

Representations of Subjective Effort Cost

Patrick S. Hogan¹, Cary D. Frydman², Joseph K. Galaro¹, Vikram S. Chib^{1,3,*}

¹ Department of Biomedical Engineering, Johns Hopkins School of Medicine, Baltimore, MD 21205, USA

² Marshall School of Business, University of Southern California, Los Angeles, CA 90089, USA

³ Kennedy Krieger Institute, Baltimore, MD 21205, USA

*Correspondence: vchib@jhu.edu

SUMMARY

How effortful an action feels critically shapes everyday decisions. Despite the importance of perceptions of effortfulness for making choices, the behavioral and neural representations of the subjective cost of physical effort are not well understood. We used functional magnetic resonance imaging (fMRI) to monitor brain activity while participants engaged in risky choices for prospective physical effort independent of reward. Behaviorally we found that participants exhibited increasing sensitivity to changes in subjective effort as objective effort levels increased. Moreover, the ventromedial prefrontal cortex (vmPFC) encoded the subjective cost of prospective effort options and integrated these representations to facilitate decisions regarding effort expenditure. These results provide insight into human decision-making by showing how neural representations of feelings of effortfulness engender choices about exertion.

INTRODUCTION

Our decisions are shaped not only by the rewards at stake, but also the effort required to obtain them. Moreover, the subjectivity of effort plays an integral role in driving choice - if a task feels very effortful one may not be willing to perform the work required to obtain reward, whereas if a task feels less effortful one may be more likely to persevere. Feelings of effortfulness have been suggested to impact everyday decisions in a variety of contexts ranging from searching for jobs (DellaVigna and Paserman, 2005) and performing labor (Abeler et al., 2011; Augenblick et al., 2015), to participating in exercise and physical activity (Sniehotta et al., 2005; Dishman, 1991). However, little is known

about the behavioral manifestations of subjective effort cost and how neural representations of these costs give rise to decisions to exert effort.

Recently subjective value signals for appetitive stimuli (e.g., money, food, consumer goods) have been found in the ventromedial prefrontal cortex (vmPFC) (Plassmann et al., 2007; Chib et al., 2009; Kable and Glimcher, 2007; Padoa-Schioppa and Assad, 2008; Bartra et al., 2013). There have also been a number of studies in animals (Hillman and Bilkey, 2012; Rudebeck et al., 2008; Walton et al., 2002) and humans (Croxson et al., 2009; Prévost et al., 2010; Kurniawan et al., 2010) that have examined how the brain makes decisions to exert physical effort for reward. From this work it has been suggested that the vmPFC encodes the stimulus value of rewards and that anterior cingulate cortex (ACC) and ventral striatum encode effort costs and effort-cost trade-offs (Rangel and Hare, 2010; Rushworth et al., 2009). These studies focused on physical effort in an objective sense and did not model individuals' subjective cost of effort. A key feature of these previous experiments was that they examined choices between effort and reward and were not designed to isolate subjective effort costs from monetary or other reward-based incentives. As such our understanding of the neurobiological basis of subjective effort cost remains incomplete.

We investigated the behavioral representations of subjective effort cost and how these effort preferences are encoded in the brain's valuation and decision-making circuitry. We hypothesized that, in a similar fashion to monetary rewards (Holt and Laury, 2002; Kahneman and Tversky, 1979), participants would exhibit subjectivity in their decisions

for prospective effort. That is, effort representations would differ from the objective amount of effort, and would instead contain a degree of subjectivity driven by an individual's perception of how effortful a task feels. Since our paradigm was designed to isolate subjective effort costs, independent of reward (i.e., not requiring a trade-off between effort and reward), we hypothesized that behavioral representations of this effort subjectivity would be encoded in vmPFC – consistent with findings in human (Kable and Glimcher, 2007; Tom et al., 2007) and animal (Padoa-Schioppa and Assad, 2008) studies investigating subjective value of appetitive and aversive stimuli.

To test our hypotheses, we recruited healthy right-handed participants ($n = 23$) who performed a novel effort-choice task while their neural activity was recorded with functional magnetic resonance imaging (fMRI). Participants were first trained to associate the force exerted on a hand-clench dynamometer with numeric effort levels between 0 and 100 (Figure 1A). These effort levels were defined relative to a participant's maximum voluntary contraction (MVC) obtained prior to the experiment. The efficacy of this association was then tested in a recall phase (Figure 1B), wherein participants produced a force consistent with an unknown effort level and then immediately reported the effort they believed to have exerted. Following the recall phase, participants were presented a series of risky choices regarding prospective effort exertion (Figure 1C). Two options were presented: the first entailed an option where the participant would either need to exert a large amount of effort (G), or no effort, both with equal probability ($P = 0.5$); in the second option, a participant would need to exert a lower effort level (S) with certainty ($P = 1$). Participants had a 4 second time window to weigh the two options against each other and

make a decision. We presented participants with one hundred and seventy choices without any feedback until the end of the experiment. This choice set was selected *a priori* to allow accurate estimation of effort subjectivity. Prior to being presented with the effort choices, participants were told that 10 of their decisions would be selected at random and realized at the end of the experiment. This was done to ensure that participants were properly incentivized on each trial. It is important to note that these effort choices did not involve monetary earnings, which allowed us to experimentally isolate effort cost representations from reward.

RESULTS

Behavioral Representations of Subjective Effort Cost

Comparison between reported and exerted effort levels during the recall phase showed a high degree of agreement, indicating that participants accurately perceived the objective effort levels (Figure 2A shows the group recall results; see Figure S1 for all individual participants).

We characterized the subjectivity of participant i 's effort choices using a subjective cost function $V_i(e) = -(-e)^{\rho_i}$, where $e \leq 0$ and V is the subjective cost of an objective effort level e . ρ_i is a participant-specific parameter that characterizes how an individual subjectively represents the effort level e . In this formulation, ρ is flexible enough to capture increasing, decreasing, or constant marginal changes in subjective effort as absolute effort levels increase (Figure 2B). The case where $\rho = 1$ indicates that a participant's subjective effort cost representation is proportional to absolute effort levels;

$\rho < 1$ indicates decreasing sensitivity to changes in subjective effort cost as effort level increases; $\rho > 1$ indicates increasing sensitivity to changes in subjective effort cost as effort level increases.

