UNIVERSITY OF CALIFORNIA

Santa Barbara

Neural Signatures of Subjective Value and Confidence in Decisions with Different Physical Costs

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy
in Psychological and Brain Sciences

by

Allison D. Shapiro

Committee in charge:
Professor Scott T. Grafton, Chair
Professor Michael Gazzaniga
Professor Barry Giesbrecht
Professor Jonathan Schooler

June 2019

The dissertation of Allison D. Shapiro is approved.

_____
Michael Gazzaniga

_____
Barry Giesbrecht

_____
Jonathan Schooler

_____
Scott T. Grafton, Committee Chair

June 2019

Neural Signatures of Subjective Value and Confidence in

Decisions with Different Physical Costs

ACKNOWLEDGMENTS

I would like to thank the members of my committee: Scott Grafton, Barry Giesbrecht, Jonathan Schooler, and Mike Gazzaniga for their support and encouragement to pursue the research questions that inspired me, and for flexibly adapting with my evolving path. They have been exceptionally generous with both their wisdom and kindness. My advisor Scott, in particular, bestowed upon me tremendous intellectual freedom, sage guidance when I most needed it, and what has amounted to a full encyclopedia of obscure trivia – delivered incrementally over the past six years.

My lab mates, collaborators, and research assistants have been absolutely invaluable. When I arrived to UCSB, Deborah Barany, Matt Cieslak, Taraz Lee, and Shiva Viswanathan took me under their wings and I was fortunate to later work alongside Michelle Marneweck, Gold Okafor, Viktoriya Babenko, Neil Dundon, James Elliott, and Lukas Volz. Thanks also to Tom Bullock, Evan Layher, Tyler Santander, and Nikki Marinsek for their generous help and collaborations. This research would not have been possible without Mario Mendoza nor my outstanding research assistants including Sydney Smith, Trinity Novak, Kayli Berlin, Julia Hsueh, Shefali Verma, Cepideh Razavi, and Alex Asturias.

Finally, thank you Spencer, Moose, Mom, Dad, and the rest of SHWAMMMIJJEET for everything else, including your wholehearted indifference to whether I became a doctor or a knitter or a permanent camper.

VITA OF ALLISON D. SHAPIRO
JUNE 2019

EDUCATION

Bachelor of Arts in Cognitive Science, University of California, Berkeley, May 2010 (honors)
Doctor of Philosophy in Psychological and Brain Sciences, University of California, Santa Barbara, June 2019 (expected)

PROFESSIONAL EMPLOYMENT

2016, 2017, 2018: Teaching Associate, Department of Psychological and Brain Sciences, University of California, Santa Barbara
2014 – 2019: Graduate Student Researcher, Department of Psychological and Brain Sciences, University of California, Santa Barbara
2013-2018: Teaching Assistant, Department of Psychological and Brain Sciences, University of California, Santa Barbara

PUBLICATIONS

Barany, D. A., Shapiro, A.D., Lee, T.G. (2015). Multivariate fMRI Approaches to Flexible Sensorimotor Maps in Parietal Cortex. *Journal of Neuroscience*, 35(34), 11763-11765.
Buxbaum, L. J., Shapiro, A. D., & Coslett, H. B. (2014). Reply: Apraxia: a gestural or a cognitive disorder? *Brain,* 138(3), e334-e334
Buxbaum, L. J., Shapiro, A. D., & Coslett, H. B. (2014). Critical brain regions for tool-related and imitative actions: a componential analysis. *Brain*, 137(7) 1971-1985.
Kalénine, S., Shapiro, A. D., Flumini, A., Borghi, A. M., & Buxbaum, L. J. (2014). Visual context modulates potentiation of grasp types during semantic object categorization. *Psychonomic Bulletin & Review*, 21(3), 645-651.
Kalénine, S., Shapiro, A. D., & Buxbaum, L. J. (2013). Dissociations of action means and outcome processing in left-hemisphere stroke. *Neuropsychologia*, 51(7), 1224-1233.

AWARDS

Graduate Student Fellow, Kavli Summer Institute in Cognitive Neuroscience, 2017
Graduate Affiliate Fellow, Interdisciplinary Humanities Center, UCSB, 2015 – 2016
Graduate Research Fellowship Honorable Mention, National Science Foundation, 2015

ABSTRACT

Neural Signatures of Subjective Value and Confidence in

Decisions with Different Physical Costs

By

Allison D. Shapiro

Two fundamental goals of decision making are to select actions that maximize rewards while minimizing costs and to have strong confidence in the accuracy of a judgment. Neural signatures of these two forms of value: the subjective value (SV) of choice alternatives and the value of the judgment (confidence), have been observed in ventromedial prefrontal cortex (vmPFC). However, the relationship between these dual value signals and their relative time courses are unknown. Furthermore, there are often several ways that one can obtain a desired reward, each entailing its own costs. Previous research investigated the neural representations of many types of decision costs with the notable exclusion one evolutionarily significant expense: cardiovascular effort. It remains unclear how physical costs, such as pain and exertion are transformed into value information to guide decision making and the extent to which this process is domain-specific versus domain-general. To test this, in Experiments 1 and 2 we recorded fMRI while participants performed a two-phase ApAv task with mixed-outcomes of monetary rewards paired with significant physical costs. In Experiment 1 costs were painful shock stimuli, in Experiment 2 costs were intervals of demanding cardiovascular exercise on

a stationary bike. Neural responses were measured during offer valuation (offer phase) and choice valuation (commit phase) and analyzed with respect to observed decision outcomes, model-estimated SV, confidence, and costs. During the offer phase, vmPFC tracked SV and decision outcomes, but it not confidence. During the commit phase, vmPFC tracked confidence, computed as the quadratic extension of SV, but it bore no significant relationship with the offer valuation itself, nor the decision. In fact, vmPFC responses from the commit phase were selective for confidence even for rejected offers, wherein confidence and SV were inversely related. Conversely, activation of the cognitive control network, including within lateral prefrontal cortex (lPFC) and dorsal anterior cingulate cortex (dACC) was associated with ambivalence, during both the offer and commit phases. None of these signals differed in magnitude between the pain and effort conditions, consistent with shared neural faculties of cost representations. However, inspection of voxel response patterns within those faculties revealed representational dissimilarity between the two types of decision costs, indicating that cost-domain-specific information is encoded throughout the decision making process. Within vmPFC in particular, the availability of cost-specific information over the course of a decision suggested qualitatively different decision processes under conditions of confidence and ambivalence. Taken together, our results reveal complementary representations in vmPFC during value-based decision making that temporally dissociate such that offer valuation (SV) emerges before decision valuation (confidence), both of which include specific information about costly consequences.

Table of Contents

Table of Contents

Chapter I: Introduction

Theoretical Overview

A fundamental goal for survival is accruing the maximum amount of appetitive resources while minimizing the aversive outcomes of those pursuits. The pressure to optimally arbitrate such cost/benefit tradeoffs has resulted in remarkable capacities for extracting value information from the environment and applying complex decision policies (Cosmides and Tooby, 1994; Montague and King-Casas, 2007). Humans, who evolved as foragers, often faced choices between approaching and avoiding opportunities to obtain resources such as food items (Hayden, 2018). Approach decisions hold the potential for rewards but come at the cost of energetic expenses and the threat of physical harm. Avoidance decisions are less risky in the short term, but increase one's reliance on the environment to afford future opportunities.

Approach-Avoidance (ApAv) decisions can be simulated in a laboratory setting. In these tasks, participants decide whether to accept (approach) or reject (avoid) individual offers of appetitive rewards paired with contingent aversive costs (Miller, 1944). Given the relevance of ApAv choices to evolved decision making capacities in humans, we selected this paradigm to measure behavioral and neural responses during value-based decision making. This enabled us to enforce true rewards and costs of physical pain and cardiovascular exercise that invoke authentic choice scenarios for the participant.

ApAv offers are one type of multi-attribute mixed-outcome (or conflict) choice. Such choices entail multiple consequences (multi-attribute) with opposing values (mixed-outcome) that induce competing objectives (conflict) to pursue rewards, avoid costs, and minimize risks (Keeney and Raiffa, 1993). No single choice satisfies all goals, instead one must make tradeoffs between them (Tversky and Shafir, 1992). The overall utility of the offer accounts

for both its positive and negative attributes. There are vast literatures of economics research specifically aimed at modeling the precise parameters in value computations revealed by human choice behavior (Coombs and Avrunin, 1977; Kahneman and Tversky, 1979; Dyer et al., 1992). Generally speaking, the utility of an offer combines weighted estimates of its reward and cost components (Neumann and Morgenstern, 1944).

The relatively recent advent of neuroeconomics brought together the fields of economics, psychology, and neuroscience, resulting in a flood of groundbreaking discoveries about the biology underlying value-based decisions (Rangel et al., 2008). Now, we are beginning to connect the dots between a widely distributed value network in the human brain, which includes ventromedial prefrontal cortex (vmPFC) and orbitofrontal cortex (OFC), the amygdala, anterior cingulate cortex and particularly its more dorsal segments (ACC / dACC), posterior cingulate cortex, anterior and posterior segments of the insula (aIC and pIC), the basal ganglia and in particular the nucleus accumbens (nAcc) (O'Doherty, 2004; Behrens et al., 2007; Rushworth and Behrens, 2008; Bartra et al., 2013)

Neural representations of rewards, reward predictions, and hedonic experiences more generally have been observed in medial prefrontal cortex (including vmPFC, OFC, and subgenual ACC [sgACC]), nAcc, the striatum, and midbrain (Ernst et al., 2004; O'Doherty, 2004; Kringelbach, 2005; Kahnt et al., 2011a). Conversely, aversive stimuli, anticipated negative outcomes, and costs elicit responses in the amygdala (Becerra et al., 2001; Yacubian, 2006), ACC (Prevost et al., 2010; Massar et al., 2015), and throughout the insula (Preuschoff et al., 2008; Kurniawan et al., 2010; Prevost et al., 2010; Kurniawan et al., 2013; Arulpragasam et al., 2018). Some nodes within the value network integrate multiple sources of value information. vmPFC, OFC and the striatum seem to encode appetitive and aversive cues

together (Brooks et al., 2010; Plassmann et al., 2010). dACC is recruited when competing value signals induce response conflict (Pochon et al., 2008; Aupperle et al., 2015), consistent with this region's more general role in managing response competition in difficult decisions (Kerns et al., 2004; Botvinick, 2007).

There are often several ways that one might reach a certain goal, and different paths promise varying types and degrees of costs. Neuroeconomic studies are beginning to show how different types of costs are accounted in value-based choices, both at the behavioral and neural levels. Previous empirical work has posed a wide array of cost stimuli in behavioral, neuroimaging, and lesion studies, including delays of rewards (Frederick et al., 2002; Kable and Glimcher, 2007), probabilistic risk (Kahneman and Tversky, 1979; Preuschoff et al., 2008; Boorman et al., 2009), pain stimuli (Park et al., 2011; Talmi, 2012), cognitive effort (Botvinick et al., 2009; Westbrook et al., 2019), physical effort in motor tasks (Treadway et al., 2012; Arulpragasam et al., 2018), and in manual grip force (Croxson et al., 2009; Kurniawan et al., 2013; Klein-Flügge et al., 2015; Hogan et al., 2018).

Premises for the Present Research

Despite the wealth of research on decision costs, the field is starkly lacking examination of one evolutionarily critical type of decision cost, cardiovascular exercise. To our knowledge, value-based choices with costs of aerobic exercise have never been implemented in a laboratory setting. Furthermore, there is little consensus on one fundamental question: are value representations domain-specific or are all types of costs and rewards represented in a domain-general manner, potentially with a common currency?

After accounting for the patterns in decision behavior that can be explained by evolutionary pressures and shared cognitive and neural capacities, there are still wide margins in individual preferences. Such individual differences have been observed in humans and non-human primates alike (Premack, 1963; Schunk and Betsch, 2006). We intuitively experience that many of our decisions are made on the basis of subjective judgments, in line with our internal beliefs and core values. This isn't entirely illusory. Every individual has a unique personal history that shapes their preferences, needs, and future goals. The relative worth of different types of rewards and the aversiveness of different types of costs depend critically on these factors. In this way, the overall value perceived in a given opportunity is, in fact, unique to each decision maker. Subjective value (SV) measures the *perceived* utility of objective value information, relative to the decision maker.

vmPFC, in particular, plays a crucial role in representing SV. This region is understood to guide decisions by integrating reward and cost attributes of choice alternatives (Basten et al., 2010; Amemori and Graybiel, 2012; Levy and Glimcher, 2012; Talmi and Pine, 2012) and automatically track the SV of items in our environment (Lebreton et al., 2009; Pessiglione and Delgado, 2015). In ApAv decision making tasks, vmPFC encodes multiple value representations including rewards, decision variables, SV, and valuation models that inform SV (Talmi et al., 2009; Park et al., 2011; Skvortsova et al., 2014; Wan et al., 2015).

Recently, it has been suggested that confidence, the extent to which one can believe that they are making the best choice, is a form of valuation in its own right (Lebreton et al., 2015). Accurate decisions are valuable decisions, and in this sense, the confidence with which one makes a decision may measure the value of that judgment. There is recent evidence that vmPFC also signals choice confidence. For example, vmPFC activation varies with confidence

about perceptual judgments (Heereman et al., 2015; Bang and Fleming, 2018; Gherman and Philiastides, 2018) and lesions lead to atypical confidence reports on general knowledge tests (Hebscher et al., 2016). These findings could be evidence that both SV and confidence emerge from shared neural computations, however only one study has tested this hypothesis directly (Lebreton et al., 2015). Moreover, it remains unclear how vmPFC manages dual signals pertaining to SV and confidence over the course of a decision. Are SV and confidence integrated into a common signal, temporally dissociated, or something else entirely?

The present work investigates behavioral and neural signatures of SV and confidence in value-based decision making with different types of costs. We recorded functional magnetic resonance imaging (fMRI) while participants decided to accept or reject real offers that entailed enduring aversive costs to receive monetary rewards. In Experiment 1 (Chapter 2), participants were offered monetary rewards in exchange for receiving painful electrical shocks. We combined neuroeconomic models of decision behavior with neuroimaging to identify when, where, and how neural activity tracks model-estimated SV and confidence. In Experiment 2 (Chapter 3), participants were offered monetary rewards in exchange for performing intervals of demanding cardiovascular exercise. We applied the same analyses from Experiment 1 on the new data collected in Experiment 2 to test whether observations about decisions involving pain costs generalized to decisions involving effort costs. Then, to test for domain-specific processing of SV and confidence with respect to decision costs, we contrasted the results from Experiments 1 and 2. In Chapter 4, data from Experiments 1 and 2 were re-analyzed with a multivariate approach to measure the representational dissimilarity (differences between patterns of activation) of effort and pain costs in voxel response patterns. In Chapter 5, we

measured individual differences the relative valuation of pain and effort costs by comparing both behavioral and neural data from the two tasks at the individual level.

Notes on Terminology

Note that in the present work, we examine multiple types of value. Offer stimuli comprise two value attributes: costs and rewards. We use "value" and "value stimuli" to refer not only to reward attributes (which have positive value), but also to cost attributes (which have negative value) and to the entire offer stimulus (which has an overall worth or SV) – all of these carry some type of value information. Moreover, in the present research, we varied the domain of cost attributes, which could influence behavioral and neural responses not only to the cost itself, but also to the reward and the overall valuation of the offer. We use "domain-specific" to refer to characteristics of behaviors and neural representations that were only observed in response to offers with either pain costs or effort costs, but not both, and can pertain to the rewards, costs, SV, or confidence of those offers. We use "domain-general" to refer to behaviors and neural representations observed across both the bike and shock task, including representations of rewards, costs, SV, and confidence. Furthermore, we occasionally use the term "common currency" to describe that value information from different domains (e.g. pain and effort) may be evaluated or represented similarly. The literature often links this term with the integration of costs and rewards into an overall utility, which is one subtype of common currency. We use this term more generally and often in reference to domain-general processing of different types of costs (Levy and Glimcher, 2012).

Chapter II: Experiment 1: Value and Confidence in Decisions with Pain Costs

Every day, we navigate a maze of choices, guided by our subjective preferences and goals. When the route forks, multiple forms of value imbue the selection of one's path. We assess the value of potential actions from their relative costs and benefits, as well as the value of our own judgment – accurate decisions are valuable decisions, regardless of the options.

Much progress has been made toward understanding how the brain resolves value-based decisions (Rangel et al., 2008; Bartra et al., 2013) and vmPFC's important role in this process. As described in Chapter 1, vmPFC is understood to integrate reward and cost of choice alternatives (Basten et al., 2010; Amemori and Graybiel, 2012; Levy and Glimcher, 2012; Talmi and Pine, 2012) and track SV (Lebreton et al., 2009; Pessiglione and Delgado, 2015). Neuroeconomic models benefit from mixed-outcome ApAv tasks, which pose realistic, consequential choice scenarios (e.g. accept or reject offers of appetitive rewards contingent on aversive costs). Such fMRI studies in humans demonstrate multiple value representations in vmPFC including rewards, decision variables, SV, and valuation models that inform SV (Talmi et al., 2009; Park et al., 2011; Skvortsova et al., 2014; Wan et al., 2015).

Recent research suggests that vmPFC also signals choice confidence. BOLD responses in vmPFC correlate with confidence about perceptual judgments (Heereman et al., 2015; Bang and Fleming, 2018; Gherman and Philiastides, 2018) and impair self-reported confidence ratings (Hebscher et al., 2016) These findings could be evidence that vmPFC performs a valuation of one's judgment, assigning high value to high confidence decisions (Lebreton et al., 2015). Confidence signals in vmPFC also accompany value-based decisions(Rolls et al., 2010; De Martino et al., 2013; Lebreton et al., 2015), which is intriguing given the relationship

between confidence and value. Self-reported confidence takes a U-shaped function with respect to first-order valuation judgments (Lebreton et al., 2009; Hebscher and Gilboa, 2016): decisions about extremely high- or low-value items elicit stronger confidence than decisions about items with neutral or ambiguous value. Similarly, vmPFC responses take a U-shaped function with respect to value in risky decision making (Schlund et al., 2016). Accordingly, Lebreton et al. (2015) operationalized confidence as the quadratic extension of value and elegantly demonstrated that vmPFC tracks modeled confidence, even in the absence of explicit ratings.

Given this literature, vmPFC should track both SV and its quadratic extension, confidence, in mixed-outcome ApAv decision making, but this has not been explicitly tested. Specifically, confidence about accept choices (typically positive SV) should increase as SV increases whereas confidence about reject choices (typically negative SV) should increase as SV decreases (Fig. 1B). It is unknown how vmPFC represents both SV and confidence in value-based choices, particularly when confidence and value are inversely related (i.e. reject choices). One possibility is that confidence evolves in parallel with decision variables (Dotan et al., 2018). Early confidence-related signals have been recorded from frontal and parietal sites (Kepecs et al., 2008; Kiani et al., 2014), including vmPFC (De Martino et al., 2013; Lebreton et al., 2015; Gherman and Philiastides, 2018). Alternatively, confidence may evolve later through retrospective metacognitive judgments or continued deliberation after choice commitment (Resulaj et al., 2009; Pleskac and Busemeyer, 2010; Moran et al., 2015; Yu et al., 2015; van den Berg et al., 2016), both of which can recruit medial prefrontal cortex (Hilgenstock et al., 2014; Fleming et al., 2018; Morales et al., 2018).

To test this, we deployed a two-phase ApAv task in which participants accepted or rejected offers of monetary rewards paired with painful shock stimuli. BOLD responses were measured during offer valuation (offer phase) and choice valuation (commit phase). Observed decision outcomes and model-based estimates of SV and confidence were used to predict neural activity in vmPFC and elsewhere during both phases of decision making.

Methods

Participants

We report the data from 28 paid volunteers that participated in the study (17 women, 27 right-handed, mean age = 21.9, sd=2.8). One additional participant completed the study but was removed from analyses due to significant susceptibility artifacts causing excessive errors in spatial normalization. No participants had a history of neurological injuries or illnesses or current daily use of psychoactive medications. All participants provided written consent in accordance with the Institutional Review Board at the University of California, Santa Barbara.

Session overview

All testing was performed on the same day. Participants first provided informed consent and were screened for disqualifying criteria. Because our experiment ApAv decisions with pain stimuli, participants next underwent a pain thresholding procedure and were familiarized with the mapping between experimental stimuli and the pain intensities they represented. Then, participants performed the decision making task while in the MRI scanner. Finally, after they had completed all experimental tasks and been removed from the scanner, participants received the pain stimuli and monetary rewards associated with the choices they made during the task.

Pain thresholding procedure

We used mild cutaneous electrical shocks as pain stimuli. Pain thresholding allowed us to control for individual differences in pain tolerance for electrical shocks such that a given experimental stimulus was associated with the same pain sensation across participants. To identify participant-specific minimum and maximum shock intensities, participants completed a pain thresholding procedure before the decision-making task, outside of the MRI scanner. Electrical shocks were administered with a constant current stimulator (Digitimer DS7A, Digitimer, Great Britain) controlled by a train generator (DG2A Train/Delay Generator, Digitimer, Great Britain). Each shock had a duration of 1 s and a frequency of 100 Hz with a 2 ms waveform. Two adhesive electrodes were placed on the back of the participant's hand approximately 1 inch above the wrist and connected to the stimulator. When a shock was administered, electric current was run between the two electrodes, causing an aversive sensation that is increasingly painful at higher levels of current.

The first stage of thresholding was a ramp-up procedure in which the experimenter delivered several shocks, each time increasing the intensity by 1 mV. The participant was instructed to report three thresholds: the lowest intensity at which they detected the shock, the lowest intensity at which the shock caused discomfort, and the intensity at which the shock became unbearably painful. The second stage was a rating procedure in which the participant received fourteen shocks with intensities between the discomfort threshold and unbearable threshold and rated the pain from each on a 0-10 scale. Next, both the ramping and rating stages were repeated to verify that we had accurately identified the participant's pain tolerance. Finally, the pain ratings and shock intensities from the second rating procedure were fit with a sigmoid function to model the relationship between shock intensity and perceived pain for this

individual (Volz et al., 2017). Based on this sigmoid function, we identified the shock intensity that predicted a pain rating of 8 out of 10. No payout trials during the decision task had costs exceeding this value.

Familiarization Procedure

Before the decision making task, participants were shown five example offer stimuli illustrating costs of 5%, 25%, 50%, 75% and 95% while the experimenter delivered a shock at the corresponding intensity relative to their discomfort and maximum pain thresholds (with 0% = minimal discomfort and 100% = maximal pain). Participants were instructed to remember the sensation associated with the example shocks and use these as points of reference when making choices during the task, but that the real offers would include shocks ranging anywhere between 0 and 100% intensity, not just at the example levels.

ApAv Task

Offer stimuli

We adapted an approach–avoidance task similar that implemented in non-human primates by Amemori and Graybiel (2012; 2015) in which participants accepted or rejected offers with mixed appetitive and aversive outcomes. On each trial, the participant was offered a certain amount of money in exchange for receiving a shock of a certain intensity. The participant chose either to accept *both* the money and the shock or reject both. Therefore, trial stimuli were deterministic mixed-outcome offers as there were no probabilistic manipulations involving chances of receiving the shock and/or money - receiving a monetary reward was contingent on also receiving its associated cost.

The offer stimuli were two horizontal bars, one had an overlaid $ symbol and illustrated the offered reward and the other had an overlaid lightning bolt symbol and illustrated the contingent cost. The width of the bars represented the amount of each attribute being offered on the current trial, with monetary rewards ranging continuously from $0.01 to $1.50 and shock intensities ranging continuously from minimal discomfort to maximal pain. Each bar was within a larger rectangular frame that illustrated the maximum possible bar width. The bars were blue and yellow, and color-attribute mappings varied between participants. The relative position of the cost and reward bars (i.e. which bar was above the other) alternated between blocks. The offer stimuli were centrally aligned and stacked just above and just below the vertical midpoint of the display. Overall, there were 189 offers, each presented as a single trial, unless the participant failed to respond, in which case the offer was repeated at the end of the experiment. Trials were split between six functional runs, with each run containing 31 or 32 trials. All participants viewed the same set of offers, presented in the same pseudo-random order that systematically covered all quadrants of the decision space.

Decision making task

On each trial, the participant first saw a fixation point in the center of the screen for either 910 ms, 1820 ms or 2730 ms, randomly varied between trials. Then, the offer stimulus was shown for 4550 ms. The participant was instructed to use this time to evaluate the offer but could not yet indicate a choice. Next, the offer remained on the screen and two response mappings appeared in left and right positions beneath the offer stimulus for 1820 ms. During this time, the participant was required to respond with a button press, indicating their commitment to either accept (approach) or reject (avoid) the offer. All stimuli remained on

the screen for the remainder of the response interval, and after the response was submitted, the response mapping corresponding to their choice was highlighted. Participants were allowed to change their response within the 1820 ms response interval. The response mappings were an upward pointing triangle representing the accept option and a downward pointing triangle representing the reject option. The participant used either the index or middle finger of their right hand to press either the left or right (respectively) button on a Cedrus LP-RH response pad transmitting through a Lumina LSC-400 controller (Lumina, Cedrus Corportation, San Pedro, CA, USA), according to the location of their preferred choice option. The left/right positions of the choice options varied from trial to trial, preventing the participant from pre-planning a motor response before the response phase. Finally, a 910 ms feedback interval followed the response interval. Feedback included the offer stimulus, which remained on the display, and text stating whether the offer was accepted or rejected. On payout trials, a small subset of all trials, an additional 9100 ms payout alert followed the decision feedback.

Because we were interested in value-processing during different phases of decision making, our analyses separately analyzed neural responses from the offer phase and the commit phase of each trial. The offer phase was the 4550 ms in which the offer was displayed but the participant was not yet able to submit a response. The commit phase was the 2730 ms comprising the response interval (1820 ms) and the feedback interval (910 ms).

Payouts

Participants did not receive the rewards and shocks for all accepted trials. Instead, during trial generation, 10 pseudorandomly selected offers were tagged as payout trials. On these trials, if the participant accepted the offer the payout alert indicated that they would

receive both the monetary reward and the shock. If they rejected the offer, the payout alert indicated they would receive neither. The payout alert also showed the cumulative number of payouts accepted so far. Participants were encouraged to assume that every trial was a payout offer and decide to accept or reject it accordingly, as they would not see the payout alert until after they'd submitted their choice. Although participants were notified whether an offer was a payout immediately after choosing to accept or reject it, the actual delivery of monetary rewards and shocks from payout trials occurred after they had completed the entire decision making task and had been removed from the MRI scanner. In pilot studies, we found that providing realtime payout feedback, relative to showing which trials were payouts at the end of a fixed-length block or at the end of the task, reduced the proportion of offers that were accepted. We believe that this may indicate that instant feedback reduces temporal discounting of future costs, relative to when that information is delayed, even though in both cases the time that the costs were actually delivered was the same (after completing the entire task).

All participants also performed an additional cost-benefit decision making task after they completed the approach-avoidance task, not reported here. Twenty-four of the participants also underwent simultaneous physiological recordings of impedance cardiography while performing the decision making tasks in the MRI scanner, impedance cardiography data is not included here and will be presented elsewhere.

Behavioral statistical analysis

Estimating SV from decision behavior

The subjective value (SV) of each offer can be represented as

$$SV = \beta_0 + \beta_r r_{offer} + \beta_s s_{offer}$$

where r_{offer} is the available reward, s_{offer} is the contingent shock, β_r and β_s describe how strongly the individual subject weights rewards and shocks, and β_0 is the individual's intercept, indicating their intrinsic motivations to pursue reward versus avoid shock (Volz et al., 2017). We separately modeled each participant's choice data with logistic regression, a specialized from of the generalized linear model, using the glm package in R. Offers to which participants failed to respond before the decision deadline were repeated at the end of the task. Choice outcomes from the second presentation were included in the dataset used for generating valuation models but these trials were excluded from the fMRI analysis. Participants' model estimates were used inversely to predict the perceived subjective value of each offer, SV, and the likelihood that it would be accepted, P(Acc) specific to each individual on each trial.

$$SV = \beta_0 + \beta_r r_{offer} + \beta_s s_{offer} = \log \frac{P(Acc)}{1 - P(Acc)}$$

$$P(Acc) = \frac{1}{1 + e^{-SV}}$$

Estimating confidence from SV

Here, we operationalize choice confidence as the extent to which one can believe they are making the best decision. In our task, given the options to either accept or reject, the best choice depends on the perceived value of the offer and consequently choice confidence is maximized when SV unambiguously points to one choice over the other. Therefore, confidence varies in a U-shaped function with respect to increasing SV such that when SV is extremely high or extremely low, one can be highly confident in choices to accept or reject the offer, respectively, but when an offer's SV is neutral there is more ambivalence about the decision.

We quantify confidence in accordance with Lebreton et al. (2015), who demonstrated that decision confidence can be estimated from the quadratic extension of perceived value in the absence of explicit confidence ratings.

$$Confidence = SV^2$$

We used these trial-by-trial estimates to test whether behavioral response times and neural activity varied with respect to choice confidence, see below. Notably, because we did not collect metacognitive self-reported confidence ratings, we do not intend to make explicit claims about the subjective experience of confidence. Previous research has found mixed results regarding how closely model-estimated confidence varies with self-reported confidence, see the discussion for a further explanation. Therefore, we primarily aim to describe neural responses that vary with model-estimated confidence.

