ARTICLE

Arezoo Pooresmaeili[1], Aurel Wannig[2], Raymond J. Dolan[1,3,4]
[1]Humboldt-Universität zu Berlin, Berlin School of Mind and Brain, Luisenstraße 56, Haus 1, 10099 Berlin, Germany
[2]Freie Schule Charlottenburg, Frankenallee 6, 14052 Berlin, Germany
[3]Wellcome Trust Center for Neuroimaging, Institute of Neurology, University College London, London WC1N 3BG, United Kingdom
[4]Max Planck- University College London Centre for Computational Psychiatry and Ageing Research, London WC1B 5EH
Correspondence should be addressed to AP, Address: Humboldt-Universität zu Berlin, Berlin School of Mind and Brain, Luisenstraße 56, Haus 1, 10099 Berlin, Germany, E-mail: arezoo.pooresmaeili@gmail.com, Phone: +49 30 2093-1794, Fax: +49 30 2093-1802
Running head: Effects of reward on estimated effort

Abstract
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 re-estimated 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.

Significance Statement
Retrospective re-evaluation 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, while 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.

"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 the **SI Text** 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 the way up the ramp were rewarded. Reward was contingent upon task difficulty (with values drawn from 6 Gaussian distributions with Means from 1.5-6.5 cents) and the strength of this contingency varied across different blocks of the experiment (Standard deviations (SD) 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, since 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 Text**).

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 backwards 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 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 if they had

seen a color change in the ball. All intervals in a trial were self-paced except for outcome reward display (2-3s).

In trials where reward information was presented before effort rating, we examined if 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, i.e. 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 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.

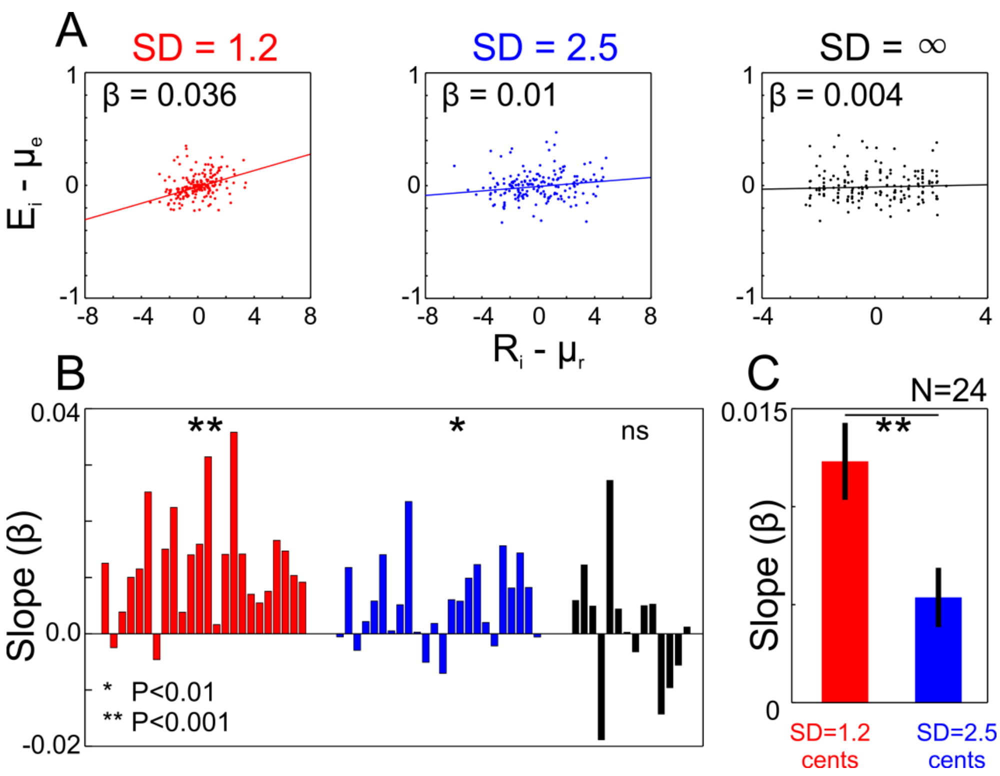

Fig. 2. (A) Relationship between trial-by-trial fluctuations in reward and retrospective estimates of expended effort in a typical subject tested with high (SD=1.2 cents, data shown in red) and low (SD=2.5 cents, data shown in blue) reward contingencies. The estimated effort E_i is normalized to each 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 if random variations of reward (SD=∞, data shown in black) impact on 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 s.e.m. Regression results are based on a robust regression analysis ("robustfit" in MATLAB with default settings) that minimizes the effect of potential outliers.