Using the behavioral data, we performed a maximum likelihood estimation procedure to characterize each participant's subjectivity of effort cost ρ and underlying consistency of choice τ . We found that participants exhibited mean parameter estimates of $\rho = 1.15$ (S.D. 0.27), $\tau = 0.23$ (S.D. 0.28). A parameter recovery procedure recovered reliable estimates of these two preference parameters, indicating that we acquired robust estimates of subjective effort cost and choice consistency (see Methods for details). A likelihood ratio test statistic indicated that the majority of participants ($n = 17$) made choices that were inconsistent with a linear subjective effort function ($\rho = 1$), and the group exhibited subjectivity parameters that were significantly greater than 1 ($t_{23} = 2.61$, $p < 0.01$). Critically, participants' effort subjectivity parameters were not significantly correlated with MVC ($\beta = -0.31$, $p = 0.15$), or measures of monetary subjective value (risk aversion: $\beta = -0.05$, $p = 0.83$; loss aversion: $\beta = -0.01$, $p = 0.98$; for details see Methods). This suggests that individuals' subjective preferences for reward and maximum strength did not have a significant impact on their perception of effort, parameterized by ρ . Together, these results reveal that participants do not make effort decisions purely based on an objective valuation of effort, and that the majority of participants instead exhibited subjectivity of effort such that larger effort levels yielded increased marginal effort costs (Figure 2C).

Neural Representations of Subjective Effort Cost

To test our neural hypothesis that subjective effort should be encoded in vmPFC, we estimated a general linear model (GLM) in SPM12 of the blood-oxygenation level dependent (BOLD) activity of the whole-brain during the choice phase. This model included a parametric modulator at the time of choice, corresponding to the decision value of the sure option relative to the gamble, $DV_{sure}(G, S) = 0.5(-G)^\rho - (-S)^\rho$. This decision value was defined by transforming the effort options using the effort subjectivity parameter ρ estimated from each individual participant's behavior (see Methods for details). This formulation allowed us isolate brain regions that encoded subjective effort cost at the time of choice.

We found that BOLD signal in vmPFC was significantly correlated with the decision value of the sure option (Figures 3A and 3B). As the decision value increased, vmPFC activity significantly increased suggesting that this region encoded the difference in subjective effort cost of the two options. The areas of vmPFC identified largely overlap those found in studies of appetitive and aversive valuation (Bartra et al., 2013).

To further investigate the components of the subjective effort cost decision, we estimated a separate GLM in which we modeled the subjective utility of the gamble and sure options as separate unorthogonalized parametric modulators at the time of choice. Consistent with the activations in vmPFC related to the decision value of the sure option, we found that activity in this region was negatively correlated with the subjective utility of the gamble option and positively correlated with that of the sure option (Figure 3C). These results

suggest that vmPFC integrates subjective effort costs of both options to facilitate decisions regarding prospective effort expenditure.

To confirm that activity in vmPFC during effort choices was best described by representing options subjectively as opposed to objectively ($\rho = 1$), we generated exceedance probability maps (EPMs) for the imaging model described above as well as a null model representing objective effort valuation ($DV_{sure}^{objective} = 0.5(-G) - (-S)$) (Rosa et al., 2010). Using these probabilistic brain maps we were able to evaluate the likelihood that areas of vmPFC better represented subjective effort costs as opposed to objective effort costs. We found a cluster of voxels in vmPFC (34 voxels, $P > 0.95$) overlapping those identified in our main imaging analysis, illustrating that vmPFC activity is best described by a subjective rather than objective model of effort costs (Figure 3D).

DISCUSSION

We used a novel effort choice paradigm in which participants made decisions regarding prospective effort exertion, independent of reward, and found that on average individual's subjective effort cost exhibited increasing marginal costs of effort. Furthermore, during choice, we found that activity in vmPFC was related to the subjectivity of the effort options as well as their difference, suggesting that vmPFC encodes the subjective costs that underlie choices involving physical effort.

Our results provide insight into an open question in behavioral neuroscience: how are effort/action costs computed and represented by the brain (Rangel and Hare, 2010;

Padoa-Schioppa, 2011; O'Doherty, 2011)? Previous studies of effort cost have focused on the trade-offs between prospective effort and reward, similar to the natural choices we make in everyday life (Croxson et al., 2009; Prévost et al., 2010; Kurniawan et al., 2010), and have suggested that in these contexts ACC encodes effort cost. However these studies were not designed to separate effort costs from reward and instead focused on the integration of both of these utilities to compute a decision. In our paradigm, however, we isolated effort irrespective of reward in order to provide a computational description of how subjective effort cost is encoded in the brain. We found that activity in vmPFC represents subjective effort cost consistent with the idea that this region encodes a general subjective value signal that subserves effort decisions, similar to the reward based decision values that have been previously reported (Chib et al., 2009; Kable and Glimcher, 2007; Tom et al., 2007). Notably, we did not find activity in ACC in our fMRI analysis of subjective effort cost. It has been proposed that the ACC activity found in previous studies could be indicative of a multiplexing node that combines action/reward values (Hayden and Platt, 2010; Shenhav et al., 2013; Klein-Flügge et al., 2016) and serves as a gateway that informs the motor system to act for reward (Cai and Padoa-Schioppa, 2012). Thus our negative result in ACC could be due to the fact that our study was designed specifically to isolate neural representations of subjective effort cost that were *independent* of (and not multiplexed with) reward. Another possibility is that ACC could encode effort in an absolute rather than a subjective sense. Consistent with this idea we found that ACC better represented objective effort costs as opposed to subjective effort costs at the time of decision (exceedance probability map reported in Figure S3). In this framework ACC could serve as a gateway between the decision and motor systems,

representing motoric parameters in an absolute sense (e.g., applied or anticipated force, metabolic cost of effort) that serve as an input for the subjective effort cost representations instantiated in vmPFC.

In this study we focused on characterizing the subjective cost of physical effort in the form of grip force. However, an individual's subjective effort costs could vary across types of effort (i.e., walking, arm movements, or even cognitive effort) in a similar fashion to how individuals exhibit different subjective values for different types of goods (Chib et al., 2009; Levy and Glimcher, 2011). Moreover, just as the subjective value of rewards can be modulated by the state of an individual, subjective costs of effort could also be influenced by state. For example, individuals having undergone physical or cognitive fatigue or training might exhibit modified representations of effort cost. Furthermore, it is possible that the subjectivity of different types of effort may exhibit similar trait-like consistency over time, as has been reported in studies of subjective valuation of money (Kable and Glimcher, 2007; Ohmura et al., 2006; Ballard and Knutson, 2009).

Characterization of subjective effort costs will provide an understanding of why some people find certain tasks to be very effortful while others complete them with ease. Such knowledge could be used to design incentive mechanisms that account for perceptions of effortfulness to maximize employees' performance. Insights into these preferences may aid in the development of more efficacious individual-specific behavioral mechanisms that enhance motivational output and effort exertion in a variety of everyday tasks.

EXPERIMENTAL PROCEDURES

Experimental Setup

Presentation of visual stimulus and acquisition of behavioral data were accomplished using custom MATLAB (<http://www.mathworks.com>) scripts implementing the PsychToolBox libraries (Brainard, 1997). During functional magnetic resonance imaging (fMRI), visual feedback was presented via a projector positioned at the back of the room. Participants viewed a reflection of the projector in a mirror attached to the scanner head coil.