We also inspected the relationship between confidence and choice *outcomes* by categorizing trials into confident accepts (AccCon), ambivalent accepts (AccAmb), confident rejects (RejCon), and ambivalent rejects (RejAmb). Each participant's trials were binned by choice outcome and then we performed a median split on *SV* for both choice bins. Accepted offers in the upper 50% SV were assigned to AccCon and the rest were assigned to AccAmb. Rejected offers in the upper 50% SV were assigned to RejAmb, and the rest were assigned to RejCer. Notably, because we observed substantial individual differences in subjective valuation, both parameterizing and categorizing the offers according to individual-specific model-based estimates allowed us to describe objectively identical offers differently for each participant, according to their own individual preferences. This was a critical feature of our study that allowed us to observe group-wide neural responses that were specific to decision outcome, SV, and confidence, regardless of the objective properties of the stimulus.

In our paradigm, P(Acc) varies as a sigmoidal function of SV while choice confidence varies as a U-shaped function of SV. Consequently, when P(Acc) is close to either 0 or 1 there is high choice confidence in either accepting or rejecting the offer, respectively and P(Acc) is close to .5, there is high ambivalence about committing to either choice. In a two-dimensional decision space in which reward values are represented on one dimension, pain values are represented on the other, and the modeled decision boundary is the vector along which P(Acc) = .5, the choice confidence associated with a reward-pain offer pair increases with its distance from the decision boundary. We present individual choice outcomes and binned choice confidence overlaid on estimated P(Acc) throughout the decision space in figure (Fig. 2).

Neuroimaging Data Acquisition and Preprocessing

Neuroanatomical ROI

We investigated task-related involvement within an a priori region of interest (ROI), vmPFC. The ROI was anatomically defined from the Harvard-Oxford cortical structural probabilistic atlas (https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/Atlases), and includes all voxels with at least 25% likelihood of being located within areas labelled Frontal Medial Cortex or Subcallosal Cortex.

MRI protocols:

Anatomical and functional MRI data were collected on a Siemens 3T Magnetom Prisma Fit with a 64-channel phased-array head and neck coil (58 channels active for functional coronal imaging). High-resolution 0.94 mm isotropic T1- (TR=2500 ms, TE=2.2 ms, FA=7°, FOV=241 mm) and T2*-weighted (TR=3200 ms, TE=570 ms, FOV=241 mm) sagittal

sequence images were acquired of the whole brain. Next, functional MRI recordings were collected while participants performed the decision making task. For each functional run, a multiband T2*-weighted echo planar gradient-echo imaging sequence sensitive to BOLD contrast was acquired (TR=910 ms, TE=32 ms, FA=52°, FOV=192 mm, multiband factor 4) provided by the Center for Magnetic Resonance Research in accordance with a current license. Each functional image consisted of 64 coronal slices acquired perpendicular to the AC-PC plane (3 mm thick; 3x3 mm in-plane resolution). Coronal orientation is necessary when acquiring simultaneous impedance cardiography to avoid artifact (Cieslak et al., 2015).

MRI pre-processing

Anatomical data was skull-stripped using Advanced Neuroimaging Tools (ANTs) brain extraction script (Avants et al., 2011). All other image pre-processing was performed with FMRIB's Software Library (FSL, www.fmrib.ox.ac.uk/fsl). The first 10 volumes of each functional run were removed to eliminate non-equilibrium effects of magnetization occurring before the start of the task. The remaining functional volumes were skull-stripped using BET (Smith, 2002) motion corrected using MCFLIRT (Jenkinson et al., 2002), spatially smoothed using a Gaussian kernel of FWHM 5mm, intensity normalized relative to the grand-mean of the entire 4D dataset by a single multiplicative factor, and underwent high-pass temporal filtering (Gaussian-weighted least-squares straight line fitting, with sigma=50.0s).

In preparation for group analyses, participants' six functional runs were registered to their anatomical image and then to the Montreal Neurological Institute (MNI) 2mm averaged 152-brain template included with FSL distributions, using FSL's linear image registration tool with 12 degrees of freedom (FLIRT; (Jenkinson and Smith, 2001; Jenkinson et al., 2002).

Refinement of the latter transformation was carried out with FSL's nonlinear registration image registration tool (FNIRT) with a 10mm warp resolution (Andersson et al., 2007a; 2007b).

Neuroimaging Statistical Analysis

fMRI analysis

FMRI data processing was carried out using FEAT (FMRI Expert Analysis Tool) Version 6.00, part of FSL (FMRIB's Software Library, www.fmrib.ox.ac.uk/fsl). Time-series statistical analysis was carried out using FSL's improved linear model (FILM) with local autocorrelation correction (Woolrich et al., 2001). We performed whole-brain statistical analyses with two general linear models (GLMs), as described below. Both GLMs had separate terms for the offer phase (off) and the choice commitment phase (com) of each trial. Offer phase regressors were time-locked to the onset of the offer and had a duration of 4.55s, during which the participant assesses the SV of the offer but cannot yet respond. Commit phase regressors were time-locked to the onset of the response mappings and had a duration of 2.73 seconds, during which the participant submits their decision about the offer and then views feedback confirming their choice. Finally, the GLMs also included a nuisance regressor for payout notifications, which were not used in analyses of interest but were intended to absorb variance in neural responses associated with subjective value (as payout notifications included an image of the payout offer) but unrelated to decision making processes. The payout regressor was time-locked to the onset of the payout notification and had a duration of 4.55 s, and the remaining 4.55 s of the payout notification screen was included with baseline activity.

Both analyses were performed at three sequential levels. First, at the run level, each participant's six runs were separately modeled to find mean within-run activity corresponding each regressor and contrast images were generated by estimating pairwise differences between conditions. Then, at the participant level, run-level data was combined (fixed effects) to find the participant's overall mean response relating to each regressor and contrast. Finally, at the group level, the participant data was combined (mixed-effects treating participant as a random effect with FSL's FLAME 1) to find the group-wide mean responses for each regressor and contrast.

We tested the results of each contrast with non-parametric permutation testing at the whole brain level with threshold-free cluster enhancement (TFCE), implemented with FSL's Randomise. This approach minimizes false positives by deriving a null distribution from the voxelwise data rather than assuming a parametric null distribution. For one-sample t tests, the distribution is created by iteratively multiplying statistical map values by 1 or -1, we performed 5,000 permutations of each contrast. TFCE detects clusters of contiguous voxels without setting an arbitrary cutoff for minimum cluster size or voxel statistic but rather summarizes the cluster-wise evidence at each voxel, against several types of cluster-forming thresholds and controls the family-wise error (FWE) rate at p=.05 (Nichols and Holmes, 2002; Winkler et al., 2014). We present figures with voxel-wise T-values from all voxels that survived whole-brain TFCE correction. Some contrasts yielded significant voxels across contiguous but widespread regions of cortex, and consequently reporting only the peak voxel of a cluster would obscure other local maxima in different anatomical regions. Therefore, we also report the coordinates, t statistics, Brodmann area, and anatomical structure labels from Automated Anatomical Labelling the of the MNI atlas for local maxima within each cluster in Table 1. Local maxima

were found with the cluster command provided with FSL and labelled with label4MRI, a freely available toolbox for R (https://github.com/yunshiuan/label4MRI). We additionally report results from sub-conditions of GLM2 restricted to and TFCE corrected only within our primary ROI of interest, vmPFC.

GLM1: Parametric analysis SV and confidence

GLM1 was a parametric statistical analysis to observe neural activity modulated by SV and confidence (CD) during the offer and commit phases of each trial:

$$Y = \beta_{off} X_{off} + \beta_{com} X_{com} + \beta_{SVoff} X_{SVoff} + \beta_{SVcom} X_{SVcom} + \beta_{CDoff} X_{CDoff} + \beta_{CDoff} X_{CDcom} + \beta_{payout} X_{payout} + \varepsilon$$

where Y is the time series of a given voxel predicted by a design matrix with one row for each time sample and one column for each of 7 trial regressors, convolved with a canonical gamma hemodynamic response function. The model included categorical terms ($\beta_{off} X_{off}$, $\beta_{com} X_{com}$) to isolate task related activity during the offer and commit trial phases (onsets and offsets described above). Parametric terms $\beta_{SVoff} X_{SVoff}$ and $\beta_{CDoff} X_{CDoff}$ were orthogonalized with respect to $\beta_{off} X_{off}$ to capture variance in neural activity during the offer phase explained by trial-by-trial SV and confidence (for which regressors were range normalized by z-score). Likewise, $\beta_{SVcom} X_{SVcom}$ and $\beta_{CDcom} X_{CDcom}$ were orthogonalized with respect to $\beta_{com} X_{com}$ and modeled neural activity related to SV and confidence during the commit phase.

GLM2: Categorical choice by confidence analysis

GLM2 estimated categorical variance in neural responses during the offer and commit phases of each trial:

$$Y = \beta_{ACoff} X_{ACoff} + \beta_{AAoff} X_{AAoff} + \beta_{RCoff} X_{RCoff} + \beta_{RAoff} X_{RAoff} + \beta_{ACcom} X_{ACcom} + \beta_{AAcom} X_{AAcom} + \beta_{RCcom} X_{RCcom} + \beta_{RAcom} X_{RAcom} + \varepsilon$$

The first four terms modeled the offer phase of confident accepts, ambivalent accepts, confident rejects, and confident accepts, respectively. The next four terms modeled the commit phase of the same conditions. The ninth regressor modeled payout notifications.

Our contrasts compared each condition with baseline, tested main effects of choice outcome (i.e. accepted offers vs. rejected offers irrespective of choice confidence) and choice confidence (i.e. high confidence choices vs. low confidence choices irrespective of choice outcome) separately during the offer and commit phases, as well as select pairwise comparisons within these conditions (AccCon vs. AccAmb, and RejCon vs. RejAmb), which were specifically inspected within our anatomical ROI, vmPFC.

Phase-specific responses in vmPFC

To further characterize the pattern of vmPFC responses during the offer and commit phases, we extracted mean parameter estimates within vmPFC for each condition at each trial phase. We tested whether vmPFC simultaneously tracks value and confidence, or if these two signals modulate vmPFC responses during different phases of decision making. Following the logic of Lebreton et al., (2015), we assumed that the function that decision phase-specific vmPFC responses take across trial conditions of ascending value (RejCon < RejAmb < AccAmb < AccCon) would be indicative of the information being processed during that trial

phase. A linear increase of vmPFC response magnitudes would be associated with value processing, whereas a quadratic function would be associated with confidence processing. If vmPFC simultaneously tracked confidence and value, model terms for both the linear and quadratic extensions of value would be necessary to fully explain the observed pattern of vmPFC activity.

To tailor the predictors in these models to individual participants, we used the mean perceived value across trials from each condition, calculated separately for each participant. SV varies substantially between participants depending on the consistency of their choice behavior. Specifically, slight differences in participants' choice consistency lead to substantial differences between their model-estimated SV predictors. Consequently, there was a large range of SV and SV^2 throughout the group and many participants' data only spanned only a portion of that range, making it difficult to draw conclusions at the group level. P(Acc) and P(Rej) (the latter is equivalent to 1-P(Acc)) are always restricted to the range of 0 to 1. Consequently, P(Acc)-P(Rej) normalizes value to range from -1 to 1 while preserving the sign of SV, with negative values predicting reject choices, positive values predicting accept choices, and values surrounding zero indicating decision ambivalence. Conversely, normalizing SV at the individual level (e.g. z-score) would not preserve the sign of raw SV. Furthermore, whereas in GLM1 we observed parametric modulation by SV, in GLM2 trials were binned according to observed choice, which is more specifically related to P(Acc)-P(Rej). Therefore, we used P(Acc)-P(Rej) as value predictors in our ROI analysis, which improved consistency for groupwide analysis while retaining the sign of SV. Previous research has used a similar approach (Lebreton et al., 2015).

Each participant's mean vmPFC parameter estimates from the four trial conditions were predicted as a function of their mean P(Acc)-P(Rej) of that condition using linear mixed effects regression, implemented with the lme4 package for R (Bates et al., 2015).

$$\text{Linear Model: } \hat{y} = b_0 + b_1(P(Acc) - P(Rej)) + \varepsilon$$

$$\text{Quadratic Model: } \hat{y} = b_0 + b_1(P(Acc) - P(Rej)) + b_2(P(Acc) - P(Rej))^2 + \varepsilon$$

The linear and quadratic extensions of P(Acc)-P(Rej) were specified as fixed effects and we included a random effect on the model intercept across subjects to account for baseline variation in vmPFC parameter estimates. Predictor values entered into the model were participants' mean P(Acc)-P(Rej) of trials from each of the four conditions. Both models were separately fit to parameter estimates of BOLD responses in vmPFC during the offer phase and during the commit phase. We report model fits as well as the results of model comparisons. Due to ambiguity in estimating denominator degrees of freedom, linear mixed model fits are not best evaluated by p-values. However, significance can be inferred from confidence intervals constructed by iteratively sampling the model posterior to estimate the likelihood of the observed parameter estimates. We ran 5,000 simulations using the posterior distributions over each parameter from the mixed models using the merTools package for R. We report significant parameters with 95% CIs that do not span zero. For interested readers, corresponding p-values estimated with Satterthwaite's method implemented in the lmerTest package for R are also provided. Model comparison (ordinary likelihood ratio test) and relative AIC and BIC values were used to determine the best fitting model for the offer phase and for the commit phase.

Behavioral Correlates of SV and Confidence

We estimated confidence as the quadratic extension of value in accordance with a similar study that validated this operationalization by demonstrating that both response times (RTs) and self-reported confidence were related by the inverse quadratic to value, and thus RTs and confidence were negatively correlated. Notably, that study found that quadratic extension of value (i.e. model-based confidence) better predicted self-reported confidence than RTs, suggesting that while RTs were a useful behavioral correlate of subjective confidence, they didn't fully explain variance in metacognitive confidence ratings (Lebreton et al., 2009; 2015).

We aimed to measure neural correlates of implicit, naturalistic experiences of confidence during decision making. Therefore, our task did not solicit metacognitive confidence ratings. Instead, we used model-based confidence (estimated in accordance to previous literature) to predict changes in BOLD responses. Given previous findings that model-based confidence better predicted confidence ratings than RTs, we did not take RTs to be a direct proxy for the subjective experience of confidence. Nonetheless, it was important to verify that model-estimated confidence mad a meaningful relationship with behavior in our task, that is, to the time it took participants to commit to a decision. Specifically, an inverse quadratic relationship between behavioral RTs and model-estimated confidence would indicate that it took participants longer to commit to decisions that were associated with lower degrees of model-estimated confidence. To test this, we fit RTs with mixed effects regression (lme4 package for R; (Bates et al., 2015)). Fixed effects specified the linear and quadratic extensions of value and a random effect on the model intercept was included to account for

baseline variation in RT.

$$\hat{y} = b_0 + b_1(P(Acc) - P(Rej)) + b_2(P(Acc) - P(Rej))^2 + \varepsilon$$

Value predictors entered into the model were generated by sorting model-estimated value into 5 equally spaced bins and the dependent measures were participants' mean RTs for each bin. The model was tested again using z-scored SV as value predictors to verify that relationship between value and RTs were consistent regardless of the method used to estimate value.

Results

Behavioral Choice Models and RTs

A separate logistic regression model was fit to each participant's decisions to accept or reject shock/reward offers made while undergoing fMRI. All participants' model fits had significantly positive reward coefficients and significantly negative shock coefficients, indicating that both offer attributes influenced SV in the intended direction, despite individual differences in their relative contribution to choice outcomes (β_r estimates: mean= .199, range = [0.064, 0.547], p values all <.001), costs (β_s estimates: mean = -0.170, range = [-0.547, -0.050], p values all <.001). Moreover, there was great variation in participants' model intercepts with ranges that spanned zero, suggesting strong individual differences in baseline tendencies to accept or reject offers (β_0 estimates: mean = 1.057, range = [-6.420, 9.750]). On average, participants tended to accept more offers than they rejected (mean=61.3%, sd=17.8%), and seemed to be engaged in the task (98.2% of all trials received responses before

the 1.8s decision deadline). Finally, linear mixed effects regression of behavioral RTs confirmed that RTs were quadratically related to SV and negatively correlated with estimated confidence, signifying that choices with low estimated confidence indeed took longer. This relationship between RT and confidence is consistent with previous research that defined confidence in a similar form (Lebreton et al., 2015).

Neuroimaging Results

GLM 1: Parametric Modulation by SV and Confidence

In GLM 1 we measured parametric modulation of BOLD responses by continuous regressors for SV and choice confidence over the course of value-based choices. During the offer phase SV correlated significantly with activation in many regions of cortex, including a network of value-related regions incorporating vmPFC, posterior cingulate cortex, orbitofrontal cortex (OFC,) the basal ganglia, posterior insula, and hippocampus, as well as other regions known to be involved with perceptual and value comparison such angular gyrus, lateral temporal cortex, and visual cortex (Fig. 3A). A comprehensive list of significant clusters for all conditions are tabulated in Table 1.

We additionally found a relatively smaller set of regions where activity negatively correlated with confidence during the offer phase including areas associated with cognitive control and conflict resolution such as dACC and right lPFC, right superior parietal lobule (SPL), premotor regions, and visual cortex (Fig. 3B). There were no voxels with activity negatively related with SV nor any voxels that were positively related with confidence during the offer phase. Notably, neural responses to neutral SV can't be easily interpreted with the SV regressor alone. Instead, the inverse of the confidence regressor can be interpreted as a measure

of choice conflict. Both the dACC and lPFC demonstrated a strong inverse relationship with confidence, consistent with previous results demonstrating their recruitment during ambivalence, conflict, uncertainty, and choice difficulty during value based decision making and other tasks (Kahnt et al., 2011b; Badre et al., 2012; Economides et al., 2014; Shenhav et al., 2014; Aupperle et al., 2015; Lopez-Persem et al., 2016; Schlund et al., 2016; Shenhav et al., 2016) . The inverse relationship between the confidence regressor and conflict-resolution and positive relationship of the regressor with SV is consistent with the idea that neural activity during the offer phase is primarily related to valuation processes.

Surprisingly, we did not observe any BOLD responses, in vmPFC or elsewhere, that varied with SV during the commit phase. Furthermore, during the commit phase, there were no regions where activity varied negatively with choice confidence. Instead, choice confidence significantly predicted activity in a cluster of voxels within vmPFC as well as other regions including posterior insula, superior temporal cortex, and premotor areas during the commit phase (Fig. 3C).

Taken together, the key findings from GLM1 regarding vmPFC were that during the offer phase vmPFC tracks SV but not confidence, and during the commit phase vmPFC tracks confidence but not SV. These results suggest that in value-based decision making, vmPFC responses do increase with both SV and confidence, however these signals do not modulate simultaneously. Rather, there may be a more dynamic process through which vmPFC is first involved in the valuation process before transitioning to signaling confidence about that valuation. This interpretation preliminarily supports the idea that vmPFC's dual roles in value-based decision making are separable if that process is examined in sequential stages. However, while SV predicts choice outcomes (i.e. offers perceived as highly valuable are likely to be

accepted), the two are not perfectly related, especially in choices that are closer to one's decision boundary. It is possible that SV is computed relatively early on, but that vmPFC responses relating to the final choice (selected with respect to SV) occur later and concurrently with choice confidence, during the commit phase. Therefore, in GLM2 we measured responses related to choice outcomes and confidence.

GLM 2: Categorical Choice by Confidence Conditions

In GLM2 we measured BOLD responses from four trial conditions: AccCon, RejCon, AccAmb, and RejAmb. These analyses mirror GLM1 except that in GLM2, value is signified as trial by trial choice outcome rather than estimated SV and confidence has been binned into high (confident) and low (ambivalent) categories. During the offer phase, we observed that a similar network of regions to those that were parametrically modulated by SV in GLM1 also showed significant contrasts for decision outcome in GLM2 such that BOLD responses were stronger preceding decisions to accept than decisions to reject (Fig. 4A). This augments our findings from GLM1, demonstrating that not only SV but also a decision variable is represented within vmPFC relatively early in the decision making process, in this case during the offer phase.

There was substantial overlap between value-related responses in GLM1 and decisions to accept in GLM2. These similarities are likely attributable to participants' relatively stable choice behavior, causing SV model estimates to be strongly predictive of choice outcomes. Consequently, neural responses associated with high estimated SV in GLM1 were also categorized as objective accept choices in GLM2. The primary exception to this rule was visual cortex, where activity tracked SV but not choice outcome, suggesting that visual cortex

activation varied with SV only insofar as it systematically related to visual properties of the offer stimuli.

Similarly, the contrast of BOLD responses from the offer phase of ambivalent versus confident choices revealed a similar network of regions as those that were negatively related to confidence during the offer phase from GLM1, with the addition of the anterior insula for GLM2 (Fig. 4B). No regions, including vmPFC, showed a stronger response during confident choices than ambivalent choices during the offer phase, nor were any regions more responsive during rejected trials than accepted trials.

During the commit phase, there were highly significant and widespread differences between confident and ambivalent choices, but no regions responded preferentially to one choice outcome (accept or reject) over the other. Critically, clusters of voxels in vmPFC as well as posterior cingulate, and an adjacent medial segment of superior parietal lobule showed a significant contrast between confident and ambivalent decisions during the commit phase, with stronger BOLD responses to confident choices (Fig. 4C, purple). The reverse contrast (ambivalent versus confident) revealed several of the same regions that responded selectively to ambivalent choices during the offer phase, such as the anterior insula, lateral PFC, and dACC (Fig. 4C, orange). This could possibly indicate sustained ambivalence-related activity that arises early in the valuation process and may still be unresolved at the time of choice commitment for highly difficult choices. The observed responses in vmPFC, selective for positive choice outcomes during the offer phase and for high confidence during the commit phase, provide additional evidence of separable stages of activity over the course of a value-based choice.

Given the robust SV- (from GLM1) and choice-related (from GLM2) responses during

the offer phase, it is possible that apparent confidence-related signals in vmPFC activity during the commit phase were in fact attributable to strong, sustained value-related responses carrying over from the offer phase. That is, relative differences in vmPFC responses to AccCon choices versus AccAmb choices could have driven the observed main effect of confidence – and moreover these effects might be better explained by value than confidence. To verify that the apparent effects of confidence could not be explained by relative differences in value, we measured contrasts of confident versus ambivalent BOLD responses separately for accepted and rejected offers. We were specifically interested in these comparisons within our a priori vmPFC ROI and therefore restricted statistical correction to this region alone (Fig. 5A). During the commit phase, we observed a cluster of voxels with significantly stronger responses during RejCon versus RejAmb trials as well as a cluster of voxels that preferred AccCon to AccAmb (Fig. 5B). Notably, because SV of RejCon trials is less than SV of RejAmb trials, this result suggests that during the commit phase, clusters of vmPFC activity are signaling choice confidence irrespective of SV. We found no voxels that demonstrated the same pattern during the offer phase, nor did we find any voxels with stronger responses to ambivalent choices during either phase.

Decision Phase-Specific Responses in vmPFC

Finally, for both phases of the decision making process, we examined the pattern of response magnitudes in vmPFC across trial conditions of increasing value, beyond relative differences between pairs of conditions and compared the fits of models with linear and quadratic terms (Fig. 5C). For the offer phase, the linear model best explained the observed pattern of vmPFC parameter estimates across the four conditions as the additional term in the

quadratic model did not improve the model fit (Linear model: AIC = 901.1, BIC = 911.98; Quadratic model: AIC = 903.1, BIC = 916.69; comparison of models: $\chi 2(1) = .0018$, $p = .966$), indicating that vmPFC activity increased linearly with P(Acc)-P(Rej) during the offer phase (signaling value). For the commit phase, the quadratic model best explained the observed pattern of vmPFC parameter estimates (Linear model: AIC = 1014.5, BIC = 1025.3, Quadratic Model: AIC = 1002, BIC = 1015.6, comparison of models: $\chi 2(1) = 14.320$, $p < .001$), indicating that vmPFC responses take a quadratic function with respect to value during the commit phase, signaling confidence. Model effect sizes are plotted in Fig. 5D and details about model fits are shown in Table 2.

In summary, vmPFC responses increased linearly across trial conditions of increasing value during the offer phase, signifying early involvement with valuation of the offer stimulus, and varied quadratically across the same conditions during the commit phase, signifying late involvement with valuation of the decision and deriving confidence. Notably, no values associated with model-estimated confidence (such as those used in GLM1) were entered into the mixed effects models. Rather, the models estimated participants' mean parameter estimates from the four trial conditions (of GLM2) from their mean P(Acc)-P(Rej) for all trials of that condition. Therefore, the plots in Figure 5c demonstrate that during the offer phase, vmPFC responses naturally take a linear function across the four trial conditions of ascending value P(Acc)-P(Rej), whereas during the commit phase, vmPFC responses naturally take a quadratic function across the same trial conditions.

Strictly speaking, because P(Acc)-P(Rej) takes a sigmoid function with respect to SV and that vmPFC increased linearly with SV in GLM1, one might predict vmPFC responses to take a logistic function with respect to P(Acc)-P(Rej). Visual inspection of individual raw

vmPFC parameter estimates (Fig. 5B) did not reveal that the underlying pattern of vmPFC parameter estimates had a logistic shape. To verify that this was the case, we repeated the linear mixed regression model, substituting SV predictors for P(Acc)-P(Rej) predictors, and found analogous statistical results. While the fMRI parameter estimates had sufficient resolution to discern between linear and quadratic functions, it is unlikely that a similar distinction could be made between two monotonic functions such as linear and logistic functions. We present results with respect to P(Acc)-P(Rej) because this measure improved the interpretability of group-wide results by controlling the range of predictor values across the group while preserving the sign of raw SV (as described above).

Behavioral RTs and Confidence

Before analysis, trials with RTs exceeding 3 SD of the overall group mean were excluded (1.56%). Additionally, trials exceeding 3 SD of the individual's mean RT for each bin of each analysis were excluded (0.98%). The model confirmed that RTs times took an inverse quadratic function with respect to value (Fixed Effects: (P(Acc)-P(Rej) estimate = -35.415, 95CI [-50.042, -20.787], SE = 7.466, t = -4.743, $p < .001$; $P(Acc)\text{-}P(Rej)^2$ estimate = -77.330, 95CI [-102.213, -52.451], SE = 12.700, t = -6.90, $p <. 001$; Random Effect SD = 88.02), indicating that participants were slower to commit to decisions associated with low model-estimated confidence (Fig. 6). Notably, in addition to the quadratic relationship, there was also a negative linear correlation between value and RTs, indicating that it took participants longer to commit to decisions about low value offers versus high value offers. This may have been caused by the observed bias to accept offers more offers than were rejected, which meant that responses to lower value offers, which were often reject choices, required

overriding the default decision to accept.

We observed statistically analogous, but somewhat weaker results when using SV as value predictors (Fixed Effects: SV estimate = -27.198, 95CI [-32.326, -22.070], SE = 2.618, t = -10.390, $p < .001$; SV^2 estimate = -7.900, 95CI [-12.230, -3.563], SE = 2.212, t = -3.570, p = .001; Random Effect SD=8.020), which is likely due to large individual differences in estimated SV, as described above. Taken together, the results of the RT analysis confirmed that there was a meaningful relationship between model-estimated confidence and behavior. Participants were slower to commit to decisions about offers we predicted would elicit ambivalence, suggesting that these choices might have been more challenging to resolve. This is not an indication that RTs measure the subjective experience of confidence, but rather that there is an empirical basis that model-estimated confidence captures aspects of our task associated with decision difficulty. Importantly, because participants varied in their subjective preferences and their decision boundaries diverged considerably from the reward=cost line (Fig. 2), increased RTs for ambivalent choices cannot be explained by difficulty in perceptual discrimination (i.e. merely determining whether costs or rewards were perceptually larger).

Discussion

We measured behavioral and neural responses during a two-phase Ap-Av decision making task with consequential mixed outcomes. Model-based SV, confidence estimates and decision variables correlated with BOLD responses throughout the cortex, including within vmPFC. We provide novel evidence that these vmPFC signals emerge during different phases

of decision making. Specifically, neural responses to offer valuation (SV) and choice determination were temporally dissociated from decision valuation (confidence), with the former dominating vmPFC prior to choice commitment, when the latter takes priority. The following sections discuss our key findings regarding neural signatures of SV and confidence as well limitations of this study and suggestions for future research.