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.3A**), 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.3A**). 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 how 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 Text). 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:

$$(1)\hat{E} = E_{\mathrm{m}} * k_{\mathrm{m}}$$

By contrast, model 2 (*Reward Only*) relies completely on reward information, while the recalled effort (E_m) is disregarded:

$$(2)\ \hat{E} = E_{\mathrm{r}} * k_{\mathrm{r}}$$

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:

$$(3)\ \hat{E} = \omega * E_{\mathrm{m}} + (1 - \omega) * E_{\mathrm{r}}$$

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

$$(4)\ \ \hat{E} = \omega_{\mathrm{m}} * E_{\mathrm{m}} + \omega_{\mathrm{r}} * E_{\mathrm{r}},$$

where $\omega_{\mathrm{m}} = 1 - \omega_{\mathrm{r}}$ and $\omega_{\mathrm{r}} = \frac{\frac{1}{\sigma_{\mathrm{r}}{}^2}}{\frac{1}{\sigma_{\mathrm{m}}{}^2} + \frac{1}{\sigma_{\mathrm{r}}{}^2}}$.

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

$$(5)\ \ \hat{E} = \frac{\frac{1}{\sigma_{\mathrm{r}}{}^2}}{\frac{1}{k\sigma_{\mathrm{m}}{}^2} + \frac{1}{\sigma_{\mathrm{r}}{}^2}} * E_{\mathrm{m}} + \left(1 - \frac{\frac{1}{\sigma_{\mathrm{r}}{}^2}}{\frac{1}{k\sigma_{\mathrm{m}}{}^2} + \frac{1}{\sigma_{\mathrm{r}}{}^2}}\right) * E_{\mathrm{r}}$$

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, see methods). BIC weights shown in **Fig. 3B** are the weight of evidence in favor of each model (*9, 10*). The average BIC weights are highest for the Bayesian weighted averaging model (for data of individual subjects see **SI Text**). 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 (see the **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 to be constant (see also Suppl. Info). **Fig. 3C** 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 accuracy in effort estimation (see **Fig.S5).**

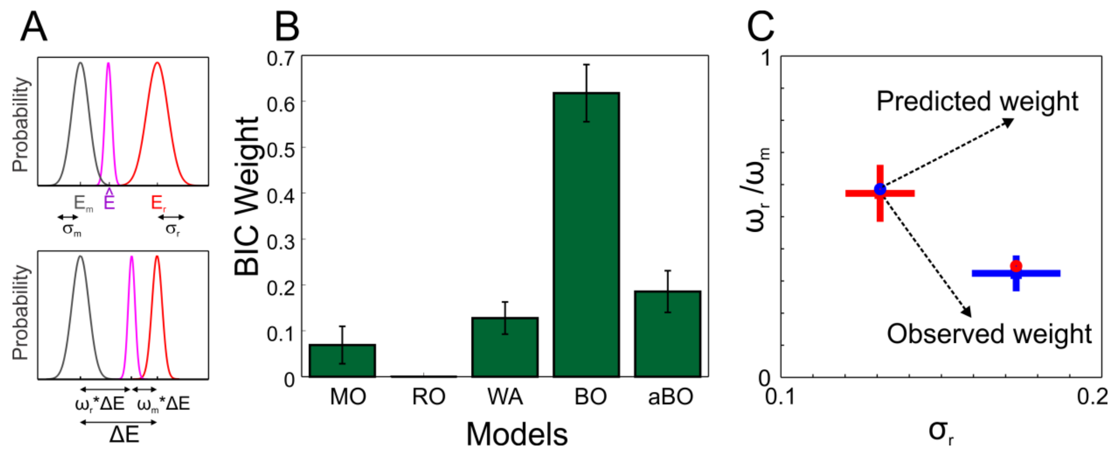

Fig. 3. (A) Two types of information can be combined to derive the estimated effort ($\hat{E}$): E_m is the recalled effort which is distorted by memory noise (σ_m); E_r is the most likely effort level given the reward magnitude in each trial, its variability σ_r being dependent on the contingency of 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 $\hat{E}$, while for high reward contingency, $\hat{E}$ is closer to E_r (upper and lower panel, respectively). In each trial, E_m and E_r differ by a certain amount ΔE (ΔE≠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 and the known reward variability σ_r of both blocks, the weight ratio ω_m/ω_r in the other block can be accurately predicted (see also **SI**).

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. While 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, since 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 sport (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 in order to make them fit with an updated knowledge base (23). A large number of other cognitive biases have been described in the past which also reflect human's 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 over-estimate 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 socio-economic norms.

