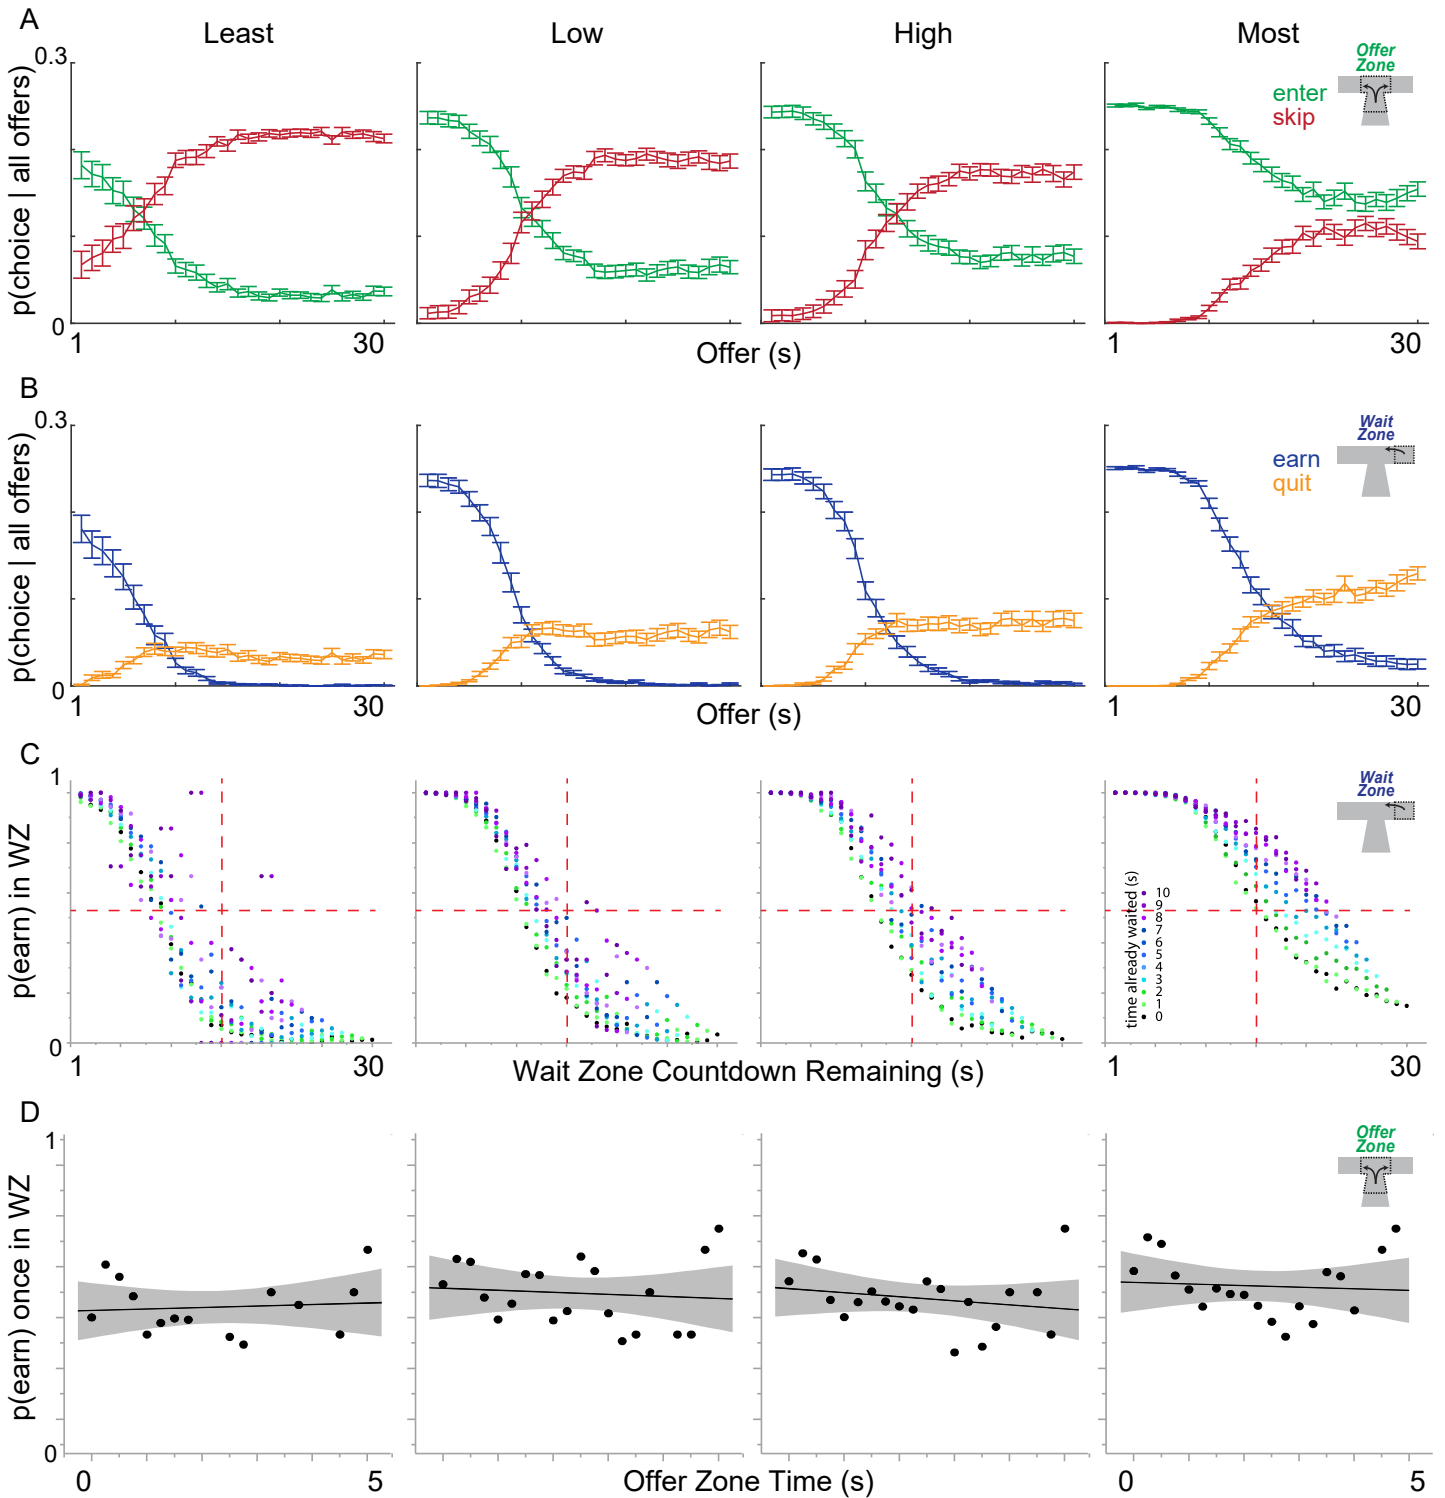


Fig. S10. Asymmetries in choices split by subjective value reassures cost discriminability in mice.

Vertical columns split data by rankings (categorized in Fig.S2). (A-B) Choice outcome probabilities in the offer zone (A) or wait zone (B) relative to all offers encountered across all restaurants for a given offer length. Because offers are randomly presented from trial to trial, differences in offer zone ability to skip vs. enter high cost offers separated by subjective value helps ensure mice are capable of discriminating tones and information uncertainty is less of a confound. (C) Sunk cost effects in the wait zone exist in all restaurants (least: $F=3.33$, $p<0.05$; low: $F=4.34$, $p<0.05$; high: $F=10.11$, $p<0.01$; most: $F=26.69$, $p<0.001$). (D) Sunk cost effects in the offer zone do not exist in all restaurants (Pearson coefficient: least: $r=0.089$, $p=0.75$; low: $r=-0.097$, $p=0.69$; high: $r=-0.207$, $p=0.38$; most: $r=-0.077$, $p=0.74$). $p(\text{earn})$ in (C-D) are relative to each restaurant's offers. (error bars ± 1 SEM, shaded region represents 95% confidence interval of correlation)



(Primary) Cohort 3
Mice (n=32)
(2016)

(Replication) Cohort 6
Mice (n=32)
(2017)

Both underwent an identical 70+ days training regimen tested in the same exact mazes