During the offer phase of each trial, participants deliberated on accepting or rejecting offers and vmPFC responses increased with estimated SV (GLM1). We complemented this finding by demonstrating that parameter estimates from an anatomical vmPFC ROI increased monotonically with P(Acc)-P(Rej) (GLM2). This is consistent with a vast literature documenting vmPFC's role in integrating items' cost and reward attributes into its overall SV (Talmi et al., 2009; Grabenhorst and Rolls, 2011; Park et al., 2011; Amemori and Graybiel, 2012). vmPFC also distinguished between decision outcomes during the offer phase, with stronger responses anticipating accept decisions than reject decisions (GLM2). These results support a goods-based models of value-based decision making, which proposes that economic choices are made between goods rather than actions (Padoa-Schioppa and Assad, 2006), and can therefore be settled prior to planning the action to submit the decision (Wunderlich et al., 2010). To our knowledge, this is the first evidence of approach-selective vmPFC responses preceding choice commitment in an economic ApAv task. One previous fMRI study that employed a strategy game requiring choices to attack or defend reported that an adjacent region, rostral ACC, preferentially responded when participants defended versus attacked and tracked the value of deploying defense strategies but not attack strategies (Wan et al., 2015). While we did not observe regional activation corresponding to reject choices, the conceptual equivalent to defending, the notion of separable but adjacent neural bases for approach and

avoidance behaviors is compelling. Further exploration of the relationship these behaviors have with valuation will be an interesting avenue for further research.

Beyond vmPFC, we observed extensive overlap between regions tracking SV in GLM1 and regions that selectively activated preceding accept decisions in GLM2. This was observed in a larger value-network including regions important for sensory association, value-comparison, and reward processing such as angular gyrus, posterior parietal cortex, lateral temporal cortex, posterior cingulate cortex, and the ventral striatum. The value network likely encodes finer-grained distinctions between SV and choice outcome than were apparent in our analyses. For example, others have suggested separable time courses for evolving value signals and decision variables (Rushworth et al., 2012). Future research aiming to find finer dissociations of SV and decision variables may benefit from novel variations of this task. Furthermore, our anatomical vmPFC ROI was selected a priori and is somewhat inclusive, combining subcallosal cortex and medial frontal cortex from the Harvard-Oxford atlas, which include structures that others have labeled medial OFC, pregenual or subgenual ACC. Thus, we cannot draw meaningful conclusions about functional specificity at a smaller scale, but we appreciate that others have made interesting discoveries on this front. For example, it has been suggested that over the course of stimulus processing, OFC is the first cortical site where reward values are assigned to appetitive stimuli, whereas later processing in vmPFC transforms value representations into choices (Grabenhorst and Rolls, 2011) that guide action selection (Rushworth et al., 2009). Future research might combine finer parcellation of prefrontal cortex with fMRI models that include terms for individual offer attributes to reveal separable patterns of responses to reward, SV, and choice outcome. However, this was beyond the scope of our study.

Remarkably, we found no regional activation that was significantly modulated by SV (GLM1) or differed by choice outcome (GLM2) during the commit phase, even when statistical correction was restricted to vmPFC (GLM2). This result was unexpected given previous findings that vmPFC representations of subjective preferences are automatically elicited by task stimuli, even when this information is task-irrelevant (Lebreton et al., 2009; Smith et al., 2010). Moreover, SV responses in vmPFC can last for the entire duration of value stimuli, even when stimulus durations long outlast the time needed to make a decision and the extra time does not change decision behavior (Sokol-Hessner et al., 2012). Together, these prior findings suggest that in our study, vmPFC should encode SV whenever the offer stimuli were visible (including during the commit phase), but we failed to find such effects. This may indicate that our two-phase task reveals the transition of vmPFC processing from offer valuation to decision valuation in such a way that other paradigms that require immediate responses, delay responses but remove task stimuli during the response, or transition to choice outcomes without providing choice feedback cannot. Alternatively, there may be enduring, but relatively weak, SV representations in vmPFC during the commit phase that did not reach statistical significance. Visual inspection of the group mean parameter estimates suggests an asymmetrical quadratic function such that accepted offers, on average, were associated with larger response magnitudes than rejected offers. Furthermore, when we modeled vmPFC parameter estimates from the commit phase, a linear-only model provided a marginal fit with the data, consistent with a weak SV effect during the commit phase. Thus, there is a hint that vmPFC represents SV, albeit weakly, through the commit phase, but this effect is markedly exceeded by confidence.

We did not find any BOLD responses that parametrically correlated with confidence during the offer phase (GLM1) nor were there any regions with selective responses for confident over ambivalent trials during the offer phase (GLM2), even when statistical correction was restricted to vmPFC. This was somewhat unexpected given the empirical and theoretical literature suggesting early emergence of confidence signals that evolve in parallel with the processing of choice stimuli (Kepecs et al., 2008; Gherman and Philiastides, 2015; Dotan et al., 2018). Instead, we found confidence-related activity emerged after SV signals. During the commit phase of each trial, while participants submitted a choice and then viewed feedback about their selection, vmPFC responses increased parametrically with decision confidence (GLM1) and similarly were stronger for confident choices than ambivalent choices (GLM2). Closer inspection revealed that vmPFC parameter estimates took a quadratic function with respect to increasing value. Therefore, during the commit phase, vmPFC responses were stronger for confident choices than ambivalent choices, even for rejected offers when high confidence is associated with low SV. This provides strong novel evidence that vmPFC encodes both value and confidence during deterministic ApAv decision making, extending previous findings of confidence representations within vmPFC during value-based decision making (Lebreton et al., 2009; De Martino et al., 2013; Schlund et al., 2016).

Others have demonstrated that model-based confidence estimates correspond closely to self-reported confidence (Lebreton et al., 2015) and the strength of the relationship between model-based confidence and vmPFC activation strongly predicts the relationship between self-reported confidence and vmPFC activation (Bang and Fleming, 2018). Because we were interested in confidence in naturalistic decision making, which tends to occur in the absence of metacognitive evaluation, we did not solicit self-reported confidence ratings. We replicated

previous findings that model-estimated confidence had an inverse quadratic relationship with RTs (Lebreton et al., 2015), such that increasing model-estimated confidence predicted faster decision commitments (Fig. 6), suggesting that model-estimated confidence captures an element of choice difficulty.

A similar study by De Martino et al. (2013), that incorporated explicit confidence ratings with an fMRI value-based choice task suggests separable neural signals corresponding to model-based and self-reported confidence and found that while vmPFC tracked unsigned value differences between choice options (which roughly correspond to our model-estimated confidence), rostrolateral PFC (rlPFC) tracked self-reported confidence. They used this basis to ground a hypothesis that rlPFC probes internal confidence signals, represented in vmPFC, and makes them available for metacognitive self report. Others have made similar claims (Fleming et al., 2012). Notably, rlPFC responses correlated with ambivalence in our task (which did not require explicit confidence ratings), suggesting that our task did not elicit similar metacognitive appraisal processes. While there is not yet broad consensus regarding the loci of implicit versus metacognitive confidence – it seems evident that decision tasks requiring metacognitive confidence judgments may recruit unique neural resources from tasks that model confidence directly from decision behavior. Nonetheless, our results support recent suggestions that implicit confidence, the valuation of one's judgment, shares a neuroanatomical basis with a variety of other valuation processes, vmPFC (Lebreton et al., 2015; Hebscher and Gilboa, 2016; Fleming et al., 2018). We add to this growing theoretical framework evidence for a dynamic process by which vmPFC shifts from valuation of external stimuli to valuation of internal value representations and the decisions they inspire.

Chapter III: Experiment 2: Value and Confidence in Decisions with Effort Costs

Chapter 2 laid a framework for empirically investigating neural representations of value in cost-benefit ApAv choices with pain as the cost. The current chapter builds upon that framework by expanding the types of cost to which participants were exposed, allowing us to identify cost-domain-general and cost-domain-specific features of human decision behavior and their corresponding neural bases.

In everyday life, there are often several ways that one might reach a certain goal, and different paths promise varying types and degrees of costs. For example, when one is hungry for a snack, they could choose to spend time and resources cooking something from scratch, or spend money on delivery from a restaurant, or spend energy (and their potential welfare) on walking into the wilderness to catch some prey. How are expected energetic expenses, cognitive demands, pain, and time transformed into value information to guide decisions? And how similar are the representations of different types of costs?

Chapter 1 outlined numerous types of effort costs that have been posed in in the laboratory, including delayed rewards, risks, pain, cognitive effort, motor tasks, and grip force. More recently, the literature has expanded to begin directly comparing neural value representations of different types of costs, side by side. For example, neuroeconomic studies have contrasted decisions with costs of delayed rewards versus cognitive effort (Massar et al., 2015) and manual grip-force effort (Rudebeck et al., 2006; Prevost et al., 2010); as well decisions requiring physical effort versus cognitive effort (Schmidt et al., 2012; Chong et al., 2017) or probabilistic risks (Burke et al., 2013).

Still, it remains unknown whether different types of costs elicit different, domain-specific neural value computations or if all types of costs are represented in a domain-general manner, potentially with a common currency. Efforts to resolve this debate yield mixed results, often times even at the level of individual studies. One straightforward example is the comparison between cognitive effort and physical effort. Even intuitively, one can see strong cases for how such cost representations might be either domain-general or domain-specific. On one hand, physical and cognitive effort are conceptually distinct and the resources required for each are minimally overlapping – many different considerations go into deciding whether you would do one hour of manual labor in exchange for a meal versus deciding whether you would complete one hour of 3-back working memory tasks for the same reward. On the other hand, if costs are all represented in distinct domain-specific modules, how can they be evaluated with respect to rewards or one another?

Empirical tests of this specific question turned up mixed results. Whereas one study concluded that both cognitive and motor systems are driven by a common "motivational node," the ventral striatum (nAcc), which represents the rewards associated with different types of effort (Schmidt et al., 2012), another reported to find an entire network of regions involved with domain-general reward devaluation, notably excluding the ventral striatum, as well as one unique domain-specific locus, the amygdala, that represents the rewards of cognitive efforts (Chong et al., 2017). The story only grows richer as additional types of costs are taken into consideration. Other studies that compared effort to risk and delay almost unanimously agree that representations of effort costs are highly domain-specific and found within dACC (see below) , while a recent report concludes that cognitive effort costs are encoded by the same domain-general network as any other type of value stimulus (Westbrook et al., 2019).

Nonetheless, as the field matures, some themes are coming into focus. dACC (and/or ACC proper) is consistently implicated as a hub for representations of both cognitive effort costs (Massar et al., 2015) and physical effort costs (Rudebeck et al., 2008; Kurniawan et al., 2010; Burke et al., 2013; Arulpragasam et al., 2018). Interestingly, it has also been demonstrated that ACC represents SV in cost-benefit choices with effort costs (Croxson et al., 2009; Prevost et al., 2010; Chong et al., 2017), even potentially supplanting vmPFC from its well-understood role in this process (Klein-Flugge et al., 2016).

Notably, studies that directly compared neural representations of effort costs against risk and delay costs have seemingly reached consensus that dACC (/ACC) is uniquely domain-specific in this regard. That is, dACC(/ACC) consistently and exclusively tracks effort costs whereas other regions such as OFC, lateral parietal, and temporal cortex do not discriminate between cost types (Massar et al., 2015); nAcc, vmPFC, and (in this case) OFC specifically encode delay costs (Prevost et al., 2010), and aIC represents risks (Rudebeck et al., 2006). However, it has also been suggested that studies reporting associations between dACC(/ACC) and effort costs of value-based decisions, such physical effort, have conflated representations of the consequences of the choices with difficulty of making the decision itself (Hogan et al., 2018).

The insula is also frequently associated with cost representations in value-based decision making, although its exact role is unclear. Several studies have found that aIC tracks physical effort costs (Prevost et al., 2010; Treadway et al., 2012; Kurniawan et al., 2013; Arulpragasam et al., 2018), others have shown that this region encodes risks (Preuschoff et al., 2008; Burke et al., 2013), rewards associated with exerting effort (Croxson et al., 2009), and SV more generally (Chong et al., 2017). pIC has been reported to represent temporally

discounted rewards (Tanaka et al., 2004; Wittmann et al., 2007), respond to pain and the expectation of pain (Singer et al., 2009) and aversive values more generally in decision making tasks (Plassmann et al., 2010). These results are generally consistent with the anatomy of the region. pIC is a sensory cortex for visceral afferents and aIC is the source of cortical efferents to the parasympathetic nervous system, i.e. visceral motor cortex.

Despite these exciting advances, the field is starkly lacking examination of one evolutionarily critical type of decision cost, cardiovascular exercise. To our knowledge, value-based choices with costs of aerobic exercise have never been implemented in a laboratory setting. Previous investigations into the neural representations of physical effort costs in decision making posed offers with difficult motor tasks, such as rapid button presses, or manual grip force. However, such tasks are not particularly representative of the decision costs humans evolved to face. When our ancestors decided what to do about foraging and prey selection, their options often involved steep costs of cardiovascular effort rather than rapid button presses to dial in a delivery order. It follows that humans should be particularly well-equipped to make such cost-benefit decisions about aerobic energetic expenses, and we were interested to know how people decide about and how the brain represents significantly physically demanding effort costs. Furthermore, we sought to understand how such choices and neural responses compare to our results from Chapter 2, in which participants faced choices with pain costs. As decisions with exercise costs had not been studied empirically, they had not been contrasted with pain costs either. In fact, to our knowledge, no previous studies have compared decisions with pain costs to decisions with any sort of effort costs.

We had participants complete two ApAv decision making tasks while undergoing fMRI. In one session, decision costs entailed significant pain stimuli, as in Chapter 2. In the

other session, decision costs were exercise intervals on a stationary bike, calibrated to challenge the individual's maximum capacity for physical exertion. We applied the same analyses as described in Chapter 2 to measure decision outcomes, SV, and confidence in choices with cardiovascular exercise costs. These analyses tested whether the results observed in Chapter 2 generalized to other types of decision costs (physical exercise) as well as to probe decision making about aerobic exercise more generally, as such a task had not been previously implemented.

In addition to testing for a replication of our original findings, we additionally sought to identify domain-specific and domain-general behaviors and value representations between decisions about pain and exercise costs.

We hypothesized that the overall pattern of results from Chapter 2 (shock version) would generally replicate here (bike version) in a domain-general manner, as we used identical experimental designs and stimuli, and because both tasks entailed significant physical costs paired with monetary rewards. While the literature strongly suggests that effort costs recruit dACC in a uniquely domain-specific manner, it was unclear exactly how this would manifest with respect to our initial findings of ambivalence-related responses in the same region. We expected that the bike task would elicit stronger responses in dACC, representing effort costs, than the shock task. Yet, because this was an entirely novel implementation of effort costs, we did not have explicit predictions about the exact nature of such domain-specific representations in dACC beyond relative magnitudes. Previous research also reported that SV is represented in ACC and *not* vmPFC in decision making about effort costs. Accordingly, we expected that value representations would be somewhat weaker in vmPFC and potentially distributed more broadly into adjacent cortical regions (dACC /ACC) in the bike task relative to the shock task.

Methods

Participants

We report data from 21 participants (13 women, 20 right-handed, mean age = 21.9, *SD*=2.5). One additional participant completed the study but was removed from analyses due to significant susceptibility artifacts in their functional neural images that caused excessive errors in spatial normalization. No participants had a history of neurological injuries or illnesses, current daily use of psychoactive medications, cardiovascular disease, nor physical injuries that would interfere with performing cardiovascular exercise. Participants provided written consent in accordance with the Institutional Review Board at the University of California, Santa Barbara.

Experiment Overview

Participants made three separate visits to the lab: once for the pain decision making task, once for a physical fitness test and once for the exercise decision making task. Seventeen participants completed the shock session first, four completed the bike version first. Experimental procedures for the pain task are detailed in Chapter 2 and an overview of both sessions is included in Figure 7. The bike experiment required an initial fitness test to identify participants' aerobic capacity, as measured by their VO_2 max, a cardio-pulmonary exercise test. On the second visit, participants were familiarized with the cost stimuli by biking several intervals at a percent of their VO_2 max on a stationary bike while viewing corresponding cost stimuli. Next, participants performed the decision making task while in the MRI scanner. Finally, after they had completed the experimental task, they returned to the stationary bike to

perform the exercise intervals they had chosen during the task and received their monetary rewards.

Session 1: Aerobic thresholding procedure (VO2 max test)

The VO_2 max test determines a person's maximal oxygen uptake during cardiovascular exercise. For the purposes of our study, this metric served as an objective benchmark of participants' capacity for physical exercise. In this way, costs in the decision making task were normalized for individual differences in physical fitness. Therefore, the VO_2 max test was analogous to the pain thresholding procedure in Chapter 2, but without any variance that might have been driven by subjective pain ratings.

Participants refrained from strenuous physical exercise for at least 48 hours before coming to the lab for a bike session. They were also instructed to eat a full meal approximately two hours before both sessions. Before performing the fitness test, we verified that participants were eligible to engage in physical exercise with the Physical Activity Readiness Questionnaire (PAR-Q) from the National Academy of Sports Medicine (www.nasm.org). Participants were then fit with electrocardiograph recording electrodes, a pulse oximeter, and a respiratory mask, all transmitting to a Vyntus CPX/ECT metabolic cart (www.vyaire.com). The respiratory mask connected with a spirometer measured respiratory gas exchanges as the participant performed the test. During the first three minutes of testing, participants rested on the stationary bike (Ergoline 900, www.vyaire.com) while baseline respiration and cardiography were measured. For the next three minutes, participants began pedaling at a constant rate with a work load determined by their reported typical exercise routines and overall physical health. Next, the resistance of the bike gradually increased until oxygen uptake

plateaued. Participants then entered a cooldown phase by pedaling for five minutes at the baseline resistance level. The entire test lasted approximately 25 minutes.

Participants also completed brief fMRI sessions immediately before and after their VO2 max test for an independent study about visual attention (data not reported here). Physiological benchmarks from the VO_2 max test were used to determine the bounding thresholds of minimally and maximally difficult exercise that defined the range of costs in the decision making task. At the start of the start of the VO2 max test, physiological signals initially increase gradually with increasing resistance load, however as the test progresses and the load further increases, heartrate and metabolic respiratory gas exchanges begin to increase nonlinearly as load increases. We identified this point of deflection through visual inspection of physiological signals measured during the VO2 max test and the resistance load at this time point was recorded as the threshold for minimally difficult exercise. The VO2 max test ends when oxygen intake eventually plateaus, signifying the individuals peak physical exertion capacity. We recorded the resistance load at which participants reached this plateau (VO2 max) as the threshold for maximum physical exertion. All costs offered during the decision making task ranged from 1 – 100% between these two thresholds.

Session 2: Cost familiarization, decision task, and decision payouts

Cost familiarization: Participants returned to the lab to complete the exercise decision making task. Before the decision making task, participants were first familiarized with cost stimuli by riding the stationary bike while viewing example offers from the decision making task. They were shown 5 example offers with costs of 5%, 25%, 50%, 75% and 95% while pedaling the stationary bike at the corresponding resistance load. The loads were calibrated

such that 1% corresponded to the recorded load for minimally strenuous exercise from the VO_2 max test, 100% was the recorded load when they reached their VO_2max (maximally strenuous), and all other values were scaled linearly between the two. Each familiarization interval was 90 seconds. Participants were informed that the duration of the familiarization intervals were longer than the costs from the decision making task, but were instructed to attend to the perceived difficulty of the exercise rather than the timing. They were additionally instructed to remember the sensation associated with the example exercise intervals and use these as points of reference when making choices during the task, but that the real offers would include exercise ranging anywhere between 0 and 100% intensity, not just at the example levels.

ApAv task: The offer stimuli, trial design, and experimental procedure were identical to the methods described in Chapter 2 for the pain version of the task. The key difference between the two sessions were the costs of the offers. In the bike version of the task, the cost stimulus (though perceptually identical to that in the shock version) represented one minute of exercise at an intensity proportionate to the size of the cost attribute stimulus, rather than the pain intensity of an electrical shock. We additionally varied which offers were tagged as payouts between sessions so that participants decisions weren't biased by memories about which offers participants would result in payouts.

Payouts: Payouts were carried out analogously to the procedures for the shock version of the study described in Chapter 2. After participants completed the decision making task in the MRI scanner, they returned to the stationary bike to perform the exercise costs from accepted payout trials and to receive the associated monetary rewards. Before the payout intervals, participants were re-fitted with a breathing mask and pulse oximeter and completed five minutes of warm-up exercise at a resistance below their minimally strenuous threshold. If

participants grew too fatigued to complete the payout intervals, the experimenter transitioned the exercise protocol to a cool down interval and the participant chose whether to complete the remaining payouts or to terminate the session and be compensated for what had been completed so far.

Statistical Analysis

Decision behavior from the bike ApAv session was analyzed with the same methods as the shock data as described in Chapter 2. Participants' choices were fit with individual logistic regression models to estimate the perceived SV, P(Acc)-P(Rej) and confidence associated with each trial offer, and trials were categorized into AccCon, AccAmb, RejCon, and RejAmb trials for contrast analyses. Likewise, the same preprocessing pipeline was applied to neuroimaging data from the bike session and the same analyses (GLM1 and GLM2 from Chapter 2, in Chapter 3 will be referred to as GLM1b and GLM2b), with the same regressors and contrasts were applied at the run level. Condition (e.g. AccCon in GLM1b) and regressor parameter estimates (e.g. SV in GLM1b), as well as contrasts between trial conditions (e.g. AccCon vs. AccAmb in GLM2b) estimated at the run level were then carried up to the second level to estimate overall statistics at the subject level, just as in Chapter 2. The novel additions to GLM1b and GLM2b are that bike runs and shock runs were both included at the subject level, allowing us to specify additional higher-order contrasts between bike and shock conditions (domain-specific: e.g. SV_{shock} vs. SV_{bike} and domain-general: e.g. SV_{shock} and SV_{bike} together), which were in turn carried up to the group level. Notably, such comparisons couldn't be made at the run-level as participants performed the bike and shock versions of the ApAv task on separate days.

In summary, GLM1b measured parametric responses to four regressors (SVoffer, SV_{commit}, $Confidence_{offer}$, and $Confidence_{commit}$) with bike costs, with shock costs, overall (domain-general; averaging across bike and shock), and finally contrasts between the two (domain-specific; e.g. SV_{offer_bike} vs. SV_{offer_shock}). GLM2b measured neural responses during both the offer and commit phases of each trial condition (AccCon, AccAmb, RejCon, and RejAmb) and select contrasts between these conditions (the same as in Chapter 2; e.g. AccCon vs. AccAmb). Each contrast was estimated for bike and shock costs separately, overall (averaging across bike and shock), and then contrasted between bike and shock conditions. Statistical outcomes of both GLMs were corrected for multiple comparisons at the whole-brain level with permutation-based TFCE, using FSL's Randomise function, as in Chapter 2. This was repeated within our vmPFC ROI and included additional sub-condition contrasts of AccCon vs. AccAmb and RejCon vs. RejAmb.

Results

The primary differences between the analyses conducted here and in Chapter 2 are that we now draw from a subset of the original group (21 instead of 28) and we include two data sets for each individual - one in which decision costs are pain stimuli and one in which decision costs are cardiovascular exercise intervals. Decision behavior was similar to that described in Chapter 2 - participants negatively valued the bike offer attributes, indicating that aerobic exercise was perceived as a proper cost. The results of the behavioral analyses are described in detail in Chapter 5. Our fMRI analyses measured neural responses separately for bike (bike-only) and shock (shock-only) data sets, both pooled together (domain-general), and contrasted the two (domain-specific). The latter comparison would identify any voxels with significantly

stronger responses to one cost versus the other. Although the analysis approach to the shock-only data is identical to what was presented in Chapter 2, here we present results from only the subset of 21 individuals who also completed the bike task.

Remarkably, we found no voxels with significant differences between the two costs in GLM1b or GLM2b, indicating that there were no domain-specific differences in the magnitude of responses correlating with any the variables that we measured. Therefore, we focus our descriptions on the results of the domain-general analyses, which incorporated neural responses from both the bike and shock tasks. Our figures additionally include the results of bike-only and shock-only analyses, which highlight two noteworthy features of the data set analyzed here. First, the effects of the reduced sample are rather apparent. Comparison of the results presented here in the shock-only analysis versus the results from Chapter 2 reveals that some effects that were robust in a sample of 28 fail to reach significance in a sample of 21 with strict statistical thresholding ($p < .05$, tfce, whole brain level). Second, in some cases there appear to be considerable differences between the bike-only and shock-only results through visual inspection of statistical maps. However, these visually apparent differences are spurious, statistically speaking, as we found no significant contrasts between the two tasks after correcting for multiple test comparisons.

GLM1b

In GLM1b we measured parametric modulation of BOLD responses by continuous regressors for SV and choice confidence during the offer and commit phases of value-based choices, as in GLM1 in Chapter 2. During the offer phase, domain-general SV correlated significantly with a wide network of regions, especially within the value network, including

vmPFC, OFC, throughout the basal ganglia including within nAcc and the caudate, posterior cingulate cortex, and the pIC (Fig. 8, left). Other robust signals were observed in lateral parietal cortex and dorsal premotor areas, and other relatively weaker but statistically significant voxels were found to track SV virtually throughout the cortex. No voxel activation was negatively correlated with SV or positively correlated with confidence during the offer phase. These results largely replicated our findings in Chapter 2, validating that SV computations in value-based decisions with exercise costs recruit highly similar neural resources as decisions about pain costs. Interestingly, ACC also seemed to track SV. Although we had not observed this effect our previous analysis, other previous studies have found similar responses in this region (Klein-Flugge et al., 2016).

BOLD responses in a partially overlapping network of cortex correlated negatively with confidence (that is, positively with ambivalence) during the offer phase. Specifically, we observed that dACC and right SPL responses increased with ambivalence, replicating our findings in Chapter 2 and extending them to decisions about exercise costs (Fig. 8, right). Responses in dACC in particular are likely best explained by the resolution of conflict evoked by valuation of neutral offers (Aupperle et al., 2015). Activation in early visual cortex also correlated with ambivalence, suggested ongoing analysis of choice stimuli during difficult decisions. Building on our findings from Chapter 2, the regions that respond both to SV and ambivalence during the offer phase are a particularly interesting case because we expect that for offers with positive SV confidence increases with SV. Thus, the pattern of responses in dACC and SPL may be specific to negatively valued items, for which ambivalence increases as value increases. Taken together, the results from the offer phase in GLM1b reinforce our

previous conclusions that task-related activation during the offer phase are largely related to processing the valuation of the offer.

During the commit phase, BOLD responses in pIC, lateral temporal cortex, and SPL correlated positively with choice confidence, mirroring our previous findings with one notable exception (Fig. 9, left). Given the results of GLM1 in Chapter 2, we predicted vmPFC would track confidence during the commit phase in GLM2b, but no voxel activation reached significance at a strict statistical threshold ($p < .05$, tfce, whole brain correction). Slightly reducing the threshold ($p < .1$, tfce, whole brain correction), did reveal task-related activity in vmPFC that correlated positively with choice confidence (Fig. 8, right). Notably, reducing the threshold only slightly changed the overall pattern of results, in most cases subtly augmenting clusters that were already significant at the stricter threshold. vmPFC was the most salient exception to this rule, as reducing the threshold revealed a new large cluster of voxel responses correlating with confidence, suggesting that we lacked sufficient power to observe a true effect at the original threshold. We did not find any neural responses that were negatively correlated with confidence during the commit phase

No regional activation was positively correlated with SV during the commit phase, replicating our original findings from Chapter 2. Instead, we observed activation that was negatively correlated with SV during the commit phase, which is a novel result relative to Chapter 2. These responses were observed throughout the insula, the amygdala, ventral mid-cingulate cortex, left angular gyrus, sensory cortex, and precuneus (Fig. 10).

GLM2b

In GLM2b we compared BOLD responses measured during accept versus reject decisions and during confident versus ambivalent decisions. During the offer phase, many regions had significantly stronger responses while evaluating offers that would be accepted versus offers that would be rejected. The network of regions showing this effect closely resembled those that tracked SV during the offer phase in GLM1b, and importantly included vmPFC (Fig. 11, left). This result additionally confirms findings from Chapter 2 that robust activation of the value network correlates both with increasing offer valuation (SV) and positive decision outcomes (accept decisions). BOLD responses within the cognitive control network, such as in dACC and lPFC, were stronger during the offer phase of ambivalent decisions than confident decisions (Fig. 11, right). These effects again suggest the recruitment of regions critical for settling response conflict during more difficult choices. Posterior parietal cortex and visual cortex showed the same effect, which may be related to amplification of processes related to perceptual and value comparison as the participants attempted to resolve ambiguous choices. As in Chapter 2, BOLD responses from the offer phase revealed no regions that responded preferentially during decisions to reject over decisions to accept, nor any regions that responded preferentially to confident choices over ambivalent choices.

During the commit phase, many regional responses showed significant contrasts between confident and ambivalent trials, but none responded differentially according to decision outcome, replicating our findings in Chapter2. dACC, lPFC, aIC, and inferior parietal cortex, responded more strongly to ambivalent trials than confident trials (Fig. 12, left, orange). Similar regions had all shown the same effect during the offer phase, and during the commit phase the activation seemed to spread and incorporate other regions important for resolving

response conflict such as aIC and wider swaths of lateral prefrontal cortex. The reverse contrast was observed in cortical regions including the pIC, superior temporal cortex, and primary sensory cortex (Fig. 12, left, purple). As in GLM1b, we had strongly predicted that vmPFC would be included in the set regions with significantly stronger responses during commitment of confident choices, vmPFC activation correlating with high confidence did not surpass statistical thresholding. We again lowered the critical threshold (to $p < .1$, tfce, whole brain correction), which modestly expanded the regions showing this contrast to now incorporate vmPFC (Fig. 12, right).