Acknowledgements:
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 a Wellcome Trust Senior Investigator Award 098362/Z/12/Z to RJD.

References

1. Pessiglione, M., Schmidt, L., Draganski, B., Kalisch, R., Lau, H., Dolan, R. J., & Frith, C. D. (2007) How the brain translates money into force: a neuroimaging study of subliminal motivation. *Science* **316**, 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. U. S. A* **110**, 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**, 15621-15630.

4. Kurniawan, I. T., Guitart-Masip, M., & Dolan, R. J. (2011) Dopamine and effort-based decision making. *Front Neurosci.* **5**, 81.

5. Wardle, M. C., Treadway, M. T., Mayo, L. M., Zald, D. H., & de Wit, H. (2011) Amping up effort: effects of d-amphetamine on human effort-based decision-making. *J. Neurosci.* **31**, 16597-16602.

6. Ernst, M. O. & Banks, M. S. (2002) Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* **415**, 429-433.

7. Alais, D. & Burr, D. (2004) The ventriloquist effect results from near-optimal bimodal integration. *Curr. Biol.* **14**, 257-262.

8. Kording, K. P. & Wolpert, D. M. (2004) Bayesian integration in sensorimotor learning. *Nature* **427**, 244-247.

9. Burnham, K. P. & Anderson, D. R. (2015) *Model selection and multimodel inference: A practical information-theoretic approach* (Springer-Verlag, New York).

10. Wagenmakers, E. J. & Farrell, S. (2004) AIC model selection using Akaike weights. *Psychon. Bull. Rev.* **11**, 192-196.

11. Lazear, E. P. (1986) Salaries and Piece Rates. *The Journal of Business* **59**.

12. Stiglitz, J. E. (1974) Alternative Theories of Wage Determination and Unemployment in LDC's: The Labor Turnover Model. *The Quarterly Journal of Economics* **88**.

13. Nalebuff, B. J. & Stiglitz, J. E. (1983) Prizes and Incentives: Towards a General Theory of Compensation and Competition. *The Bell Journal of Economics* **14**.

14. Van-Yperen, N. W. & Duda, J. L. (1999) Goal orientations, beliefs about success, and performance improvement among young elite Dutch soccer players. *Scandinavian Journal of Medicine & Science in Sports* **9**, 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. *Journal of Cross-Cultural Psychology* **26**, 537-553.

16. Frieze, I. H. & Snyder, H. N. (1980) Children's beliefs about the causes of success and failure in school settings. *J. Educ Psychol.* **72**, 186-196.

17. Tollefson, N. (2000) Classroom Applications of Cognitive Theories of Motivation. *Educational Psychology Review* **12**, 63-83.

18. Livengood, J. (1992) Students' motivational goals and beliefs about effort and ability as they relate to college academic success. *Res High Educ* **33**, 247-261.

19. Fleming, S. M., Dolan, R. J., & Frith, C. D. (2012) Metacognition: computation, biology and function. *Philos. Trans. R. Soc. Lond B Biol. Sci.* **367**, 1280-1286.

20. Siegrist, J. (1996) Adverse health effects of high-effort/low-reward conditions. *J. Occup. Health Psychol.* **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**, 1317-1327.

22. Roese, N. J. & Vohs, K. D. (2012) Hindsight Bias. *Perspectives on Psychological Science* **7**, 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**, 566-581.

24. Kahneman, D., Slovic, P., & Tversky, A. (1982) *Judgment under uncertainty: Heuristics and biases.* (Cambridge University Press, New York).

25. Geisler, W. S. & Kersten, D. (2002) Illusions, perception and Bayes. *Nat. Neurosci.* **5**, 508-510.

26. Trommershauser, 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**, 27-34.

28. Marshall, J. A., Trimmer, P. C., Houston, A. I., & McNamara, J. M. (2013) On evolutionary explanations of cognitive biases. *Trends Ecol. Evol.* **28**, 469-473.

29. Fennell, J. & Baddeley, R. (2012) Uncertainty plus prior equals rational bias: an intuitive Bayesian probability weighting function. *Psychol. Rev.* **119**, 878-887.

30. Trimmer, P. C., Houston, A. I., Marshall, J. A., Mendl, M. T., Paul, E. S., & McNamara, J. M. (2011) Decision-making under uncertainty: biases and Bayesians. *Anim Cogn* **14**, 465-476.

31. Bach, D. R. & Dolan, R. J. (2012) Knowing how much you don't know: a neural organization of uncertainty estimates. *Nat. Rev. Neurosci.* **13**, 572-586.

