

# Chronic and Acute Stress Promote Overexploitation in Serial Decision Making

Jennifer K. Lenow,<sup>1</sup> Sara M. Constantino,<sup>1</sup>  Nathaniel D. Daw,<sup>3</sup> and Elizabeth A. Phelps<sup>1,2,4</sup>

<sup>1</sup>Department of Psychology and <sup>2</sup>Center for Neural Science, New York University, New York, New York 10003, <sup>3</sup>Princeton Neuroscience Institute and Department of Psychology, Princeton University, Princeton, New Jersey 08544, and <sup>4</sup>Nathan Kline Institute for Psychiatric Research, Orangeburg, New York 10962

Many decisions that humans make resemble foraging problems in which a currently available, known option must be weighed against an unknown alternative option. In such foraging decisions, the quality of the overall environment can be used as a proxy for estimating the value of future unknown options against which current prospects are compared. We hypothesized that such foraging-like decisions would be characteristically sensitive to stress, a physiological response that tracks biologically relevant changes in environmental context. Specifically, we hypothesized that stress would lead to more exploitative foraging behavior. To test this, we investigated how acute and chronic stress, as measured by changes in cortisol in response to an acute stress manipulation and subjective scores on a questionnaire assessing recent chronic stress, relate to performance in a virtual sequential foraging task. We found that both types of stress bias human decision makers toward overexploiting current options relative to an optimal policy. These findings suggest a possible computational role of stress in decision making in which stress biases judgments of environmental quality.

**Key words:** decision making; foraging; stress

## Significance Statement

Many of the most biologically relevant decisions that we make are foraging-like decisions about whether to stay with a current option or search the environment for a potentially better one. In the current study, we found that both acute physiological and chronic subjective stress are associated with greater overexploitation or staying at current options for longer than is optimal. These results suggest a domain-general way in which stress might bias foraging decisions through changing one's appraisal of the overall quality of the environment. These novel findings not only have implications for understanding how this important class of foraging decisions might be biologically implemented, but also for understanding the computational role of stress in behavior and cognition more broadly.

## Introduction

A key challenge that animals have faced throughout evolution is how to solve foraging-like problems (e.g., when to give up on a current patch of food in search of a better one). Today, humans still confront problems that are fundamentally analogous to these foraging situations (e.g., when to give up on a current job). Notably, these foraging decisions require us to choose between a

current known option and a future unknown option, raising a dilemma as to how these two quantities can be compared (or, principally, even estimated). Optimal foraging theory has provided an elegant solution to this class of patch-foraging decisions, the marginal value theorem (MVT), which prescribes that an animal should leave the current option when they expect on average to find a better alternative in the environment (Charnov, 1976). Indeed, intuitively, what would seem like a good option (e.g., job offer), relative to potential future options, depends on the overall quality of the environment (e.g., job market).

The average reward rate of the environment serves as the optimal leaving threshold because it effectively sets the opportunity cost of time spent exploiting the current option. When the instantaneous reward rate of the current, depleting option falls below this level, an animal's time would be better spent doing something else. The MVT thus predicts that animals should adjust their patch-leaving thresholds according to the average reward in the environment, for example, exploiting patches for

Received Nov. 22, 2016; revised March 16, 2017; accepted April 10, 2017.

Author contributions: J.K.L., S.M.C., N.D.D., and E.A.P. designed research; J.K.L. performed research; J.K.L. and S.M.C. contributed unpublished reagents/analytic tools; J.K.L. and N.D.D. analyzed data; J.K.L., S.M.C., N.D.D., and E.A.P. wrote the paper.

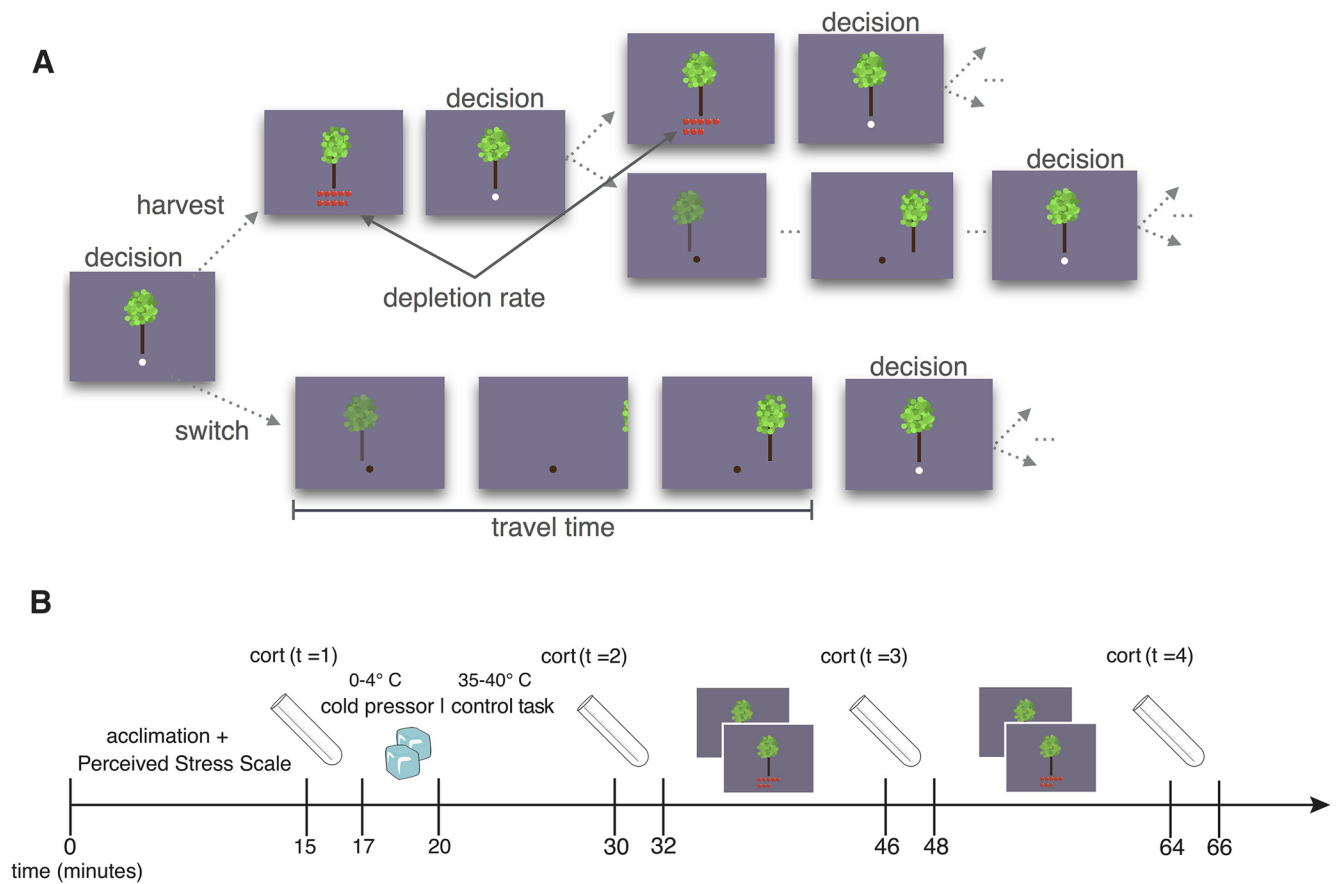
This work was supported by National Institutes of Health (Grant R01 AG039283 to E.A.P.) and the National Science Foundation (National Science Foundation Graduate Research Fellowship Program Grant to J.K.L.).