Finally, we measured contrasts BOLD responses during the offer and commit phases of confident versus ambivalent trials within our priori vmPFC ROI. As in Chapter 2, during the offer phase, there were no significant differences in vmPFC responses between confident and ambivalent trials of either decision type. However, during the commit phase, we observed clusters of voxels showing significantly stronger responses for confident accepts and confident rejects versus ambivalent accepts and ambivalent rejects, respectively (Fig. 13).

Discussion

We measured behavioral and neural responses during the same two-phase mixed-outcome Ap-Av decision making task presented in Chapter 2, except decision costs now included cardiovascular exercise on a stationary bike. No prior studies had used exercise costs to study value-based decision making. The results presented here validate the use of the methodology for future research and suggest that domain-general valuation mechanisms are at work during decisions about physical pain costs as decisions about physical exercise costs.

Overall, analysis of the bike and shock data sets revealed remarkably similar neural response patterns, to the extent that we found no significant contrasts between the bike and shock task for any of the variables that we measured in GLM1b and GLM2b. That is, we cannot confirm hypotheses predicting domain-specific valuation or decision processes between our two tasks. Instead, our domain-general analyses revealed highly similar results to those presented in Chapter 2, in which decision costs were only pain stimuli. We additionally observed novel activation in ventral mid-cingulate cortex, the amygdala, aIC and pIC, precuneus, and inferior parietal cortex that was inversely related with SV, suggesting that these regions tracked the aversiveness of offer stimuli. We organize our discussion of key results around the regions showing the most theoretically interesting results and that are used in ROI analyses in Chapters 4 and 5.

vmPFC

We observed confirmatory evidence that vmPFC signals relating to SV and confidence emerge during different phases of value-based decision making such that processing of offer valuation (SV) and decision variables arise during the offer phase, prior to neural signals relating to decision confidence, which arise during the commit phase. While late confidence signals in vmPFC were somewhat weaker than we initially observed in Chapter 2, we take the overall pattern of results as preliminary evidence that such value and confidence representations are domain-general across pain and exercise costs.

During the offer phase, vmPFC responses increased parametrically with SV in GLM1b and were stronger for accept decisions than reject decisions in GLM2b, as in GLM1 in Chapter 2. The fact that we did not observe significant contrasts between the two cost conditions,

suggests that SV and decision outcomes are represented in a domain-general fashion for offers with both exercise and pain costs. Moreover, such early, domain-general encoding of decision outcomes in vmPFC is further evidence that our results support a goods-based model of value representations in vmPFC and OFC. This model posits that economic choices are made on the basis of the relative SV of one's options, which is encoded in abstract (domain-general) value representations in vmPFC and OFC (Padoa-Schioppa, 2011). Furthermore, decision variables are computed prior to and irrespective of the planning of actions required to execute the choice. (Padoa-Schioppa and Assad, 2006).

The majority of studies that employed decision tasks with effort costs report that ACC / dACC, and *not* vmPFC, encodes SV, our data did not replicate this effect. In our task, vmPFC *did* track SV during decisions about both pain and effort costs. Despite the broader consensus, our observation of effort-related SV signals in vmPFC was not entirely unprecedented. Earlier studies that employed physical effort costs of manual grip force and challenging motor tasks also found that vmPFC tracked SV (Treadway et al., 2012; Arulpragasam et al., 2018; Hogan et al., 2018). However, it has been suggested that vmPFC responses in decisions about physical effort are specifically related to reward maximization, rather than representing the overall SV of the offer, which is encoded by ACC instead (Klein-Flügge et al., 2015). Closer inspection of within-ROI response patterns in Chapters 4 and 5 attempted test these competing hypotheses more directly.

During the commit phase, we did not find a statistically significant parametric relationship between vmPFC responses and model-estimated confidence when correcting for multiple comparisons at the whole brain level (GLM1b). Similarly, contrasts between the commit phase of confident versus ambivalent trials did not reveal decision-specific activation

in vmPFC (GLM2b) when corrected at the whole-brain level. Reducing the whole-brain threshold changed little throughout the rest of the cortex, but revealed large clusters of voxel activation in vmPFC that were positively correlated with decision confidence. Thus, as an exploratory study, we can make a reasonable case that vmPFC is likely to be involved in representing confidence.

Finally, to verify that apparent late confidence signals in vmPFC couldn't be explained by strong SV signals (from positively valued items only) carrying over from the offer phase, we separately contrasted AccCon vs. AccAmb and RejCon vs. RejAmb within in vmPFC ROI. Note that for reject decisions, which are associated with negative value, confidence *decreases* as value increases, whereas the opposite is true for accept decisions. Both contrasts revealed clusters within vmPFC with significantly stronger responses associated with high confidence decisions, indicating that during the commit phase, vmPFC responses relate to decision confidence, irrespective of SV. This is a key replication of our original novel finding of temporally dissociated signals corresponding to SV and decision confidence within vmPFC (Chapter 2). With this result, we can extend our conclusion that implicit confidence (valuation of one's judgment) shares a neuroanatomical basis with valuation processes in vmPFC: These complementary representations are domain-general insofar as they are commonly represented in decisions with different costs. Furthermore, the dynamic process by which vmPFC shifts from early processes related to offer valuation to later processes related to decision valuation, is likewise domain-general – as we observed it across both types of costs.

ACC / dACC

ACC / dACC correlated with ambivalence during both the offer (GLM1b and GLM2b) and commit phases (GLM2b), suggesting that the pattern of results observed in Chapter 2 was indicative of domain-general resolution of response conflict during difficult decisions. This is consistent with numerous, diverse observations of dACC's role in decision making under ambivalence, conflict, and uncertainty (Kahnt et al., 2011a; Badre et al., 2012; Economides et al., 2014; Aupperle et al., 2015).

During the offer phase, we observed that clusters of activation centered in vmPFC but extending into adjacent dACC that were positively correlated with SV (GLM1b) and were stronger for accept decisions than reject decisions (GLM2). Notably, these extensions into dACC were not observed in Chapter 2. Although contrasts between the bike-only and shock-only data sets were not significant, by visual inspection it seems that bike-only SV and accept-selective activation is more dorsal and posterior than that observed in the shock-only analysis. While we can't confirm these signals are cost-domain-specific, the basic finding of SV signals in dACC are consistent with other previous reports (Croxson et al., 2009; Prevost et al., 2010; Chong et al., 2017).

We additionally found that responses in ventral mid-cingulate cortex were inversely related with SV, which we had not observed in Chapter 2. It is tempting to deduce that this effect might have been primarily driven by the bike task, given that this region and adjacent dACC are consistently identified as domain-specific modules for processing effort costs and the SV of effortful offers (Rudebeck et al., 2006; Prevost et al., 2010; Burke et al., 2013; Massar et al., 2015). However, when the bike and shock data were analyzed separately, we

found that neither had significant activation in ACC, dACC, or mid-cingulate, nor were there significant contrasts between the two.

Previous claims about domain-specificity of effort costs ACC / dACC were made on the basis of empirical comparisons between decisions with effort costs versus decisions with risk or delay costs. When we compared physical effort costs with *physical pain costs* we did not find evidence of domain-specific effort representations in ACC / dACC, but rather domain-general activation in more inferior and posterior portions of cingulate cortex that was negatively correlated with SV. It is possible that this region is encoding something about enduring aversive events, which include both pain and intense exercise, but not risks or delays, per se. In fact, ACC responds robustly to pain, among other aversive stimuli (Shackman et al., 2011). In this light, we might describe our observations as domain-specific cost representations (-SV) insofar as pain and exercise belong to the same domain of 'endurance of aversive events.' This is a testable hypothesis. Future research may compare pain, physical effort, and risk or delay costs to determine if cost representations in dACC are truly domain-specific, and if so – specific to what? Still, this interpretation does not explain why our analysis of the bike-only data failed to replicate numerous accounts that dACC tracks effort costs in value-based decision making.

An alternative explanation is that dACC is in fact involved in a process other than computing effort costs. It was recently pointed out that many influential decision making studies inadvertently conflate decision costs with decision difficulty, leading to unfounded assumptions that dACC encodes negative value representations (Shenhav et al., 2014; 2016). Specifically, task designs often failed to orthogonalize SV with respect to choice difficulty by systematically pairing difficult decisions with lesser value increments. Controlling for this

revealed that dACC only tracks negative value when it is confounded with decision difficulty (Shenhav et al., 2014; 2016). Another study directly tested this hypothesis in decision making with effort costs and found that responses in dACC were explained by choice difficulty rather than by negative valuation of effort costs (Hogan et al., 2018). Our design is well-suited for dissociating negative SV from decision difficulty. Ambivalence takes an inverse U-shaped function with respect to SV such that there are ambivalent decisions with both positive and negative SV. Therefore, having appropriately controlled for potential confounds of ambivalence and effort costs, we might have circumvented the potential for spurious correlations between effort costs and activation in dACC.

aIC and pIC

The posterior insula is the primary sensory cortex for visceral afferents and associated with somatic representations whereas the anterior insula is the source of parasypmathetic efferents to the viscera (Tsakiris and Haggard, 2003; Taggart et al., 2016) Thus, activations in these areas provide a window into potential connections between somatic visceral representations and decision making. With that in mind, we found interesting patterns of responses in pIC and aIC that seem to roughly mirror those in vmPFC and dACC, respectively. During the offer phase, pIC activation parametrically correlated with SV (GLM1b) and responded more strongly to accept decisions than reject decisions (GLM2b). During the commit phase, pIC responses increased parametrically with confidence (GLM1b) and were stronger in categorically binned confident versus ambivalent trials (GLM2b). This is the same pattern that we observed in vmPFC, if not a bit stronger – as late confidence signals in pIC survived whole-brain statistical correction at our original statistical threshold. Our robust

results in pIC are novel as the literature is largely silent on the potential contributions of this region to value-based decision making beyond its role in responding to aversive stimuli. There is one crucial dissimilarity between pIC and vmPFC, pIC responses in the commit phase correlate negatively with SV. This fascinating comparison demonstrates subtle dissociations between how the two regions differentially represent value and confidence over the course of a decision. Earlier in the trial both regions respond positively with increasing value, but later while vmPFC hones in on confidence, pIC signals both confidence and negative value. Chapters 4 and 5 will more closely examine pIC's role in in our task as well as potential domain-specific response patterns.

We did not observe significant involvement of aIC during the offer phase. Instead, during the commit phase, BOLD responses in aIC were negatively correlated with SV in GLM1b and were stronger when participants committed to ambivalent versus confident decisions. Thus, whereas dACC seemed to be exclusively engaged by decision ambivalence during the commit phase, aIC encoded both ambivalence and negative value, mirroring the relationship between pIC and vmPFC.

The pattern of responses we observed in aIC agree, to some extent, with other similar studies. aIC is frequently co-reported with dACC in representing negative value of effort costs (Kurniawan et al., 2010; Prevost et al., 2010; Treadway et al., 2012; Arulpragasam et al., 2018) and other negative decision outcomes such as risk (Preuschoff et al., 2008; Burke et al., 2013). Given the discussion above regarding the often misreported relationship between dACC and decision costs, perhaps aIC responses are truer neural signatures of negative value. Conversely, aIC has also been observed to represent the rewards associated with exerting effort (Croxson et al., 2009) and positive SV more generally (Chong et al., 2017), suggesting that there is

deeper complexity to value representations in aIC, which will be explored in the following chapters.

Amygdala

The amygdala tracked negative value during the commit phase (GLM1b), consistent with this region's known responsivity to aversive stimuli. Research is beginning to explain how the amygdala's role in emotional response regulation interacts with circuits for value-based decision making, steering behavior away from aversive stimuli and even participating in the valuation process (Hsu et al., 2005; De Martino et al., 2006; Basten et al., 2010; Jenison et al., 2011; Zangemeister et al., 2016). Interestingly, the amygdala has also been observed to track rewards associated with physical effort in a domain-specific manner (i.e., it did not similarly track the rewards of cognitive effort; (Chong et al., 2017)). In our task, BOLD responses that were significantly stronger for accepted trials than rejected trials spread marginally into the amygdala, however it is unclear whether this is indicative of a true positive value representation in the amygdala or an artifact of spatial smoothing. To explore this further and to get a finer-grained look at response patterns within this region associated with our variables of interest, we include the amygdala in our ROI analyses in Chapters 4 and 5.

Nucleus Accumbens

Finally, nAcc is well-understood to respond to rewards. Accordingly, in our task, nAcc tracked SV (GLM1B) and engaged during the offer phase of decisions to accept (GLM2b). We included nAcc in our ROI analyses to understand if this region differentially encoded rewards according to the costs with which it was paired. While the magnitudes of such differences weren't statistically significant after correction at the whole-brain level, it was possible that

they were represented within patterns of activation or only apparent with finer-grained comparisons of choice variables.

Taken together the results of Chapter 3 validated a novel methodology for implementing cardiovascular effort costs in empirical studies of cost-benefit decision making, replicated and generalized the results of Chapter 2, and presented evidence that neural representations of SV and confidence are domain-general with respect to costs. Although we did not observe differences in response magnitudes between decisions about exercise and pain costs, it is possible that these costs elicit unique patterns of voxel activation within the set of ROIS highlighted in the discussion for showing theoretically interesting results. In Chapter 4, we test this hypothesis by measuring representational dissimilarity in voxel response patterns in vmPFC, ACC / dACC, aIC, pIC, amygdala, and nAcc.

Chapter IV: Representational Dissimilarity of Pain and Effort Costs

Chapter 3 compared neural responses recorded while people made decisions about offers with effort costs (bike version) versus pain costs (shock version). We specifically looked for differences in how effort and pain costs might influence BOLD responses associated with decision outcomes, SV, and confidence during two phases of decision making. The goal was to determine which, if any, of these three types of value representations are similar (domain-general) or unique (domain-specific) when associated with different types of costs. The magnitudes of neural responses from the bike and shock tasks were statistically indistinguishable for all three variables, suggesting that neural processing of decision costs is domain-general. Here, we test this conclusion by applying a different analysis that measured representational dissimilarity in voxel response patterns rather than the relative magnitudes of responses.

The decision tasks that participants performed in Chapter 3 were identical in all ways (experimental session schema, thresholding procedures, task design, stimuli, trial order, etc.) except the *meaning* of the cost stimulus. In the bike version costs were cardiovascular effort, whereas in the shock version costs were physical pain. While pain and effort are conceptually different, there are reasons to believe that differences in neural responses during these two tasks may be more nuanced than what can be detected by comparing the magnitudes of BOLD responses alone. First, both costs are corporeal in nature – that is, rather than other costs often implemented in the laboratory such as cognitive effort, delay, or probabilistic risk of reward loss; exercise and pain are physical, tangible costs. Therefore, we might expect similar neural mechanisms to respond to the anticipation of both costs. Second, if the pain and effort costs

we posed *are* encoded by overlapping cortical regions, we may not necessarily expect that what varies between the two is the strength of the overall response. Instead, pain and effort costs may elicit different patterns of responses across voxels within the same region. Although the overall amplitude of the responses associated with each cost were comparable, we hypothesized that this may have occluded differences in the underlying patterns of voxelwise activation.

Because we suspected that our initial analyses undermined our attempts to observe cost-domain-specific processing, we maintained our initial hypotheses posed in Chapter 3, but took a multivariate analysis approach. Differences between patterns of activation can be detected by multivariate approaches to analysis of neuroimaging data, such as representational dissimilarity analysis (RDA), which is more often referred to as representational similarity analysis (RSA) (Kriegeskorte, 2008; Walther et al., 2015a). Some RSA/RDA analyses attempt to measure similarities between patterns of activation in a manner akin to correlation analysis, resulting in representational similarity matrices (RSMs). We were interested in the converse, and sought to measure differences between patterns of activation, and therefore computed representational *dissimilarity* matrices (RDMs) comprising Euclidean distances between multivoxel patterns of activation.

GLM2b was recomputed to measure neural responses during AccCon, AccAmb, RejCon, and RejAmb trials, but this time preprocessing and modeling steps were tailored to a secondary analysis with RDA. Furthermore, to increase our ability to observe subtle differences in voxel response patterns, we measured neural responses on a more precise timescale. In Chapters 2 and 3, we looked at neural responses averaged across the offer (4.55s) and commit (3.27s) phases. Here, we measure neural responses during each of 9 individual

TRs (.91s each) over the course of the trial, beginning with the fixation cue. We predicted that this approach would reveal cost-domain-specific patterns in voxel responses, confirming the hypotheses originally put forth in Chapter 3.

In addition to comparisons between conditions with different costs, we also measured all pairwise distances between trials that varied in decision outcome, confidence, and cost. Figure 16 illustrates non-specific predictions about the relative differences between conditions, irrespective of the ROI in which they measured. It is primarily intended to illustrate the comparisons under consideration. Identical trial conditions (diagonal that runs bottom left to top right) will have a distance of 0 by definition. We predicted that pairs of trials that vary on every dimension (different costs, decisions, and confidence) would have the greatest distances (diagonal that runs top left to bottom right). Beyond this basic prediction, specifics of how costs, decision outcomes, and confidence would influence representational similarity would depend on the ROI under consideration as described in Chapter 3, as well as the time at which the distance was recorded. Generally speaking, we predicted that distances between pairs of trial conditions would increase with the number of dimensions on which they varied. Furthermore, we expected that the time at which patterns of voxel activation begin to differentiate between trial conditions would be indicative of the type of information being processed there as well as the evolution of the neural representation of that decision content. For example, if a region was involved in representing different costs, then it would show highly similar patterns of activation at the start of all trials. Then, as the decision was deliberated, the patterns of activation between bike trials and shock trials would gradually begin to separate until the choice was settled or until confidence about the choice was estimated. Conversely, a

region that does not encode decision costs would show lesser differentiation in its voxel response patterns between bike and shock trials.

Methods

GLM2b (FIR version for RDA)

Pre-processing of functional imaging data and distance estimation were conducted at the individual participant level until the final stage of group-wide statistical comparisons. GLM2b was re-estimated at the run-level; this time functional data was not spatially smoothed and a finite impulse response (FIR) model was applied. Neural responses were sampled in each of nine consecutive .91s time bins during each trial, beginning with the final .91s of the fixation cue before the offer stimulus appeared. Then, for each individual, for each of the 12 functional runs (6 bike and 6 shock), parameter estimates (PEs) and model residuals from the four trial conditions (AccCon, AccAmb, RejCon, RejAmb) during nine time intervals were extracted from a priori regions of interest.

ROIs

Individuals' parameter estimates from the four trial conditions of the shock and bike tasks (choice X confidence X decision phase X cost = 16 conditions per participant) were extracted from a priori anatomical ROIs: vmPFC, ACC / dACC, amygdala, nAcc, aIC, and pIC, as well as a control ROI, primary auditory cortex (A1). ROIs were anatomically defined from the Harvard-Oxford cortical and subcortical structural probabilistic atlases (https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/Atlases), thresholded to include only voxels that exceeded 50% probability of being located within each structure and then divided into left-

and right-hemisphere components. Most ROIs correspond directly to labels on the Harvard-Oxford atlas, but our vmPFC ROI combines Harvard-Oxford's Frontal Medial Cortex and Subcallosal Cortex (as in Chapter 2) and ACC / dACC combines Cingulate Gyrus anterior division and paracingulate gyrus. aIC and pIC ROIs were created by masking Harvard-Oxford's Insular Cortex with insula subdivisions labeled on the Brainnetome atlas (http://atlas.brainnetome.org): aIC includes ventral agranular, dorsal agranular, and dorsal dysgranular insular cortex from Brainnetome and pIC includes hypergranular, ventral granular and dorsal granular insular cortex. See Figure 14 for an illustration of all ROIs.

Estimating Representational Dissimilarity of Trial Conditions

RDA measures the dissimilarity of patterns in voxel-wise responses. Here, we compared patterns of responses observed in our ROIs during AccCon, AccAmb, RejCon, and RejAmb trial conditions of GLM2b (FIR version). Conceptually, if an ROI has n voxels, each trial condition can be described as a point in n-dimensional space, with its coordinates determined by voxel-wise parameter estimates from the corresponding condition. The dissimilarity between two conditions is then the distance between their respective coordinates, or their locations in that n-dimensional space. We quantified these relationships with de-noised and normalized Euclidean distances.

An ROI's response to a given trial condition was represented as a vector of n parameter estimates (one for each voxel). This was performed at the level of a single FIR time bin, within single ROI, for individual participants. Voxels with missing data from any run, likely caused by signal artifacts during imaging or during spatial normalization to template space, were removed from further analysis - conceptually speaking, this eliminated one dimension (per

missing voxel) from the ROI space. Then, in order to reduce the influence of noisy voxels on distance estimates, each voxel's PEs were noise normalized by weighting each PE with the unbiased within-run variance (Σ). Σ was calculated by multiplying the matrix of model residuals (t timepoints by n voxels in the ROI) by its transpose and dividing by t . ((Walther et al., 2015b) Equation 4). Therefore, the square root of the diagonal of Σ represents the standard deviation of each voxel's residuals across time points from that run. Along that diagonal, larger values correspond to greater amounts of overall variance, therefore when it is divided into the vector of PEs from the corresponding voxel ((Walther et al., 2015b) Equation 5), noisier voxels are down-weighted, resulting in univariate noise reduction.

Next, we calculated *all* pairwise Euclidean distances between the de-noised voxel responses from different trial conditions (separately for each FIR time bin, for each ROI, for each participant). Pairs of trial conditions could vary in decision outcome, confidence, cost, and the experimental run in which they were observed. We computed distances between *all* trial conditions *across runs* rather than measuring only within-run distances, which is another popular method for RDA analyses. Computing distances between runs allowed us to do three crucial things: First, this allowed us to take distances between cost conditions which were never presented in the same run. Second, it allowed us to cross validate distance estimates by taking the mean distance of two conditions computed across all runs. The can eliminate spurious mutivariate similarity based on within run temporal correlation of the residual error (Walther et al., 2015b). Third, this allowed us to estimate the baseline variance of each condition by computing the distance between a given condition and itself in all other runs, and taking the average of all within-condition-between-run pairwise comparisons. Although we initially

computed within-run distances, these were excluded from group-wide analyses in the results presented, as not all distances could be computed at the within-run level.

We took the average Euclidean distances between trial conditions across runs, d(A, B), and normalized them with respect to the baseline between-run variance inherent in the two conditions themselves, d (A, A) and d(B, B). The number of trials per condition could not be counterbalanced across runs by design because they depended on the participants' choice behavior. Cases in which there were only one or two trials of a given condition in a given run, lead to extremely large distance estimates for all comparisons including that condition, including when comparing that same condition to itself in another run (Fig. 15). Therefore, the baseline distances could be used to attenuate this bias without discarding any data. Normalized distances estimates were calculated as:

$$d(A,B)^{*} = \frac{|\, d(A,B)^2 - (d(A,A) \times d(B,B))\,|}{d(A,A) \times d(B,B)}$$

Normalization made it such that d(A, B)* reflected the relative distance between two conditions, beyond variance inherent in the conditions themselves and because d(A,B), d(A,A), and d(B,B) were calculated separately for different ROIs, d(A, B)* also controlled for ROI size as larger ROIs lead to larger d(A, B) values prior to normalization.

Emergence of pattern dissimilarity between costs

In order to visualize the overall pattern and time course of representational similarity between conditions, we plotted group mean d(A, B)* for all comparisons between conditions. (Fig. 16). However, our RDA analyses of interest specifically concerned between-cross comparisons, as we had already observed differences in the magnitude of responses to different trial conditions with the same costs.

We aimed to identify whether and when over the course of a trial, our ROIs revealed cost-specific representations in their patterns of activation. To test this, we isolated all d(A, B)* cases in which A and B only differed by cost (e.g. d(AccConBike, AccConShock)* and compared their representational similarity at sequential time points after the onset of the offer relative to baseline (during the fixation cue). Separate linear mixed effects models, implemented in the lme4 package for R (Bates et al., 2015) predicted representational dissimilarity as a function of the time bin in which they were observed. There were five models, one for each of the trial conditions (AccCon, AccAmb, RejCon, RejAmb), and one for overall dissimilarity collapsed across trial condition. Each took the form:

$$\hat{y} = b_0 + b_1 FIR_{bin} + \varepsilon$$

where $\hat{y}$ predicts d(A, B)*, the fixed effect *FIRbin* is a categorical label for the time bin of the neural response, and we included a random effect on the model intercept across subjects to account for baseline individual differences overall pattern similarity. Because we were only interested in comparing each *FIRbin* to baseline, we only measured select contrasts of each post-cue time bin (FIR 2-9) against the baseline (FIR 1) and report significant effects corrected for multiple comparisons with the Dunnett one-to-many test using the glht and confint functions from the multcomp package for R (Hothorn et al., 2008).

Results

We measured the representational dissimilarity of voxel response patterns during different types of value-based decisions. Neural responses were modeled at nine time points over the course of a decision to examine when representations of task-relevant information were available in our ROIs. Although the primary focus of this analysis was to identify cost-

specific representations (i.e., representational dissimilarity between decisions with bike and shock costs), we also explored representational dissimilarity of all combinations of decision variables over the course of a trial to understand the general time course of these dissociations.

Overview of all pairwise distances

Figure 17 illustrates the pairwise distances between all possible combinations of decision outcomes, confidence, and costs in 6 ROIs and the control region. Each column is a time point (FIR bin) with the offer onset occurring in the second time bin. Note that because an FIR model was implemented, the neural responses may be offset from the marked trial events. Darker colors indicate smaller distances whereas brighter colors are larger distances. The dark diagonal that runs from the bottom left to top right of each FIR bin is the identity comparison – a condition compared to itself – where distances are always zero.

Figure 17 does not illustrate statistical comparisons, but instead depicts a descriptive overview of all observed distances, which is useful for demonstrating general trends in the time course of representational dissimilarity and highlights salient local comparisons. During the earliest time bins, all regions exhibit relatively small distances between conditions, suggesting that distinctions between different task conditions have not yet emerged. Gradually, over the course of the trial (typically beginning around FIR bin 3, ~ 1.8s after trial onset), distances between voxel response patterns begin to grow and by the latest time bins (around FIR bin 8, ~6.4s after trial onset), most response distances were returning towards their baseline values. Note that the control region (bottom row) exhibits patterns in representational dissimilarity that are generally similar to those in the network of value ROIs, indicating that there are global dissimilarities in voxel activation patterns beyond our network of value ROIs. aIC and ACC /

dACC have the longest sustained response pattern differences (see the 8th and 9th FIR bins for each of these ROIs), which are apparent both when comparing across different costs and different choices with the same cost.

Across all ROIs, the largest distances are found along the diagonal that runs bottom right to top left, where pairs of trials differ on every dimension, and particularly comparisons of bike rejects with shock accepts (the top left and bottom right corners of the center squares). The smallest distances are typically between trial conditions that differ only by decision confidence (adjacent to the identity diagonal). Within-cost comparisons (such as bike AccCon versus bike RejCon bike) tend to be smaller for the bike task (bottom left) than the shock task (top right). The control ROI, A1, differs from the value ROIs in that in A1 distances are smaller overall and less stable over the course of the trial, suggesting that the control region isn't encoding value information (which differs between trial conditions) as precisely as the other ROIs.

Cost-Domain-Specific Representations, Across All Conditions

We next analyzed distances between pairs of trials that varied solely by the type of cost in the offer (e.g. ignoring trials with different costs *and* different choices, etc.). Larger distances indicate greater representational dissimilarity and imply that a region is differentially encoding information about the two types of costs. We assumed the logic that a meaningful difference between cost representations could be inferred if there were significantly greater distances between voxel response patterns when the offer was on the display versus before it appeared (during the cue period). Therefore, we compared the representational dissimilarity of decision costs, collapsed across all trial conditions (AccCon, AccAmb, RejCon, and RejAmb), at each

time point (FIR bins 2-9) after the offer appeared against the baseline (FIR bin 1), within each of our ROIs.

Distances were compared by fitting linear mixed effects models to normalized Euclidean distances between voxel response patterns from the bike and shock tasks. Separate models were fit for each ROI. The fixed effects predictor was FIR bin (categorical) and the model included a random intercept term to account for individual differences in overall pattern dissimilarity within each ROI. We statistically compared FIR bins 2-9 against baseline, FIR bin 1, with Dunnett's one-to-many test. All relevant statistics can be found in Table 3, here we summarize the key findings.

Cost distances in the control region, A1, did not differ before and after the offer appeared on the display. That is, no distances in FIR bins 2-8 differed significantly from baseline (FIR bin 1). This is consistent with our assumption that patterns of voxel activation in the control region would not encode meaningful differences between effort and pain costs. However, all value-network ROIs did differ from baseline in at least one FIR bin, post offer onset Fig. 18). Distances in ACC / dACC, aIC, and pIC differed from baseline during FIR bins 3 - 7. There were significant distances from baseline in the amygdala during FIR bins 3, 5, and 6; in vmPFC during FIR bins 5 and 6; and in nAcc during FIR bin 5.

Cost-Domain-Specific Representations, Within Trial Conditions

Our second ROI analysis tested whether and when there were meaningful distances between pain and effort costs that were specific to the trial condition. For example, how similarly are pain and effort costs represented during ambivalent rejects? We again only analyzed distances between pairs of trials that varied solely by the type of cost in the offer,

however this time it was done separately for each of the four trial conditions (e.g. AccConBike versus AccConShock). We fit the same linear mixed models as in the previous analysis, except this time separately for each of the four trial conditions, for each ROI. All relevant statistics can be found in Table 4, and we summarize the key findings here.

