Receipt of reward leads to altered estimation of effort

Arezoo Pooresmaeili[a,1], Aurel Wannig[b], and Raymond J. Dolan[a,c,d]

[a]Berlin School of Mind and Brain, Humboldt-Universität zu Berlin, 10099 Berlin, Germany; [b]Freie Schule Charlottenburg, 14052 Berlin, Germany; [c]Wellcome Trust Center for Neuroimaging, Institute of Neurology, University College London, London WC1N 3BG, United Kingdom; and [d]Max Planck University College London Centre for Computational Psychiatry and Ageing Research, London WC1B 5EH, United Kingdom

Edited by Paul W. Glimcher, New York University, New York, NY, and accepted by the Editorial Board September 8, 2015 (received for review April 21, 2015)

Effort and reward jointly shape many human decisions. Errors in predicting the required effort needed for a task can lead to suboptimal behavior. Here, we show that effort estimations can be biased when retrospectively reestimated following receipt of a rewarding outcome. These biases depend on the contingency between reward and task difficulty and are stronger for highly contingent rewards. Strikingly, the observed pattern accords with predictions from Bayesian cue integration, indicating humans deploy an adaptive and rational strategy to deal with inconsistencies between the efforts they expend and the ensuing rewards.

effort | reward | retrospective | Bayesian | cue integration

> Aye, and I saw Sisyphus in violent torment, seeking to raise a monstrous stone with both his hands.
>
> Homer, Book XI of *The Odyssey*

The adage "it was well worth the effort" highlights an assumed interdependency between attainment of reward and retrospective effort assignment. Despite its ubiquity we know little about the nature of these retrospective effort estimations. Previous studies have focused on the interaction of effort and reward as costs and benefits when a choice is to be made (1–5). Whether receipt of reward influences a retrospective estimation of effort is not known. Intuitively, we assume to have immediate and unbiased access to internal representations of how much "effort" we expended in an endeavor. Here, we demonstrate that retrospective estimation of effort is strongly affected by the amount of monetary reward attained and, as such, is profoundly biased. This bias adheres to established principles of Bayesian cue integration (6–8) and, on this basis, is not irrational.

In a behavioral experiment, participants pressed two buttons on a keyboard to push a ball up a virtual ramp and rated their experienced physical effort in each trial (Fig. 1; see also *SI Materials and Methods* for additional information regarding the task). The ball rolled back by a constant amount on each frame of the display, hence simulating a gravity force that varied so as to manipulate task difficulty ($n = 6$ difficulty levels, adjusted individually for each participant). Successful trials where participants managed to push the ball all of the way up the ramp were rewarded. Reward was contingent upon task difficulty (with values drawn from six Gaussian distributions with means from 1.5 to 6.5 cents) and the strength of this contingency varied across different blocks of the experiment (SDs of 1.2 or 2.5 cents). Additionally, we included a control experiment in which reward receipt was unrelated to task difficulty (SD = ∞).

The reward information was presented either before or after the rating of effort (in 90% and 10% of trials, respectively). Trials in which reward was shown after the estimation of effort served as a reference, because here subjects are not influenced by preceding reward information. Participants were instructed to pay attention to all information presented in a trial, including a brief color change of the ball (50% of the trials), which they needed to detect on each trial. This manipulation was implemented to distract subjects from the true purpose of experiment, discouraging ad hoc strategies that might link effort and reward (also see *SI Materials and Methods*). Twenty-six participants (15 females, age 20–39 y, mean: 27.07 ± 5.1 y) took part in our main experiment. Two participants were excluded from the analysis because in our debriefing they mentioned they had not paid attention to the reward magnitudes during the experiment. Fourteen participants (nine females, age 21–35 y, mean: 27.5 ± 4.1 y) participated in the control experiment. Participants gave oral and written consent for their attendance. The study was approved by the local ethics committee of Berlin Charité University Hospital.