Varied number and gender of experimenters

85% free feeding weight food-deprivation

90% free feeding weight food-deprivation

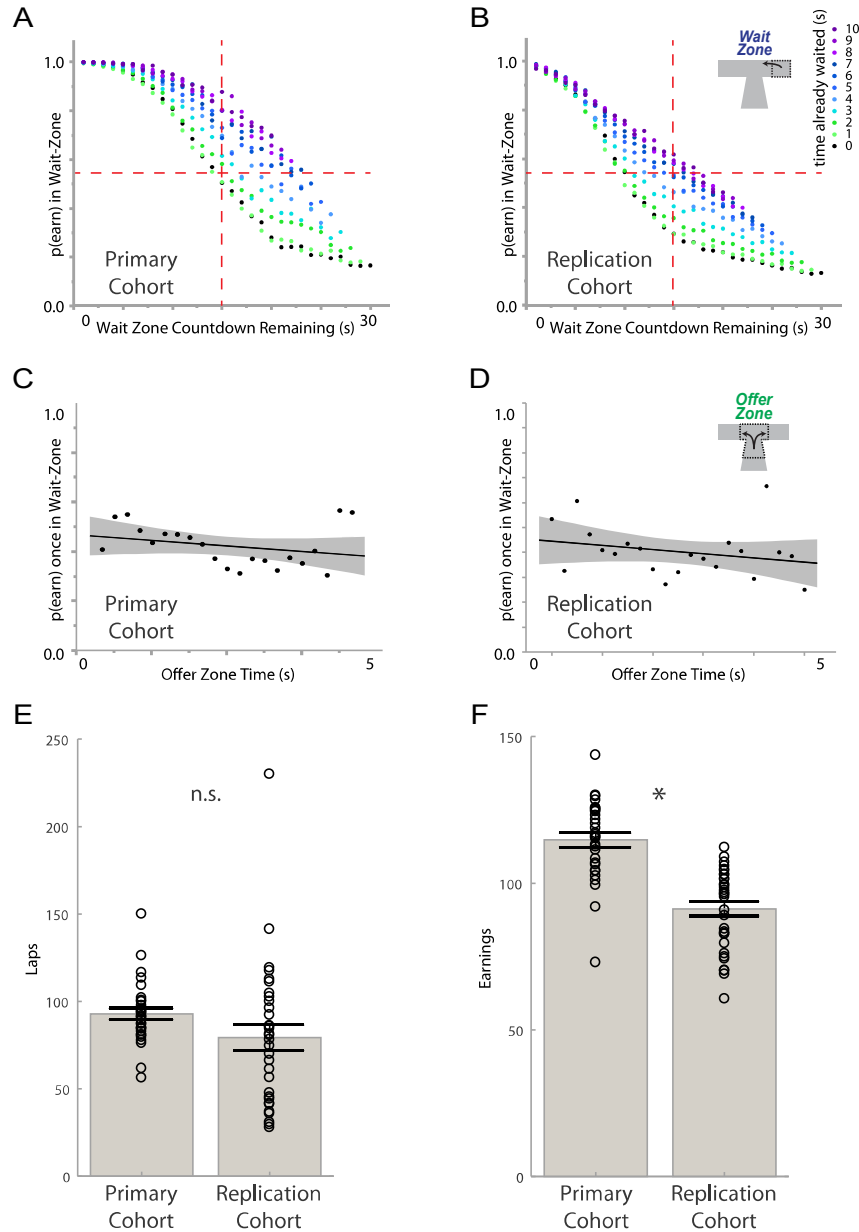


Fig. S11. Replication cohort of mice varying environmental factors.

An additional 32 mice were tested on the main variant of the task (offer zone and wait zone). Experimenters were varied, and these mice were trained in the exact same apparatuses as the original cohort under the exact same training protocol. Additionally, these mice were intentionally food deprived to a lesser extreme. Despite these different environmental factors, mice still demonstrated a robust sunk cost effect in the wait zone (B, $F=26.56$, $p<0.01$, see original cohort in A for comparison) but no effect as a function of time spent in the offer zone (D, $r=-0.254$, $p=0.27$, see C for comparison). Environmental factors, while not affecting locomotor abilities nor number of offers encountered by not altering number of laps run (E, $F=2.86$, $p=0.10$), did result in a reduction in average number of pellets earned on the task (F, $F=46.79$, $*p<0.0001$). This is reflected in an overall left shift of the economic budget curves in panel B (vertical red dashed line indicates 15s, horizontal red line indicates 0.5 p(earn)), reflecting an alteration in reward value and willingness to wait. (open circles in (E-F) represent individual animals, error bars ± 1 SEM, shaded region represents 95% confidence interval of correlation, not significant, n.s.)