The control region, A1, again did not differ from baseline in any trial condition, during any FIR bin. All value ROIs did differ from baseline during at least one FIR bin of at least one trial condition (Fig. 19). In the RejCon condition, there were cost-specific differences in voxel response patterns in ACC / dACC (FIR bins 3, 5) and vmPFC (FIR bins 5, 7). For the RejAmb condition, there were cost-specific differences in ACC / dACC (FIR bins 3-7), vmPFC (FIR bins 3, 5-6), the amygdala (FIR bin 6), and aIC (FIR bin 5). In the AccAmb condition, there were cost-specific differces in vmPFC (FIR bins 4-5, 7), the amygdala (FIR bins 3, 5), aIC (FIR bin 7), and pIC (FIR bin 7). Finally, for the AccCon condition, there were cost-specific differences in ACC (FIR bins 5-6), in vmPFC (FIR bins 5-6), in the amygdala (FIR bins 5-6), and nAcc (FIR bin 7

Discussion

Although the magnitudes of BOLD responses elicited by effort and pain costs did not differ in Chapter 3, here we observed unique patterns of voxelwise activation between the two costs in all of our ROIs, except the control region, A1. This consistent with the hypothesis that our value ROIs encode domain-specific representations of effort and pain costs.

Cost distances were measured at successive time points after the offer appeared on the display and compared against baseline (when the fixation cue was on the display) to identify when domain-specific voxel response patterns began to emerge over the trial. This was done

once for distances collapsed across all trial conditions (AccCon, AccAmb, RejCon, RejAmb), and then once analyzing each trial condition separately. The former revealed that all value ROIs differed from baseline during at least one FIR bin after the onset of the offer, confirming that all ROIs were encoding cost-specific information, irrespective of other choice variables. The differences (collapsed across trial conditions) began the earliest and were longest sustained in aIC, pIC, and ACC / dACC. In the amygdala, cost differences onset and offset relatively earlier, whereas in vmPFC they onset and offset relatively later, and in nAcc distances were only observed in one FIR bin, at the temporal midpoint of the trial. Beyond these subtle differences, the most interesting and theoretically relevant results were observed in analyses of individual trial conditions. Below, we discuss domain-specific cost representations observed in each of our ROIs during different trial conditions.

vmPFC

vmPFC was the only ROI that encoded cost-domain-specific representations in all four trial conditions, suggesting that information about the identity of the cost is included in neural computations of SV. Moreover, there is an intriguing structure to the relative times at which cost-domain-specific information was available in vmPFC. Specifically, during ambivalent choices, when offer valuation is most challenging, vmPFC encodes unique representations of pain and effort costs early in the decision making process. During confident choices, such differences don't emerge until relatively later.

It is noteworthy that the temporal dissociation of cost representations between ambivalent and confident trials in vmPFC corresponds with the temporal dissociation of SV/decision outcome and confidence signals in the same region. In Chapters 2 and 3 we

observed that during the offer phase of each trial (relatively early), vmPFC tracked the SV of the offer stimulus and the decision outcome, whereas during the commit phase (relatively late), vmPFC tracked decision confidence – or the value of one's judgment. Here we show that during ambivalent choices, when offer valuation requires more careful consideration, patterns of voxel activation in vmPFC distinguished between the types of costs earlier in the trial, whereas during confident choices those differences emerged relatively later. Taken together, the similar time courses of these representations suggest that domain-specific cost information co-occurs with and may even be factored into vmPFC representations of SV, decision outcomes, and decision confidence. That is, vmPFC may represent both decision outcomes and choice confidence in a cost-domain-specific manner but at different time points over the course of a decision. We discuss this hypothesis more thoroughly in the general discussion.

ACC / dACC

ACC / dACC encoded cost-domain-specific representations during RejCon, RejAmb, and AccCon trial conditions, confirming numerous previous accounts that this region distinguishes between effort and other types of costs (Rudebeck et al., 2006; Prevost et al., 2010; Burke et al., 2013; Massar et al., 2015). The largest and longest-sustained differences in representational disimilarity were observed in the RejAmb condition, suggesting that cost-specific representations become most critical in highly deliberated decisions that ultimately result in rejecting the offer. In Chapters 1 and 2 we observed ambivalence related responses in ACC / dACC during both the offer and commit phases, consistent with this region's known role in managing response competition and conflict resolution. However, the magnitude of ambivalence signals did not differ between the bike task and the shock task. Here, the timing

of cost distances in the RejAmb condition lines up well with the offer and commit phases of GLM1(b) and GLM2(b), and we find that differences between cost representations were available within ACC / dACC when measured in distances between voxel response patterns rather than response magnitudes.

Amygdala

There were cost-domain-specific voxel response patterns in the amygdala during RejAmb, AccAmb, and AccCon trials. Given the amygdala's known role in processing aversive stimuli, it was somewhat unexpected that we found no differences in cost representations in the RejCon condition in which decision costs are the most salient. However, we did observe relatively later cost-specific differences in the amygdala during RejAmb trials – which may be related to our finding in GLM2b that BOLD responses in this region were inversely correlated with SV. That is, as individuals carefully evaluate offers with borderline negative SV, the costs are particularly salient and the amygdala encodes cost-domain-specific representations of pain and effort. Our other finding, that the amygdala encodes cost-specific voxel responses patterns during AccAmb and AccCon, is consistent with a recent report that the amygdala tracks the rewards of effort costs – although in that study, the effort costs of mention were cognitive efforts (Chong et al., 2017).

aIC,

Cost-specific voxel response patterns were observed in aIC for both ambivalent conditions, but only during one time point for each. In GLM2b (Chapter 3) BOLD responses in aIC were stronger during the commit phase of ambivalent trials versus confident trials,

suggesting sustained conflict resolution processes. Again, the relative magnitudes of these contrasts did not differ between the bike and shock tasks. However, in the RDA analysis using the same conditions, the FIR bins during which we observed cost-specific patterns in voxel activation, which were relatively later in the trial, likely overlap in time with the commit phase in GLM2b. Taken together, the RDA and GLM2b results suggest that cost-specific patterns of activation (RDA) co-occurred with amplification of ambivalence signals (GLM2b) within aIC. That is, not only did aIC preferentially respond to ambivalent decisions over confident decisions, but ambivalence-related activation in this region seems to be cost-domain-specific.

Interestingly, cost-specific patterns of activation arose relatively later for ambivalent *accepts* than ambivalent *rejects*. Given that aIC sends projects efferent inputs to the parasympathetic nervous system, cost-specific representations in aIC while the participant accepts an offer may encode cost-specific preparative information – that is, what the body needs to do in order to face pain or effort costs. Future research may explore this hypothesis more directly. We collected simultaneous physiological recordings. Incorporating this physiological data with the neural data presented here may reveal a link between the emergence cost-specific patterns of activation in aIC and bodily physiological responses.

pIC

pIC demonstrated only one instance of cost-specific neural response patterns – in the AccAmb condition, relatively late in the trial (FIR bin 7). This was unexpected given pIC's known role in receiving and processing visceral afferents (Craig, 2002; 2009)– which suggested that the processing of prospective pain and effort costs could also be encoded in unique representations. Moreover, in GLM2b we observed highly similar activation in pIC as

we observed in vmPFC, suggesting that the two regions were tracking similar information. However, here we found that whereas vmPFC shows cost-domain-specific patterns of voxel responses in *all* conditions pIC shows few. Taken together, it seems that whereas both vmPFC and pIC track SV, decision outcomes, and confidence – in vmPFC these choice variables are encoded in a cost-domain-specific manner, whereas in pIC they are almost exclusively encoded in a cost-domain-general manner.

Nucleus Accumbens

nAcc demonstrated only one difference in response pattern similarity between pain and effort costs, during FIR bin 7 of AccCon trials. It is unsurprising that nAcc selectively encodes cost-specific information with respect to offers with the largest rewards. nAcc has long been understood to respond to reward stimuli (Ikemoto and Panksepp, 1999), but as the field progresses we are understanding more about how costs modulate those responses. nAcc is critical for motivating efforts to obtain rewards (Phillips et al., 2006; Botvinick et al., 2009; Talmi and Pine, 2012) and it is also sensitive to the anticipation and avoidance of pain (Jensen et al., 2003; Shackman et al., 2011). One of its key associations in this process is responding to the prospective prediction of rewards (Knutson et al., 2001). Such anticipatory reward responses may explain our observation that nAcc encodes cost specific information about AccCon trials – as on these trials the participant expects a relatively large reward but one that is accompanied by a task. Above, we hypothesized that late cost-domain-specific confidence signals in vmPFC may be driven by predictions about the consequences of one's choices. Here, we extend the same idea to our results in nAcc. Cost-domain-specific responses during the late

stages of AccCon trials suggest strong predictions about large rewards, that are specific to the nature of the contingent cost.

Limitations

Our experimental paradigm had relatively brief inter-trial intervals (ITIs) of .91 – 2.7s which maximized the number of samples we collected, but can make it difficult to interpret the results of fine-grained temporal analyses. We extended our FIR time bins to last until the offset of the feedback stimulus (the end of the trial), when we could be sure that any BOLD responses recorded would be associated with the stimulus on the current trial. Given the delay of the hemodynamic response, it is likely that neural responses occurring towards the end of the trial were eliminated from analysis. We observed an overall trend through which distances between costs gradually increased from their smallest values at baseline, peaked, and then decreased back towards their initial baseline near the end of the trial. In fact, we observed no distances that were significantly different from baseline from the final two FIR bins (8 and 9), suggesting that most of the differences in cost representations had occurred relatively early in the trial and resolved by the point at which we stopped measuring them. Nonetheless, it is possible that there were differences in cost representations very late in the trial that we did not measure. Future variations of our paradigm that incorporate longer ITIs could test for the presence of differences in voxel response patterns with later onsets.

We also performed the bike and shock versions of the decision making task on separate days. This meant that we could never make within-run or within-session comparisons between bike and shock trials. Therefore, any observations that we make about cost-specific neural responses may also be influenced by differences between the shock and the bike session, such

as overall arousal and mood, physical fatigue or residual effects following familiarization with the shock and bike costs (that preceded the experimental task), or the order in which participants completed the sessions. Future research may seek to integrate both types of cost stimuli into individual runs or sessions, although this could pose significant logistical challenges.

Chapter V: Individual Differences in Preferences and Value Transformations

The analyses conducted in Chapters 2-4 explored behavioral and neural signatures of SV at the group level. Correlating model-based value estimates with neural responses allowed us to control for individual differences in how a given cost-reward pair is perceived. For example, an individual who is intrinsically motivated to exercise may not weigh bike costs as strongly as another individual who doesn't share this motivation, even when controlling for the individuals' physical capacities for exercise. Estimating SV separately for each individual and for each task enabled us to isolate behavioral and neural responses with respect to *perceived* value, regardless of the underlying features of the stimulus. This methodology was useful for highlighting the shared properties of subjective value representations across cost types and individuals and irrespective of the objective value attributes from which they were derived. However, it is unclear whether there are likewise individual differences in *relative* preferences, or if there is some extrinsic or systematic relationship between individuals' preferences and neural valuation mechanisms for pain versus effort costs.

In the present chapter, we use behavioral and neural measures to look more closely at *how* individuals arrived at their choices in both tasks and the extent to which these mechanisms differ between individuals. We specifically aim to determine whether individuals vary in the types of information they prioritize when making decisions about different types of costs. In Chapters 3 and 4 we investigated whether neural representations of value were cost-domain-specific or cost-domain general, which led to somewhat mixed results. One reason this might occur is if there were large individual differences in the relative decision strategies applied between the two tasks. For example, if some individuals prioritized costs in the bike task and

rewards in the shock task whereas others showed the reverse pattern, meaningful cost-domain-specific signals could be washed out by taking an average across the group. Here, we inspect individual differences in both behavioral and neuroimaging recordings to test for individual differences in cost-domain-specific preferences, valuation strategies, and neural value transformations.

Individual differences in preferences

In the academic literature and in everyday life, humans and even non-human primates exhibit remarkably different preferences and choice behaviors when given full freedom to choose subjectively (Premack, 1963; Schunk and Betsch, 2006). The profound individual differences in choice behavior is strong evidence that perceived value is informed not only by objective value information, but also by the individuals' unique experiences, preferences, motivations. However, it is unclear whether individuals likewise vary in the extent to which they use consistent decision strategies across various domains.

One suggestion is that individuals are largely domain-general in their application of decision strategies: while there may be individual differences in preferred decision strategies, individuals are largely consistent with themselves (Betsch, 2011). An example of this principle could be that participants vary in the relative weights and rewards they attribute to costs and rewards, however any one individual always attributes the same reward and cost weights regardless of what they entail (i.e. an individual would give equal priority to bike costs and shock costs).

A second suggestion is that there are individual differences in decision strategies and individuals exhibit domain-specific decision strategies, but all participants make the same

relative adjustments between domains. An example of this principle could be that individuals assign different weights to exercise costs, but all individuals perceive effort costs to be twice as aversive as pain costs. While this may not fit one's intuitions about subjective preferences more generally, within the constraints of a controlled experiment with identical offer stimuli for both costs, it was quite possible that individuals employed a general-purpose valuation heuristic in both tasks or that they used similar value transformations between bike and shock costs.

A third suggestion is that there are individual differences in both decision strategies *and* the domain-specific application of those strategies. An example of this principle could be that individuals vary in the weight they assign to pain costs, and the weights they assign to effort costs vary independently of that. Previous research about has reported individual differences in domain-specific preferences for decision strategies (Pachur and Spaar, 2015). That study asked participants how they make different types of preference-based decisions such as choosing a dress or choosing a doctor. They found that individuals varied not only in whether they would rely on intuitive versus deliberative strategies within choice domains, but that individuals varied in the extent to which their decision strategies changed between tasks. Another study concerning preferences for risky decision making similarly reported that individuals' preferences for risk-taking "was highly domain-specific, i.e. not consistently risk-averse or consistently risk-seeking across all content domains" (Weber et al., 2002).

To test these competing hypotheses, we correlated behavioral model parameters related to individuals' decision boundaries, as well as the proportion of offers they accepted in both the bike and the shock tasks. If individuals' decision strategies are domain-general (e.g. same decisions in both the bike and shock tasks) or if individuals' strategies are domain-specific but

all individuals make the same adjustment between domains (e.g. participants equally adjust the slope of their decision boundary between the bike and shock tasks), then there should be strong correlations between model parameters and decision outcomes at the group level. If there are individual differences both in the relationship between decision strategies on the two tasks, then we should observe no correlation between model correlations or decision outcomes.

Individual differences in neural value transformations

Beyond the behavioral choice data, another way to measure individual differences in the application of valuation strategies is to identify where and how strongly value information is being represented in the brain at the individual level. Here, we aimed to understand whether individuals differ with respect to the decision strategies applied at the neural level. We adopted the following logic: Given that individuals' valuation models diverge both from the objective utility of offer stimuli and from one another, at some point during neural value processing, the objective features of offer stimuli must be transformed into a different subjective value. If the extent to which a brain region's response to a value stimulus changes between the bike and shock tasks predicts the extent to which behavior in response to the stimulus changes between the two tasks, this would suggest that responses in this region track value transformations from objective stimulus features to subjective value information that guides behavior. For example, if for a given individual, a particular brain region were twice as responsive to costs in the bike task than the shock task *and* the participant's behavioral choice models reveal that they weighted bike costs twice as strongly as shock costs, we would have reason to believe that this brain region was encoding information about the transformation of the objective value of a stimulus into its subjective utility.

Our previous analyses parameterized offers by decision outcomes and model-estimated SV, therefore they could not detect such objective to subjective value transformations because we were dealing only with the outputs of those transformations (as estimated by behavioral models). Furthermore, because our previous analyses only parameterized SV we could only identify regions that tracked offer's *overall* subjective valuation, accounting for both its cost and reward attributes. If a module in the value network is specifically involved in the computation or representation of SV, then it would not differentially respond to offers with qualitatively different cost attributes (i.e. pain and exercise) with different objective utility (i.e. 40% vs. 80%), so long as they had equitable overall subjective value. These approaches were also insensitive to individual differences in value transformations, which we additionally aimed to explore here.

To identify potential sites of such value transformations, we compared the relative neural responses to objective cost and rewards in the bike and shock tasks with the relative subjective valuation of bike and shock offers identified from decision behavior. Specifically, we correlated model weights from analysis of neural data (tracking objective value) with model weights from the analysis of behavioral data (tracking subjective value). We predicted that the amplitude of neural responses to objective value information would be correlated with the influence of those value attributes on decision behavior in regions involved in value transformation. In this way, individual differences in the strength of neural responses to the objective features of value stimuli would predict individual differences in the influence of those attributes on behavior.

To test this, we built a third model for analysis of our fMRI dataset that tested for neural responses tracking the objective cost and reward attribute of offer stimuli (irrespective of their

SV or decision outcomes). Then, we extracted those parameter estimates from several a priori anatomical ROIs within the value network and correlated them with cost and reward weights from behavioral valuation models. We were particularly interested in the relative transformations of bike and shock value information, and therefore focused our analyses on comparisons between the two.

Methods

Group-wide Cost Preferences in Decision Behavior

In order to quantify preferences between exercise and pain costs, we examined individuals' decision outcomes and their individual valuation model parameters. We first tested for group-wide pattern in the bike versus shock task by comparing group means for overall proportion of offers accepted as well as model terms for the slope and intercept of the decision boundaries with paired sample t-tests, Bonferroni corrected for (3) multiple comparisons. Before running statistics on the behavioral comparisons, one participant was removed from analyses of model coefficients as their decision boundary intercept and slope coefficients were beyond three SD of the group mean.

Behavior-Behavior Correlations (Individual Differences in Cost Preferences)

While group mean differences between decision behavior and modeled preferences from the bike and shock tasks could reveal overall biases, they might also obscure more nuanced variation at the individual level. We were particularly interested in individual differences in subjective preferences, and specifically whether they could be revealed in our tasks. Notably, because the offer stimuli and trial set were identical in the bike and shock

versions of the task, differences in decision behavior could only be explained by varying subjective valuation of objectively equal pain and exercise costs. The way that participants computed the relative values of bike and shock costs might have been driven by subjective preferences, determined by their unique histories, biases, and internal goals. Conversely, if there were no ostensible differences in valuation strategies or decision behavior between individuals, this may be indicative of a more universal value transformation between pain and exercise costs that is insensitive to subjective preferences.

To test this, we ran three correlations with the same variables (percent of shock offers accepted versus percent of bike offers accepted, decision bound slope in the shock task versus the bike task, and the decision bound intercept in the shock task versus bike task). Confidence intervals were corrected for multiple (3) comparisons. If there were systematic value transformations at the group level, then there should be a strong correlation between observed preferences on the bike and shock tasks. However, if preferences for pain versus exercise costs are truly subjective, then decision behavior and model terms in one cost task should not predict decision behavior on the other at the group level. We additionally provide comparison plots of all individuals' complete choice sets on the bike and shock versions of the task (as in Ch. 2).

GLM3: Parametric analysis of objective rewards and costs

Whereas GLM1(b) and GLM2(b) modeled BOLD responses associated with participants' subjective preferences and decision outcomes, GLM3 measured neural activity modulated by the objective quantities of rewards and costs offered on each trial:

$$Y = \beta_{off} X_{off} + \beta_{com} X_{com} + \beta_{Roff} X_{Roff} + \beta_{Rcom} X_{Rcom} + \beta_{Coff} X_{Coff} + \beta_{Ccom} X_{Ccom} + \beta_{payout} X_{payout} + \varepsilon$$

This analysis is analogous to GLM1(b), described in full detail in Chapters 2 and 3, except the baseline decision phase regressors were parametrically modulated by the reward ($\beta_{Roff}X_{Roff}, \beta_{Rcom}X_{Rcom}$) and costs ($\beta_{Coff}X_{Coff}, \beta_{Ccom}X_{Ccom}$) offered on each trial, irrespective of how the participant valued or decided about them. Notably, because the stimuli in the bike and shock versions of the ApAv task were perceptually identical, significant differences in neural responses correlating with each between the bike and shock tasks would indicate cost-context-specific processing. That is, the neural response to a given offer stimulus was influenced by the type of cost under consideration. This could manifest both in the representation of the cost itself, as well as in the representation of the reward. While both reward stimuli and the meaning of the reward were identical in the bike and shock tasks (rewards always ranged \$0.01-\$1.50), how the reward is perceived and neurally represented may be differentially modulated by the cost with which it is paired.

Anatomical ROIs

Subject-level parameter estimates for the reward and cost regressors of GLM3 were extracted from the same ROIs used in Chapter 4 (vmPFC, ACC, amygdala, aIC, and pIC) and we additionally included the nucleus accumbens (nAcc), also extracted from the Harvard-Oxford atlas (thresholded at 50%) because we predicted this region would be particularly responsive to objective reward stimuli.

Brain-Behavior correlations (Individual Differences in Value Computations)

To test if the strength of neural responses to objective value information from the bike and shock tasks predicted differences in decision behavior, we extracted cost and reward

parameter estimates from the offer phase in GLM3 ($\beta_{shock_{Rew}}, \beta_{shock_{Cost}}, \beta_{bike_{Rew}}, \beta_{shock_{Cost}}$; note the difference in notation from the description in GLM3 above, these terms now represent parameter estimates from the offer phase) and cost and reward weights from individuals decision behavior models ($b_{shock_{Rew}}, b_{shock_{Cost}}, b_{bike_{Rew}}, b_{shock_{Cost}}$). Specifically, for each individual we calculated relative differences in weights from the neural data

$$\frac{\beta_{bike_{Rew}} - \beta_{shock_{Rew}}}{\beta_{bike_{Rew}} + \beta_{shock_{Rew}}} \text{ and } \frac{\beta_{bike_{Cost}} - \beta_{shock_{Cost}}}{\beta_{bike_{Cost}} + \beta_{shock_{Cost}}}$$

with relative differences in weights from the behavioral data

$$(\frac{b_{bike_{Rew}} - b_{shock_{Rew}}}{b_{bike_{Rew}} + b_{shock_{Rew}}} \text{ and } \frac{b_{bike_{Cost}} - b_{shock_{Cost}}}{b_{bike_{Cost}} + b_{shock_{Cost}}}$$

and then correlated the analogous pairs (brain costs with behavior costs; brain rewards with behavior rewards), correcting for multiple comparisons. We predicted that if these regions played a universal and uniform role in transforming objective value of offer stimuli into subjective value representations that inform decision behavior, then there should be a strong correlation between the selected brain and behavioral model terms. However, if individuals differ not only in their relative preferences between bike and shock costs, but also in the neural bases of those preferences, then there should not be a relationship between the brain and behavior terms from the two tasks.

Results

Behavior-Behavior Correlations (Individual Differences in Cost Preferences)

Overall, the exercise costs seemed to be perceived as more aversive than the bike costs. Fewer bike offers were accepted (*M* = .532, *SD* = .111*)* than shock offers (*M* = .618*, SD* =

.104; t(20) = -2.816, p = 0.011, CI_{adj}[-0.165, -0.006]. On average, the intercept of the decision boundary in the shock task (M = -8.166, SD = 21.412) was greater than boundary intercepts in the bike task (M = -21.053, SD = 37.813), but not significantly so (t(19) = -1.585, p = 1.29, CI_{adj}[-34.226, 8.451]) and there was far more variance in the bike task. The boundary intercept is often an index of one's baseline tendency to accept offers relative to rejecting them, therefore this is a second demonstration that participants were more inclined to accept offers with pain costs than they were to accept offers with exercise costs. Decision boundaries were steeper in the bike task (M = 1.36, SD = .546) than the shock task (M = .937, SD = .427; t(19) = 3.99, p = .0009, CI_{adj}[0.144, 0.699]), indicating that the influence of costs relative to rewards was greater in the bike task. Notably, in the shock task the slope of the decision boundary was < 1, suggesting that rewards were more influential than pain costs in participants' decisions. However, in the bike task, the slope of the decision boundary was < 1, suggesting that effort costs were more influential than rewards

Beyond group trends, we specifically wanted to know the extent to which individual differences in preferences for exercise versus pain costs were apparent in their decision behavior and valuation strategies. Although we observed wide variation in preferences within a task (Chapter 2), it was unclear whether these differences reflected cost-independent (i.e. common-currency) valuation strategies or if there were even further diversity in cost-specific preferences. Visual inspection of comparison plots did suggest that there were substantial individual differences in the extent to which individuals adapted their decision strategies to the cost context, suggesting that while some individuals preferences were cost-domain-specific, others were cost-domain-general (Figs. 20 and 21).

To investigate this quantitatively, we correlated individuals' accept rates, and decision boundary slopes and intercepts from the bike task with those from the shock task. There was no meaningful relationship between the proportion of offers a participant accepted in the bike task and the proportion of offers accepted in the shock task (r = .158, p = .494, CI_{adj}[-.384, .619], nor were their decision boundary intercepts significantly correlated (r = .35, p = .130, CI_{adj}[-.212, .738]). However, there was a modestly significant relationship between the slopes of the decision boundaries from the two tasks, after correcting for multiple comparisons (r = .55, p = .0112, CI_{adj}[.040, .834]) (Fig. 22).

Brain-Behavior correlations (Individual Differences in Value Computations)

We found no significant correlations between relative influence of value attributes on decision behavior between the bike and shock tasks and the relative influence of value attributes on neural responses in our anatomical ROIs (Figure 23 for scatter plots and statistics). In fact, only the correlation between $\frac{b_{bike_{Cost}} - b_{shock_{Cost}}}{b_{bike_{Cost}} + b_{shock_{Cost}}}$ and $\frac{\beta_{bike_{Cost}} - \beta_{shock_{Cost}}}{\beta_{bike_{Cost}} + \beta_{shock_{Cost}}}$ from the posterior insula was significant before correction. The general trend of our results showed little relationship between brain and behavior reward weights, however there were more promising patterns in the cost data. The relative differences in cost parameter estimates from the bike and shock tasks in posterior insula, nAcc, and the amygdala showed positive relationships with the relative cost weights in behavioral data from the bike and shock tasks (post. insula r = .526; nAcc r = .395; amygdala r = .397), however these effects failed to reach significance after correcting for multiple comparisons.

Discussion

While individuals acceptance rates and decision boundary intercepts in the bike and shock tasks were unrelated, there was a meaningful relationship between individuals' decision boundary slopes in the two tasks. That is, a participants' decision boundary slope in one task significantly predicted the boundary slope in the other. This may be evidence of a common tendency across individuals to preserve similar cost/reward exchange rates across different types of costs but to adjust one's baseline acceptance rate in a context-dependent manner, according to the subjective averseness of type of cost one is offered.

We did not find strong relationships between behavioral and neural cost transformations between the bike and shock tasks. This was unpredicted given the strong theoretical bases we had for selecting task-relevant ROIs, many of which were recruited for various aspects of the decision making task, as demonstrated in GLM1b and GLM2b. With only twenty-one participants, and in some cases even fewer observations (after removing outlier cases > 3SD outside the group mean), our data was rather noisy, making it difficult to draw definitive conclusions from the analysis.

Still, we believe that the introduction of a novel approach for investigating value transformations with analysis of neuroimaging data will prove to be fruitful in future research. Furthermore, as an exploratory analysis, our results hint at very promising results, demonstrating relationships between neural value transformations in key ROIs and the behaviors they might inspire, opening a couple interesting avenues for further investigation. First, is possible that the suggestions of relationships between neural and behavioral cost representations that failed to reach statistical significance were indicative of true but underpowered effects. First, the observation that differences in reward weights on behavior

between the bike and shock tasks were ostensibly unrelated to differences in neural responses to objective rewards in the value network suggests that changing the type of cost did not influence the neural transformation of objective rewards into subjective reward value in these regions. Second, we observed a tentative relationship between behavioral cost representations and neural cost representations in regions that are sensitive to aversive stimuli, the amygdala and posterior insula, and cost representations. Specifically, as the difference between the relative influences of bike shocks and shock costs grew, so did the difference between the relative neural responses to the objective cost stimuli. In other words, stronger responses in these regions were associated with stronger influences of cost stimuli on decision behavior. Surprisingly, we also observed this pattern in nAcc, which is typically understood to respond primarily to reward stimuli. However, we are not the first to observe such an effect. A previous study found that nAcc responded to aversive conditioning, increasing activation in anticipation of pain stimuli (Jensen et al., 2003), even in the absence of rewards or any possibility of evading the pain stimulus. Other studies have reported that responses in this region track rewards discounted by different types of costs, including effort costs (Botvinick et al., 2009; Croxson et al., 2009; Prevost et al., 2010). It is possible that pattern we see with respect to costs is only one half of the equation. Regardless of what's driving this effect, it may be important to update conceptions about this region's role in value-based decision making.