The authors declare no competing financial interests.

Correspondence should be addressed to Elizabeth A. Phelps, New York University, 6 Washington Place, Room 863, New York, NY 10001. E-mail: liz.phelps@nyu.edu.

DOI:10.1523/JNEUROSCI.3618-16.2017

Copyright © 2017 the authors 0270-6474/17/375681-09\$15.00/0



**Figure 1.** Task design and experimental procedures. **A**, Schematic of fixed-duration foraging task. In this task, subjects decide sequentially whether to harvest a currently displayed tree or travel to a different tree. Trees yield fewer apples with each successive harvest according to a randomly drawn depletion rate. New trees have an unharvested supply of apples, but entail a travel time (waiting) cost, the length of which varies across two environment types (with long and short travel times). **B**, Timeline of experimental procedures, which include the collection of questionnaire measurements, an experimental stress or control manipulation, the periodic collection of saliva (cortisol) measurements, and the behavioral task.

longer when the environment is scarcer. This optimal strategy, which relies on an index of the overall quality of the environment, has sufficiently described foraging behavior across diverse species (Cowie, 1977; Pyke, 1980; Stephens and Krebs, 1986; Cassini et al., 1993; Hayden et al., 2011) and, in humans, across behavioral domains (Hills et al., 2008, 2012; Wilke et al., 2009).

However, despite a steadily growing interest in this class of decisions within neurophysiology and cognitive neuroscience (Hayden et al., 2011; Kolling et al., 2012, 2016; Wikenheiser et al., 2013), we still know very little about the biological mechanisms by which these highly conserved, context-dependent foraging decisions are made. Specifically, we do not know how the brain dynamically tracks this critical decision variable of overall environment quality. There are, however, other biologically conserved mechanisms in place, such as stress, that could possibly serve such a role in tracking environmental information. Stress recruits biological systems that monitor and respond to changes in the environment (specifically those that pose a threat to homeostasis), orchestrating a cascade of hormonal, neurophysiological, and behavioral adjustments to do so (Goldstein and McEwen, 2002; Arnsten, 2009; Joëls and Baram, 2009). In particular, the longer timescale stress response, as measured by hypothalamic-pituitary-adrenal (HPA) axis activity, is a highly plausible candidate signal for indexing the overall quality of the environment, with physiological properties able to provide contextual information exactly of the sort (and with the appropriate timescale) needed for making optimal foraging decisions. To the

extent that stress signals aversive information about one's environment, it should accordingly be associated with a lower perceived average reward rate and consequently greater exploitation of current options (as if foraging in a scarcer environment).

To explore this hypothesis in humans, in the current study, we used a virtual patch-foraging task to determine how acute physiological and chronic subjective forms of stress might modulate the threshold at which subjects decide to give up on current diminishing options in order to forage for better options—this decision being fundamentally related to how we appraise the quality of our decision-making environments and estimate the opportunity cost of time within these environments. Specifically, we hypothesized that subjects under stress would demonstrate greater exploitation of, or lower reward thresholds for leaving, current options.

## Materials and Methods

**Procedure.** To investigate the relationship between stress and foraging behavior (Fig. 1b), we measured and operationalized stress in two independent ways: according to changes in cortisol in response to an acute stress manipulation and according to subjective responses to a questionnaire about chronic perceived life stress. We then used these estimated levels of acute physiological and chronic subjective stress to predict individual differences in foraging behavior as characterized by performance on a computerized patch-foraging task.

**Subjects.** Subjects were recruited from the general New York University (NYU) community via posted flyers and online Cloud-based subject pool software (Sona Systems). Eligibility criteria were the following: sub-

jects must be at least 18 years of age and speak English as their first language and subjects must not be taking any psychoactive medications, have any major medical disorders, or have any irregularities in heart rate or blood pressure. The NYU Committee on Activities Involving Human Subjects approved all recruiting instruments and experimental protocols. Subjects were paid \$15 per hour and earned up to an additional \$10 based on performance on the task. After removing outliers, the control group and stress group were composed of 29 and 36 subjects, respectively. The proportion of females in the stress group was 0.64 and in the control group 0.60; there was no relationship between sex and group identity ( $t_{(64)} = -0.15, p = 0.88$ ). There was furthermore no relationship between age and group identity ( $t_{(64)} = 0.01, p = 1$ ). The mean age across all subjects was 22.7 years ( $SD = 4.1$ ).

**Foraging task.** Subjects participated in a sequential patch-foraging task (highly similar to that used in prior work (Constantino and Daw, 2015), in which they spent a fixed amount of time in a series of 4 orchards (7 min each) of 2 environment types (counterbalanced in ABAB/BABA block design; see Figure 1a for a graphical depiction of the task sequence). Subjects were told that their goal should be to harvest as many apples as possible during this fixed time because total apples earned would be converted into money at the end of the experiment at an exchange rate of approximately one cent per apple. On each trial, subjects decided whether to harvest the current tree or move to a different tree. The only systematic difference between the environment types was the distance between trees—that is, the time that it took to move virtually between trees. The travel time (i.e., the time elapsed between when the apples from the previous harvest appeared to when the apples from the first harvest at a new tree appeared) was 6 s in short travel time orchards and 12 s in long travel time orchards. The harvest time (i.e., the time elapsed between successive harvests of a single tree) was 3 s in all trees and orchards. The difference in the travel time affected the overall reward rate of the environment and, equivalently, the opportunity cost of time spent exploiting the current tree because in longer travel time environments more time was necessarily spent traveling between trees rather than earning apples.