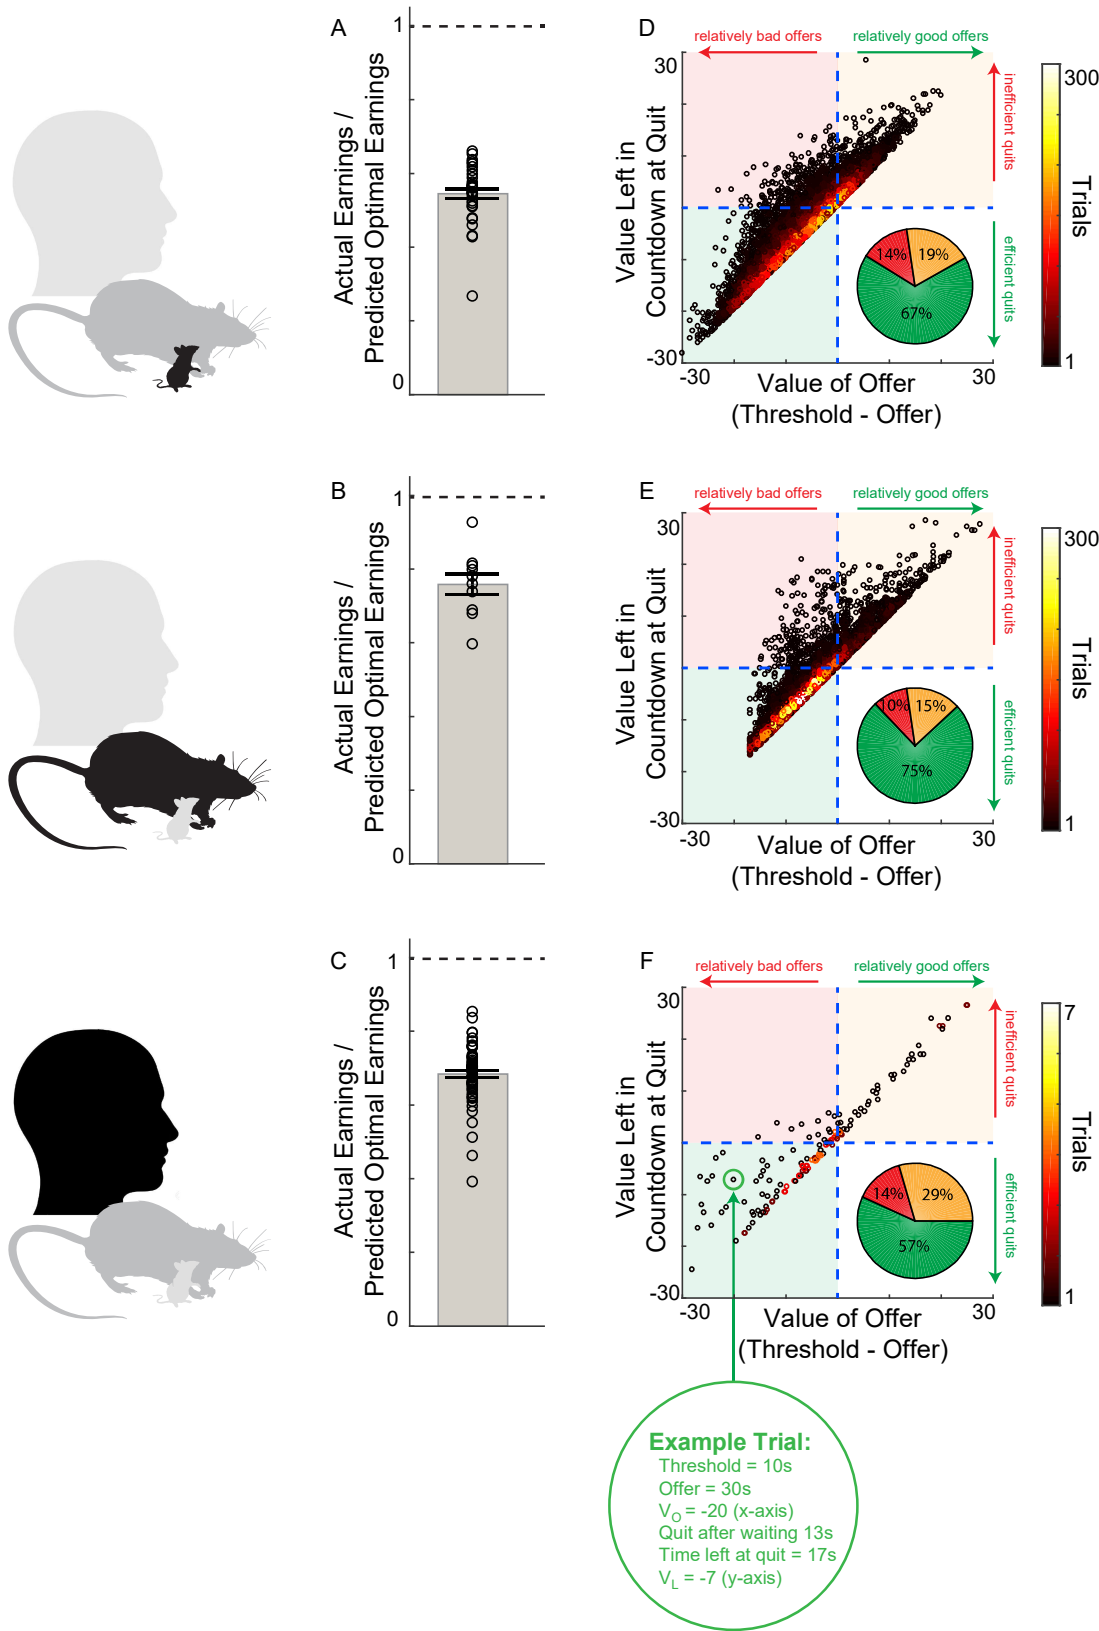


Fig. S12. Modeling sub-optimality and economic efficiency across species.

Proportion of total rewards mice (A), rats (B), and humans (C) actually earned relative to model-estimated maximal predicted earnings taking into account individual differences in behavioral performance and subjective valuation preferences (see Methods and Supplemental Discussion). Observations are normalized to perfectly optimal earnings in our model (1.0, e.g., using minimal reaction times and no quits based on each subject's behavior and restaurant-specific thresholds). Mice, rats, and humans all behaved significantly sub-optimally according to this model (t-test comparing means against 1.0, mice: $t=-31.962$, $p<0.0001$; rats: $t=-8.43$, $p<0.0001$; humans: $t=-20.29$, $p<0.0001$). (open circles in (A-C) represent individual animals, error bars ± 1 SEM). (D-F) Pooled data across all subjects to visualize economic