32. Brainard, D. H. (1997) The Psychophysics Toolbox. *Spat. Vis.* **10**, 433-436.

33. Pelli, D. G. (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat. Vis.* **10**, 437-442.

34. Burnham K.P. & Anderson D.R. (2004) Multimodel Inference: Understanding AIC and BIC in Model Selection. *Sociological Methods Research* **33**, 261-304.

35. Wagenmakers, E. J. & Farrell, S. (2004) AIC model selection using Akaike weights. *Psychon. Bull. Rev.* **11**, 192-196.

Supporting Text

Participants. 26 participants (15 females, age: 20-39, mean: 27.07±5.1) took part in our main experiment. Two participants were excluded from the analysis since in our debriefing they mentioned they had not paid attention to the reward magnitudes during the experiment. 14 participants (9 females, age: 21-35, mean: 27.5±4.1) 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.

Stimuli and task. Stimuli were produced with MATLAB and the Psychophysics Toolbox (1, 2). Each trial consisted of a stimulus (ball and ramp), a reward and an effort rating display all shown on a black background (see **Fig.1**). The stimulus display contained the ball (radius: 3 visual degrees), initially at the starting part of the ramp (ramp length: 19 visual degrees; both ball and ramp had a light gray color). The ball was displaced up the ramp with consecutive alternate key presses (left and right arrow keys) until it reached the upper plateau. Each key press resulted in a constant amount of displacement (0.87 visual degrees per key press) and was counteracted by a gravity force of variable strength that displaced the ball backwards. To determine the levels of gravity force, at the beginning of the experiment we asked each subject to push the ball up the ramp by pressing both keys alternately and consecutively as fast as they possibly could. 90% of the gravity force necessary to counteract the maximum number of key presses in a limited time (10s) determined the maximum gravity force used in the experiment. Based on this individualized estimate, 6 equally spaced gravity levels were defined and used in the experiment. A trial was aborted if key presses did not occur fast enough (max pause allowed: 2s). If participants were able to successfully push the ball all the way up, they received a monetary reward, with the amount contingent on task difficulty (gravity force level). Reward magnitude was defined based on 6 Gaussians with means of 1.5, 2.5, 3.5, 4.5, 5.5 and 6.5 cents, using two different standard deviations (SD of 1.2 or 2.5 cents) which determined the contingency between reward magnitude and task difficulty. Reward display consisted of a pie chart that depicted subjects' reward as a proportion of maximum reward possible and a number which showed the reward in Arabic numerals. Effort rating display consisted of a slider, and participants were instructed to set the slider at a position which represents their experienced effort during a trial proportionate to the maximum effort they had ever experienced during the experiment. In 90% of trials, the reward display was shown immediately after the stimulus display, while in 10% of trials it occurred after the rating of effort.

The main experiment was done in two separate blocks in which different reward contingencies (SD of 1.2 or 2.5 cents) were used. Participants did not receive any instruction regarding the strength of contingency between task difficulty and reward, and were only told that a change in the strength of this relationship would occur across blocks. The order of these blocks was counterbalanced across subjects. Due to the random assignment of this order (i.e. whether the contingency of the first block was high or low), the direction of change across blocks (from weak to strong or vice versa) was also unknown to the experimenters. Therefore, participants had to experience the relationship between task difficulties and reward by themselves, in the absence of any prior knowledge regarding the strength of any contingencies. Subjects first performed 12 training trials before the start of each block so that they were acquainted with the task and the reward-effort relationship. These trials were not included in our analyses. Each block consisted of 6 smaller mini-blocks, each consisting of 36 trials. Participants could take a pause and rest between the mini-blocks. In the control experiment, rewards were randomly chosen and varied between 0.5 and 5 cents, without any relationship to the difficulty levels.

In pilot experiments, some participants reported that in order to work out a relationship between task difficulty and reward they had adopted different ad-hoc strategies. Since task difficulty and reward were the only parameters that varied across trials, participants had presumably focused exclusively on relating the two. In the main experiment, to distract subjects from its true purpose we introduced a second task into the main paradigm. Thus, subjects were also asked to report if they had seen a brief (21ms or 3 frames) color change on the ball (to green, red or blue) at the end of each trial, which occurred in 50% of the trials. We debriefed the participants after the experiments and note that none reported that the main question of the study involved effort and reward relationships, nor did they now report using ad-hoc strategies.

Computational modeling. Our models are based on two independent effort estimates: An estimate purely based on the subjective memory of the effort spent in a trial (E_m), which is unaffected by reward information, and an estimate purely based on the information conveyed by reward magnitude (E_r).