On each trial, the subject made the decision via a key press whether to harvest the currently displayed tree (down arrow) or move to a different tree (right arrow). If the subject decided to harvest, s/he would then be shown the apples successfully harvested after a short harvest time delay. If s/he instead decided to move to a different tree, s/he would wait a period of time (the duration of which depended on the type of environment) before arriving to a new, unharvested tree. The subject faced the same stay-or-switch decision at this new tree. If subjects did not make a response in the allotted 1 s response window, a warning message was displayed and a timeout was incurred before the subject was afforded another response opportunity. The average number of warnings per subject was 2.0 ( $SD = 2.8$ ), or 0.04% of a subject's average number of decisions, indicating that subjects overall remained engaged in the task and responded within the allotted time. Other than through these warnings, subjects could not affect the speed of events by responding more quickly or slowly because the reaction time was subtracted from the programmed harvest or travel delay. See Figure 1a for schematic of the task.

Each tree's initial supply of apples was randomly drawn from a Gaussian distribution with a mean of 10 and  $SD$  of 1. The depletion rate for each successive harvest of a tree was randomly drawn from a Beta distribution with parameters  $\alpha = 14.9$  and  $\beta = 2.0$ . These parameters were set such that the mean rate at which harvesting depleted a tree's supply of apples was 0.88 with an  $SD$  of 0.07 (i.e., harvesting a tree repeatedly yielded, on average, 88% of the apples received on the previous harvest). Subjects were informed that trees would vary in terms of their quality and depletion rate (i.e., some trees would be richer or poorer than others and some trees would deplete slower or faster than others), but that the trees varied in the same way across all orchards. Subjects were instructed that the only factor that might change across orchards would be the time it took to travel between trees. While performing the task, subjects were informed when the environment changed by a message, a brief break, and a change in background color indicating a new block (where colors were counterbalanced across blocks and environment types). Before the

stress/control manipulation, subjects received task instructions and performed a short practice version of the task in which they briefly visited each type of environment (although with different travel delays from those in the experiment) to familiarize themselves with task parameters and response windows. This task preexposure minimized the possibility of stress interfering with task comprehension or strategy selection.

**Stress induction.** To induce an endogenous acute stress response, we used the well validated cold pressor task, in which subjects assigned to the stress group were asked to submerge their nondominant arm in ice-cold water (0–4°C) for 3 min (Lovallo, 1975). This method has been shown to induce subjective, physiological, and neuroendocrine responses that are characteristic of mild-to-moderate levels of stress (Velasco et al., 1997). In a matched control task, subjects submerged their nondominant arm in warm (35–40°C) water for 3 min. All experiments were conducted between the hours of 12:00 and 6:00 P.M. to control for diurnal variability in cortisol levels.

**Saliva samples.** To index subjects' physiological stress responses, we collected saliva samples throughout different time points to assay for changes in cortisol. We selected cortisol as our primary index of acute stress, a hormone which is synthesized and released via the HPA axis and which crosses the blood–brain barrier, binding to receptors throughout the brain and thus producing behaviorally relevant neurobiological effects (Arnsten, 2009; Sandi and Haller, 2015); however, we acknowledge the limitations of using only a single physiological metric to characterize the complete, multicomponent stress response. Salivary samples were collected via absorbent oral swabs throughout the experiment to assess stress hormones at different times before and after the stress/control manipulation, following standard protocol (Lighthall et al., 2013; Raio et al., 2013). Specifically, a baseline (time = 1) sample (after acclimation and before the stress/control manipulation) was collected between 10 and 15 min after subjects' arrival. Ten minutes after subjects either performed the cold pressor or matched control task (time = 2), another saliva sample was collected; this 10 min delay between the stressor and poststressor sample (immediately preceding the task) was enacted to account for the temporal delay of the cortisol response (mediated by the slower timescale HPA axis) and so that the peak of the response would occur during the task. Immediately after the second sample, subjects began the foraging task, completing the first 2 blocks before they provided another saliva sample midtask (time = 3), ~25 min after the stress/control manipulation. After completion of the second half of the task (time = 4), subjects provided a final saliva sample ~40 min after the stress/control manipulation. Salivary concentrations of cortisol, which are highly correlated with blood serum concentrations, provide a reliable index of levels of unbound cortisol that are available to bind to glucocorticoid receptors in neural tissue (Kirschbaum et al., 1994). Saliva samples were stored in sterile tubes and upon collection immediately refrigerated at –20°C. Samples were analyzed using Salimetrics Testing Services. See Figure 1b for the experiment protocol.

**Stress measures.** Once subjects acclimated to the laboratory environment, they completed the Perceived Stress Scale (PSS). The PSS is a self-report, 10-item psychological questionnaire that assesses global (i.e., domain-general) and chronic (i.e., across the past month) levels of subjectively experienced stress (Cohen et al., 1983). Subjects completed this questionnaire before learning their group assignment or undergoing any group-specific procedures; therefore, chronic stress (PSS) scores were by design independent of the acute stress manipulation. Statistically, using linear regression, we found that PSS scores were unrelated to both basal cortisol levels ( $t_{(64)} = -0.66, p = 0.51$ ) and cortisol responses ( $t_{(64)} = -0.91, p = 0.37$ ).

We quantified acute stress responses using salivary cortisol measurements. Such measurements have repeatedly been found to demonstrate positive skewness, perhaps due to inherent nonlinearity in cortisol secretory dynamics and measurement error (Miller et al., 2013). To better approximate a Gaussian distribution of cortisol values across subjects, we log-transformed the cortisol data. To compute a per-subject physiological stress index that represented a subject's average cortisol response throughout the duration of task performance, we calculated the baseline-corrected mean of log-transformed cortisol values across time points 2, 3, and 4. This procedure yielded a log cortisol  $\Delta$  value for each subject,

representing the average change in cortisol from baseline over the task. Below is the formula that we used to derive this index, where  $\log \text{cort}$  is the (natural) log-transformed cortisol concentration:

$$\log \text{cort } \Delta = \frac{\log \text{cort}_{t=2} + \log \text{cort}_{t=3} + \log \text{cort}_{t=4}}{3} - \log \text{cort}_{t=1}$$

This continuous measure, which varied across subjects and reflected individually and physiologically defined acute stress responses, was used in all primary analyses. This approach, which has been used successfully in other stress and decision-making studies (Lighthall et al., 2013; Otto et al., 2013), allowed us to leverage the between-subject variance in physiological stress responses. Notwithstanding the large individual variation, we also observed a groupwise difference in log cortisol  $\Delta$  values ( $t_{(64)} = -4.58, p < 0.00001$ ), suggesting the overall efficacy of our stress manipulation. Reported in raw units (micrograms per deciliter), the means and SDs of cortisol concentrations for the control group at each time point are as follows: time point 1,  $M = 0.30, SD = 0.15$ ; time point 2,  $M = 0.27, SD = 0.15$ ; time point 3,  $M = 0.24, SD = 0.10$ ; and time point 4,  $M = 0.22, SD = 0.09$ . For the stress group, these values are as follows: time point 1,  $M = 0.24, SD = 0.13$ ; time point 2,  $M = 0.34, SD = 0.22$ ; time point 3,  $M = 0.47, SD = 0.37$ ; and time point 4,  $M = 0.33, SD = 0.24$ .