In trials where reward information was presented before effort rating, we examined whether reward magnitude influenced the effort estimation (Fig. 2). We measured the regression slope between trial-by-trial variations in reward ($R_i - \mu_r$, with R_i being the reward on each trial in cents and μ_r being the mean reward of each difficulty level) and estimated effort ($E_i - \mu_e$, with E_i being the estimated effort on each trial and μ_e being the mean estimated effort of each difficulty level; see also Fig. S1). In both blocks with different reward contingencies (red vs. blue bars in Fig. 2), there was a significant relationship between reward variation and estimated effort (Wilcoxon sign rank test, $P = 0.00094$ for SD = 1.2 cents and $P = 0.002$ for SD = 2.5 cents). This effect of reward on estimated effort was stronger when reward variance was smaller, that is, when reward was highly contingent upon the task difficulty (mean slopes of 0.012 and 0.005 for SD = 1.2 and SD = 2.5 cents, respectively; Wilcoxon sign rank test for the difference of both slopes, $P = 0.004$). This result was highly consistent across subjects (Fig. S2). We observed the same pattern of results when reward magnitude was balanced across blocks using a stratification method (Fig. S3). In a control experiment, where reward was randomly varied and unrelated to the task difficulty (SD = ∞), regression slopes did not differ from zero (Wilcoxon sign rank test, $P = 0.15$; for individual data see Fig. S4). These results indicate that reward influences effort

Significance

Retrospective reevaluation of effort is a pervasive aspect of everyday life, such as when we assess our professional satisfaction after knowing the ensuing outcomes. Previous studies have focused on the interaction of effort and reward when a choice is to be made, whereas retrospective interactions have been largely ignored. Here we show that humans revise their estimation of effort after receiving a reward. When rewarded more than average, subjects tend to overestimate their effort, with a converse effect observed for low rewards. The size of this bias depends strongly on the contingency between reward magnitude and task difficulty and is dynamically adjusted when changes occur in these contingencies. These results reveal a sophisticated mechanism to cope with reward–effort inconsistencies.

Author contributions: A.P. and A.W. designed research; A.P. performed research; A.P. and A.W. analyzed data; and A.P., A.W., and R.J.D. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. P.W.G. is a guest editor invited by the Editorial Board.

Freely available online through the PNAS open access option.

[1]To whom correspondence should be addressed. Email: arezoo.pooresmaeili@gmail.com.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1507527112/-/DCSupplemental.

Fig. 1. Participants were asked to move a ball up a ramp by engaging in fast, alternating key presses. A gravity force was simulated, displacing the ball backward by a constant amount on each display frame. We used six levels of task difficulty, corresponding to the amount of ball displacement per time frame. After the ball was successfully pushed all of the way to the top of the ramp, participants received a monetary reward, where reward amount was contingent upon task difficulty. The strength of this contingency was varied in two separate blocks. Reward receipt information was either displayed before (90%) or after the rating of effort (10% of trials, not shown here). Subjects rated their effort by shifting the position of a sliding bar. At the end of each trial, they were asked to indicate whether they had seen a color change in the ball. All intervals in a trial were self-paced except for outcome reward display, which was in view for 2–3 s.

estimation only when there is a reliable relationship with task difficulty, and hence a reliable relationship with the true exertion level subjects expend while pushing the ball.

We next compared our behavioral data to predictions arising out of five separate computational models (Fig. 3). At the time when participants estimate their effort, the requirement was to retrospectively recall their true effort level. This recalled effort, E_m (gray distribution in Fig. 3*A*), is predictive of the true effort but is corrupted by noise, as reflected in the variance σ_m^2. When reward is correlated with task difficulty, reward magnitude yields by itself an independent effort estimation E_r with uncertainty σ_r^2 (red distribution in Fig. 3*A*). In each trial, E_r and E_m can differ by a certain amount ($\Delta E \neq 0$). The models we consider are based on distinct ways in which a conflict between these two informational sources (E_r and E_m) might be resolved.