efficiency of quit decisions. This was measured by calculating two "value" metrics (offer value defined as the difference between each subject's restaurant-specific normalized threshold and offer, value left defined as the difference between threshold and the remaining countdown at the moment of quitting) (see Methods and Supplemental Discussion). Dashed blue lines indicate 0 value for both metrics, where either the offer = threshold (vertical) or the time remaining in the countdown when quitting = threshold (horizontal). In (F), a single example trial is circled in green and described in detail for demonstration purposes. Pie charts in (D-F) quantify percentage of all quits that fall into each of the 3 color-coded quadrants of the scatter plot (green = smart quits; red/orange = bad quits).

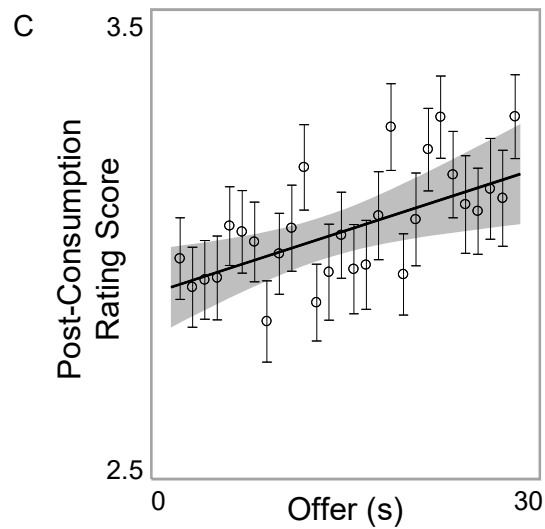

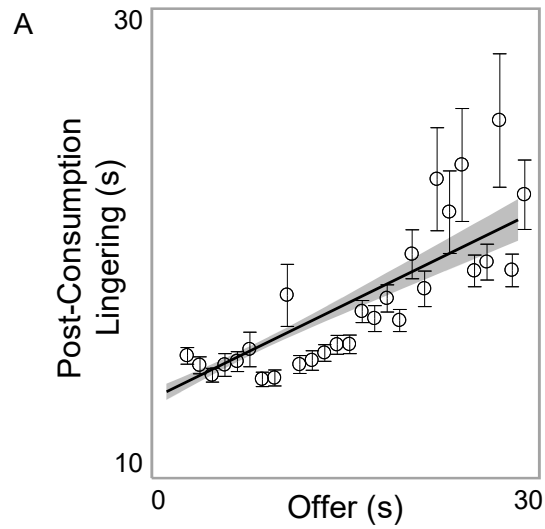


Fig. S13. Offer cost and post-consumption valuations.