An MRI compatible hand clench dynamometer (TSD121B-MRI, BIOPAC Systems, Inc., Goleta, CA) was used to record grip force effort. During experiments, signals from this sensor were sent to our custom designed software for visual real-time feedback of participants' effort exertion. Effort exertion was performed while participants held the force transducer in their right hand with arm extended while lying in the supine position.

To record participants' choices we used an MRI compatible multiple button-press response box (Cedrus RB-830, Cedrus Corp., San Pedro, CA).

Experiment Participants

All participants were right handed, and were prescreened to exclude those with prior history of neurological or psychiatric illness. The Johns Hopkins School of Medicine Institutional Review Board approved this study, and all participants gave informed consent.

Twenty-eight healthy participants participated in the experiment (mean age, 22 years; age range, 18-31 years; 13 females). Of these participants, a total of five were excluded because they did not generate salient associations between effort levels and applied effort ($n = 2$), and/or they had highly imprecise estimates of effort cost parameters ($n = 4$).

Effort Paradigm

Prior to the experiment, participants were informed that they would receive a fixed show-up fee of \$30. It was made clear that this fee did not, in any way, depend on performance or behavior over the course of the experiment.

The experiment began with acquiring participants' maximum voluntary contraction (MVC) by selecting the maximum force achieved over the course of three consecutive repetitions on the hand clench dynamometer. During these repetitions participants did not have knowledge about the subsequent experimental phases, and were instructed and verbally encouraged to squeeze with their maximum force.

Next, participants performed an association phase in which they were trained to associate effort levels (defined relative to MVC) with the force they exerted against the hand dynamometer. The effort levels ranged from 0 to 100, corresponding to no exertion and a force equal to 80% of a participant's MVC, respectively. A single training block consisted of five trials of training for each target level, where the targets varied from 10 to 80 in

increments of 10 and training blocks were presented in a randomized order. A single trial of a training block began with the numeric display of the target effort level (2 s), followed by an effort task with visual-feedback in the form of a black vertical bar, similar in design to a thermometer, which filled white the harder participants gripped the dynamometer (4 s). The bottom and top of this effort gauge represented effort levels 0 and 100, respectively. Participants were instructed to reach the target zone (defined as ± 5 effort levels of the target number and visually represented as a rectangle alongside the effort gauge) as fast as possible and maintain their force within the target zone for as long as possible over the course of 4 seconds. Visual indication of the target zone was colored green if the effort produced was within the target zone, and red otherwise. At the end of the effort task, if individuals were within the target zone for more than 2/3 of the total frames observed during squeezing, the trial was counted a success. To minimize participants' fatigue, a fixation cross (2-5 s) separated the trials within a training block and 60 seconds of quiet rest were provided between training blocks.

Following the association phase, we performed a recall phase to test if participants successfully developed an association between the effort levels and the effort exerted. Participants were tested on each of the previously trained effort levels (10 to 80, increments of 10), six times per level, presented in a random order. Each recall trial consisted of the display of a black horizontal bar that participants were instructed to completely fill by gripping the transducer – turning the force-feedback from red to green once the target effort level was reached. For the recall phase, the full bar did not correspond to effort level 100 as in the previous phase, but instead was representative of

the target effort level being tested in a particular trial. Participants were instructed to reach the target zone as fast as possible, to maintain their produced force as long as possible, and to get a sense of what effort level they were gripping during exertion (4 s). Following this exertion task, participants were presented a number line (from 0 to 100) and told to select the effort level they believed they had just gripped. Selection was accomplished by using two push-buttons to move a cursor left and right along the number line, and a third button to enter their believed effort level. Participants had a limited amount of time to make this effort assessment (4 s), and if no effort level was selected within the allotted time the trial was considered missed. No feedback was given to participants as to the accuracy of their selection.

Finally, during the choice phase of the experiment we scanned participants' brains with fMRI while they were presented with a series of effort gambles and the choices from these gambles were used to characterize the subjectivity of individuals' effort costs. A single effort gamble consisted of choosing between two options shown on the screen under a time constraint (4 s): one option promised exerting a low amount of force (S) with absolute certainty (known as the "sure" option); whereas the other entailed taking a risk which could result in either high exertion (G) or no exertion, with equal probability (known as the "flip" option). The effort levels were presented using the 0 to 100 scale that participants were trained on during the association phase. Participants made their choice by pressing one of two buttons on a hand-held button-box with their right and with either their first or second digits. Gambles were not resolved following choice. Effort gambles (170 in total) were presented consecutively in a randomized order.

At the end of the choice phase, the computer selected 10 of the trials at random to be implemented. The outcomes of the selected trials, and only those trials, were implemented. In this way, participants did not have to worry about spreading their effort exertion over all of their trials. Instead, they were able to treat each individual trial as if it were the only decision that counted.

Monetary Subjective Value Task

After the effort paradigm, participants were endowed with \$25 and were presented with a set of mixed monetary gambles ; Sokol-Hessner et al., 2012; Chib et al., 2014) for the purpose of estimating subjectivity parameters for monetary rewards – risk aversion and loss aversion. This monetary task was conducted independent of the main effort experiment.

Effort Choice Values

The effort amounts were chosen to accommodate a range of effort sensitivity from decreasing to increasing marginal sensitivity as effort level increases. If we consider for a single effort gamble the ratio Γ of the worst possible outcome (choosing flip and having to exert that effort level) to the sure option, a set of gambles can be generated based on this ratio. If the sure value is more than half the flip ($\Gamma < 2$), participants will choose to flip because it will have the lower expected exertion of the two options. Varying Γ over a range associated with increasing marginal utility, and applying this criterion to a set of sure values (which at the maximal Γ will yield a flip amount of 100 or less), a set of effort options can be generated. Since we reasoned that participants would primarily exhibit

increasing marginal utility for effort we chose a range of $\Gamma \in [1.75, 2.75]$, as $\Gamma < 2$ represents decisions where the flip option is objectively more advantageous than the sure option. We reasoned that participant's point of indifference would lie within a range of gambles in which the sure option was objectively better than the flip option, and variability in the point of indifference between participants would occur as the result of differences in subjective preferences for effort. In our gamble set, sure values ranged from 5-35 in increments of 3.25 and were multiplied by the variable Γ to generate in total 100 unique effort gambles, all with effort levels below 100. To span a more complete set of Γ , additional gambles were designed in a similar method described above, except that the ratio of flip to sure was halved and then once more multiplied by the sure values ($\Gamma_{1/2} \in [0.88, 1.38]$). 30 of these 100 additional effort gambles resulted in trivial (G,S) pairings with the flip values less than sure values ($\Gamma < 1$), and were thus excluded. The end result was 170 unique effort gambles. The complete set of these gambles is shown in Supplementary Table 1.