In all models, E_m is computed based on the trials where reward was presented after the estimation of effort (i.e., where the estimation of effort is unaffected by reward information). E_r is computed based on the posterior probability distribution of task difficulty given an obtained reward (for details see *SI Materials and Methods*). Model 1 (memory only) assumes that participants solely rely on the recalled effort (E_m) and completely ignore reward information. Hence, the effort estimate $\hat{E}$ is equal to E_m multiplied by a scaling factor k_m:

$$\hat{E} = E_m K_m. \quad [1]$$

By contrast, model 2 (reward only) relies completely on reward information, whereas the recalled effort (E_m) is disregarded:

$$\hat{E} = E_r K_r. \quad [2]$$

Model 3 assumes that both E_m and E_r contribute to effort estimation with $\hat{E}$ being a weighted average (WA) of E_m and E_r:

$$\hat{E} = \omega E_m + (1-\omega)E_r. \quad [3]$$

Models 1–3 all assume that information regarding each signal's variance is not explicitly exploited by the participants. However, model 4 (Bayes optimal, BO) assumes that a Bayesian optimal strategy is used by the participants where signals are weighted based on their respective reliability or inverse variance:

$$\hat{E} = \omega_m E_m + \omega_r E_r, \quad [4]$$

where $\omega_m = 1 - \omega_r$ and $\omega_r = 1/\sigma_r^2/(1/\sigma_m^2 + 1/\sigma_r^2)$.

Similar reliability-based Bayesian models have been used previously to explain integration of sensory cues during perceptual decision making (6–8). The variances σ_m^2 and σ_r^2 are derived from the data and reward probability distributions. It is, however, debatable whether σ_m^2 is indeed inferred correctly using the trials where reward was presented after the rating of effort. Instead, the true variance σ_m^2 might be a multiple of the variance in these trials. Therefore, model 5 (adapted Bayes optimal, aBO) is a modified version of model 4, assuming that the variance of E_m (σ_m^2) is scaled by a free parameter k ($\widehat{\sigma_m}^2 = k\sigma_m^2$):

$$\hat{E} = \frac{\frac{1}{\sigma_r^2}}{\frac{1}{k\sigma_m^2}+\frac{1}{\sigma_r^2}} E_m + \left(1 - \frac{\frac{1}{\sigma_r^2}}{\frac{1}{k\sigma_m^2}+\frac{1}{\sigma_r^2}}\right) E_r. \quad [5]$$

We evaluated these models by computing their maximum likelihood fits to the trial-by-trial data of individual subjects, measuring the quality of fits by Bayesian information criterion (BIC, *SI Materials and Methods*). BIC weights shown in Fig. 3*B* are the weight of evidence in favor of each model (9, 10). The average BIC weights are highest for the Bayesian weighted

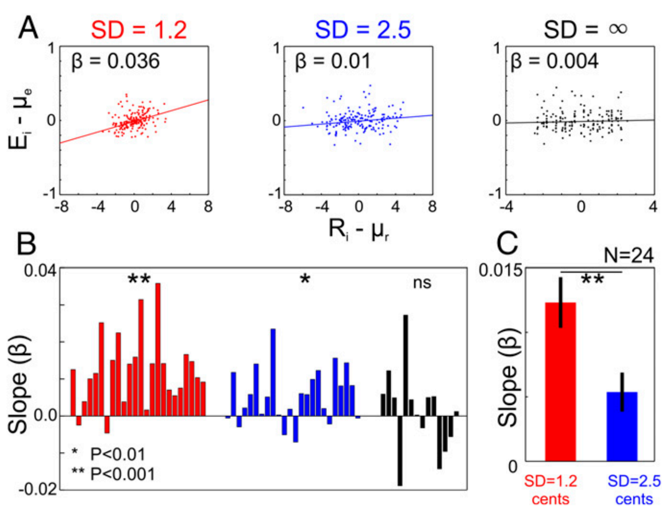

Fig. 2. (*A*) Relationship between trial-by-trial fluctuations in reward and retrospective estimates of expended effort in a typical subject tested with high (SD = 1.2 cents, data shown in red) and low (SD = 2.5 cents, data shown in blue) reward contingencies. The estimated effort E_i is normalized to each subject's maximum estimated effort in the whole experiment, and the mean effort μ_i of each difficulty level is subtracted. The same procedure is implemented in relation to rewards R_i and their means μ_r. The slope β of the linear regression is larger for the higher reward contingency. In a control experiment, we tested whether random variations of reward (SD = ∞, data shown in black) affect estimates of effort. Here, slopes did not significantly differ from zero (P = 0.15; Wilcoxon sign rank test.) (*B*) Regression slopes of all individual subjects in both experiments (colors as in *A*). Average slopes differ from zero only when reward is contingent on task difficulty (P levels from Wilcoxon sign rank test). (*C*) Average regression slopes across subjects are higher when reward is more contingent on task difficulty than when less contingent (P = 0.004, Wilcoxon sign rank test). Error bars denote SEM. Regression results are based on a robust regression analysis ("robustfit" in MATLAB with default settings) that minimizes the effect of potential outliers.