Alternatively, it is possible that our results were veridical and there isn't a universal value transformation mechanism, at least not within our ROIs. This may indicate that there are significant individual differences in the neural mechanisms recruited for value transformations in value-based decision making. For example, individuals who are motivated by exercise and weary of pain costs may exhibit different patterns neural cost transformations from individuals

who are equally averse to both. Another possibility is that, neural cost transformations are more complex than linear gain functions, as computed in our behavioral models, and therefore the true relationship was undermined by a linear correlation. More sophisticated models of neural value transformations, perhaps drawing non-linear discounting functions documented in neuroeconomic literature will provide a better fit with the observed data.

Chapter VI: General Discussion

Summary of Key Results

The present body of work investigated how the human brain resolves cost/benefit choices. While many forms of value imbue such choices, we paid particular attention to three: SV, confidence, and costs. In Experiment 1, participants underwent fMRI while performing an ApAv decision making task that offered monetary rewards in exchange for enduring painful electrical shocks. In Experiment 2, participants performed the same task, but with a novel implementation of physical costs. This was the first study to use intervals of cardiovascular exercise as cost stimuli in a value-based decision making task. Behavioral and neuroimaging data from Experiments 1 and 2 were analyzed separately (Chapters 2 and 3). Then, these two studies were considered together and in contrast with one another (Chapter 3) in terms of the magnitude of brain activation. Then we reconsidered potential differences of the costs in these experiments by quantifying representational dissimilarity of voxel response patterns (Chapter 4), and characterized individual differences with respect to preferences and neural value transformations (Chapter 5).

Our key novel finding was that value-based decision making entails a dynamic process by which vmPFC shifts from valuation of external stimuli (SV) to valuation of internal signals (confidence). We also found that the magnitudes of neural signals tracking SV, confidence, and decision outcomes did not differ between choices with pain costs and choices with effort costs, consistent with shared neural faculties for value computation. On the other hand, within those faculties, a more sensitive analysis using representational dissimilarity of voxel response patterns between the two costs suggested cost-domain-specific value processing. Specifically,

vmPFC encodes costs differently earlier on in ambivalent decisions than confident decisions, suggesting that early scrutiny of value information is exclusive to conditions of ambivalence.

Analysis of behavior at the individual level revealed a correlation between the relative weights participants assigned to costs and rewards in the bike and shock tasks but no correlation between their general tendencies for approach behavior (proportion accepted offers and decision boundary intercept), suggesting domain-general valuation strategies and domain-specific approach criteria.

Summary of Theoretical Implications

The overall series of results within vmPFC suggest qualitatively different decision processes for ambivalent versus confident decisions. When the environment affords easy choices, careful valuation of the opportunity may be initially bypassed, but value information is later examined to make high-fidelity predictions about the outcomes of a decision, amounting to strong choice confidence. When the environment presents more ambiguous opportunities, value information is promptly scrutinized in detail, but predictions about decision outcomes remain relatively poor, amounting to choice ambivalence.

Discussion

In Experiment 1 (Chapter 2), participants underwent fMRI while performing an ApAv decision making task that offered monetary rewards in exchange for enduring painful electrical shocks. This study was motivated by recent research reporting that vmPFC, a region well understood to mediate SV computations in cost-benefit opportunities (Park et al., 2011; Talmi and Pine, 2012), also tracked decision confidence in both perceptual (Heereman et al., 2015;

Bang and Fleming, 2018; Gherman and Philiastides, 2018) and value-based decisions (De Martino et al., 2013; Lebreton et al., 2015). This led to a useful reframing of the nature of confidence signals. It was suggested that confidence, like SV, is a valuation in its own right measuring the extent to which one can believe that they are making the best choice (Lebreton et al., 2015). Given that accurate decisions are valuable, vmPFC may perform valuations of internal judgments similarly to how it performs valuations of external stimuli.

The results of Experiment 1 showed that while participants evaluated the offer (offer phase), vmPFC tracked SV and decision outcomes but not confidence. When participants submitted a decision (commit phase), vmPFC tracked confidence but not SV or decision outcomes. In Experiment 2, we found that this generalized to a different type of decision cost, cardiovascular exercise, indicating that vmPFC carries similar SV and confidence signals across different domains of value stimuli. These results provide evidence supporting the recent hypothesis that confidence signals in vmPFC are valuations of judgments (Lebreton et al., 2015) and generalize the hypothesis to two new types of decision making tasks.

Note that the matter of timing isn't a trivial one. There has been extensive debate in the literature concerning when confidence emerges over the course of decision making. Some have proposed that confidence evolves concurrently with a decision and therefore emerges relatively early (Kepecs et al., 2008; Lebreton et al., 2009; Chib et al., 2012; Kiani et al., 2014; Gherman and Philiastides, 2015; Dotan et al., 2018; Gherman and Philiastides, 2018), others have reported evidence suggesting that confidence lags behind decision variables and value estimates, emerging closer to the time of choice commitment or even later (Resulaj et al., 2009; Pleskac and Busemeyer, 2010; Hilgenstock et al., 2014; Moran et al., 2015; Yu et al., 2015; Fleming et al., 2018; Morales et al., 2018). Our results decidedly endorse the latter argument,

however given the sizable amount of evidence in support of each hypothesis, it seems worthwhile to consider exactly what is being measured.

We found it useful to conceptualize confidence as a valuation process. Other conceptions of confidence, often described in studies of perceptual decision making, concern degrees of certainty about information in the environment and the reliability of perceptual signals (Kiani and Shadlen, 2009). Note that the two descriptions of confidence mentioned are not mutually exclusive nor are they necessarily different in intended meaning. For both, confidence is the extent to which one can believe they're making the best choice (Pouget et al., 2016; Lebreton et al., 2018), but each takes a slightly different perspective. In perceptual decision making tasks, one is attempting to discern a ground truth about the environment. The ultimate goal is to get an accurate reading of that environment, irrespective of oneself. In these cases, one's sense of confidence is largely dependent on the quality or reliability of information to which one has access – the extent to which they can believe they have accurately perceived the environment. In subjective value-based decision making tasks, in which there are no "correct" answers nor latent optimal strategies, one is attempting to discern a value with respect to their own subjective needs and preferences. The ultimate goal is to maximize future value for oneself. In these cases, one's sense of confidence is largely dependent on predicted outcomes with respect to oneself – the extent to which they can believe that their choice promises a satisfying outcome. With these considerations in mind, it becomes clear that defining confidence in a general way (e.g. the extent to which one can believe they're making the best choice) can take different meanings with respect to the task at hand. This has important implications for how we might interpret results regarding timing of confidence.

In Chapter 4, we measured the representational dissimilarity of voxel response patterns associated with different decision costs during confident and ambivalent decisions to accept or reject offers. In this analysis, vmPFC was the only ROI that encoded cost-domain-specific representations in all four trial conditions, indicating that value representations consistently encoded information about the specific cost consequence of the current offer. Importantly, ambivalent and confident decisions showed systematically different patterns in the timing at which cost information was available.

Offer valuation is obviously a major contributor to the factors that can challenge a choice and lead to ambivalence. Notably, for our task ambivalence is not driven by other contributors, such as a perceptual uncertainty. Instead, the overall SV of the offer under consideration approaches 0 – it is hard to know whether it amounts to a net loss or gain. Due to the difficulty of estimating SV during ambivalent choices, it is likewise difficult to predict the consequences of accepting such an offer with respect to one's future satisfaction. On the other hand, offer valuation is relatively easy during confident choices. The overall SV of the offer under consideration approaches extremely negative and positive values and consequently, one can accurately predict how this will influence one's future satisfaction.

During ambivalent choices, vmPFC encoded unique representations of pain and effort costs *early* in the decision making process. During confident choices, such differences didn't emerge until relatively later. This aligned with results from Chapters 2 and 3, showing that vmPFC first tracked value information about the offer early in the trial, and then tracked information about confidence later in the trial. The similar time courses of these representations suggest that domain-specific cost information informs vmPFC representations of SV, decision outcomes, and confidence. That is, vmPFC may represent both decision outcomes and choice

confidence in a cost-domain-specific manner but at different time points over the course of a decision.

While this is still somewhat indirect evidence of such a relationship, we offer the following interpretations that we believe are consistent with our previous observations and the relevant literature: For ambivalent choices, the offer stimuli are closer to the decision boundary and require careful consideration of exactly what accepting the offer would entail. On these trials, the SV of the offer and ultimately the participant's decision cannot be determined by simple cost-domain-general heuristics (e.g. choose accept if the reward is X times larger than the cost). Instead, cost-domain-specific computations of SV and the decision outcome are computed in real time.

For confident choices, the offer stimuli are farther from the decision boundary and may not require precise valuation in order for the participant to make a decision. The participant could instead rely on automated responses, quick estimations, or heuristics (Korn and Bach, 2019) such as the example given above. These strategies may be cost-domain-general. Participants are more likely to repeat confident decisions in both the bike and shock tasks than they are to repeat ambivalent decisions (details about decision behavior are provided in Chapter 5). Alternatively, decision strategies for confident choices could be set towards the beginning of the task, after which the participant doesn't need to invoke representations of the type of cost under consideration. Note that because we collected the bike and shock data in separate sessions, the participant could use the same strategies throughout a session. In either case, cost-domain-specific representations would not be critical for valuation of the offer stimulus on confident trials, or at least not to the extent that they are required for ambivalent decisions.

We observed that cost-domain-specific representations emerge relatively later (closer to choice commitment) on confident trials. As described in previous chapters, it has been suggested that confidence signals in vmPFC are related to the valuation of one's judgment – assessing how strongly one can believe they've made the right decision (Lebreton et al., 2015). If on confident trials participants took a valuation shortcut to arrive at a decision, bypassing cost-specific information, cost representations could become crucial for assessing confidence insofar as they guide predicting the specific consequences of one's decision. That is, to assess the quality of a decision, a representation of what specifically that decision will entail is needed. In our task, decisions entail costs of pain or effort. Confidence signals may thus invoke cost-specific representations that had been bypassed in the original determination of the decision in order to assess its quality.

These interpretations refine the hypothesis that confidence is a valuation of one's judgment by suggesting that such a mechanism is particularly important for predicting the consequences of one's actions. Accurate decisions are valuable decisions because they facilitate the learning of associations between actions and outcomes, allow us to adapt dynamically to what's ahead, and to maximize the likelihood of future satisfaction.

References

Amemori K-I, Graybiel AM (2012) Localized microstimulation of primate pregenual cingulate cortex induces negative decision-making. Nat Neurosci 15:776–785.

Amemori KI, Amemori S, Graybiel AM (2015) Motivation and Affective Judgments Differentially Recruit Neurons in the Primate Dorsolateral Prefrontal and Anterior Cingulate Cortex. Journal of Neuroscience 35:1939–1953.

Andersson JLR, Jenkinson M, Smith SM (2007a) Non-linear optimisation. FMRIB technical report.

Andersson JLR, Jenkinson M, Smith SM (2007b) Non-linear registration, aka Spatial normalisation.

Arulpragasam AR, Cooper JA, Nuutinen MR, Treadway MT (2018) Corticoinsular circuits encode subjective value expectation and violation for effortful goal-directed behavior. Proc Natl Acad Sci USA 115:E5233–E5242.

Aupperle RL, Melrose AJ, Francisco A, Paulus MP, Stein MB (2015) Neural substrates of approach-avoidance conflict decision-making. Hum Brain Mapp 36:449–462.

Avants BB, Tustison NJ, Song G, Cook PA, Klein A, Gee JC (2011) A reproducible evaluation of ANTs similarity metric performance in brain image registration. NeuroImage 54:2033–2044.

Badre D, Doll BB, Long NM, Frank MJ (2012) Rostrolateral prefrontal cortex and individual differences in uncertainty-driven exploration. Neuron 73:595–607.

Bang D, Fleming SM (2018) Distinct encoding of decision confidence in human medial prefrontal cortex. Proc Natl Acad Sci USA 115:6082–6087.

Bartra O, McGuire JT, Kable JW (2013) The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. NeuroImage 76:412–427.

Basten U, Biele G, Heekeren HR, Fiebach CJ (2010) How the brain integrates costs and benefits during decision making. Proc Natl Acad Sci USA 107:21767–21772.

Bates D, Mächler M, Bolker BM, Walker SC (2015) Fitting linear mixed-effects models using lme4. Journal of Statistical Software 67.

Becerra L, Breiter HC, Wise R, Gonzalez RG, Borsook D (2001) Reward circuitry activation by noxious thermal stimuli. Neuron 32:927–946.

Behrens TEJ, Woolrich MW, Walton ME, Rushworth MFS (2007) Learning the value of information in an uncertain world. Nat Neurosci 10:1214–1221.

Betsch C (2011) Chronic preferences for intuition and deliberation in decision making: Lessons learned about intuition from an individual differences approach. In: Intuition: In Judgment and Decision Making, 1st ed., pp 231–248. Psychology Press.

Boorman ED, Behrens TEJ, Woolrich MW, Rushworth MFS (2009) How Green Is the Grass on the Other Side? Frontopolar Cortex and the Evidence in Favor of Alternative Courses of Action. Neuron 62:733–743.

Botvinick MM (2007) Conflict monitoring and decision making: Reconciling two perspectives on anterior cingulate function. Cognitive, Affective, & Behavioral Neuroscience 7:356–366.

Botvinick MM, Huffstetler S, McGuire JT (2009) Effort discounting in human nucleus accumbens. Cognitive, Affective, & Behavioral Neuroscience 9:16–27.

Brooks AM, Pammi VSC, Noussair C, Capra CM, Engelmann JB, Berns GS (2010) From Bad to Worse: Striatal Coding of the Relative Value of Painful Decisions. Front Neurosci 4.

Burke CJ, Brunger C, Kahnt T, Park SQ, Tobler PN (2013) Neural Integration of Risk and Effort Costs by the Frontal Pole: Only upon Request. Journal of Neuroscience 33:1706–1713.

Chib VS, De Martino B, Shimojo S, O'Doherty JP (2012) Neural Mechanisms Underlying Paradoxical Performance for Monetary Incentives Are Driven by Loss Aversion. Neuron 74:582–594.

Chong TTJ, Apps M, Giehl K, Sillence A, Grima LL, Husain M (2017) Neurocomputational mechanisms underlying subjective valuation of effort costs. PLoS Biol 15:e1002598.

Cieslak M, Ryan WS, Macy A, Kelsey RM, Cornick JE, Verket M, Blascovich J, Grafton S (2015) Simultaneous acquisition of functional magnetic resonance images and impedance cardiography. Psychophysiol 52:481–488.

Coombs CH, Avrunin GS (1977) Single-peaked Functions and the Theory of Preference. Psychological Review 84:216–230.

Cosmides L, Tooby J (1994) Better than Rational: Evolutionary Psychology and the Invisible Hand. The American Economic Review 84:327–332.

Craig AD (2002) How do you feel? Interoception: the sense of the physiological condition of the body. Nat Rev Neurosci 3:655–666.

Craig ADB (2009) How do you feel--now? The anterior insula and human awareness. Nature Publishing Group 10:59–70.

Croxson PL, Walton ME, O'Reilly JX, Behrens TEJ, Rushworth MFS (2009) Effort-Based Cost-Benefit Valuation and the Human Brain. Journal of Neuroscience 29:4531–4541.

De Martino B, Fleming SM, Garrett N, Dolan RJ (2013) Confidence in value-based choice. Nat Neurosci 16:105–110.

De Martino B, Kumaran D, Ben Seymour, Dolan RJ (2006) Frames, Biases, and Rational Decision-Making in the Human Brain. Science 313:684–687.

Dotan D, Meyniel F, Dehaene S (2018) On-line confidence monitoring during decision making. Cognition 171:112–121.

Dyer JS, Fishburn PC, Steuer RE, Wallenius J, Zionts S (1992) Multiple Criteria Decision Making, Multiattribute Utility Theory: The Next Ten Years. Management Science 38:645–654.

Economides M, Guitart-Masip M, Kurth-Nelson Z, Dolan RJ (2014) Anterior cingulate cortex instigates adaptive switches in choice by integrating immediate and delayed components of value in ventromedial prefrontal cortex. Journal of Neuroscience 34:3340–3349.

Ernst M, Nelson EE, McClure EB, Monk CS, Munson S, Eshel N, Zarahn E, Leibenluft E, Zametkin A, Towbin K, Blair J, Charney D, Pine DS (2004) Choice selection and reward anticipation: an fMRI study. Neuropsychologia 42:1585–1597.

Fleming SM, Huijgen J, Dolan RJ (2012) Prefrontal contributions to metacognition in perceptual decision making. Journal of Neuroscience 32:6117–6125.

Fleming SM, van der Putten EJ, Daw ND (2018) Neural mediators of changes of mind about perceptual decisions. Nat Neurosci 21:617–624.

Frederick S, Loewenstein G, O'Donoghue T (2002) Time discounting and time preference: A critical review. Journal of Economic Literature 40:351–401.

Gherman S, Philiastides MG (2015) Neural representations of confidence emerge from the process of decision formation during perceptual choices. NeuroImage 106:134–143.

Gherman S, Philiastides MG (2018) Human VMPFC encodes early signatures of confidence in perceptual decisions. eLife 7.

Grabenhorst F, Rolls ET (2011) Value, pleasure and choice in the ventral prefrontal cortex. Trends in Cognitive Sciences 15:56–67.

Hayden BY (2018) Economic choice: the foraging perspective. Current Opinion in Behavioral Sciences 24:1–6.

Hebscher M, Barkan-Abramski M, Goldsmith M, Aharon-Peretz J, Gilboa A (2016) Memory, Decision-Making, and the Ventromedial Prefrontal Cortex (vmPFC): The Roles of Subcallosal and Posterior Orbitofrontal Cortices in Monitoring and Control Processes. Cereb Cortex 26:4590–4601.

Hebscher M, Gilboa A (2016) A boost of confidence: The role of the ventromedial prefrontal cortex in memory, decision-making, and schemas. Neuropsychologia 90:46–58.

Heereman J, Walter H, Heekeren HR (2015) A task-independent neural representation of subjective certainty in visual perception. Front Hum Neurosci 9:195–12.

Hilgenstock R, Weiss T, Witte OW (2014) You'd Better Think Twice: Post-Decision Perceptual Confidence. NeuroImage 99:323–331.

Hogan PS, Galaro JK, Chib VS (2018) Roles of Ventromedial Prefrontal Cortex and Anterior Cingulate in Subjective Valuation of Prospective Effort. Cerebral Cortex 101:470.

Hothorn T, Bretz F, Westfall P (2008) Simultaneous Inference in General Parametric Models. Biom J 50:346–363.

Hsu M, Bhatt M, Adolphs R, Tranel D, Camerer CF (2005) Neural systems responding to degrees of uncertainty in human decision-making. Science 310:1680–1683.

Ikemoto S, Panksepp J (1999) The role of nucleus accumbens dopamine in motivated behavior: A unifying interpretation with special reference to reward-seeking. Brain Research Reviews 31:6–41.

Jenison RL, Rangel A, Oya H, Kawasaki H, Howard MA (2011) Value Encoding in Single Neurons in the Human Amygdala during Decision Making. Journal of Neuroscience 31:331–338.

Jenkinson M, Bannister P, Brady M, Smith S (2002) Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images. NeuroImage 17:825–841.

Jenkinson M, Smith S (2001) A global optimisation method for robust affine registration of brain images. Med Image Anal 5:143–156.

Jensen J, McIntosh AR, Crawley AP, Mikulis DJ, Remington G, Kapur S (2003) Direct activation of the ventral striatum in anticipation of aversive stimuli. Neuron 40:1251–1257.

Kable JW, Glimcher PW (2007) The neural correlates of subjective value during intertemporal choice. Nat Neurosci 10:1625–1633.

Kahneman D, Tversky A (1979) Prospect Theory: An Analysis of Decision under Risk. Econometrica 47:263–292.

Kahnt T, Heinzle J, Park SQ, Haynes J-D (2011a) Decoding different roles for vmPFC and dlPFC in multi-attribute decision making. NeuroImage 56:709–715.

Kahnt T, Heinzle J, Park SQ, Haynes J-D (2011b) Decoding different roles for vmPFC and dlPFC in multi-attribute decision making. NeuroImage 56:709–715.

Keeney RL, Raiffa H (1993) Decisions with Multiple Objectives. Cambridge University Press.

Kepecs A, Uchida N, Zariwala HA, Mainen ZF (2008) Neural correlates, computation and behavioural impact of decision confidence. Nature 455:227–231.

Kerns JG, Cohen JD, MacDonald AW, Cho RY (2004) Anterior cingulate conflict monitoring and adjustments in control. Science 303:1023–1026.

Kiani R, Corthell L, Shadlen MN (2014) Choice Certainty Is Informed by Both Evidence and Decision Time. Neuron 84:1329–1342.

Kiani R, Shadlen MN (2009) Representation of confidence associated with a decision by neurons in the parietal cortex. Science 324:759–764.

Klein-Flugge MC, Kennerley SW, Friston K, Bestmann S (2016) Neural Signatures of Value Comparison in Human Cingulate Cortex during Decisions Requiring an Effort-Reward Trade-off. Journal of Neuroscience 36:10002–10015.

Klein-Flügge MC, Kennerley SW, Saraiva AC, Penny WD, Bestmann S (2015) Behavioral Modeling of Human Choices Reveals Dissociable Effects of Physical Effort and Temporal Delay on Reward Devaluation Torres-Oviedo G, ed. PLoS Comput Biol 11:e1004116–e1004131.

Knutson B, Adams CM, Fong GW, Hommer D (2001) Anticipation of increasing monetary reward selectively recruits nucleus accumbens. Journal of Neuroscience 21:RC159.

Korn CW, Bach DR (2019) Minimizing threat via heuristic and optimal policies recruits hippocampus and medial prefrontal cortex. Nature Human Behaviour:1–17.

Kriegeskorte N (2008) Representational similarity analysis – connecting the branches of systems neuroscience. Front Syst Neurosci:1–28.

Kringelbach ML (2005) The human orbitofrontal cortex: linking reward to hedonic experience. Nat Rev Neurosci 6:691–702.

Kurniawan IT, Guitart-Masip M, Dayan P, Dolan RJ (2013) Effort and Valuation in the Brain: The Effects of Anticipation and Execution. Journal of Neuroscience 33:6160–6169.

Kurniawan IT, Seymour B, Talmi D, Yoshida W, Chater N, Dolan RJ (2010) Choosing to Make an Effort: The Role of Striatum in Signaling Physical Effort of a Chosen Action. J Neurophysiol 104:313–321.

Lebreton M, Abitbol R, Daunizeau J, Pessiglione M (2015) Automatic integration of confidence in the brain valuation signal. Nat Neurosci 18:1159–1167.

Lebreton M, Jorge S, Michel V, Thirion B, Pessiglione M (2009) An Automatic Valuation

System in the Human Brain: Evidence from Functional Neuroimaging. Neuron 64:431–439.

Lebreton M, Langdon S, Slieker MJ, Nooitgedacht JS, Goudriaan AE, Denys D, van Holst RJ, Luigjes J (2018) Two sides of the same coin: Monetary incentives concurrently improve and bias confidence judgments. Sci Adv 4:eaaq0668.

Levy DJ, Glimcher PW (2012) The root of all value: a neural common currency for choice. Current Opinion in Neurobiology 22:1027–1038.

Lopez-Persem A, Domenech P, Pessiglione M (2016) How prior preferences determine decision-making frames and biases in the human brain. eLife 5:2308.

Massar SAA, Libedinsky C, Weiyan C, Huettel SA, Chee MWL (2015) Separate and overlapping brain areas encode subjective value during delay and effort discounting. NeuroImage 120:104–113.

Miller NE (1944) Experimental studies of conflict. In: Personality and the behavior disorders, pp 431–465. Personality and the behavior disorders.

Montague PR, King-Casas B (2007) Efficient statistics, common currencies and the problem of reward-harvesting. Trends in Cognitive Sciences 11:514–519.

Morales J, Lau H, Fleming SM (2018) Domain-General and Domain-Specific Patterns of Activity Supporting Metacognition in Human Prefrontal Cortex. Journal of Neuroscience 38:3534–3546.

Moran R, Teodorescu AR, Usher M (2015) Post choice information integration as a causal determinant of confidence: Novel data and a computational account. Cognitive Psychology 78:99–147.

Neumann von J, Morgenstern O (1944) *Theory of Games and Economic Behavior*.

Nichols TE, Holmes AP (2002) Nonparametric permutation tests for functional neuroimaging: a primer with examples. Hum Brain Mapp 15:1–25.

O'Doherty JP (2004) Reward representations and reward-related learning in the human brain: insights from neuroimaging. Current Opinion in Neurobiology 14:769–776.

Pachur T, Spaar M (2015) Domain-specific preferences for intuition and deliberation in decision making. Journal of Applied Research in Memory and Cognition 4:303–311.

Padoa-Schioppa C (2011) Neurobiology of Economic Choice: A Good-Based Model. Annu Rev Neurosci 34:333–359.

Padoa-Schioppa C, Assad JA (2006) Neurons in the orbitofrontal cortex encode economic value. Nature 441:223–226.

Park SQ, Kahnt T, Rieskamp J, Heekeren HR (2011) Neurobiology of Value Integration: When Value Impacts Valuation. Journal of Neuroscience 31:9307–9314.

Pessiglione M, Delgado MR (2015) The good, the bad and the brain: neural correlates of appetitive and aversive values underlying decision making. Current Opinion in Behavioral Sciences 5:78–84.

Phillips PEM, Walton ME, Jhou TC (2006) Calculating utility: preclinical evidence for cost–benefit analysis by mesolimbic dopamine. Psychopharmacology 191:483–495.

Plassmann H, O'Doherty JP, Rangel A (2010) Appetitive and Aversive Goal Values Are Encoded in the Medial Orbitofrontal Cortex at the Time of Decision Making. Journal of Neuroscience 30:10799–10808.

Pleskac TJ, Busemeyer JR (2010) Two-stage dynamic signal detection: a theory of choice, decision time, and confidence. Psychological Review 117:864–901.

Pochon JB, Riis J, Sanfey AG, Nystrom LE, Cohen JD (2008) Functional Imaging of Decision Conflict. Journal of Neuroscience 28:3468–3473.

Pouget A, Drugowitsch J, Kepecs A (2016) Confidence and certainty: distinct probabilistic quantities for different goals. Nat Neurosci 19:366–374.

Premack D (1963) Rate differential reinforcement in monkey manipulation. J Exp Anal Behav 6:81–89.

Preuschoff K, Quartz SR, Bossaerts P (2008) Human Insula Activation Reflects Risk Prediction Errors As Well As Risk. Journal of Neuroscience 28:2745–2752.

Prevost C, Pessiglione M, Metereau E, Clery-Melin ML, Dreher JC (2010) Separate Valuation Subsystems for Delay and Effort Decision Costs. Journal of Neuroscience 30:14080–14090.

Rangel A, Camerer C, Montague PR (2008) A framework for studying the neurobiology of value-based decision making. Nat Rev Neurosci 9:545–556.

Resulaj A, Kiani R, Wolpert DM, Shadlen MN (2009) Changes of mind in decision-making. Nature 461:263–266.

Rolls ET, Grabenhorst F, Deco G (2010) Choice, difficulty, and confidence in the brain. NeuroImage 53:694–706.

Rudebeck PH, Behrens TE, Kennerley SW, Baxter MG, Buckley MJ, Walton ME, Rushworth MFS (2008) Frontal Cortex Subregions Play Distinct Roles in Choices between Actions and Stimuli. Journal of Neuroscience 28:13775–13785.

Rudebeck PH, Walton ME, Smyth AN, Bannerman DM, Rushworth MFS (2006) Separate neural pathways process different decision costs. Nat Neurosci 9:1161–1168.

Rushworth MF, Kolling N, Sallet J, Mars RB (2012) Valuation and decision-making in frontal cortex: one or many serial or parallel systems? Current Opinion in Neurobiology 22:946–955.

Rushworth MF, Mars RB, Summerfield C (2009) General mechanisms for making decisions? Current Opinion in Neurobiology 19:75–83.

Rushworth MFS, Behrens TEJ (2008) Choice, uncertainty and value in prefrontal and cingulate cortex. Nat Neurosci 11:389–397.

Schlund MW, Brewer AT, Magee SK, Richman DM, Solomon S, Ludlum M, Dymond S (2016) The tipping point: Value differences and parallel dorsal-ventral frontal circuits gating human approach-avoidance behavior. NeuroImage 136:94–105.

Schmidt L, Lebreton M, Cléry-Melin M-L, Daunizeau J, Pessiglione M (2012) Neural Mechanisms Underlying Motivation of Mental Versus Physical Effort. PLoS Biol 10:e1001266.

Schunk D, Betsch C (2006) Explaining heterogeneity in utility functions by individual differences in decision modes. Journal of Economic Psychology 27:386–401.

Shackman AJ, Salomons TV, Slagter HA, Fox AS, Winter JJ, Davidson RJ (2011) The integration of negative affect, pain and cognitive control in the cingulate cortex. :1–14.

Shenhav A, Straccia MA, Botvinick MM, Cohen JD (2016) Dorsal anterior cingulate and ventromedial prefrontal cortex have inverse roles in both foraging and economic choice. Cognitive, Affective, & Behavioral Neuroscience 16:1127–1139.

Shenhav A, Straccia MA, Cohen JD, Botvinick MM (2014) Anterior cingulate engagement in a foraging context reflects choice difficulty, not foraging value. Nat Neurosci 17:1249–1254.