**Regression analysis.** In a mixed-effects linear regression analysis (using the `lm4` package in R programming language, with  $p$ -values estimated using Satterthwaite approximation implemented in the `lmerTest` package), we regressed estimated tree-leaving thresholds on the following explanatory variables: environment type (i.e., long or short travel delay), log cortisol  $\Delta$  (representing subject's baseline-corrected change in cortisol in response to the experimental manipulation), subjective chronic stress (as measured by the PSS), group status, and an intercept term. The dependent variable was comprised of each subject's tree-level exit thresholds, which were estimated as the average of the last two rewards observed before a decision to leave a tree. To account for within-subject repeated observations, we included subject-specific random effects for intercept and environment type. Specifically, we took the intercept and environment type terms as random effects that were allowed to vary from subject to subject. Finally, we included the categorical group assignment as a fixed-effect regressor to account for any marginal effects of group that were not captured by log cortisol  $\Delta$  (our main *a priori* variable of interest with regard to the acute stress manipulation). Although group and log cortisol  $\Delta$  are positively correlated ( $r = 0.47, p < 0.0001$ ), this correlation is arguably not high enough to preclude their simultaneous inclusion. That is, because the two regressors share only 22% of their variance, each variable (after adjusting for their shared variance) has sufficient variance that is unaccounted for, thus allowing for the independent estimation of both effects. Imperfectly correlated predictors, furthermore, result in less precision (i.e., inflated SEs) without biasing coefficient estimates.

For the analysis examining subjects' deviations from optimal behavior, we performed the same mixed-effects linear regression analysis, adjusting only the dependent variable to instead reflect for each subject, environment type, and tree, her deviation from optimal behavior (by subtracting the optimal threshold in each environment type from every subject's tree-level exit threshold in that environment). PSS and log cortisol  $\Delta$  scores were first transformed into  $z$ -scores. Group and block type were both coded as  $-1/1$ .

**Outlier removal.** One participant (in the control group) was removed from analyses due to an abnormally high baseline cortisol concentration that was  $>2.3$  SDs from the mean, falling in the 99<sup>th</sup> percentile of the sample's distribution of cortisol baseline scores. Before data collection, we defined a procedure for determining outliers based upon task performance. We set this criterion to exclude people who adopted a decision strategy so qualitatively different from other subjects that we would not be able to compare them appropriately within the same decision theoretic framework (i.e., they did not perform the task sufficiently well). Specifically, we excluded participants whose average number of harvests per tree was 2.3 SDs either above or below the sample mean (i.e., in the bottom or top 1% of the sample distribution). Qualitatively, these extremely high or low thresholds suggest fixed (i.e., insensitive to reward) decision strategies that are akin, respectively, to leaving trees almost im-

mediately upon arrival (e.g., after on average one harvest per tree) or harvesting trees down to nearly the last apple (e.g., perseverating on trees). This procedure identified three outliers (one in the control group and two in the stress group), who on average harvested trees 1, 15, and 16 times, respectively. These outliers were excluded from all analyses.

**Optimal analysis.** The intuition for the MVT strategy is that, when making these kinds of stay-or-switch foraging decisions, one should weigh the value of the current option against some measure of the average quality of all other options in the environment. The MVT formalizes this intuition by showing that, in this class of decision tasks, the optimal reward rate-maximizing strategy is to leave a diminishing resource whenever the expected reward from one more exploitation falls below the long-run average reward obtained in the environment. This strategy would thus prescribe lower thresholds for leaving in scarcer environments such as those with longer interpatch travel times (where the average reward is lower because a larger proportion of the agent's allotted time must necessarily be spent traveling between trees rather than collecting apples). See Constantino and Daw (2015) for a full proof of the MVT as it applies to this specific (stochastic, discretized) patch-foraging task. According to the MVT, the optimal decision rule (i.e., the decision rule that maximizes the long-run average reward rate) is to leave when the expected number of apples from one more harvest is smaller than the number of apples that you would expect on average in a given environment. More formally, this inequality can be stated as follows:

$$\mathbb{E}_{S_{i+1}}[r_{i+1}] < \rho h$$


where  $\mathbb{E}_{S_{i+1}}[r_{i+1}]$  is the expected reward of harvesting at state  $i + 1$  and  $\rho h$  is the opportunity cost of harvesting, given by the average reward in the environment,  $\rho$ , multiplied by the harvest time,  $h$ . In this task, the expected reward of harvesting at the next state,  $\mathbb{E}_{S_{i+1}}[r_{i+1}]$ , is more specifically defined as  $\kappa s_i$ , the reward at the current state,  $s_i$ , multiplied by the mean depletion rate,  $\kappa$ . The optimal threshold for leaving is thus given in terms of  $\rho$ , the average reward of the environment, or, in other words, the opportunity cost of continuing to harvest the current tree.

Because the MVT guarantees that the optimum is given by some threshold,  $\rho h$ , we used numerical optimization to find the optimal threshold for each environment in our task, given the task-general and environment-specific parameters (see "Foraging task"). Specifically, we searched over all possible exit thresholds and, for each threshold, computed the (probabilistically) expected reward over time under that exit policy in order to numerically find the threshold that maximizes this quantity. This method returns optimal exit thresholds for each environment type in units of reward. For the short travel time environment, this threshold is 6.52; for the long travel time environment, it is 5.31.

## Results

### MVT foraging

To test the extent to which patch-foraging behavior was sensitive to environmental quality in a manner consistent with an MVT policy, we determined the reward thresholds at which subjects decided to leave current diminishing trees in pursuit of new trees in the environment. We estimated subjects' tree-leaving thresholds by averaging the last two rewards observed before an exit decision, which were taken to reflect the lower and upper bounds on subjects' true internal continuous-valued reward thresholds. That is, we assumed that a subject's true threshold for leaving a tree is some factor smaller than the second-to-last reward and larger than the last reward and is therefore bounded by these two rewards. For consistency, we excluded trees that were harvested only once before an exit decision because these observations provide only a lower bound on subjects' thresholds. Across all subjects, we eliminated 190 such decisions (constituting only 0.05% of the data). This procedure enabled us to construct, for each subject, tree-by-tree exit thresholds for each environment type. On average, this yielded 57 (SD = 11) observations per subject.



**Figure 2.** Mean exit thresholds per environment type. Gray bars indicate optimal thresholds for each environment type (long or short) as given by the MVT. Open circles indicate per-subject mean exit thresholds for each environment type. Diamonds indicate mean exit thresholds across all subjects for each environment type, with 95% confidence intervals. Overall, subjects demonstrated sensitivity to changes in the environment quality by adjusting their thresholds adaptively in a manner consistent with the MVT.