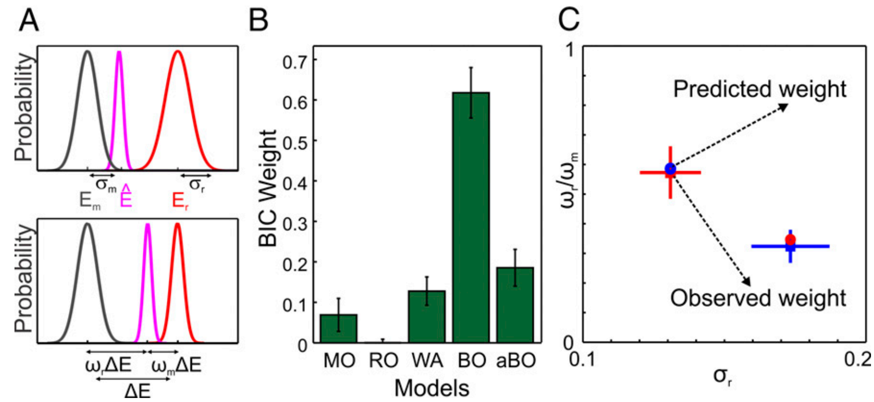

Fig. 3. (*A*) Two types of information can be combined to derive the estimated effort ($\hat{E}$): E_m is the recalled effort that is distorted by memory noise (σ_m); E_r is the most likely effort level given the reward magnitude in each trial, its variability σ_r being dependent on the contingency of reward and task difficulty. According to Bayes optimal models, the influence of each signal on $\hat{E}$ depends on its reliability. In blocks with low contingency, σ_r is large, and E_r, has a weaker influence on the estimate $\hat{E}$, whereas for high reward contingency $\hat{E}$ is closer to E_r (*Upper* and *Lower*, respectively). In each trial, E_m and E_r differ by a certain amount ΔE ($\Delta E \neq 0$). (*B*) Model comparison: BIC weights indicate the weight of evidence in favor of each model. (*C*) Using the weight ratio ω_m/ω_r derived from the data in one block (squares) and the known reward variability σ_r of both blocks, the weight ratio ω_m/ω_r in the other block (circles) can be accurately predicted (see also *Supporting Information*). Error bars denote SEM.

averaging model (see also Table S1). Importantly, in most subjects the model that merely relied on the memorized effort alone (without assuming any reward influence) performed considerably worse in terms of explaining the data (Table S2).

Model 4 (BO) also holds that the ratio of weights ω_m/ω_r derived from the data in one block should be predictive of the ratio of weights in the other block, respectively, given the known variances σ_r^2 of the reward signal in each block, and assuming the uncertainty of memory σ_m^2 to be constant (see also *Supporting Information*). Fig. 3*C* shows that this prediction provides a good match to the data ($P > 0.5$, Wilcoxon sign rank, for comparison of observed and predicted ω_m/ω_r). The reliance of participants on the reward signal can thus be accurately predicted by its variance. Finally, the variance of joint estimates was on average smaller than the variance of each signal alone, supporting that the Bayes optimal strategy participants use for signal integration improves their general precision in effort estimation (Fig. S5).

In this study, we show that after receiving reward information, humans revise their estimation of effort required for its attainment. The strength of correlation between task difficulty and obtained reward had a profound influence on retrospective effort estimates. Importantly, participants were adept at adjusting their estimation of effort when reward contingencies changed within the same experimental session.

We note that in our experiments more difficult trials required larger number of key presses and therefore took more time to complete than was the case for easy trials, entailing a decreased reward density per unit time. Therefore, participants' estimation of their effort might also be influenced by their estimation of trial time or perceived reward density. Although we cannot rule out a contribution from this factor, we would suggest that a dependency on time might be an inherent feature of effort estimations, because highly demanding tasks usually entail longer realization times.