E_m was derived from a Gaussian distribution with mean μ_{m} and variance σ_{m}^2, which were computed based on the trials in which estimation of effort was unaffected by reward (i.e. where effort was rated after receiving reward). Effort estimates of these unbiased trials will be referred to as E_u. The average μ_{m} at a difficulty level d_j was calculated as the median of a participant's rated efforts E_u in that difficulty level. The variance σ_m^2 across trials t was calculated as:

(1) $\sigma_m^2 = \frac{1}{n}\sum_{t=1}^{n}\left(E_{u_t} - \mu_m\right)^2$

To estimate E_m in a given trial, a random value was then drawn from a Gaussian distribution with mean $\mu_m$and variance σ_m^2.

E_r is the most likely effort level given a certain reward. It was derived from a Gaussian distribution with mean μ_r and variance σ_r^2. To infer μ_r, we computed the most probable difficulty level $\hat{d}$ from all levels d_j, given a reward r, which can be obtained maximizing the posterior probability function

(2) $P(d_j|r) = P(r|d_j) * P(d_j)/P(r)$,

where $P(d_j)$ and $P(r)$ are the probability of each difficulty level and each reward magnitude, respectively, and $P(r|d_j)$ is the probability of a given reward at each difficulty level. The average estimate μ_r of all trials with a given $\hat{d}$ was computed as the median of estimated efforts E_u at that difficulty level. The variance σ_r^2across all trials t was approximated as:

(3) $\sigma_r^2 = \frac{1}{n}\sum_{t=1}^{n}(\mu_r - \mu_m)^2$

To estimate E_r in a given trial, a random value was then drawn from a Gaussian distribution with mean μ_r and variance σ_r^2.

In model 4 (BO), for each block, the corresponding σ_r^2 was used; σ_m^2 was assumed to remain unchanged across blocks. Model 5 (aBO) challenges the assumption that the variance σ_m^2 is correctly reflected in the trials where participants rate their effort before receiving reward: The variance in these trials is influenced by motor noise, which might not be incorporated in σ_m. Moreover, at the time when subjects retrospectively evaluate their effort, the variance of their estimates might be affected by unknown factors (such as increased memory noise). Therefore, in model *aBO*, a free parameter (k) is used to scale σ_m^2. All other aspects of model 5 remained the same as in model 4. In all models, we only included trials with a conflict between E_m and E_r, i.e. where the estimate of task difficulty $\hat{d}$ provided by reward differed from the actual difficulty level. All trials of the two blocks (reward SD of 1.2 and 2.5) were modeled at once.

Model Comparison. Model one, two, three and five each had 1 free parameter: k_m, k_r, ω and k respectively. Model 4 (BO) had no free parameter. The fit of each model (n=5) to the data was evaluated by computing log-likelihood and BIC:

(4) $BIC = -2LL + \ln(n) * m$

where LL is the model log-likelihood, and m is the number of free parameters. The individual BIC values contain arbitrary constants and are very much affected by sample size. We therefore rescaled BIC to obtain ∆BICs (3, 4), where:

(5) $\Delta BIC = BIC_i - BIC_{min}$

and BIC_{min} is the minimum of the N different BIC_i values. This transformation forces the best model to have $\Delta BIC = 0$, while the rest of the models have positive values. The Bayesian weights are then derived from ΔBIC:

$$(6) \quad \omega_i = \frac{\exp(-0.5\Delta BIC_i)}{\sum_{n-1}^{N}(-0.5\Delta BIC_n)}$$

The Bayesian weights (ω_i) of all the models in a set sum up to 1 and indicate the probability for each model to be the best model for the data.

Prediction of Bayesian weights using reward contingencies

When the memory signal E_m and the reward signal E_r yield different effort estimates ($\Delta E \neq 0$), participants integrate both signals into an estimate $\hat{E}$, using the weights ω_m and ω_r, respectively. The ratio of these weights can be directly derived from the data (see also **Fig. 3A**):

$$(7) \quad \frac{\omega_r}{\omega_m} = \frac{|\hat{E}-E_m|\,\Delta E}{|\hat{E}-E_r|\,\Delta E}$$

Model 4 also holds that the weights of the reward and memory signal are inversely proportional to their relative variance:

$$(8) \quad \frac{\omega_r}{\omega_m} = \frac{{\sigma_m}^2}{{\sigma_r}^2}$$

Assuming that the memory variance ${\sigma_m}^2$ remains the same in the two blocks with different reward contingencies, we can compute the variances ${\sigma_{r1}}^2$ and ${\sigma_{r2}}^2$ as:

$${\sigma_{r1}}^2 * \frac{\omega_{r1}}{\omega_{m1}} = {\sigma_{r2}}^2 * \frac{\omega_{r2}}{\omega_{m2}}$$

,

$$(9) \quad \frac{\omega_{r1}}{\omega_{m1}} = \frac{{\sigma_{r1}}^2}{{\sigma_{r2}}^2} * \frac{\omega_{r2}}{\omega_{m2}} \qquad .$$

The variances ${\sigma_{r1}}^2$ and ${\sigma_{r2}}^2$ can be inferred from the known contingencies between reward and task difficulty (eq. 2). Thus, eq. 9 involves a direct prediction of the ratio of weights in one block from the ratio of weights in the other block.