MRI Protocol

A 3 Tesla Philips Achieva Quasar X-series MRI scanner and radio frequency coil was used for all the MR scanning sessions. High resolution structural images were collected using a standard MPRAGE pulse sequence, providing full brain coverage at a resolution of 3 mm x 3 mm x 2 mm. Functional images were collected at an angle of 30° from the anterior commissure-posterior commissure (AC-PC) axis, which reduced signal dropout in the orbitofrontal cortex (Deichmann et al., 2003). Forty-eight slices were acquired at a resolution of 3 mm x 3 mm x 2 mm, providing whole brain coverage. An echo-planar

imaging (FE EPI) pulse sequence was used (TR = 2800 ms, TE = 30 ms, FOV = 240, flip angle = 70°).

Data Analysis

We used a 2 parameter model to estimate participants' decisions during the choice phase, and their underlying subjective effort cost functions. We performed a parametric analysis via a stochastic choice model to estimate participants' subjective effort sensitivity. Using a form similar to that used to model monetary utility (Tversky and Kahneman, 1992), we represented a participant's cost function $V(e)$ for effort e as a power function of the form:

$$V(e) = -(-e)^\rho, \quad e \leq 0$$

In this definition of effort cost, the effort level e is defined as negative, with the interpretation being that force production is perceived as a loss. The parameter ρ represents sensitivity to changes in subjective effort cost as the effort level increases. A large ρ represents a higher sensitivity to increases in absolute effort level. $\rho = 1$ means that there are no changes in sensitivity with increases in effort level – absolute effort levels are objectively represented.

Representing the effort levels as prospective costs, and assuming participants combine probabilities and utilities linearly the decision value of choosing the sure option can be written as follows:

$$DV_{sure}(G, S) = Value(sure) - Value(gamble)$$

$$DV_{sure}(G, S) = -(-S)^\rho - (-0.5(-G)^\rho)$$

$$DV_{sure}(G, S) = 0.5(-G)^\rho - (-S)^\rho$$

This we take to be the formal definition of the decision value of the sure option.

We then assume that the probability that a participant chooses the sure option for the k^{th} trial is given by the softmax function:

$$P_k(G, S) = 1/[1 + \exp(-\tau DV(G, S))]$$

where τ is a temperature parameter representing the stochasticity of a participant's choice ($\tau = 0$ corresponds to random choice).

We used maximum likelihood to estimate parameters rho and tau for each participant, using 170 trials of effort choices (G, S) with participant response $y \in \{0, 1\}$. Here, $y = 1$ indicates that the participant chose to forego the gamble in favor of the sure option. This estimation was performed by maximizing the likelihood function:

$$\sum_{k=1}^{170} y_i \log(P_k(G, S)) + (1 - y_i) \log(1 - P_k(G, S))$$

The significance of the ρ parameter was obtained by performing a likelihood ratio test compared to a null model ($\rho = 1$ which implies that subjective and objective effort costs coincide).

As described in a number of previous studies (Sokol-Hessner et al., 2009; Frydman et al., 2011; Sokol-Hessner et al., 2012; Chib et al., 2014), a separate maximum likelihood procedure was used to estimate subjectivity parameters for monetary reward (risk aversion: $\alpha = 0.81$ (0.30), loss aversion: $\lambda = 1.69$ (1.32)) using the choice data obtained from the monetary choice task. Of the twenty-three included participants in the effort gamble experiment, four were excluded from this monetary analysis on the basis of inconsistent choices ($n = 2$) and parameter estimates outside two standard deviations from the mean ($n = 2$). Exclusion from the analysis on these grounds was independent of the effort gamble experiment, and visa-versa.

Parameter Recovery Procedure

We performed a parameter recovery procedure to evaluate the effectiveness of our parameter extraction. For this procedure we injected noise into each participant's decision process (as modeled by the subjective effort cost model) to create a group of hypothetical participants based on each single participant, and generated distributions of these parameters. Using these distributions we were able to determine if our initially estimated model parameters, for a given participant, fell within the range of those estimated for our hypothetical group. If they did, this would indicate that we were able to recover reliable

estimates of subjective effort preferences using our set of gambles and estimation procedure.

We began the parameter recovery procedure by passing each participant's model parameters (ρ and τ) through the softmax function defined above, for the full set of 170 effort gambles. This operation resulted in the probability of accepting the sure option P_k , for each effort gamble, and a new set of choices were simulated by comparing these probabilities to a random variable X_k , drawn from a uniform distribution on the interval (0,1). For a given effort gamble a choice C_k is simulated using the following rule:

$$\forall k \begin{cases} C_k = 1, & \text{if } X_k < P_k \\ C_k = 0, & \text{otherwise} \end{cases}$$

A value of $C_k = 1$ indicates rejection of an effort gamble in favor of the sure option. In following this rule, a higher probability of choosing the sure option, using the behaviorally estimated parameters, yielded a greater chance of a simulated acceptance of the sure bet. This procedure resulted in a set of simulated choices for the effort gambles (essentially the choices of a hypothetical participant). We repeated this procedure 1000 times for each participant's model parameters to generate a set of 1000 hypothetical participants corresponding to each actual participant. Estimating subjective effort preference parameters for each hypothetical participant allowed us to generate distributions of model parameters. We used these distributions to assess how well we were able to recover parameters that matched our initially estimated parameters.

Image Processing and fMRI Statistical Analysis

The SPM12 software package was used to analyze the fMRI data (Wellcome Trust Centre for Neuroimaging, Institute of Neurology; London, UK). A slice-timing correction was applied to the functional images to adjust for the fact that different slices within each image were acquired at slightly different time-points. Images were corrected for participant motion, spatially transformed to match a standard echo-planar imaging template brain, and smoothed using a 3D Gaussian kernel (8 mm FWHM) to account for anatomical differences between participants.

To examine regions of the brain that encode participants' subjective effort costs, we estimated participant-specific (first level) general linear models (GLM) for the effort choice phase of the experiment. This GLM included an event based condition corresponding to the appearance of the effort choice screen, and parametric modulators corresponding to a participant's decision value of the effort choice $DV(G,S)$ presented, and the participant's choice between the effort gamble and the sure effort option. Trials with missing responses were modeled as separate nuisance regressors. In addition, regressors modeling the head motion as derived from the affine part of the realignment procedure were included in the model.

Using this model we were able to test brain areas in which activity was related to participants' decision value of choosing an effort gamble, and their underlying subjective effort cost representations. This was done by creating contrasts with the aforementioned parametric modulator for decision value at the time of gamble presentation.

We created a separate GLM to investigate how the vmPFC separately represents the gamble and sure effort options. This GLM included an event based condition corresponding to the appearance of the effort choice screen, and separate unorthogonalized parametric modulators for the utility of the risky G^p and S^p safe effort options. Regressors modeling the head motion as derived from the affine part of the realignment procedure were also included in this model. With this model we were able to test how vmPFC actively represented the individual components of the effort choice options.