How do these findings extend to real-life situations? In many instances, the effort we expend is closely tethered to consequential outcomes. For example, in the context of performance-based pay employees are remunerated in proportion to the degree to which a task is accomplished (11–13). Accordingly, there is a strong prior in many societies that rewards (wages) are contingent upon effort (labor invested). Equally, in fields such as sports (14, 15) and education (16–18), there is a common belief in a contingency between effort and success. In our study we show that such a prior is deployed by humans when they retrospectively evaluate effort. However, a direct comparison of the impact of real-life contingencies and the contingencies used in our paradigm is still missing. Moreover, in our task design and description participants may have paid more attention to reward, and the information conveyed by it, than is the case in real life. Future studies are needed to reveal whether, and to what extent, our findings generalize to other situations including real-life scenarios.

A question arises as to whether a flexible estimation of past effort serves a functional role. We suggest that rethinking one's efforts after receiving rewards constitutes a metacognitive ability (19) to negotiate uncertainties in effort–reward relationships. In everyday life, rewards are usually correlated with effort, but the strength of this correlation may change. The here reported acute adjustment of effort estimation can greatly aid goal-directed behavior when decisions are based on online monitoring of environmental factors. Indeed, failure of such mechanisms might lead to occupational disorders such as burn-out syndrome that are thought to be related to perception of an effort–reward imbalance (20, 21).

The distorted effort rating observed in the current study also bears resemblance to hindsight bias (22), reflecting a tendency to change recalled probability estimations once outcomes are known. As in a hindsight bias, retrospective change in effort estimation might reflect a general tendency to reshape memory contents to make them fit with an updated knowledge base (23). A large number of other cognitive biases have been described in the past that also reflect humans' deviation from rationality (24). Similarly, perceptual decisions are prone to deviate from the veridical as seen in phenomena such as sensory illusions (7, 25). Recent theoretical work has suggested that these "biases" reflect humans' ability to deal with the uncertainties in the world using the probabilistic structure of the environment (26, 27). Therefore, seemingly erroneous judgments are not only very common in the course of evolution (28) but are also optimal and rational (29–31). The finding that the impact of rewards on retrospectively evaluated effort conforms to a Bayesian rule of cue integration is in line with these previous studies, extending them to a domain with relevance to most aspects of our daily life.

We have shown that human subjects either under- or overestimate their past effort when rewards are smaller or larger than average, respectively. Whether a similar tendency influences normative beliefs, for example regarding the distribution of wealth in society, is a potentially important avenue of further investigation. For example, individuals with higher incomes who are exposed to greater-than-average rewards might have an inflated perception of the effort they expended to acquire their wealth. Conversely, those with low income might have the opposite perception. One might also speculate that such a biased perception of effort could contribute to stabilizing inequality. Indeed, the increasing equality gaps in many societies reinforce the importance of gaining insight into the complex interplay between retrospective assignments and the emergence of socioeconomic norms.

ACKNOWLEDGMENTS. We thank Alexandra Klein, Rafaela Wahl, and Valerie Keller for their help with the data collection, and Ulf Toelch for his advice on the modeling analysis. This work was performed while Ray Dolan was a Visiting Einstein Fellow at the Humboldt-Universität, Berlin School of Mind and Brain and was supported by Wellcome Trust Senior Investigator Award 098362/Z/12/Z (to R.J.D.).