Overall, across both groups, exit thresholds were near the optimal values predicted by foraging theory (Fig. 2). That is, qualitatively, most individuals adaptively adjusted their exit thresholds to the blockwise changes in environmental average reward rate in the direction predicted by optimal foraging theory. Specifically, subjects adopted lower thresholds in lower quality (longer travel time) environments and higher thresholds in richer (shorter travel time) environments and this adjustment was significant at the group level ( $t_{(65)} = 9.20, p < 0.00001$ ), upholding the cardinal qualitative prediction of the MVT.

**MVT foraging and stress**

We used a mixed-effects linear regression analysis to study the relationship between our measures of stress and exit thresholds in the foraging task. Although group (stress/control) status alone was not a significant predictor of behavior (see also Otto et al., 2013), across all subjects, both acute physiological stress (as measured by log cortisol  $\Delta$ , expressed here in z-scores) and perceived chronic stress (as measured by PSS scores, expressed here as z-scores) predicted lower exit thresholds (respectively,  $\beta = -0.37, t_{(61)} = -2.06, p < 0.05$ ;  $\beta = -0.49, t_{(61)} = -2.91, p < 0.01$ ; Table 1). That is, subjects who demonstrated higher indices of stress were more likely to harvest trees for longer despite diminishing returns. Because exit thresholds effectively indicate the valuation of the broader environment from which new trees are drawn, this result is consistent with the hypothesis that stress leads to lower estimates of environmental quality and, consequently, a reluctance to seek new trees.

To quantitatively test the extent to which subjects’ foraging behavior was consistent with the optimal strategy given by the

**Table 1. Linear mixed-effects regression on mean exit thresholds across all subjects**

	$\beta$ estimate (SE)	t-value
Coefficient (intercept)	4.89 (0.26)	18.76***
Environment type	-0.6112 (0.07)	-9.11***
Group	0.7232 (0.37)	1.94
PSS	-0.4809 (0.17)	-2.90**
Log cortisol $\Delta$	-0.3675 (0.18)	-2.04*

PSS and log cortisol  $\Delta$  scores were transformed to z-scores.

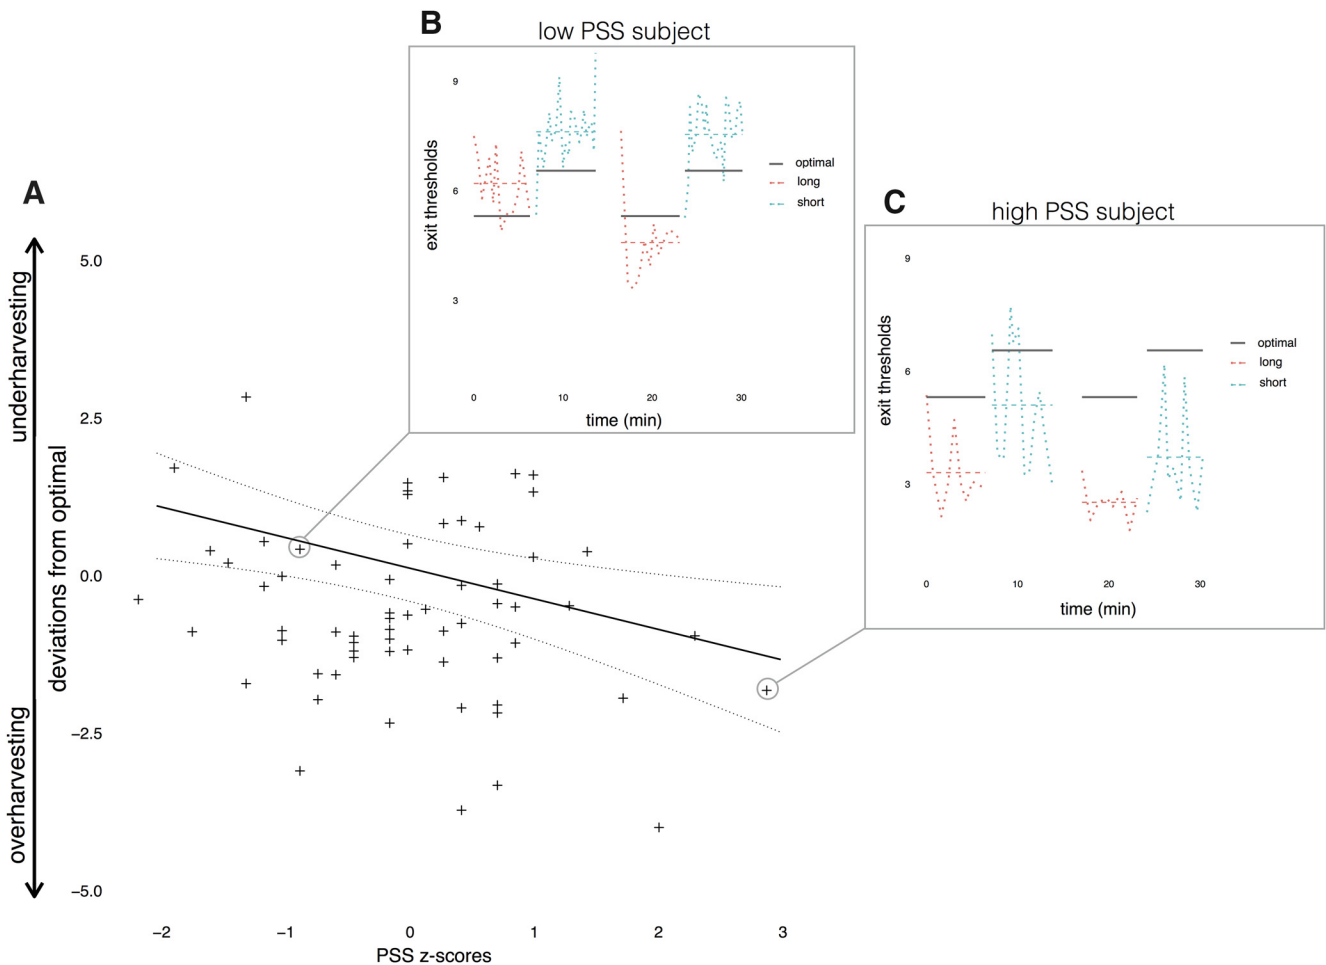
\*\*\* $p < 0.001$ ; \*\* $p < 0.01$ ; \* $p < 0.05$ .