We analyzed the vmPFC activity reported in Figure 3 within an *a priori* region of interest (ROI) defined from an extensive meta-analysis of studies examining decision values of appetitive and aversive stimuli (5 mm radius sphere centered at Montreal Neurological Institute coordinates [2,46,-8]) (Bartra et al., 2013). The statistics reported in Figure 3A were small volume family-wise error corrected $p < 0.05$. To create the bar plots in Figure 3B, 3C we regressed our design matrix on a representative time course in this ROI, calculated as the first eigenvariate.

Bayesian Model Selection

To determine if a subjective effort cost representation better accounted for neural activity in vmPFC than an objective representation, we performed a Bayesian model selection analysis (Rosa et al., 2010). We began by creating an additional GLM that was identical to our original expected decision value model $DV(G,S)$, except in this model the

parametric modulator corresponded to the difference in expected objective value of the effort options presented $DV_{sure}^{objective} = 0.5(-G) - (-S)$. This GLM captured the null choice model (objective valuation of effort; $\rho = 1$).

We used the first level Bayesian estimation procedure in SPM12 to compute voxel-wise whole-brain log-model evidence maps for every participant and each model. To model inference at the group level we applied a random effects approach at every voxel of the log evidence data within an anatomical mask of vmPFC. We used this data to create exceedance probability maps (EPM) that allowed us to test which representation of effort cost, subjective or objective, was more likely to describe activity in vmPFC. The EPMs shown in Figure 3D and Figure S3 were generated using anatomical masks for vmPFC and ACC, respectively.

ACKNOWLEDGEMENTS

We thank D. McNamee and V. Stuphorn for their insightful comments on this study. This work was supported by the Eunice Kennedy Shriver National Institute Of Child Health & Human Development of the National Institutes of Health under Award Number K12HD073945 to V.S.C. J.G. was supported by the National Defense Science and Engineering Graduate Fellowship. The authors declare no competing financial interests.

REFERENCES

- Abeler, J., Falk, A., Goette, L., and Huffman, D. (2011). Reference Points and Effort Provision. *Am. Econ. Rev.* 101, 470–492.
- Augenblick, N., Niederle, M., and Sprenger, C. (2015). Working Over Time: Dynamic Inconsistency in Real Effort Tasks*. *Q. J. Econ.* qjv020.
- Ballard, K., and Knutson, B. (2009). Dissociable neural representations of future reward magnitude and delay during temporal discounting. *NeuroImage* 45, 143–150.
- Bartra, O., McGuire, J.T., and Kable, J.W. (2013). The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage* 76, 412–427.
- Brainard, D. (1997). The Psychophysics Toolbox. *Spat. Vis.* 433–436.
- Cai, X., and Padoa-Schioppa, C. (2012). Neuronal Encoding of Subjective Value in Dorsal and Ventral Anterior Cingulate Cortex. *J. Neurosci.* 32, 3791–3808.
- Chib, V.S., Rangel, A., Shimojo, S., and O'Doherty, J.P. (2009). Evidence for a Common Representation of Decision Values for Dissimilar Goods in Human Ventromedial Prefrontal Cortex. *J. Neurosci.* 29, 12315–12320.
- Chib, V.S., Shimojo, S., and O'Doherty, J.P. (2014). The Effects of Incentive Framing on Performance Decrements for Large Monetary Outcomes: Behavioral and Neural Mechanisms. *J. Neurosci.* 34, 14833–14844.
- Croxson, P.L., Walton, M.E., O'Reilly, J.X., Behrens, T.E.J., and Rushworth, M.F.S. (2009). Effort-Based Cost–Benefit Valuation and the Human Brain. *J. Neurosci.* 29, 4531–4541.
- Deichmann, R., Gottfried, J.A., Hutton, C., and Turner, R. (2003). Optimized EPI for fMRI studies of the orbitofrontal cortex. *Neuroimage* 19, 430–441.
- DellaVigna, S., and Paserman, M.D. (2005). Job Search and Impatience. *J. Labor Econ.* 23, 527–588.
- Dishman, R.K. (1991). Increasing and maintaining exercise and physical activity. *Behav. Ther.* 22, 345–378.
- Frydman, C., Camerer, C., Bossaerts, P., and Rangel, A. (2011). MAOA-L carriers are better at making optimal financial decisions under risk. *Proc. R. Soc. Lond. B Biol. Sci.* 278, 2053–2059.
- Hayden, B.Y., and Platt, M.L. (2010). Neurons in Anterior Cingulate Cortex Multiplex Information about Reward and Action. *J. Neurosci.* 30, 3339–3346.

- Hillman, K.L., and Bilkey, D.K. (2012). Neural encoding of competitive effort in the anterior cingulate cortex. *Nat. Neurosci.* *15*, 1290–1297.
- Holt, C.A., and Laury, S.K. (2002). Risk Aversion and Incentive Effects. *Am. Econ. Rev.* *92*, 1644–1655.
- Kable, J.W., and Glimcher, P.W. (2007). The neural correlates of subjective value during intertemporal choice. *Nat. Neurosci.* *10*, 1625–1633.
- Kahneman, D., and Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica* *47*, 263–291.
- Klein-Flügge, M.C., Kennerley, S.W., Friston, K., and Bestmann, S. (2016). Neural Signatures of Value Comparison in Human Cingulate Cortex during Decisions Requiring an Effort-Reward Trade-off. *J. Neurosci.* *36*, 10002–10015.
- Kurniawan, I.T., Seymour, B., Talmi, D., Yoshida, W., Chater, N., and Dolan, R.J. (2010). Choosing to Make an Effort: The Role of Striatum in Signaling Physical Effort of a Chosen Action. *J. Neurophysiol.* *104*, 313–321.
- Levy, D.J., and Glimcher, P.W. (2011). Comparing Apples and Oranges: Using Reward-Specific and Reward-General Subjective Value Representation in the Brain. *J. Neurosci.* *31*, 14693–14707.
- O’Doherty, J.P. (2011). Contributions of the ventromedial prefrontal cortex to goal-directed action selection. *Ann. N. Y. Acad. Sci.* *1239*, 118–129.
- Ohmura, Y., Takahashi, T., Kitamura, N., and Wehr, P. (2006). Three-month stability of delay and probability discounting measures. *Exp. Clin. Psychopharmacol.* *14*, 318–328.
- Padoa-Schioppa, C. (2011). Neurobiology of economic choice: a good-based model. *Annu. Rev. Neurosci.* *34*, 333–359.
- Padoa-Schioppa, C., and Assad, J.A. (2008). The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nat. Neurosci.* *11*, 95–102.
- Plassmann, H., O’Doherty, J., and Rangel, A. (2007). Orbitofrontal Cortex Encodes Willingness to Pay in Everyday Economic Transactions. *J. Neurosci.* *27*, 9984–9988.
- Prévost, C., Pessiglione, M., Météreau, E., Cléry-Melin, M.-L., and Dreher, J.-C. (2010). Separate Valuation Subsystems for Delay and Effort Decision Costs. *J. Neurosci.* *30*, 14080–14090.
- Rangel, A., and Hare, T. (2010). Neural computations associated with goal-directed choice. *Curr. Opin. Neurobiol.* *20*, 262–270.
- Rosa, M.J., Bestmann, S., Harrison, L., and Penny, W. (2010). Bayesian model selection maps for group studies. *NeuroImage* *49*, 217–224.