Reference List

1. Brainard, D. H. (1997) The Psychophysics Toolbox. *Spat. Vis.* **10**, 433-436.

2. Pelli, D. G. (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat. Vis.* **10**, 437-442.

3. Burnham K.P. & Anderson D.R. (2004) Multimodel Inference: Understanding AIC and BIC in Model Selection. *Sociological Methods Research* **33**, 261-304.

4. Wagenmakers, E. J. & Farrell, S. (2004) AIC model selection using Akaike weights. *Psychon. Bull. Rev.* **11**, 192-196.

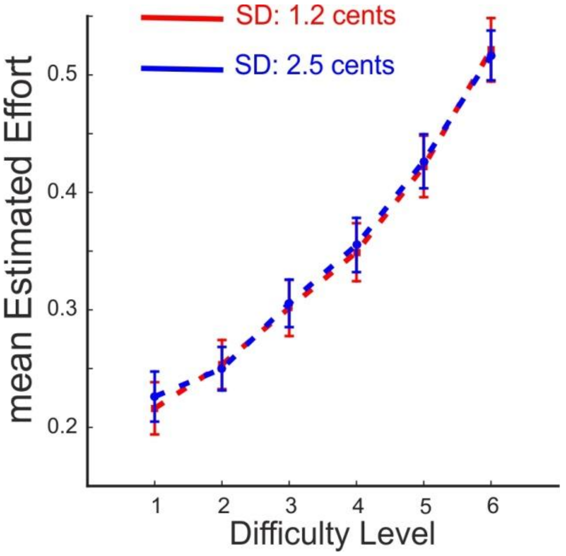

Fig. S1. Mean effort (μ_i) for each of the six difficulty levels *i*. μ_i increases monotonically with task difficulty. In our analysis, to assess the effect of reward independently of task difficulty, μ_i is subtracted from participants' rated efforts (see also **Fig. 2A**). The resulting effort fluctuations (E_i - μ_i) are therefore only influenced by variations of reward (also see Figure 2 and Supplementary Figure 3).