MVT, we compared subjects’ thresholds to optimal thresholds (the thresholds that maximize the average reward rate for the two environments; Charnov, 1976). Subtracting these optimal, rate-maximizing thresholds from subjects’ empirical thresholds yielded tree-by-tree and environment-by-environment deviation-from-optimal scores. Negative values of these deviation scores indicate overharvesting relative to an optimal policy (i.e., exercising a lower threshold or staying longer at a tree than optimal), whereas positive values indicate underharvesting. We reestimated our regression model with these deviation scores as the dependent measure. Because this merely shifts the dependent measure by a constant, the results with regard to overexploitation and stress were unchanged (Figs. 3, 4, Table 2). However, we found that the effect of environment type was no longer significant ( $p = 0.81$ ); that is, representing subjects’ exit thresholds relative to optimal nearly completely accounted for the previously highly significant effect of environment type on exit thresholds (previously,  $p < 0.00001$ ), consistent with the observation that changes in behavior between environments resulted from optimal adjustments. Furthermore, this suggests that there were no differences in the extent of overharvesting relative to optimal between environment types.

Finally, the significantly negative intercept in this regression reflects a mean bias away from optimal: a tendency to overharvest for a subject with average PSS score and average log cortisol  $\Delta$ . To determine whether this bias persisted after accounting for the effects attributed to stress (i.e., whether this tendency would be predicted for a subject who demonstrated no acute or chronic stress whatsoever), we again reestimated the regression, expressing PSS and log cortisol  $\Delta$  instead in their natural units (so that zero on each scale indicates, respectively, a complete absence of reported perceived life stress and no change in cortisol from baseline). When completely accounting for these effects of stress (i.e., holding these stress levels constant at their natural zero points), the intercept was no longer significant ( $p = 0.8$ ; Table 3). This suggests that, after accounting for the effects of stress, there was no detectable population-level bias away from optimal behavior and, conversely, that the overharvesting that we detected in our population was explainable in terms of chronic and acute stress. Although we exercise caution in interpreting these null results, the lack of significant differences between optimal and observed behavior offers additional support that our travel time manipulation was successful in modulating subjects’ foraging behavior in a way that is not detectably different from the optimal strategy and that stress was associated with an overly exploitative shift away from optimal behavior.

**Discussion**

We found that both acute physiological and chronic psychological stress bias human decision makers toward overexploiting current, known resources despite diminishing returns. Such overexploitation may be an adaptive response to situations that




**Figure 3.** Relationship between PSS scores and foraging behavior. **A**, Deviation from optimal scores conditional on the mixed-effects linear regression plotted as a function of PSS scores across all subjects. Note that positive values of this variable indicate underharvesting (i.e., desisting sooner than optimal) and negative values indicate overharvesting (i.e., desisting later than optimal). The regression line is computed from the group-level intercept and PSS fixed effects. Dashed lines indicate 2 SEs estimated from the group-level mixed-effects linear regression. **B**, Tree-by-tree data for a subject with a relatively low PSS score (based upon a median split). Solid lines indicate exit thresholds per tree as a function of time ( $y$ -axis), block (discontinuities between lines), and environment type (color). Dashed lines indicate subjects' average exit thresholds per block. Gray bars indicate optimal thresholds per environment type. **C**, Tree-by-tree threshold data for a subject with a relatively high PSS score. Subjects in **B** and **C** had comparable log cortisol  $\Delta$  values. Note that, although both subjects are sensitive to environmental changes in the direction predicted by the MVT, the subject in **C** was biased toward suboptimally overexploiting trees.

pose a threat to homeostasis (to the extent that stress confers veridical information about one's environment). In particular, in optimal foraging theory, increased exploitative behavior is appropriate when there is a lower appraisal of overall environment quality or, equivalently, a decreased opportunity cost of time. Such general evaluations of environmental quality are of particular importance in serial decision tasks such as foraging, in which one must infer the expected value of seeking new options that have not yet been experienced. Our results thus suggest that stress biases subjects' appraisal of environmental quality, the decision variable that most directly governs patch-leaving decisions in computational model fits to this task (Constantino and Daw, 2015).

This possible computational role of stress in evaluating environmental quality accords with its more general role in mediating a coordinated bodily response to a challenging context. Indeed, it seems intuitive that stress would bias us toward appraising our environments more negatively: stressful environments are defined, almost necessarily, as aversive and threatening. However, when the circumstances of the stressor and the circumstances of the decision making problem are disjoint, as in our laboratory

task and as is often also the case in everyday life, these overly exploitative and pessimistic biases could become maladaptive or even (as in related computational views of learned helplessness and depression) pathological (Huys et al., 2015). That is, to the extent the stressor is independent of, or incidental to, the decision-making context (as opposed to being meaningfully embedded in, and integral to, the context), this bias proves no longer adaptive. Future research should investigate how changing the instrumental relationship between agents' affective and external contexts might mediate the role of emotion in decision making. A better understanding of the development of such potentially maladaptive biases might bear important implications for disorders that are characterized by inappropriately exploitative behavior in which stress often plays a critical role.

In addition to the hypothesized effect on the opportunity cost of time, there are other (interrelated) computational decision variables through which stress could bias foraging behavior, among which the current study cannot clearly adjudicate, but which are not mutually exclusive and could even act synergistically. For instance, our stress findings could be explained through stress modulating parameters of the objective function



**Figure 4.** Relationship between log cortisol  $\Delta$  and foraging behavior. Deviation from optimal scores conditional on the mixed-effects linear regression is plotted as a function of log cortisol  $\Delta$  across all subjects. The regression line is computed from the group-level intercept and log cortisol  $\Delta$  fixed effects. Dashed lines indicate 2 SEs estimated from the group-level mixed-effects linear regression.

**Table 2. Linear mixed-effects regression on deviation from optimal scores across all subjects**

	$\beta$ estimate (SE)	t-value
Coefficient (intercept)	-1.02 (0.26)	-3.92***
Environment type	-0.01 (0.07)	-0.08
Group	0.72 (0.37)	1.94
PSS	-0.48 (0.17)	-2.90**
Log cortisol $\Delta$	-0.37 (0.18)	-2.04*

PSS and log cortisol  $\Delta$  scores were transformed to z-scores.

\*\*\* $p < 0.001$ ; \*\* $p < 0.01$ ; \* $p < 0.05$ .

**Table 3. Linear mixed-effects regression on deviation from optimal scores across all subjects**

	$\beta$ estimate (SE)	t-value
Coefficient (intercept)	0.09 (0.43)	0.21
Environment type	-0.01 (0.07)	-0.08
Group	0.72 (0.37)	1.94
PSS	-0.07 (0.03)	-2.90**
Log cortisol $\Delta$	-0.74 (0.36)	-2.04*

PSS and log cortisol  $\Delta$  scores are in natural (untransformed) units.