Rudebeck, P.H., Behrens, T.E., Kennerley, S.W., Baxter, M.G., Buckley, M.J., Walton, M.E., and Rushworth, M.F.S. (2008). Frontal Cortex Subregions Play Distinct Roles in Choices between Actions and Stimuli. *J. Neurosci.* 28, 13775–13785.

Rushworth, M.F., Mars, R.B., and Summerfield, C. (2009). General mechanisms for making decisions? *Curr. Opin. Neurobiol.* 19, 75–83.

Shenhav, A., Botvinick, M.M., and Cohen, J.D. (2013). The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron* 79, 217–240.

Sniehotta, F.F., Scholz, U., and Schwarzer, R. (2005). Bridging the intention–behaviour gap: Planning, self-efficacy, and action control in the adoption and maintenance of physical exercise. *Psychol. Health* 20, 143–160.

Sokol-Hessner, P., Hsu, M., Curley, N.G., Delgado, M.R., Camerer, C.F., and Phelps, E.A. (2009). Thinking like a trader selectively reduces individuals' loss aversion. *Proc. Natl. Acad. Sci.* 106, 5035–5040.

Sokol-Hessner, P., Camerer, C.F., and Phelps, E.A. (2012). Emotion regulation reduces loss aversion and decreases amygdala responses to losses. *Soc. Cogn. Affect. Neurosci.* nss002.

Tom, S.M., Fox, C.R., Trepel, C., and Poldrack, R.A. (2007). The Neural Basis of Loss Aversion in Decision-Making Under Risk. *Science* 315, 515–518.

Tversky, A., and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *J. Risk Uncertain.* 5, 297–323.

Walton, M.E., Bannerman, D.M., and Rushworth, M.F.S. (2002). The Role of Rat Medial Frontal Cortex in Effort-Based Decision Making. *J. Neurosci.* 22, 10996–11003.

FIGURE LEGENDS

Figure 1. Experimental paradigm. (A) Association phase; participants were trained to associate numeric effort levels with force exerted on a hand-clench dynamometer. Numeric effort levels ranged from 0 (no force) to 100 (80% of maximum grip force). A training block consisted of five trials at each target level of effort. Each trial began with presentation of the target effort level, followed by an effortful grip with real-time visual feedback of the exerted force represented as a bar that increased height with increased exertion. A green visual cue was also displayed within which participants were instructed to maintain their exerted effort. Feedback of success or failure was provided at the end of each trial. (B) Recall phase; participants were instructed to fill a horizontal bar by gripping the transducer. In this phase, on each trial, the full bar corresponded to a different target effort level that was unknown to participants. Successfully achieving the effort target resulted in the bar turning from red to green. Following exertion, participants used push-buttons to move a cursor along a 0-100 number line to select the effort level they believed they had squeezed. No feedback was provided as to the accuracy of participants' reported effort levels. (C) During the choice phase participants were presented a series of risky gambles which involved choosing between two options under a time constraint: exerting a low amount of effort with certainty ("Sure"), or taking a gamble that could result in either a higher level of exertion or no exertion with equal probability ("Flip"). Gambles were not realized following choice. At the end of the choices phase, to ensure participants revealed their true preferences for effort, 10 choices were randomly selected and played out such that any effort required would need to be exerted before they left the experiment.

Figure 2. Behavioral representations of subjective effort cost. (A) Recall results showing the mean and standard error across all participants of the reported effort levels plotted against the values tested. The dashed line indicates perfect recall of exerted effort. (B) The function used to model the subjective cost of effort in a choice. This function has the form $V(e) = -(-e)^\rho$. Each curve represents the absolute value of an individual's hypothetical cost function for effort. The dashed line represents objective valuation of effort ($\rho = 1$), with curves above this line representing that an incremental change in the effort level results in a greater subjective cost of that effort for higher effort levels. (C) Estimated ρ coefficients at the participant level. Asterisks indicate a significant difference ($p < 0.05$) from the null case of objective valuation ($\rho = 1$) using a likelihood ratio test statistic.

Figure 3. vmPFC encodes subjective effort cost. (A) A region of vmPFC in which activity was positively correlated with the decision value of the sure option at the time of choice (Montreal Neurological Institute (MNI) coordinates $(x, y, z) = [0, 52, -2]$). (B) Beta values in vmPFC were positively correlated with the decision value of the sure option. (C) ROI analyses illustrating vmPFC activity related to the value of the gamble and value of the sure options at the time of choice. (D) Exceedance probability map (EPM) resulting from the Bayesian model comparison of objective and subjective models. This map indicates, within an anatomical mask of vmPFC, that subjective effort cost best describes activity in this region.