(A-C) Hedonic valuation metrics measured immediately after consuming rewards as a function of offer length in mice (A), rats (B), and humans (C). Rodents (A-B) lingered at the reward site after consuming rewards before leaving for the next trial at the next restaurant (see Supplementary Discussion). Humans (C) rated videos on a scale from 1-4 (4=most enjoyed) immediately after viewing earned rewards. Post-consumption valuations positively correlated with offer cost in mice (A, Pearson coefficient $r = 0.737$, $P < 0.001$), rats (B, $r = 0.733$, $P < 0.001$), and humans (C, $r = 0.473$, $P < 0.05$). (error bars ± 1 SEM, shaded region represents 95% confidence interval of correlation).

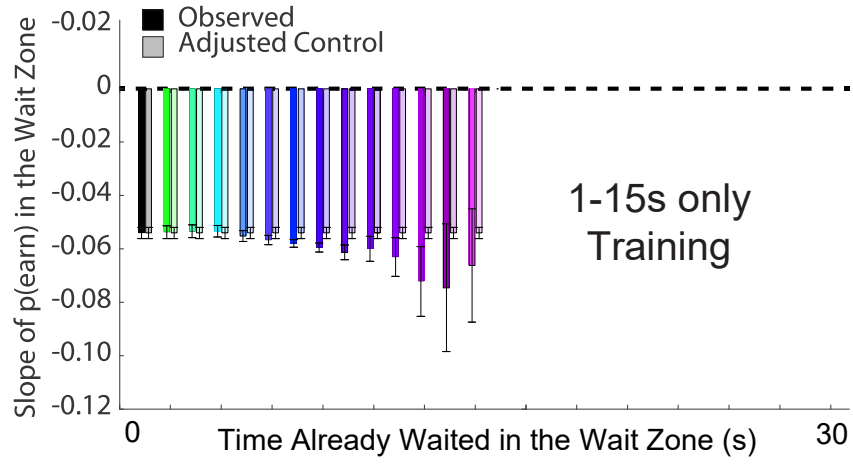


Fig. S14. Wait zone sunk cost analysis in mice in a reward-rich environment.

Before being exposed to the full range of 1-30s offers in the Restaurant Row task, mice were trained on 1-15s offers where reward cost was relatively inexpensive (see Methods, supplementary text S6). Slope of linear regressions plotted with ± 1 SEM. There were no significant differences ($P > 0.05$) between sunk costs conditions, controlling for adjusted data distributions previously described.

	Cohort 1	Cohort 2	Cohort 3	Cohort 4	Cohort 5	Cohort 6
Date	Spring 2014	Spring 2014	Fall 2016	Spring 2017	Spring 2017	Fall 2017
Species	Rat	Human	Mouse	Rat	Human	Mouse
Breed	Brown-Norway	undergraduates	C57BL6J	Fisher Brown-Norway	undergraduates	C57BL6J
Sample Size & Sex	22 (M) & 0 (F)	4 (M) & 13 (F)	32 (M) & 0 (F)	4 (M) & 6 (F)	24 (M) & 41 (F)	32 (M) & 0 (F)
Age	8-12 months	19.63 years (mean)	13 weeks	6-10 months	20.23 years (mean)	13 weeks
Task Variant	wait zone only	wait phase only	offer zone + wait zone	offer zone + wait zone	offer phase + wait phase	offer zone + wait zone
Experimenters & Gender	1 (M) & 1 (F)	3 (M) & 5 (F)	2 (M) & 3 (F)	2 (M) & 2 (F)	0 (M) & 6 (F)	3 (M) & 3 (F)
Length of Training	20+ days	5 minutes	70+ days	20+ days	5 minutes	70+ days
Food Deprivation	>80% free weight	N / A	>80% free weight	>85% free weight	N / A	>90% free weight

Table S1. Cohort information in chronological order of data collection.

The sequence of experiments that took place across labs inspired and informed subsequent experiments. The similarities and differences across species in fact reflect a strength that our work capitalizes on, demonstrating robustness of main effects despite these variations as well as harnessing interesting differences in certain effects because of between-cohort differences. The approach and design we used makes this naturalistic foraging task easily expandable to a wide variety of populations, including different ages or patient populations in humans, as well as other strains of rodents or other species beyond rodents.