\*\*\* $p < 0.001$ ; \*\* $p < 0.01$ ; \* $p < 0.05$ .

itself (e.g., preferences about time or risk), indirectly leading to the phenomenon of overexploitation. In particular, greater discounting of future rewards could potentially explain increased exploitation of current options. The optimal analysis assumes, however, that subjects maximize the undiscounted rate of reward—that is, that rewards received later are not inherently less valuable except insofar as delays decelerate the rate of receiving

rewards. This undiscounted reward has traditionally been argued to be the relevant currency in foraging theory (Charnov, 1976) and it corresponds to maximizing monetary income in the current task where subjects received their total compensation in a single payment immediately after the task regardless of their choices. Furthermore, delay amounts typically used in standard intertemporal choice tasks are expressed in units of days as opposed to seconds, resulting in the estimation of discount parameters that operate on a distinct, much longer timescale than could explain the current findings. In other words, time preferences observed in humans are often with regard to much longer intervals of time than seconds (e.g., days, weeks) and would thus not predict preferences over a span of 6–12 s as would be necessary to explain the current results.

It is also possible that the overexploitation that we observed could be partially mediated by a change in risk preferences. For concreteness, consider diminishing marginal utility (i.e., curvature in the function relating magnitudes of money to their subjective utility, which, in standard economic accounts, produces risk-averse choices). Such magnitude distortion would render new, unharvested trees less subjectively valuable, thereby biasing subjects to harvest at individual trees for longer. The

relationship between diminishing marginal utility (or risk preferences) and opportunity cost, however, is likely bidirectional. Because we assume that subjects learn the environment’s average reward rate through experience with reward over time, risk aversion would lead to learning a lower subjective opportunity cost of time. Furthermore, to the extent that the average reward rate serves as a reference point against which (risky) prospects are compared, within a prospect theory framework, a lower perceived average reward could also promote risk aversion (Cools et al., 2011). In trying to disambiguate this possibly insidious role of risk, the burgeoning literature on stress and risk offers little interpretative guidance because the reported effects of stress on risk preferences are extremely directionally mixed (Starcke et al., 2008; Porcelli and Delgado, 2009; Lighthall et al., 2012; Mather and Lighthall, 2012; Kandasamy et al., 2014; van den Bos et al., 2013; Sokol-Hessner et al., 2016). Indeed, these inconsistencies might be expected if effects on risk sensitivity are secondary to effects on reference points, leading to greater contextual dependence for results on risk preferences per se.

Finally, although not necessarily inconsistent with the accounts above, stress could bias perceptions of environment quality through cognitive appraisal mechanisms. Negative affective states such as sadness or fear have been shown to have carry-over effects on appraisal, leading to more negative judgments in emotionally irrelevant domains (Lerner and Keltner, 2000). To the extent that stress induces such appraisal tendencies, this negative evaluative bias, or pessimism, could account for the lower perceived quality of the foraging environments in our task. In fact, it

is possible that psychological-level accounts of mood and appraisal can actually be more precisely understood in terms of computational-level mechanisms that entail assessing the average reward within environments (Huys et al., 2015; Eldar et al., 2016). Further research at behavioral, computational, and neural levels is needed to more specifically identify the mechanism(s) underlying the effects of stress on foraging behavior and to fully characterize an information processing account of stress effects on decision making.

Based on extensive empirical and theoretical work in behavioral ecology and more recently psychology and neuroscience, however, we argue that the perceived opportunity cost of time plays not only a fundamental role in governing patch-leaving behavior, but also in a range of other decision domains, including but not limited to those discussed above. Indeed, shifts in the opportunity cost of time can theoretically account for changes in risk and intertemporal preferences, habitual versus deliberative control, vigor, and exploration, among other decision effects (Cools et al., 2011; Boureau et al., 2015). For instance, a change in perceived opportunity costs could bias the speed–accuracy tradeoffs that are thought to control whether the brain invokes habits versus model-based deliberation (Keramati et al., 2011). The logic by which scarcer reward environments (and stress as a signal of them) drives sticking with a diminishing tree in our task would be expected to extend more generally to increased stickiness (or, at the extreme, perseveration), both motor and cognitive (Niv et al., 2007). This potentially diverse range of decision effects highlights not only the inherently confounding relationship among such interrelated decision variables and the difficulty in experimentally disassociating them but also the potentially unifying role of opportunity cost in accounting for them within a single theoretical framework. Accordingly, the current study may in fact help provide a route toward a more fundamental, unifying account of the effects of stress on other decision tasks through biasing evaluation of one's environmental context.

More broadly, the current study's findings contribute to the mounting evidence that both acute and prolonged states of negative affect can lead to cognitive and behavioral biases that can be suboptimal or maladaptive in inappropriate contexts (Haushofer and Fehr, 2014). For instance, our study suggests that it could be the case that appraising the quality of one's environment (e.g., a job market) more negatively might bias a real-world decision maker under stress toward exploiting suboptimal options (e.g., staying in a job with diminishing returns on time and energy investments) rather than seeking better options. Suboptimal routines of behavior might become increasingly difficult to override because they reinforce negative beliefs about the environment, highlighting the intimately reciprocal relationship between the actions that we select upon our environment (based upon our potentially biased information) and the information that we acquire about our environments (based upon our actions). Such findings thus have important real-world implications because we frequently confront these kinds of persist-or-desist foraging problems in the everyday world, often in a state of psychological and physiological stress. Furthermore, individuals at a higher risk of experiencing daily stress (such as those who are economically impoverished, trait anxious, or discriminated against) and individuals who experience pathological states of stress (such as those with anxiety, stress, and mood disorders) could be especially vulnerable to these effects. To the extent that the current study's task probes a domain-general decision-making mechanism, our findings might offer a new lens with which to understand stress-related cog-

native and behavioral deficits and to develop behavioral interventions to mitigate them.