Figure 1

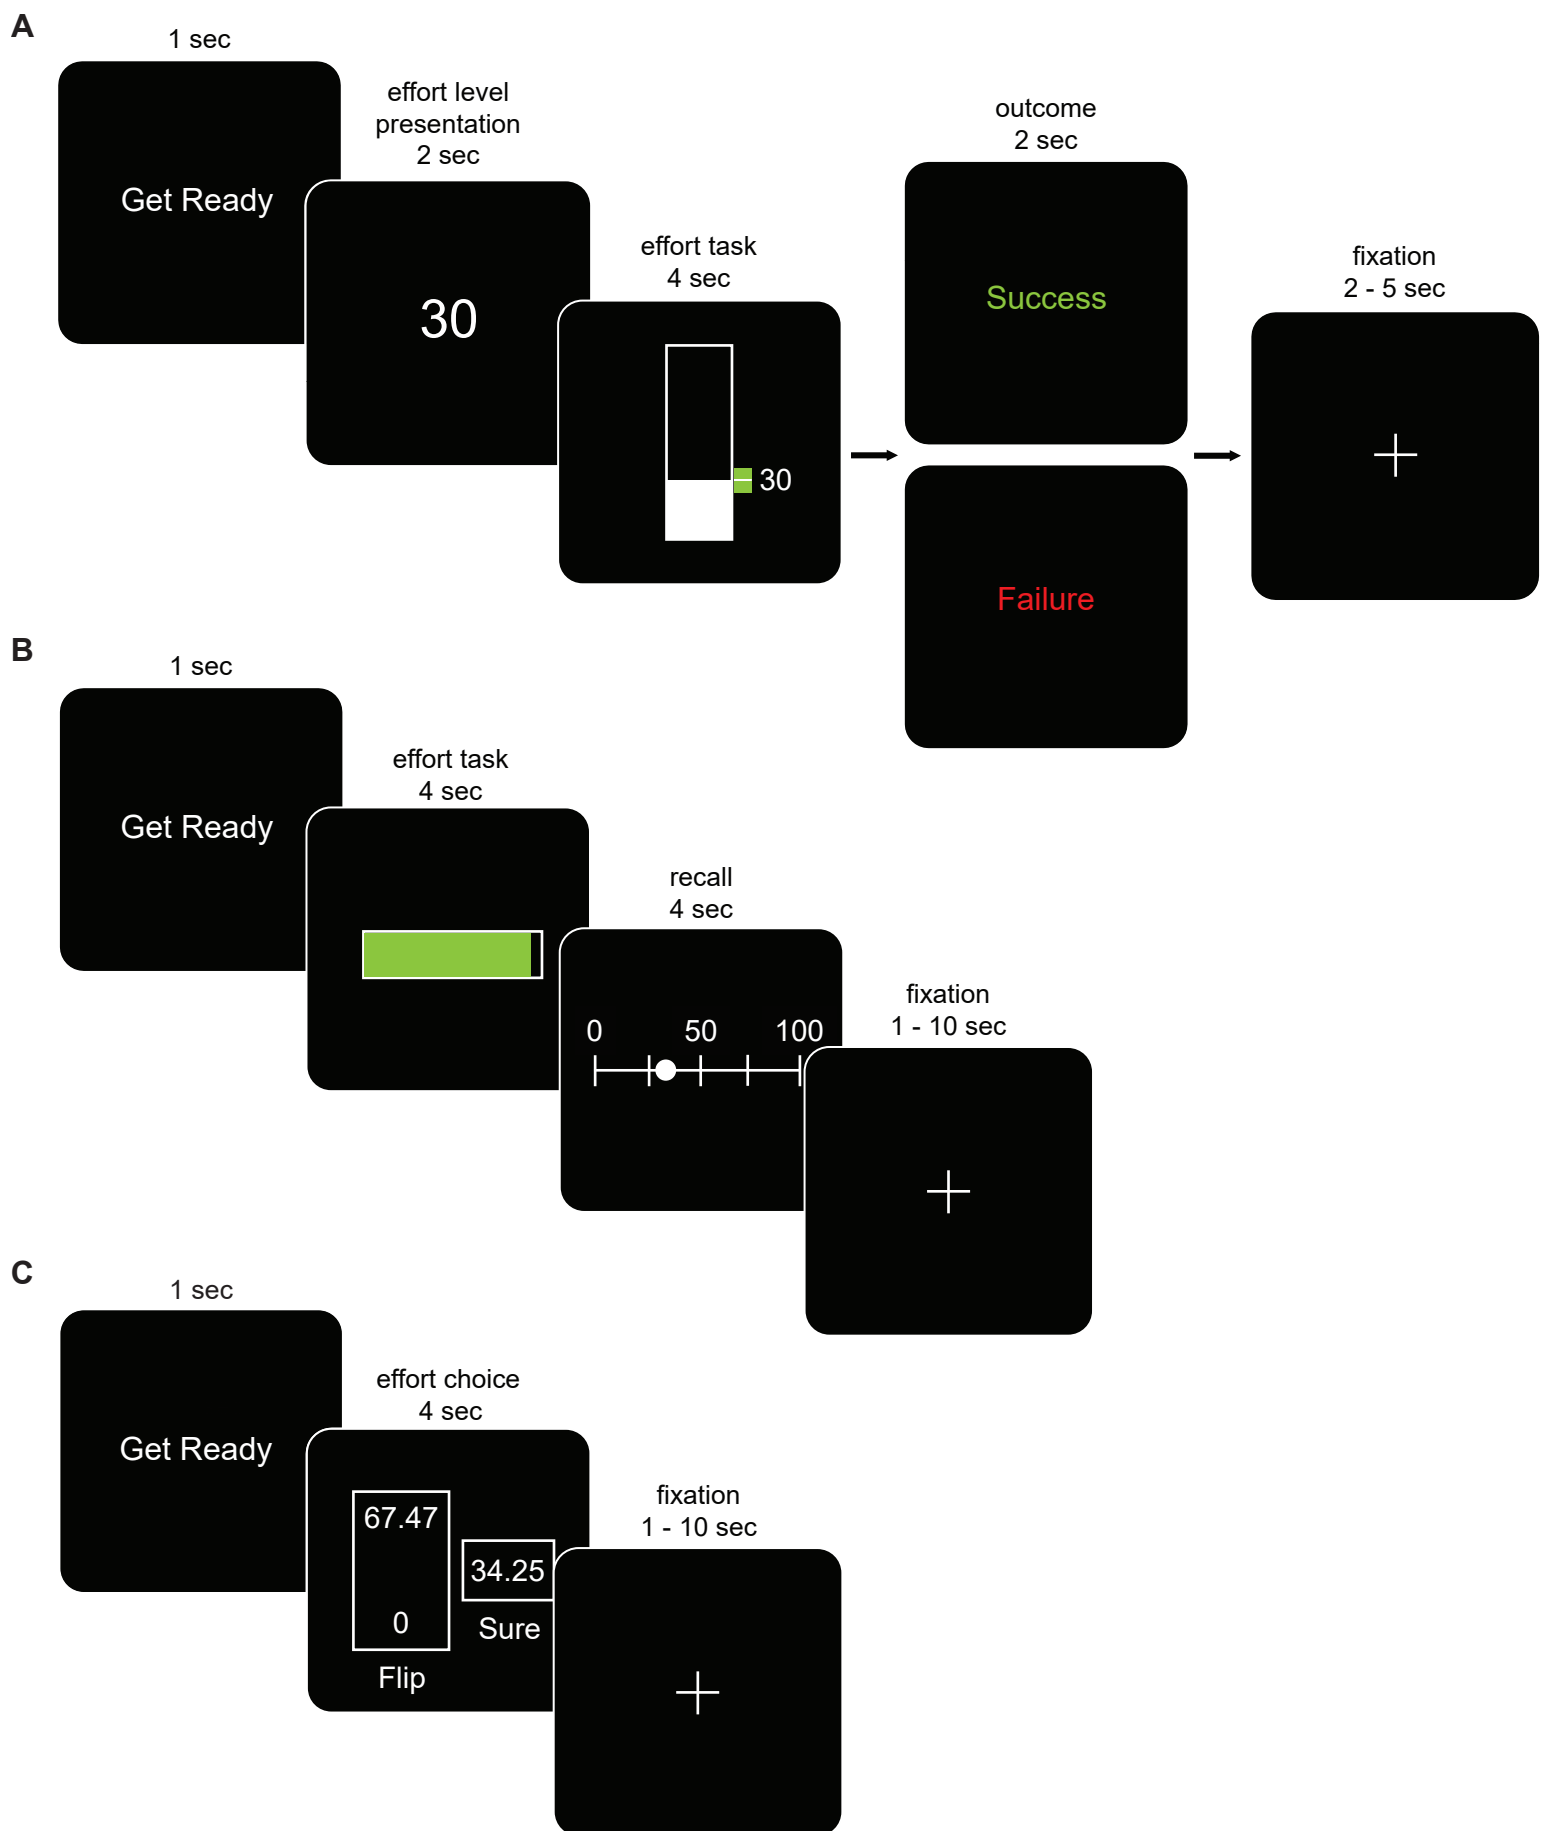