1. Pessiglione M, et al. (2007) How the brain translates money into force: A neuroimaging study of subliminal motivation. *Science* 316(5826):904–906.
2. Meyniel F, Sergent C, Rigoux L, Daunizeau J, Pessiglione M (2013) Neurocomputational account of how the human brain decides when to have a break. *Proc Natl Acad Sci USA* 110(7):2641–2646.
3. Skvortsova V, Palminteri S, Pessiglione M (2014) Learning to minimize efforts versus maximizing rewards: Computational principles and neural correlates. *J Neurosci* 34(47):15621–15630.
4. Kurniawan IT, Guitart-Masip M, Dolan RJ (2011) Dopamine and effort-based decision making. *Front Neurosci* 5:81.
5. Wardle MC, Treadway MT, Mayo LM, Zald DH, de Wit H (2011) Amping up effort: Effects of d-amphetamine on human effort-based decision-making. *J Neurosci* 31(46): 16597–16602.
6. Ernst MO, Banks MS (2002) Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415(6870):429–433.
7. Alais D, Burr D (2004) The ventriloquist effect results from near-optimal bimodal integration. *Curr Biol* 14(3):257–262.
8. Körding KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. *Nature* 427(6971):244–247.
9. Burnham KP, Anderson DR (2015) *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach* (Springer, New York).
10. Wagenmakers EJ, Farrell S (2004) AIC model selection using Akaike weights. *Psychon Bull Rev* 11(1):192–196.
11. Lazear EP (1986) Salaries and piece rates. *J Bus* 59(3):405–431.
12. Stiglitz JE (1974) Alternative theories of wage determination and unemployment in LDC's: The labor turnover model. *Q J Econ* 88(2):194–227.
13. Nalebuff BJ, Stiglitz JE (1983) Prizes and incentives: Towards a general theory of compensation and competition. *Bell J Econ* 14(1):21–43.
14. Van-Yperen NW, Duda JL (1999) Goal orientations, beliefs about success, and performance improvement among young elite Dutch soccer players. *Scand J Med Sci Sports* 9(6):358–364.
15. Si G, Rethorst S, Willimczik K (1995) Causal attribution perception in sports achievement: A cross-cultural study on attributional concepts in Germany and China. *J Cross Cult Psychol* 26(5):537–553.
16. Frieze IH, Snyder HN (1980) Children's beliefs about the causes of success and failure in school settings. *J Educ Psychol* 72(2):186–196.
17. Tollefson N (2000) Classroom applications of cognitive theories of motivation. *Educ Psychol Rev* 12(1):63–83.
18. Livengood J (1992) Students' motivational goals and beliefs about effort and ability as they relate to college academic success. *Res Higher Educ* 33(2):247–261.
19. Fleming SM, Dolan RJ, Frith CD (2012) Metacognition: Computation, biology and function. *Philos Trans R Soc Lond B Biol Sci* 367(1594):1280–1286.
20. Siegrist J (1996) Adverse health effects of high-effort/low-reward conditions. *J Occup Health Psychol* 1(1):27–41.
21. de Jonge J, Bosma H, Peter R, Siegrist J (2000) Job strain, effort-reward imbalance and employee well-being: A large-scale cross-sectional study. *Soc Sci Med* 50(9): 1317–1327.
22. Roese NJ, Vohs KD (2012) Hindsight bias. *Perspect Psychol Sci* 7(5):411–426.
23. Hoffrage U, Hertwig R, Gigerenzer G (2000) Hindsight bias: A by-product of knowledge updating? *J Exp Psychol Learn Mem Cogn* 26(3):566–581.
24. Kahneman D, Slovic P, Tversky A (1982) *Judgment Under Uncertainty: Heuristics and Biases* (Cambridge Univ Press, New York).
25. Geisler WS, Kersten D (2002) Illusions, perception and Bayes. *Nat Neurosci* 5(6): 508–510.
26. Trommershäuser J (2009) Biases and optimality of sensory-motor and cognitive decisions. *Prog Brain Res* 174:267–278.
27. Summerfield C, Tsetsos K (2015) Do humans make good decisions? *Trends Cogn Sci* 19(1):27–34.
28. Marshall JA, Trimmer PC, Houston AI, McNamara JM (2013) On evolutionary explanations of cognitive biases. *Trends Ecol Evol* 28(8):469–473.
29. Fennell J, Baddeley R (2012) Uncertainty plus prior equals rational bias: An intuitive Bayesian probability weighting function. *Psychol Rev* 119(4):878–887.
30. Trimmer PC, et al. (2011) Decision-making under uncertainty: Biases and Bayesians. *Anim Cogn* 14(4):465–476.
31. Bach DR, Dolan RJ (2012) Knowing how much you don't know: A neural organization of uncertainty estimates. *Nat Rev Neurosci* 13(8):572–586.
32. Brainard DH (1997) The Psychophysics Toolbox. *Spat Vis* 10(4):433–436.
33. Pelli DG (1997) The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spat Vis* 10(4):437–442.
34. Burnham KP, Anderson DR (2004) Multimodel inference: Understanding AIC and BIC in model selection. *Sociol Methods Res* 33(2):261–304.