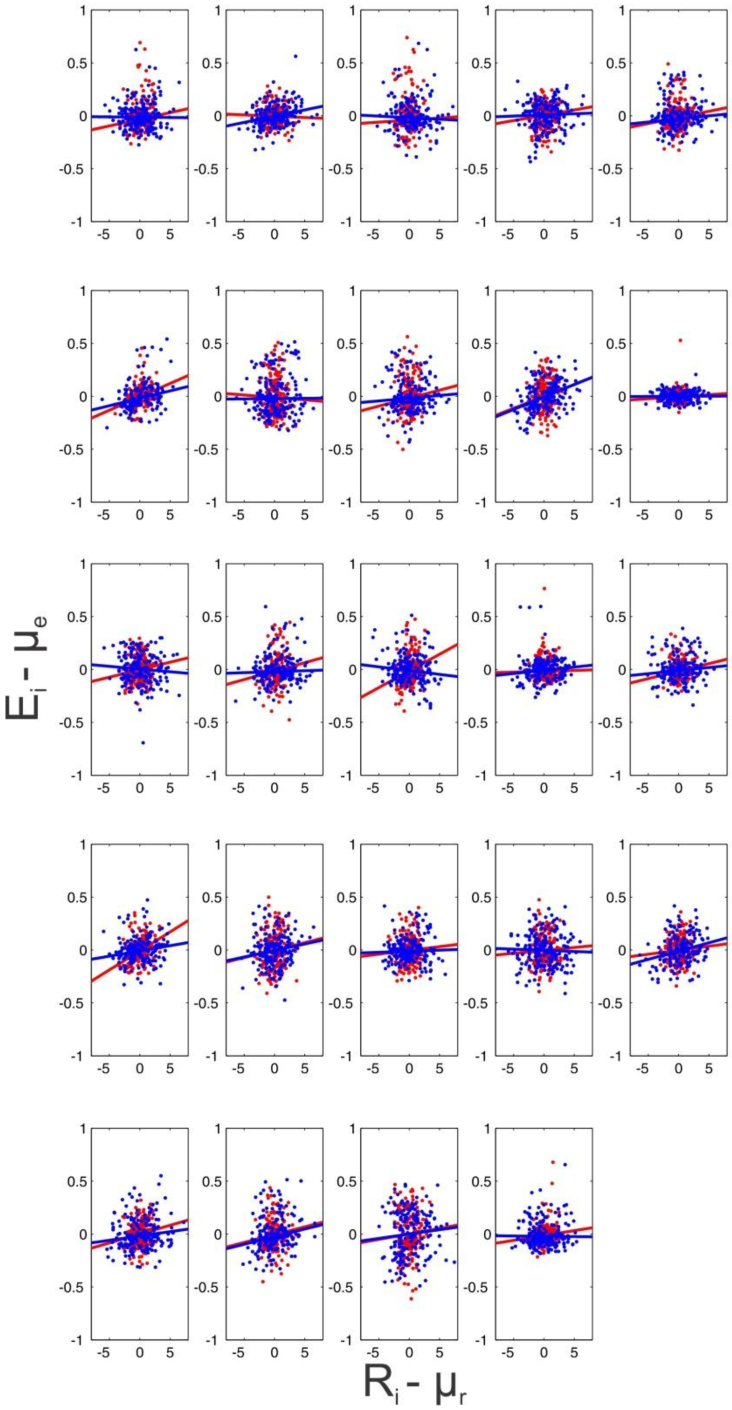

Fig. S2. Correlation between variations in rewards and estimated effort in individual subjects. Data in red corresponds to the block with stronger contingency between reward and task difficulty (SD=1.2), while data in blue corresponds to the block with a weaker contingency (SD=2.5)

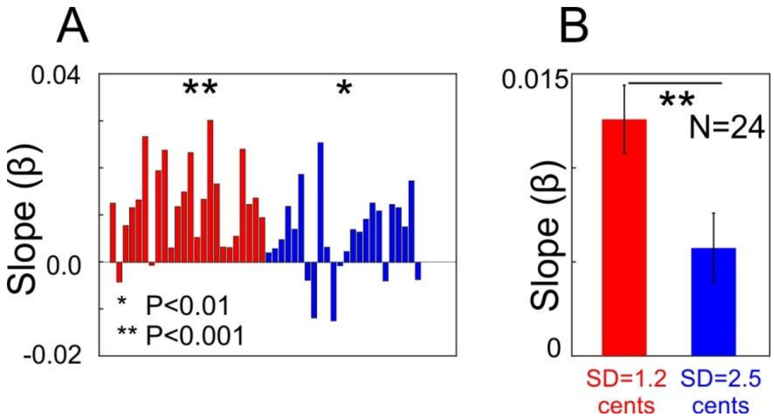

Fig. S3. Regression slopes of individual subjects (A) and average regression slopes (B) for stratified data (colors as in Fig 2). Difference in contingency between reward and task difficulty entailed that reward spread could differ across blocks (compare reward spreads in Fig. 2A), which might in turn affect regression slopes. We therefore used a stratification method to balance reward magnitudes across blocks, and checked if the observed slope differences do still uphold. Stratification was done by defining 6 equally spaced bins between minimum and maximum reward (R_i - μ_r). We randomly removed surplus trials until trial number was the same for the two blocks in every bin. We recomputed the regression slopes for the stratified data. As shown in this figure, all our results did also hold for the stratified data (see also Figure 2B and 2C). In all our regression analyses (in the main text as well as in Supplementary Material), we used a robust regression analysis ('robustfit' in MATLAB with the default bisquare weighing function). This was done to minimize the contribution of potential outliers on the regression slopes.