Figure 2

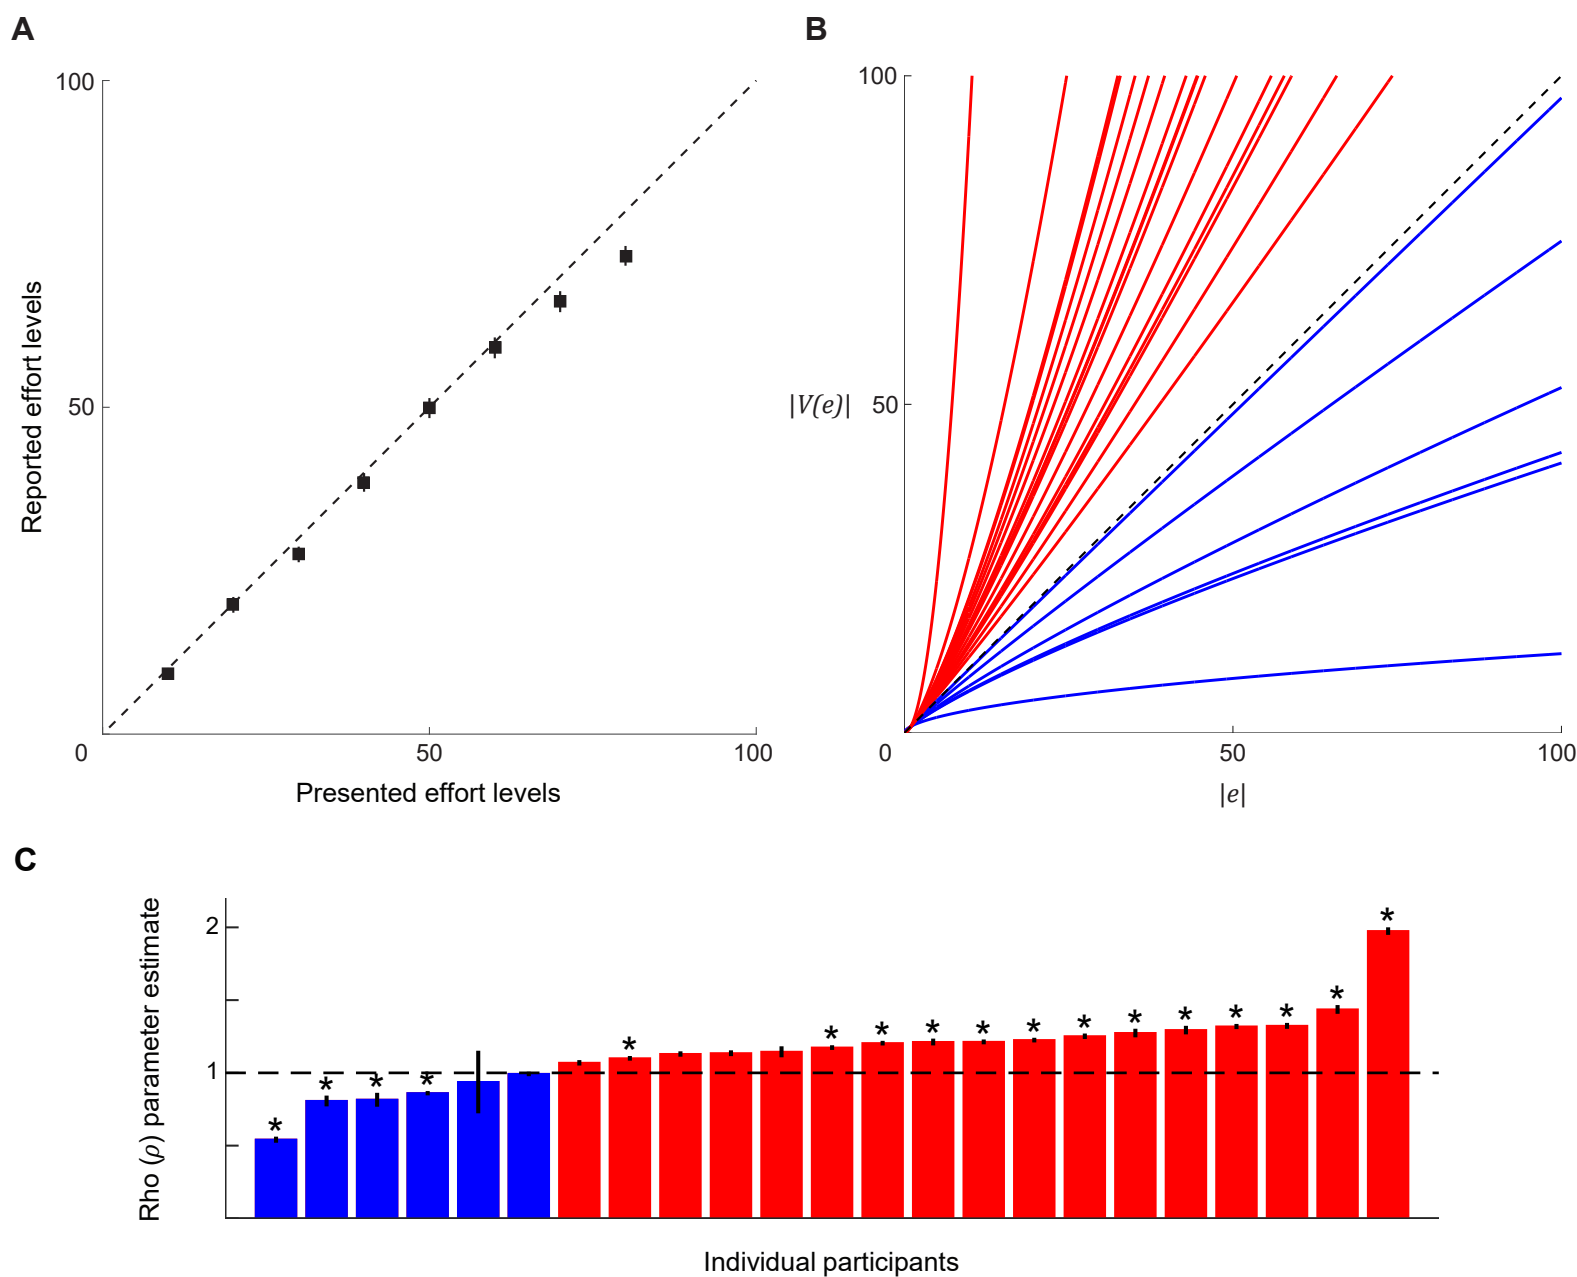


Figure 3