Singer T, Critchley HD, Preuschoff K (2009) A common role of insula in feelings, empathy and uncertainty. Trends in Cognitive Sciences 13:334–340.

Skvortsova V, Palminteri S, Pessiglione M (2014) Learning To Minimize Efforts versus Maximizing Rewards: Computational Principles and Neural Correlates. Journal of Neuroscience 34:15621–15630.

Smith DV, Hayden BY, Truong T-K, Song AW, Platt ML, Huettel SA (2010) Distinct value signals in anterior and posterior ventromedial prefrontal cortex. Journal of Neuroscience 30:2490–2495.

Smith SM (2002) Fast robust automated brain extraction. Hum Brain Mapp 17:143–155.

Sokol-Hessner P, Hutcherson C, Hare T, Rangel A (2012) Decision value computation in DLPFC and VMPFC adjusts to the available decision time. Eur J Neurosci 35:1065–1074.

Taggart P, Critchley H, van Duijvendoden S, Lambiase PD (2016) Significance of neuro-cardiac control mechanisms governed by higher regions of the brain. Autonomic Neuroscience: Basic and Clinical 199:54–65.

Talmi D (2012) How costs influence decision values for mixed outcomes. :1–21.

Talmi D, Dayan P, Kiebel SJ, Frith CD, Dolan RJ (2009) How Humans Integrate the Prospects of Pain and Reward during Choice. Journal of Neuroscience 29:14617–14626.

Talmi D, Pine A (2012) How Costs Influence Decision Values for Mixed Outcomes. Front Neurosci 6.

Tanaka SC, Doya K, Okada G, Ueda K, Okamoto Y, Yamawaki S (2004) Prediction of immediate and future rewards differentially recruits cortico-basal ganglia loops. Nat Neurosci 7:887–893.

Treadway MT, Buckholtz JW, Cowan RL, Woodward ND, Li R, Ansari MS, Baldwin RM, Schwartzman AN, Kessler RM, Zald DH (2012) Dopaminergic Mechanisms of Individual Differences in Human Effort-Based Decision-Making. Journal of Neuroscience 32:6170–6176.

Tsakiris M, Haggard P (2003) Awareness of somatic events associated with a voluntary action. Experimental Brain Research 149:439–446.

Tversky A, Shafir E (1992) Choice under Conflict: The Dynamics of Deferred Decision. Psychological Science 3:358–361.

van den Berg R, Anandalingam K, Zylberberg A, Kiani R, Shadlen MN, Wolpert DM (2016) A common mechanism underlies changes of mind about decisions and confidence. eLife 5:e12192.

Volz LJ, Welborn BL, Gobel MS, Gazzaniga MS, Grafton ST (2017) Harm to self outweighs benefit to others in moral decision making. Proc Natl Acad Sci USA 114:7963–7968.

Walther A, Nili H, Ejaz N, Alink A, Kriegeskorte N, Diedrichsen J (2015a) Reliability of dissimilarity measures for multi-voxel pattern analysis. NeuroImage:1–13.

Walther A, Nili H, Ejaz N, Alink A, Kriegeskorte N, Diedrichsen J (2015b) Reliability of dissimilarity measures for multi-voxel pattern analysis. NeuroImage:1–13.

Wan X, Cheng K, Tanaka K (2015) Neural encoding of opposing strategy values in anterior and posterior cingulate cortex. Nat Neurosci 18:752–759.

Weber EU, Blais A-RE, Betz NE (2002) A domain-specific risk-attitude scale: measuring risk perceptions and risk behaviors. J Behav Decis Making 15:263–290.

Westbrook A, Lamichhane B, Braver T (2019) The Subjective Value of Cognitive Effort is Encoded by a Domain-General Valuation Network. Journal of Neuroscience.

Winkler AM, Ridgway GR, Webster MA, Smith SM, Nichols TE (2014) Permutation inference for the general linear model. NeuroImage 92:381–397.

Wittmann M, Leland DS, Paulus MP (2007) Time and decision making: differential contribution of the posterior insular cortex and the striatum during a delay discounting task. Experimental Brain Research 179:643–653.

Woolrich MW, Ripley BD, Brady M, Smith SM (2001) Temporal Autocorrelation in Univariate Linear Modeling of FMRI Data. NeuroImage 14:1370–1386.

Wunderlich K, Rangel A, O'Doherty JP (2010) Economic choices can be made using only stimulus values. Proc Natl Acad Sci USA 107:15005–15010.

Yacubian J (2006) Dissociable Systems for Gain- and Loss-Related Value Predictions and Errors of Prediction in the Human Brain. Journal of Neuroscience 26:9530–9537.

Yu S, Pleskac TJ, Zeigenfuse MD (2015) Dynamics of postdecisional processing of confidence. Journal of Experimental Psychology: General 144:489–510.

Zangemeister L, Grabenhorst F, Schultz W (2016) Neural Basis for Economic Saving Strategies in Human Amygdala-Prefrontal Reward Circuits. Current Biology 26:3004–3013.

Tables and Figures

Figure 1: Experimental task and variables of interest

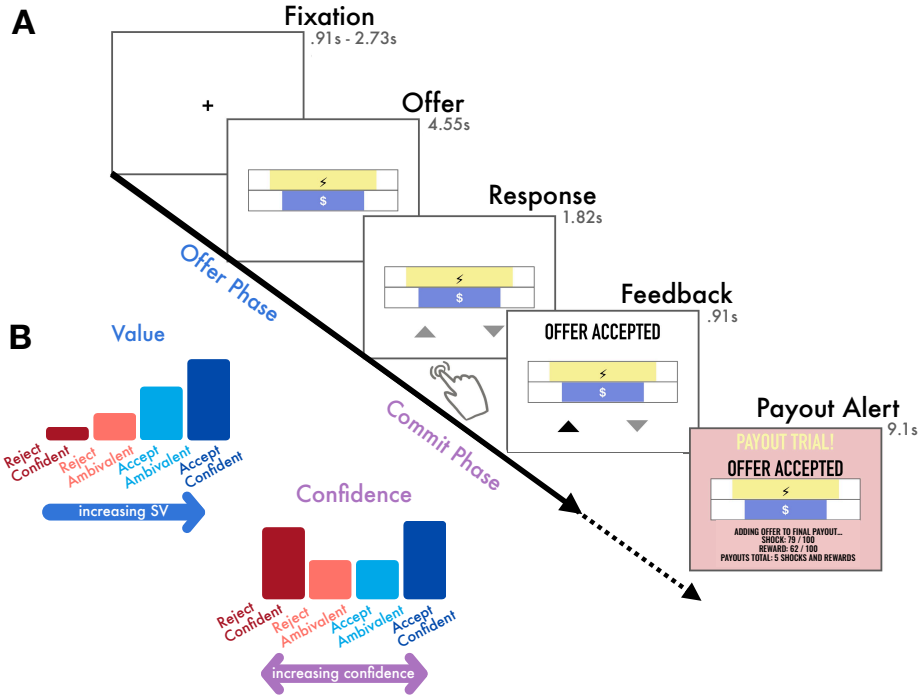

(A) Experimental Task: On each trial, the participant was offered one monetary reward ($ bar width represents amount, ranging continuously $0.01-$1.50) contingent on enduring 1 painful shock (lightning bolt bar width represents pain intensity, ranging continuously from minimally to maximally painful). Participants were instructed to use the offer phase to evaluate the offer and to decide if they would accept or reject it, but they were not yet able to respond. During the commit phase, response mappings appeared and participants made a left or right button press according to the location of the triangle representing their choice (up triangle = accept, down triangle = reject), which varied randomly between trials to prevent preparation of motor responses during the offer phase. After submitting a response, the corresponding triangle was highlighted. Finally, feedback indicating whether the offer was accepted or rejected was added. On payout trials (10 random trials of 189 total), a payout alert followed the feedback. If the participant had accepted the offer, they would receive the monetary reward and also endure the shock at the end of the task, otherwise the participant would receive neither. **(B)** Illustration of variables of interest: We were interested in two types of value inherent in economic decision making: the perceived value of the offer stimulus (SV) accounting for its cost and reward attributes, and the value of one's judgment (confidence), which measures the extent to which one believes they are making the best decision. For accept decisions (blue bars), there is a positive relationship between SV and confidence: as SV increases, one becomes increasingly confident that accepting the offer is the best decision. However, for reject decisions (red bars), there is an inverse relationship between SV and confidence: one becomes increasingly confident about rejecting offers as SV *decreases.* We tested whether they are evaluated simultaneously or during different phases of decision making.

Figure 2: Observed decisions and model-based estimates of confidence and value

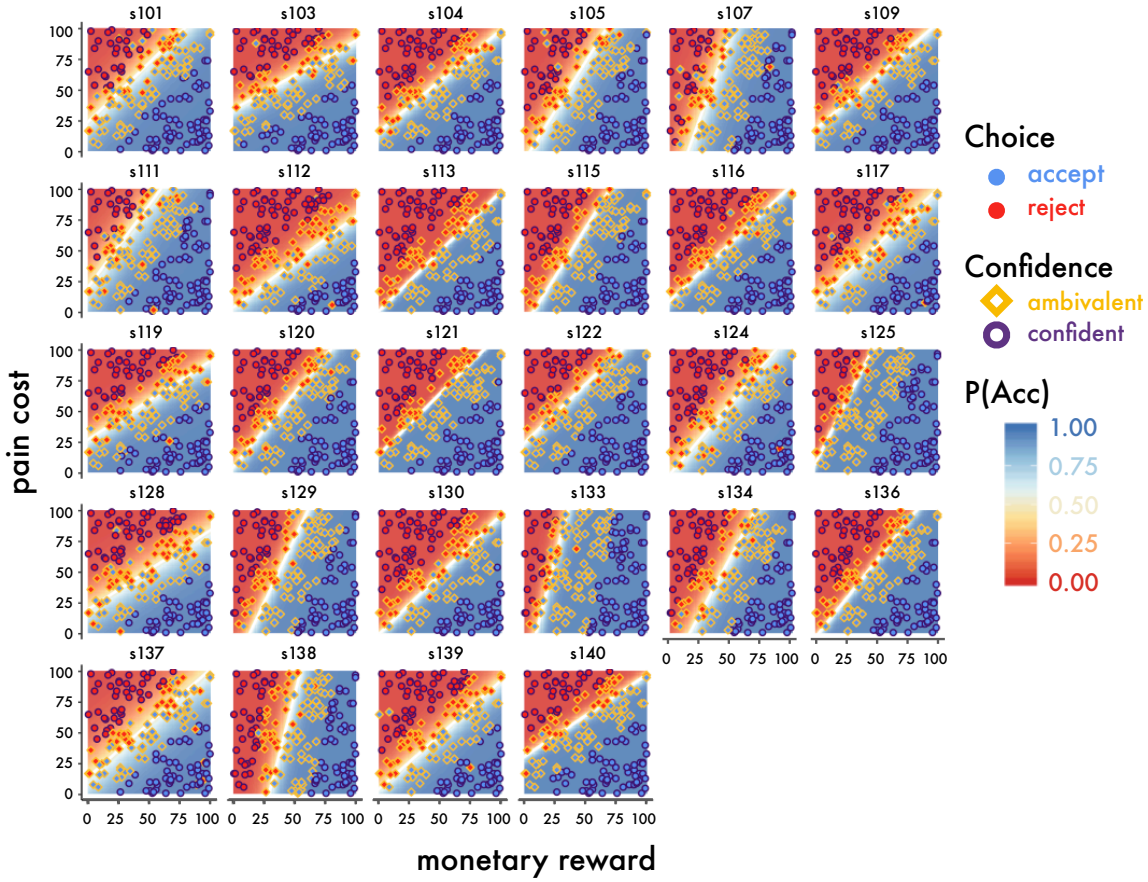

Individuals' choice outcomes are plotted over the decision space, which is shaded according to model-based P(Acc). Each coordinate is a possible offer, with the x-dimension representing percent maximum reward and the y-dimension representing the percent maximum pain cost. Observed choices are overlaid points, the filled color represents the observed decision (red = reject, blue = accept). The point outline represents model-based confidence (yellow = ambivalent, purple = confident). Decision boundaries are overlaid in white. There were substantial individual differences in choice behavior and model-estimated value, as well as differences in choice consistency (indicated by the width of the band of neutral color surrounding the decision boundary). Note that decision boundaries diverged considerably from objective perceptual equality (where reward = cost, not displayed).

Figure 3: GLM1: Parametric correlation with SV and –Confidence

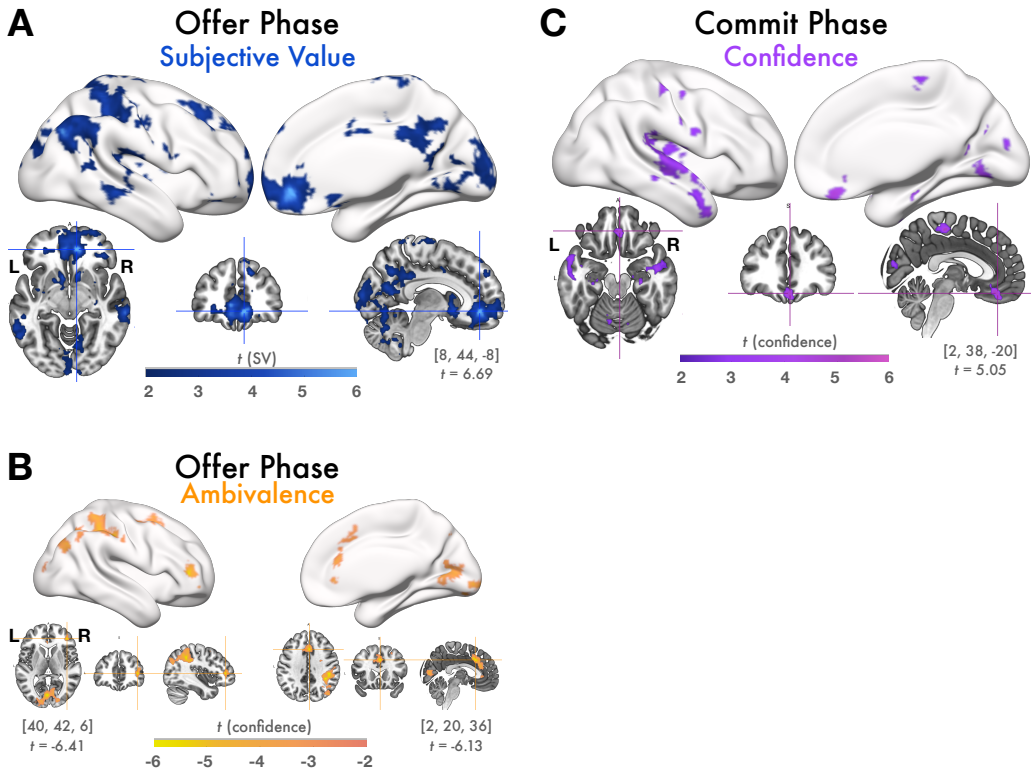

During the offer phase, vmPFC tracked SV but not confidence and during the commit phase, vmPFC tracks confidence but not SV. Figures display all statistically significant results at p<.05 TFCE-corrected at the whole-brain level. **(A)** BOLD responses during the offer phase that correlated positively with model-estimated SV. During the offer phase, a large cluster of voxels in vmPFC tracked SV while participants evaluated the offer. Slice images show local peak activation in vmPFC, along with the MNI coordinates and t-value of the local maximum in this cluster. There was also significant activation throughout the value network including within posterior cingulate cortex, the basal ganglia, insula, and hippocampus; regions involved with value-comparison such as angular gyrus and lateral temporal cortex; and visual cortex. No voxels correlated positively with SV during the commit phase, suggesting that value-responses, particularly in vmPFC, emerge relatively early in the decision making process. **(B)** BOLD responses positively correlated with ambivalence (inversely correlated with confidence) during the offer phase. Areas involved with cognitive control and response competition, such as lPFC and dACC tracked ambivalence while participants deliberated accepting or rejecting the offer. Slice images show local maxima in regions of theoretical interest (lPFC and dACC). No voxels in vmPFC or elsewhere responded positively with choice confidence during the offer phase, suggesting that neural responses corresponding to high confidence emerge relatively later than SV and ambivalence. **(C)** BOLD responses correlated positively with model-estimated confidence during the commit phase. During the commit phase, while participants submitted a response and viewed feedback about their decision, a cluster of voxels in vmPFC tracked decision confidence, as well as regions including posterior insula and lateral temporal cortex. Slices images show peak activation in vmPFC.

Figure 4: GLM2: Contrasts of decisions and confidence conditions

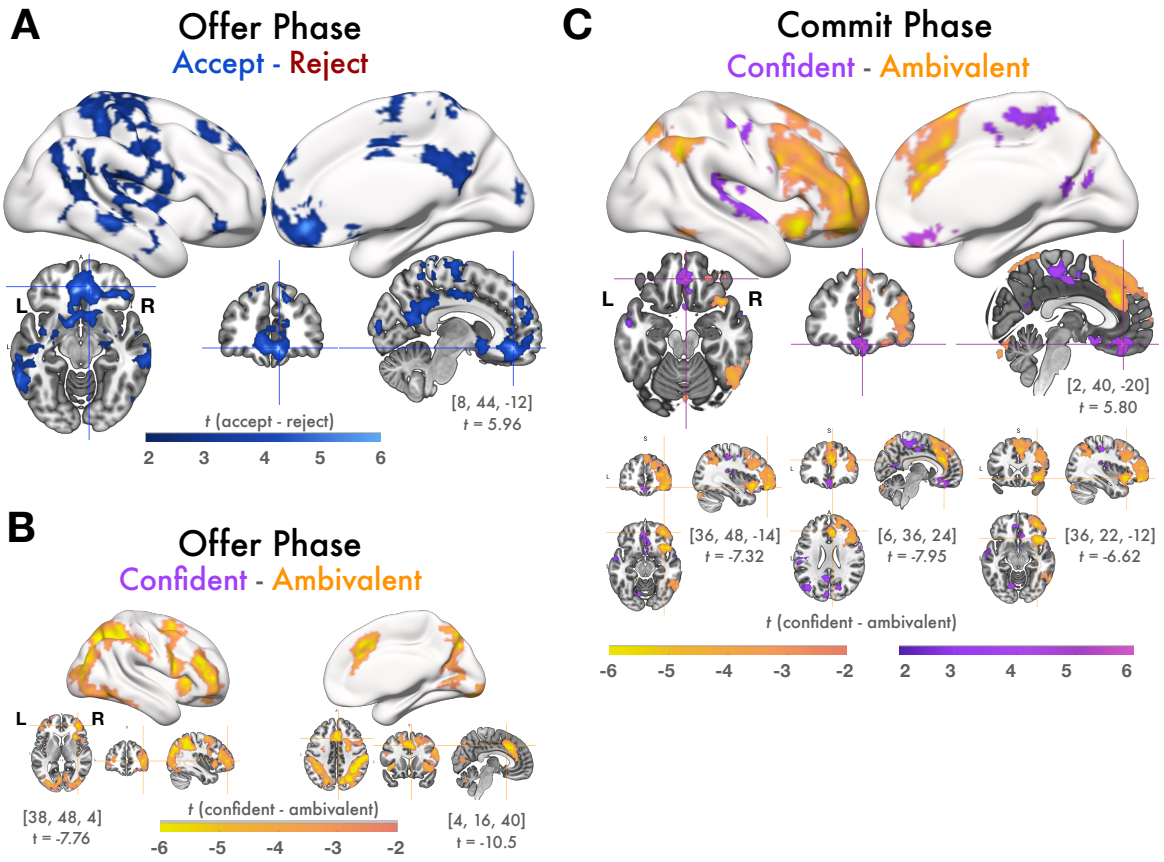

vmPFC encodes decision variables during the offer phase but not the commit phase, and encodes decision confidence during the commit phase, but not the offer phase. Figures display all statistically significant results at p<.05 TFCE-corrected at the whole-brain level. **(A)** Contrasts of response magnitudes from the offer phase of accept trials > reject trials: During the offer phase, clusters of voxels in vmPFC and other regions in the value network, many of which were also parametrically correlated with SV, had stronger response magnitudes preceding accept decisions than reject decisions. Slices images show peak coordinates in vmPFC. No voxels exhibited this contrast during the commit phase, nor were there any regions with stronger responses during rejected trials than accepted trials during either phase. **(B)** Contrasts of response magnitudes from the offer phase of confident trials < ambivalent trials. Slice images illustrate peak activation in lPFC and dACC. **(C)** Contrasts of response magnitudes from the commit phase of confident > ambivalent trials (purple) were observed in vmPFC were stronger for confident trials than ambivalent trials. No regions showed the same contrast during the offer phase. Clusters of voxels in lateral OFC, the anterior insula, and dACC showed the reverse pattern, with stronger response magnitudes during the commit phase of ambivalent trials than confident trials. Note that, several regions (dACC, lPFC, lateral OFC, and anterior insula) show the confident < ambivalent contrast during both the offer and commit phases. Slice images show peak coordinates for confident > ambivalent in vmPFC and peak coordinates for confident < ambivalent in lateral OFC, dACC, and anterior insula.

Figure 5: vmPFC ROI analysis

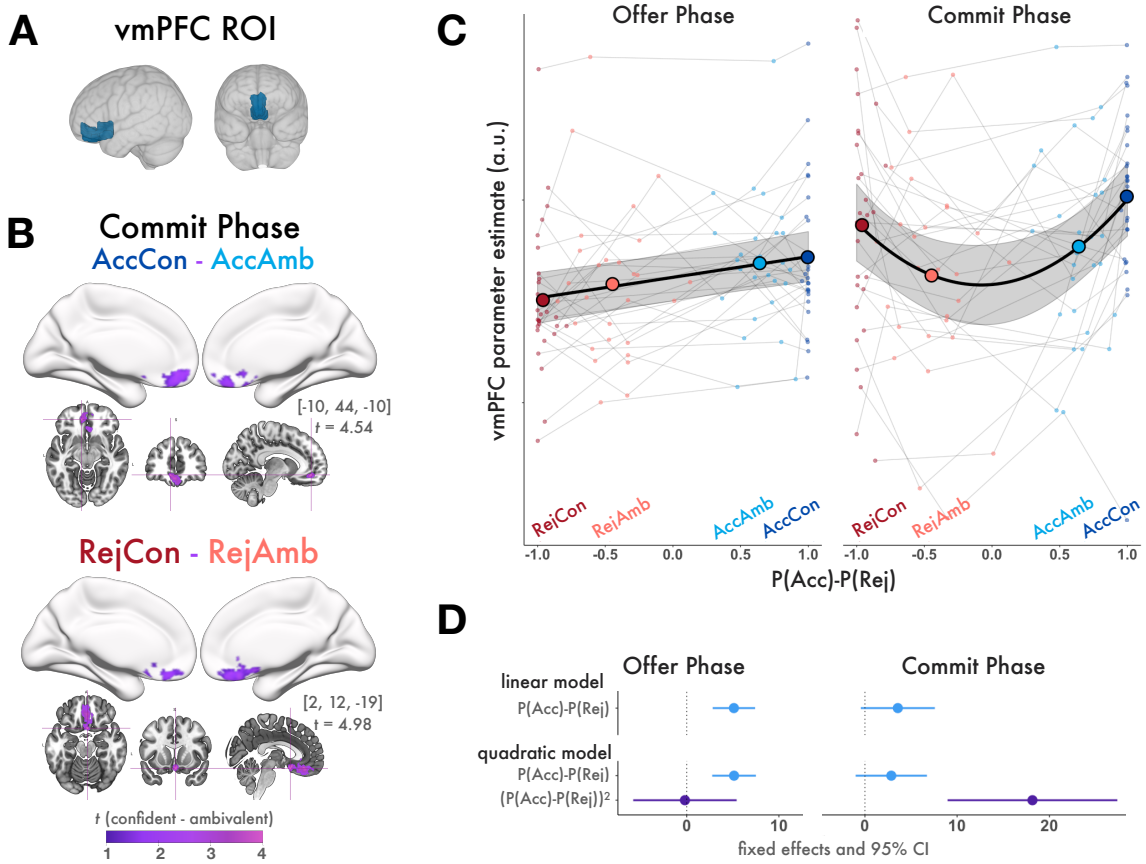

(A) anatomical vmPFC ROI **(B)** AccCon > AccAmb and RejCon > RejAmb contrasts of BOLD responses during the commit phase. Figures display all statistically significant results at p<.05 TFCE-corrected within the vmPFC ROI. During the commit phase, vmPFC responses are stronger for high confidence choices, even for reject choices when confidence is inversely related with SV. **(C)** vmPFC Parameter estimates from the offer (left) and commit phase (right). Small points are individuals' raw vmPFC parameter estimates with respect to their mean P(Acc)-P(Rej) of each trial condition. Large points are group mean vmPFC parameter estimates with respect to the group mean P(Acc)-P(Rej) of each condition. The fit line indicates regression predictions and 95% CI from the winning model for that decision phase. During the offer phase, vmPFC responses increased linearly across trial conditions with increasing value. During the commit phase, they increased quadratically. **(D)** Comparison of mixed effects regression models predicting vmPFC parameter estimates (GLM2) from the and quadratic (bottom) extensions of value. Model estimates and 95% confidence intervals show fixed effects of linear and quadratic model terms, confidence intervals not spanning 0 are considered significant. vmPFC parameter estimates from the offer phase (left) were best fit with a linear model, indicating that during the offer phase, vmPFC tracks value. vmPFC parameter estimates from the commit phase (right) were best fit with a quadratic model, indicating that during the commit phase vmPFC tracks confidence.

Figure 6: RT analysis

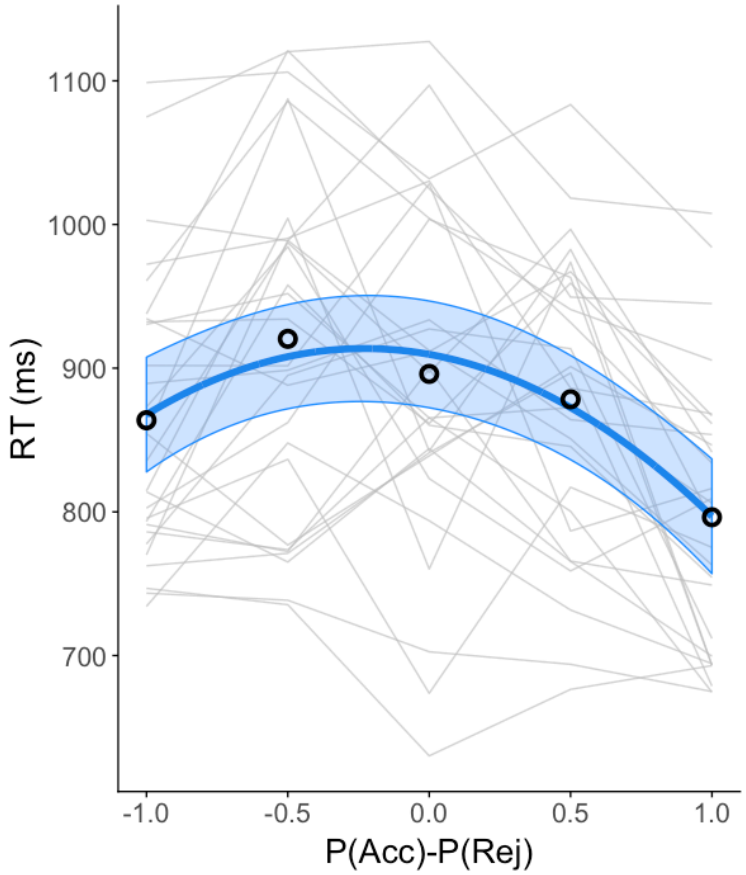

Individual and group mean RTs (small lines) are plotted with respect to value (P(Acc)-P(Rej)). Fit lines represent model predicted RTs and 95% CIs. RTs take an inverse quadratic function with respect to model-estimated value. The time it took participants to commit to decisions was negatively correlated with model-estimated confidence (the quadratic extension of value).