## References

- Arnsten AF (2009) Stress signalling pathways that impair prefrontal cortex structure and function. *Nat Rev Neurosci* 10:410–422. [CrossRef Medline](#)
- Boureau YL, Sokol-Hessner P, Daw ND (2015) Deciding how to decide: self-control and meta-decision making. *Trends Cogn Sci* 19:700–710. [CrossRef Medline](#)
- Cassini MH, Lichtenstein G, Ongay JP, Kacelnik A (1993) Foraging behaviour in guinea pigs: further tests of the marginal value theorem. *Behav Processes* 29:99–112. [CrossRef Medline](#)
- Charnov EL (1976) Optimal foraging, the marginal value theorem. *Theor Popul Biol* 9:129–136. [CrossRef Medline](#)
- Cohen S, Kamarck T, Mermelstein R (1983) A global measure of perceived stress. *J Health Soc Behav* 24:385–396. [CrossRef Medline](#)
- Constantino SM, Daw ND (2015) Learning the opportunity cost of time in a patch-foraging task. *Cogn Affect Behav Neurosci* 15:837–853. [CrossRef Medline](#)
- Cools R, Nakamura K, Daw ND (2011) Serotonin and dopamine: unifying affective, motivational, and decision functions. *Neuropsychopharmacology* 36:98–113. [CrossRef Medline](#)
- Cowie RJ (1977) Optimal foraging in great tits (*Parus major*). *Nature* 268:137–139. [CrossRef](#)
- Eldar E, Rutledge RB, Dolan RJ, Niv Y (2016) Mood as representation of momentum. *Trends Cogn Sci* 20:15–24. [CrossRef Medline](#)
- Goldstein DS, McEwen B (2002) Allostasis, homeostats, and the nature of stress. *Stress* 5:55–58. [CrossRef Medline](#)
- Haushofer J, Fehr E (2014) On the psychology of poverty. *Science* 344:862–867. [CrossRef Medline](#)
- Hayden BY, Pearson JM, Platt ML (2011) Neuronal basis of sequential foraging decisions in a patchy environment. *Nat Neurosci* 14:933–939. [CrossRef Medline](#)
- Hills TT, Todd PM, Goldstone RL (2008) Search in external and internal spaces - Evidence for generalized cognitive search processes. *Psychol Sci* 19:802–808. [CrossRef Medline](#)
- Hills TT, Jones MN, Todd PM (2012) Optimal foraging in semantic memory. *Psychol Rev* 119:431–440. [CrossRef Medline](#)
- Huys QJ, Daw ND, Dayan P (2015) Depression: a decision-theoretic analysis. *Annu Rev Neurosci* 38:1–23. [CrossRef Medline](#)
- Joëls M, Baram TZ (2009) The neuro-symphony of stress. *Nat Rev Neurosci* 10:459–466. [CrossRef Medline](#)
- Kandasamy N, Hardy B, Page L, Schaffner M, Graggaber J, Powlson AS, Fletcher PC, Gurnell M, Coates J (2014) Cortisol shifts financial risk preferences. *Proc Natl Acad Sci U S A* 111:3608–3613. [CrossRef Medline](#)
- Keramati M, Dezfouli A, Piray P (2011) Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLoS Comput Biol* 7:e1002055. [CrossRef Medline](#)
- Kirschbaum C, Hellhammer DH (1994) Salivary cortisol in psychoneuroendocrine research: recent developments and applications. *Psychoneuroendocrinology* 19:313–333. [CrossRef Medline](#)
- Kolling N, Behrens TE, Mars RB, Rushworth MF (2012) Neural mechanisms of foraging. *Science* 336:95–98. [CrossRef Medline](#)
- Kolling N, Wittmann MK, Behrens TE, Boorman ED, Mars RB, Rushworth MF (2016) Value, search, persistence and model updating in anterior cingulate cortex. *Nat Neurosci* 19:1280–1285. [CrossRef Medline](#)
- Lerner JS, Keltner D (2000) Beyond valence: toward a model of emotion-specific influences on judgement and choice. *Cognition & Emotion* 14:473–493. [CrossRef](#)
- Lighthall NR, Sakaki M, Vasunilashorn S, Nga L, Somayajula S, Chen EY, Samii N, Mather M (2012) Gender differences in reward-related decision processing under stress. *Soc Cogn Affect Neurosci* 7:476–484. [CrossRef Medline](#)
- Lighthall NR, Gorlick MA, Schoeke A, Frank MJ, Mather M (2013) Stress modulates reinforcement learning in younger and older adults. *Psychol Aging* 28:35–46. [CrossRef Medline](#)
- Lovall W (1975) The cold pressor test and autonomic function: a review and integration. *Psychophysiology* 12:268–282. [CrossRef Medline](#)
- Mather M, Lighthall NR (2012) Both risk and reward are processed differently in decisions made under stress. *Curr Dir Psychol Sci* 21:36–41. [CrossRef Medline](#)
- Miller R, Plessow F, Rauh M, Gröschl M, Kirschbaum C (2013) Comparison



- of salivary cortisol as measured by different immunoassays and tandem mass spectrometry. *Psychoneuroendocrinology* 38:50–57. [CrossRef Medline](#)
- Niv Y, Daw ND, Joel D, Dayan P (2007) Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology* 191:507–520. [CrossRef Medline](#)
- Otto AR, Raio CM, Chiang A, Phelps EA, Daw ND (2013) Working-memory capacity protects model-based learning from stress. *Proc Natl Acad Sci U S A* 110:20941–20946. [CrossRef Medline](#)
- Porcelli AJ, Delgado MR (2009) Acute stress modulates risk taking in financial decision making. *Psychol Sci* 20:278–283. [CrossRef Medline](#)
- Pyke GH (1980) Optimal foraging in bumblebees: calculation of net rate of energy intake and optimal patch choice. *Theor Popul Biol* 17:232–246. [CrossRef Medline](#)
- Raio CM, O'Rederu TA, Palazzolo L, Shurick AA, Phelps EA (2013) Cognitive emotion regulation fails the stress test. *Proc Natl Acad Sci U S A* 110:15139–15144. [CrossRef Medline](#)
- Sandi C, Haller J (2015) Stress and the social brain: behavioural effects and neurobiological mechanisms. *Nat Rev Neurosci* 16:290–304. [CrossRef Medline](#)
- Sokol-Hessner P, Raio CM, Gottesman SP, Lackovic SF, Phelps EA (2016) Acute stress does not affect risky monetary decision making. *Neurobiol Stress* 5:19–25. [CrossRef Medline](#)
- Starcke K, Wolf OT, Markowitsch HJ, Brand M (2008) Anticipatory stress influences decision making under explicit risk conditions. *Behav Neurosci* 122:1352–1360. [CrossRef Medline](#)
- Stephens DW, Krebs JR (1986) *Foraging theory*. Princeton, NJ: Princeton University.
- van den Bos R, Taris R, Scheppink B, de Haan L, Verster JC (2013) Salivary cortisol and alpha-amylase levels during an assessment procedure correlate differently with risk-taking measures in male and female police recruits. *Front Behav Neurosci* 7:219. [CrossRef Medline](#)
- Velasco M, Gómez J, Blanco M, Rodríguez I (1997) The cold pressor test: pharmacological and therapeutic aspects. *Am J Ther* 4:34–38. [CrossRef Medline](#)
- Wikenheiser AM, Stephens DW, Redish AD (2013) Subjective costs drive overly patient foraging strategies in rats on an intertemporal foraging task. *Proc Natl Acad Sci U S A* 110:8308–8313. [CrossRef Medline](#)
- Wilke A, Hutchinson JM, Todd PM, Czienskowski U (2009) Fishing for the right words: decision rules for human foraging behavior in internal search tasks. *Cogn Sci* 33:497–529. [CrossRef Medline](#)