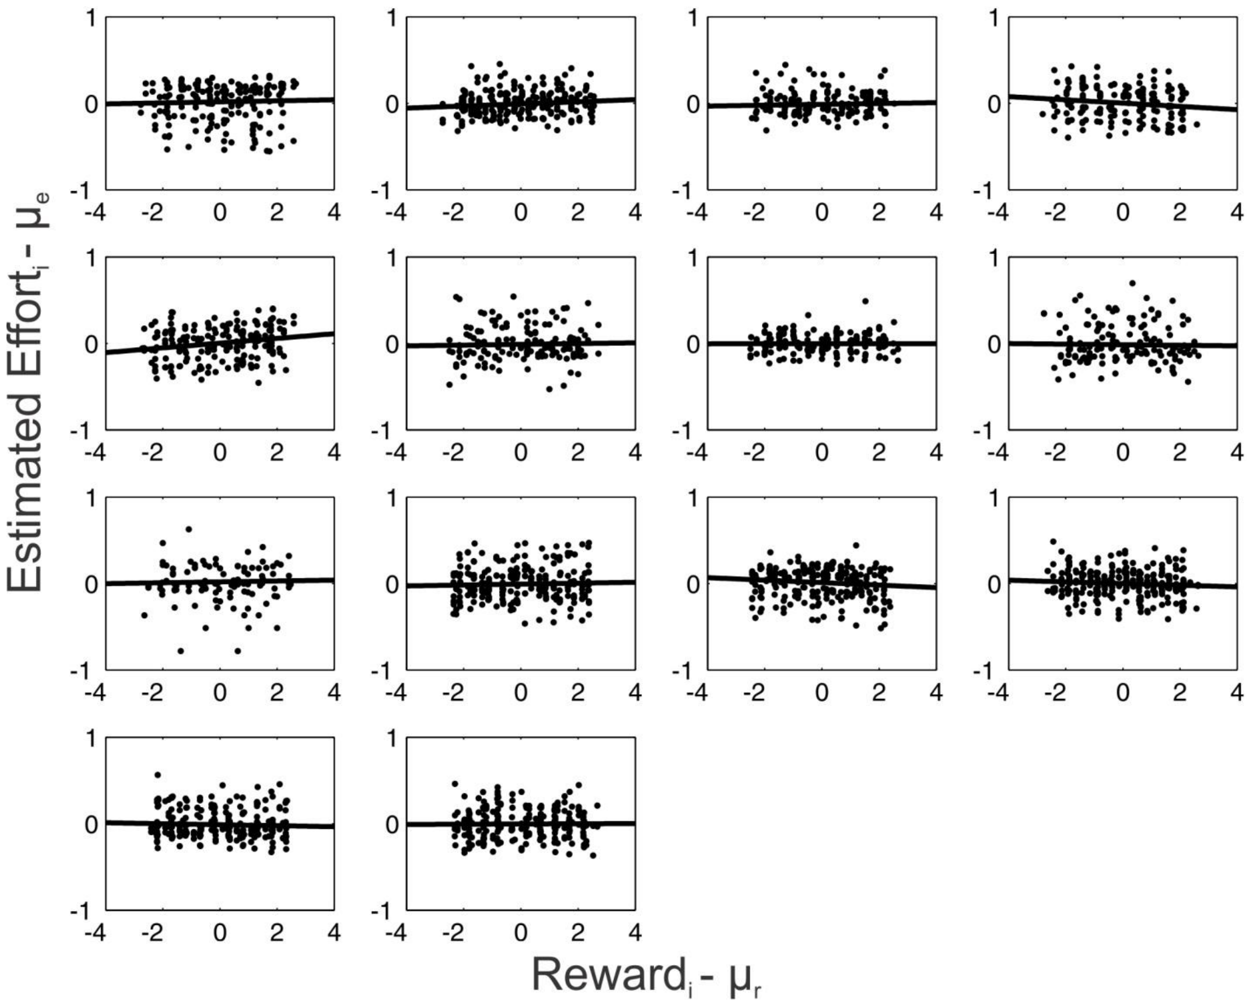

Fig. S4. Rewards that are not contingent on task difficulty (SD= ∞) have no effect on estimated effort.

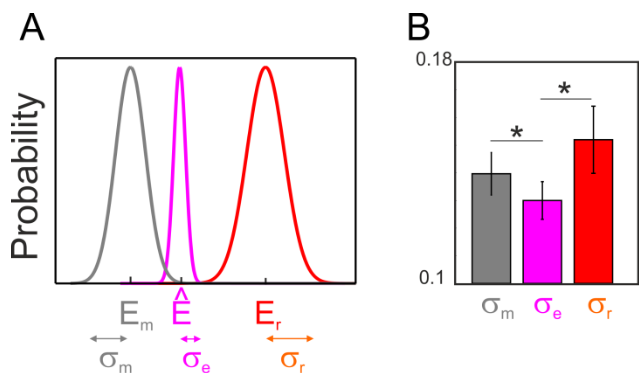

Fig. S5. **A)** Estimated effort ($\hat{E}$), shown in magenta is computed based on E_m (shown in gray) and E_r (shown in red). Each of these estimates has its corresponding uncertainty reflected in the standard deviation of the distribution (σ_e, σ_m and σ_r, respectively). Bayesian optimal integration predicts σ_e to be smaller than both σ_m and σ_r. **B)** Standard deviations σ_e, σ_m and σ_r, averaged across all subjects and both reward contingency blocks. σ_e is smaller than σ_m ($p=0.003$) and σ_r ($p=0.01$, Wilcoxon Sign rank test).