Figure 7: Schematic of experimental session procedures for Experiments 1 and 2

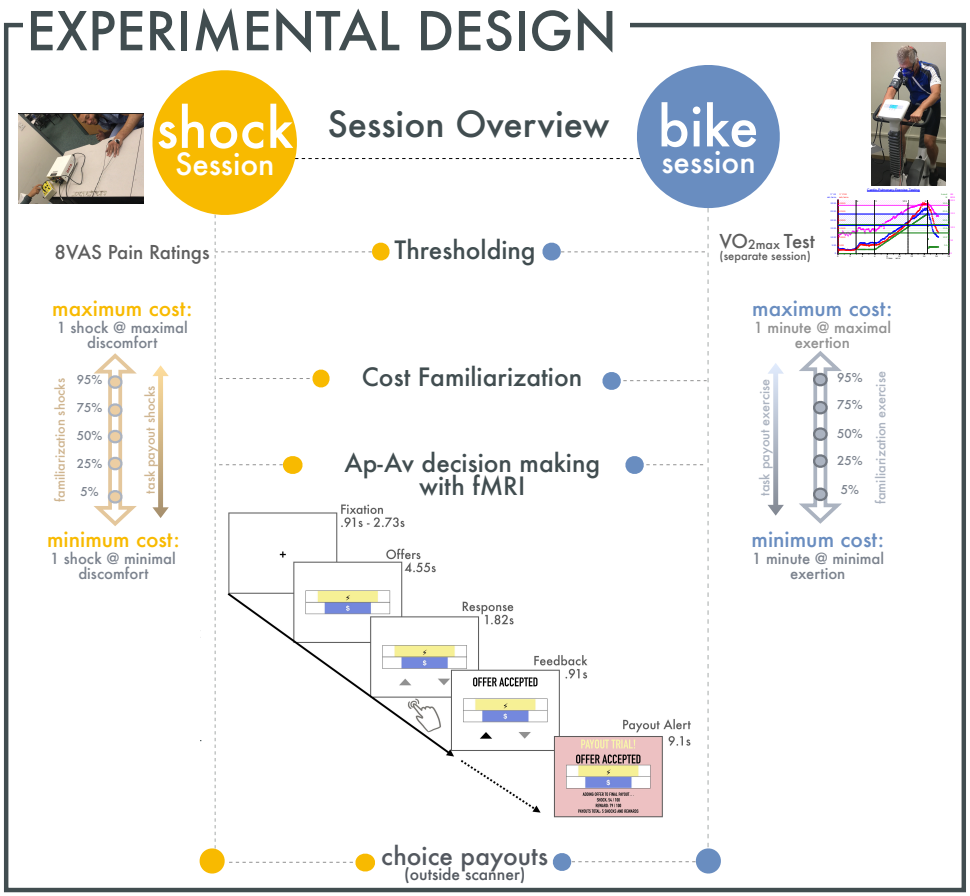

Both the bike and the shock task followed nearly identical procedures. The main difference between the two was that the bike task required cost thresholding to be completed on a separate day, prior to the experimental task. Both pain and effort costs were calibrated to the individual's capacity, with shocks ranging from minimally painful to maximally painful and bike intervals ranging from minimally effortful to maximally effortful. Decision costs ranged continuously between these values and varied independently of decision rewards. The task design, stimuli, and trial order were identical for the two tasks, however the trials randomly selected as payouts differed.

Figure 8: GLM1b offer phase: parametric correlation with SV and –confidence

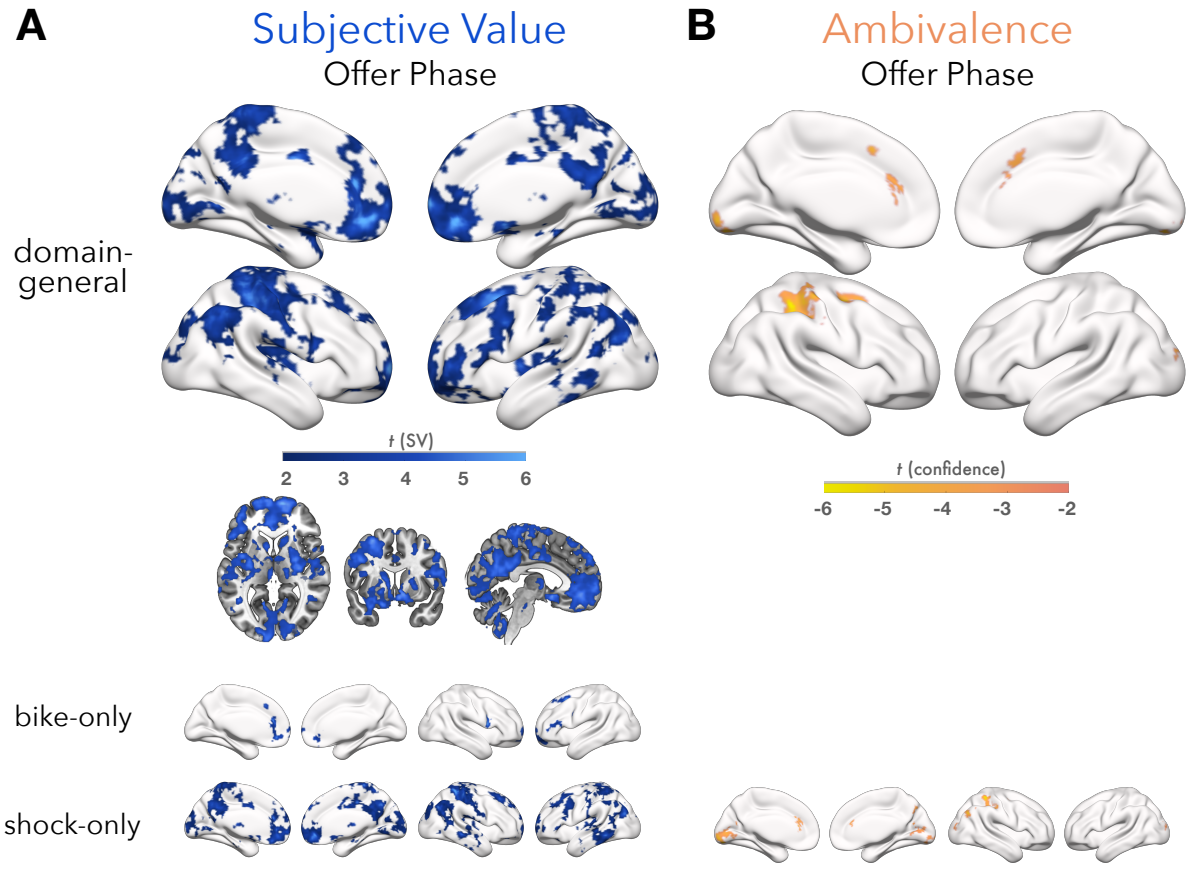

During the offer phase, vmPFC tracked SV but not confidence. dACC tracked ambivalence. The top panel shows analysis of BOLD responses from the offer phase, combining data from both the bike and shock tasks (domain-general). The bottom panel is the results of each task separately. Figures display all statistically significant results at $p<.05$ TFCE-corrected at the whole-brain level. **(A)** BOLD responses during the offer phase that correlated positively with model-estimated SV. During the offer phase, a large cluster of voxels in vmPFC tracked SV while participants evaluated the offer. Slice images show local activation throughout the value network including within vmPFC, pIC, nAc and other structures in the basal ganglia, OFC, lPFC, and posterior cingulate cortex. **(B)** BOLD responses positively correlated with ambivalence (inversely correlated with confidence) during the offer phase. dACC tracked ambivalence while participants deliberated accepting or rejecting the offer. No voxels in vmPFC or elsewhere responded positively with choice confidence during the offer phase, suggesting that neural responses corresponding to high confidence emerge relatively later than SV and ambivalence.

Figure 9: GLM1b commit phase: parametric correlation with confidence

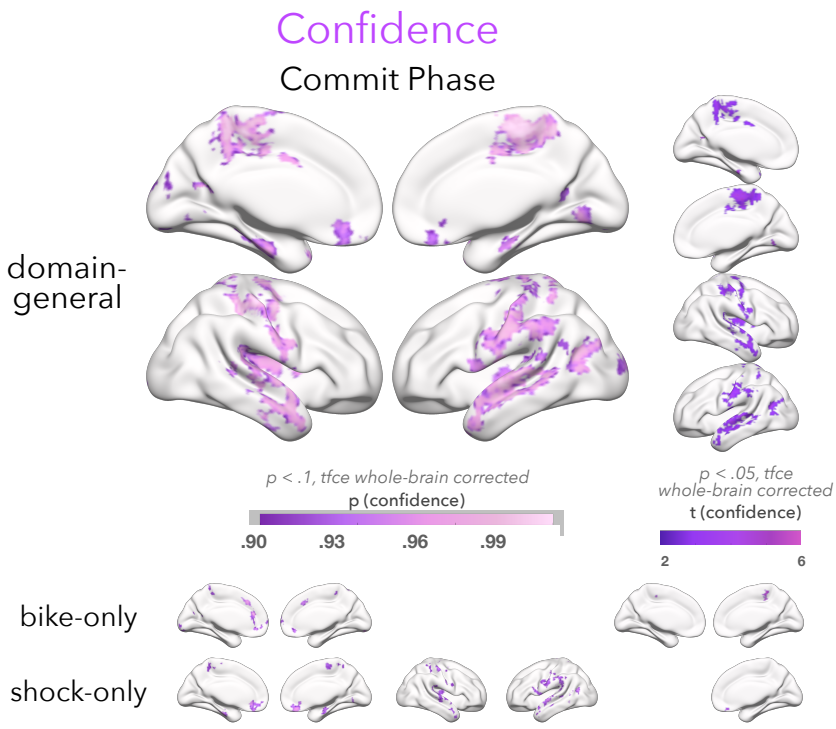

BOLD responses correlated positively with model-estimated confidence during the commit phase. On the left are voxels passing the reduced statistical threshold of $p < .1$, tfce whole-brain corrected. On the right are voxels passing the original statistical threshold of $p < .05$, tfce whole-brain corrected. Top panels are combined analyses (domain-general), bottom panels are from the bike- and shock-only analyses. During the commit phase, while participants submitted a response and viewed feedback about their decision, a cluster of voxels in vmPFC tracked choice confidence, as well as local activation in pIC, lateral temporal cortex and inferior parietal cortex, and precuneus. No voxels correlated positively with SV during the commit phase, suggesting that positive value signals within in vmPFC, emerge and then fade relatively early in the decision making process.

Figure 10: GLM1b commit phase: parametric correlation with -SV

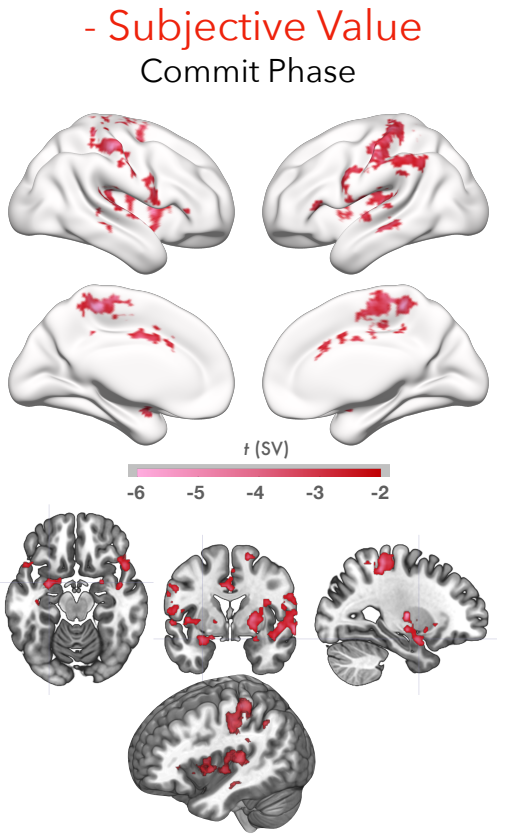

BOLD responses correlated negatively with model-estimated SV during the commit phase. Figures display all statistically significant results at p<.05 TFCE-corrected at the whole-brain level, from the domain-general analysis only. Clusters of voxels in the amygdala, inferior ACC and mid-cingulate, aIC, pIC, precuneus, and lateral parietal cortex responded inversely to subjective value. Slice images show local activation in the amygdala and rendered cortex with cutout shows activation throughout the insula.

Figure 11: GLM2b offer phase: contrasts of decisions and confidence conditions

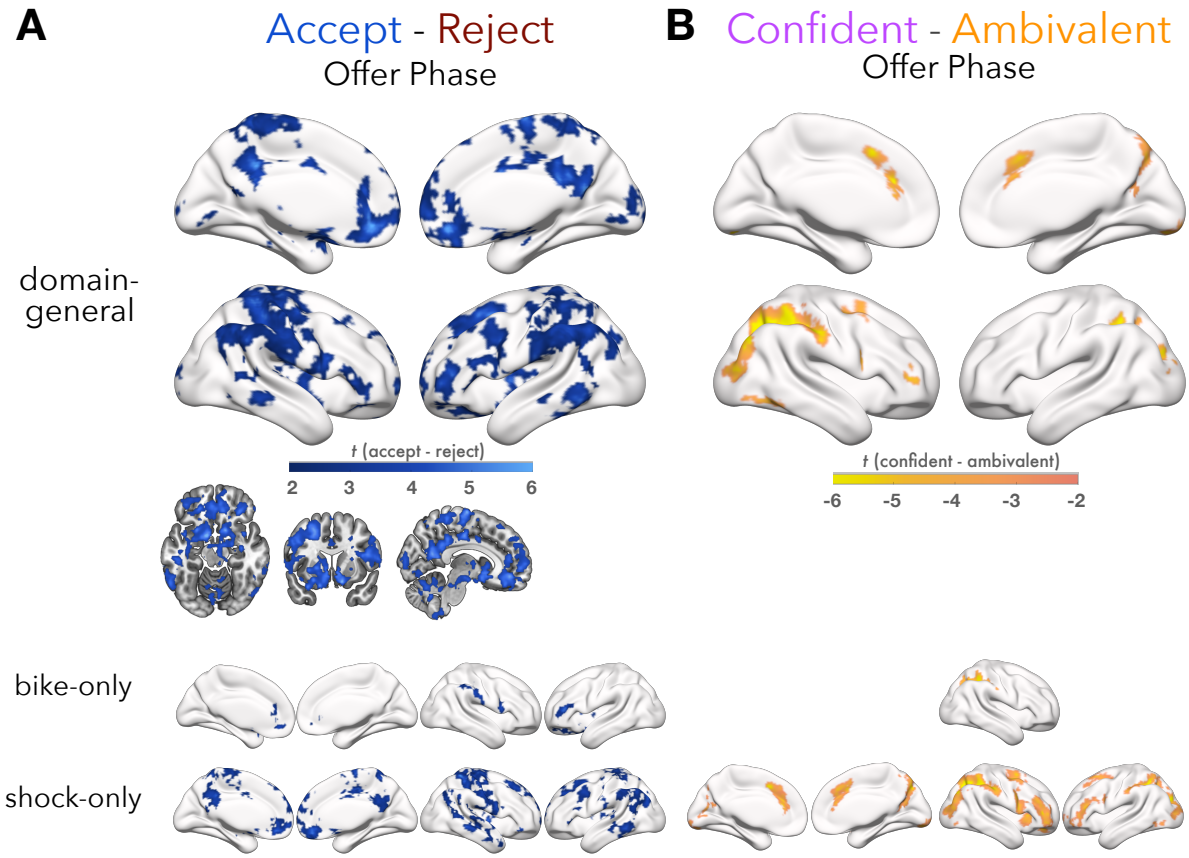

During the offer phase, vmPFC encodes decision variables and dACC and rlPFC respond more strongly to ambivalent trials than confident trials. Top panel is combined analysis, bottom panel is separate bike and shock analyses. Figures display all statistically significant results at p<.05 TFCE-corrected at the whole-brain level. **(A)** Contrasts of response magnitudes from the offer phase of accept trials > reject trials: During the offer phase, clusters of voxels in vmPFC and other regions in the value network, many of which were also parametrically correlated with SV, had stronger response magnitudes preceding accept decisions than reject decisions. Slice images show local activation in the value network including within vmPFC, pIC, nAc and throughout the basal ganglia, OFC, lPFC, and posterior cingulate cortex. No voxels exhibited this contrast during the commit phase, nor were there any regions with stronger responses during rejected trials than accepted trials during either phase. **(B)** Contrasts of response magnitudes from the offer phase of confident trials < ambivalent trials. Slice images illustrate peak activation in lPFC and dACC.

Figure 12: GLM2b commit phase: contrasts of decisions and confidence conditions

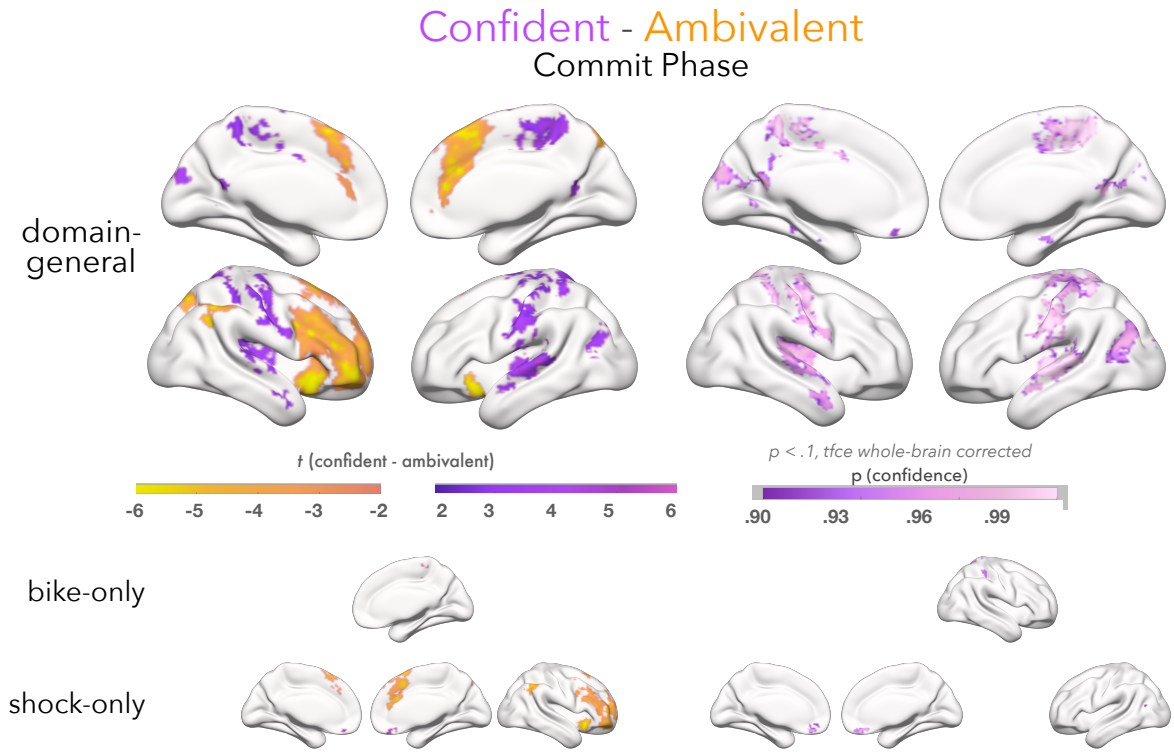

Left: Commit phase contrasts of confident > ambivalent (purple) and confident < ambivalent (orange) significant at original statistical threshold of $p < .05$ tfce, whole brain corrected. Right: Voxels passing reduced statistical threshold of $p < .1$ tfce, whole brain corrected for the confident > ambivalent contrast only.

Figure 13: GLM2b commit phase: vmPFC ROI analysis

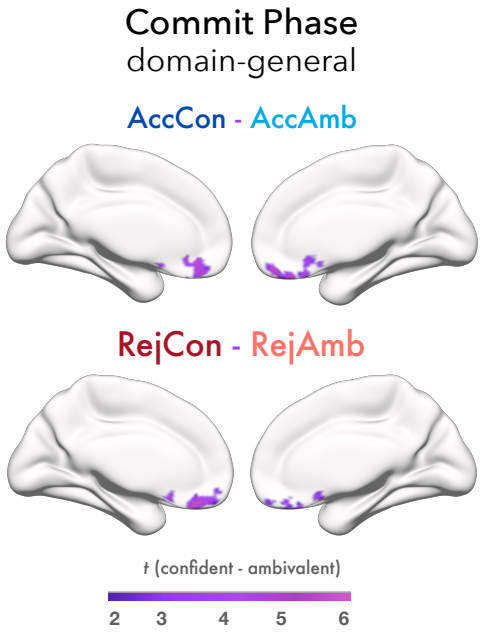

AccCon > AccAmb and RejCon > RejAmb contrasts of BOLD responses during the commit phase. Figures display all statistically significant results at p<.05 TFCE-corrected within the vmPFC ROI. During the commit phase, vmPFC responses are stronger for high confidence choices, even for reject choices when confidence is inversely related with SV.

Figure 14: Illustration of ROIs

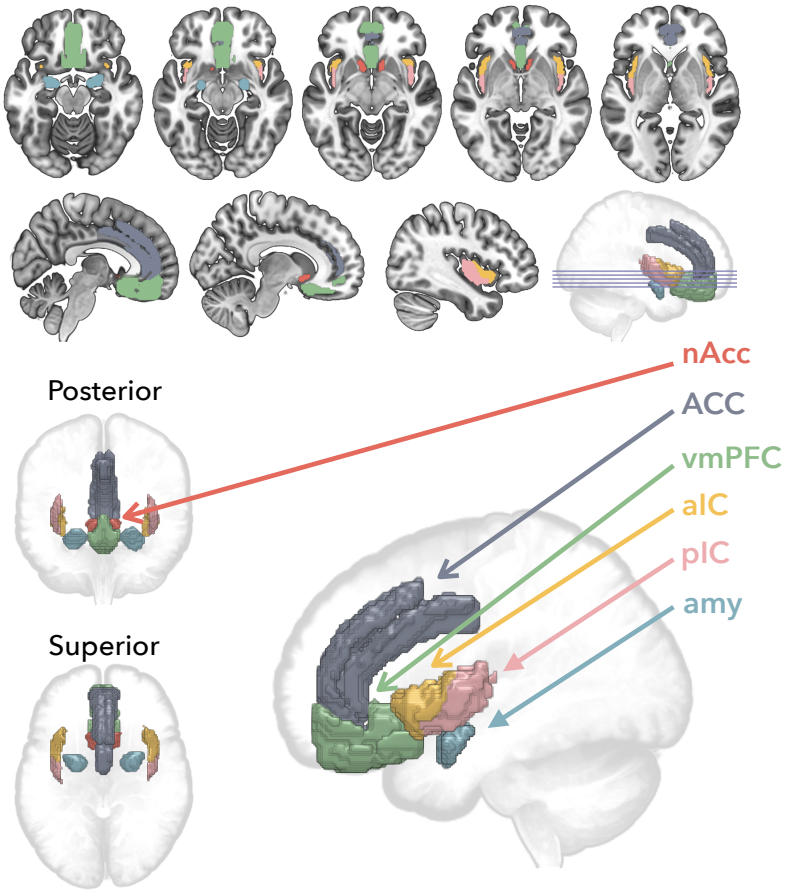

Figure 15: Representational dissimilarity distance normalization

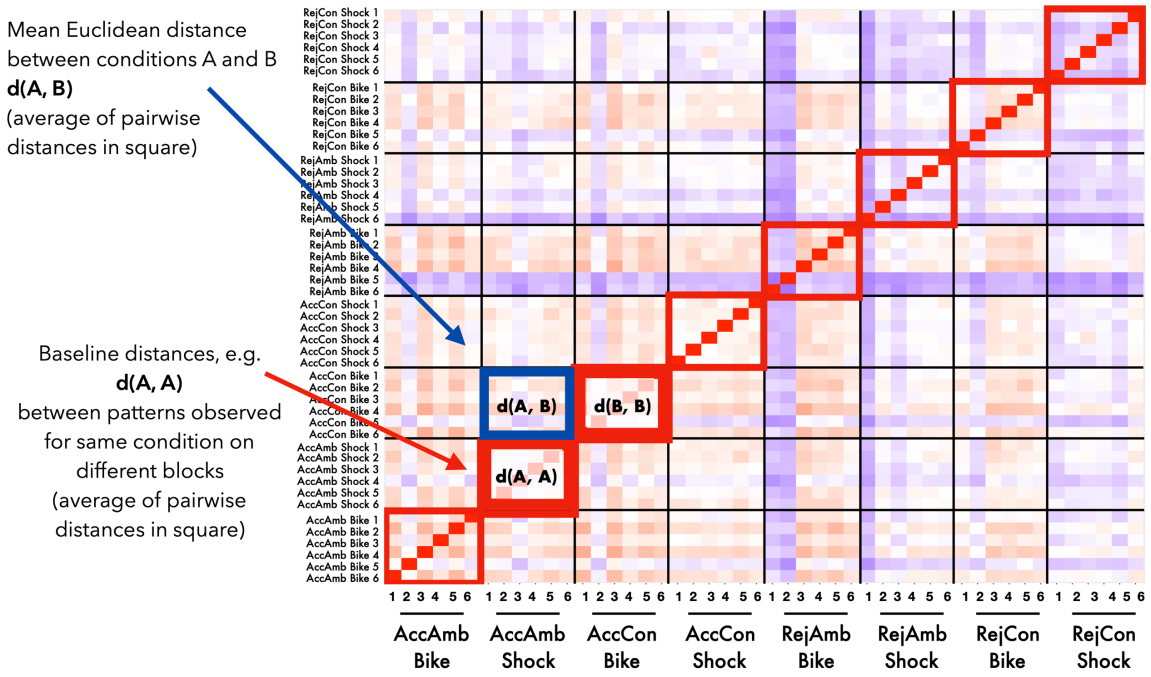

Normalized Distance:

$$\mathbf{d(A, B)^*} = \frac{|\, d(A, B)^2 - (d(A, A)*d(B, B))\,|}{(d(A, A)*d(B, B))}$$

Example data from a single participant, during single FIR, within a single ROI. Voxel response pattern distances were estimated for all pairwise comparisons of observations (2 Choice X 2 Confidence X 2 Cost X 6 runs). Red diagonal illustrates identity condition – comparison of a single observation to itself yields 0 distance. Surrounding the red diagonal are baseline distances, the distance between a condition and itself in another run. Gridlines illustrate individual trial conditions. The mean distance between a pair of conditions (average across pixels within a grid square) is the cross-validated distance. Then normalization was applied to control for between-run variance within a trial condition.

Figure 16: Illustration of representational dissimilarity comparisons of interest

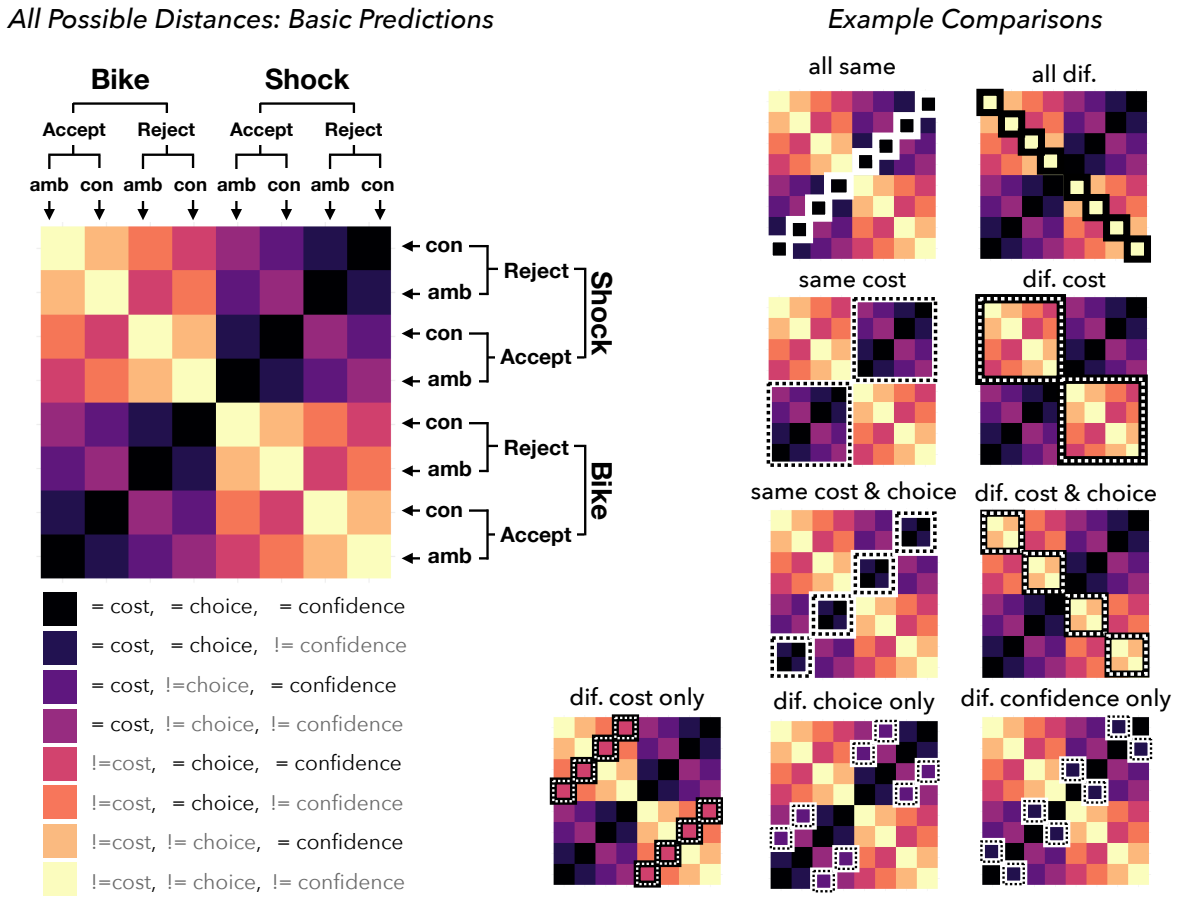

Plots contain basic predictions about relative distances, not data. This figure illustrates comparisons of interest to help guide interpretations of actual data presented in Figure 17. Further explanation is provided in the main body text.

Figure 17: Pairwise representational dissimilarity across all trial conditions

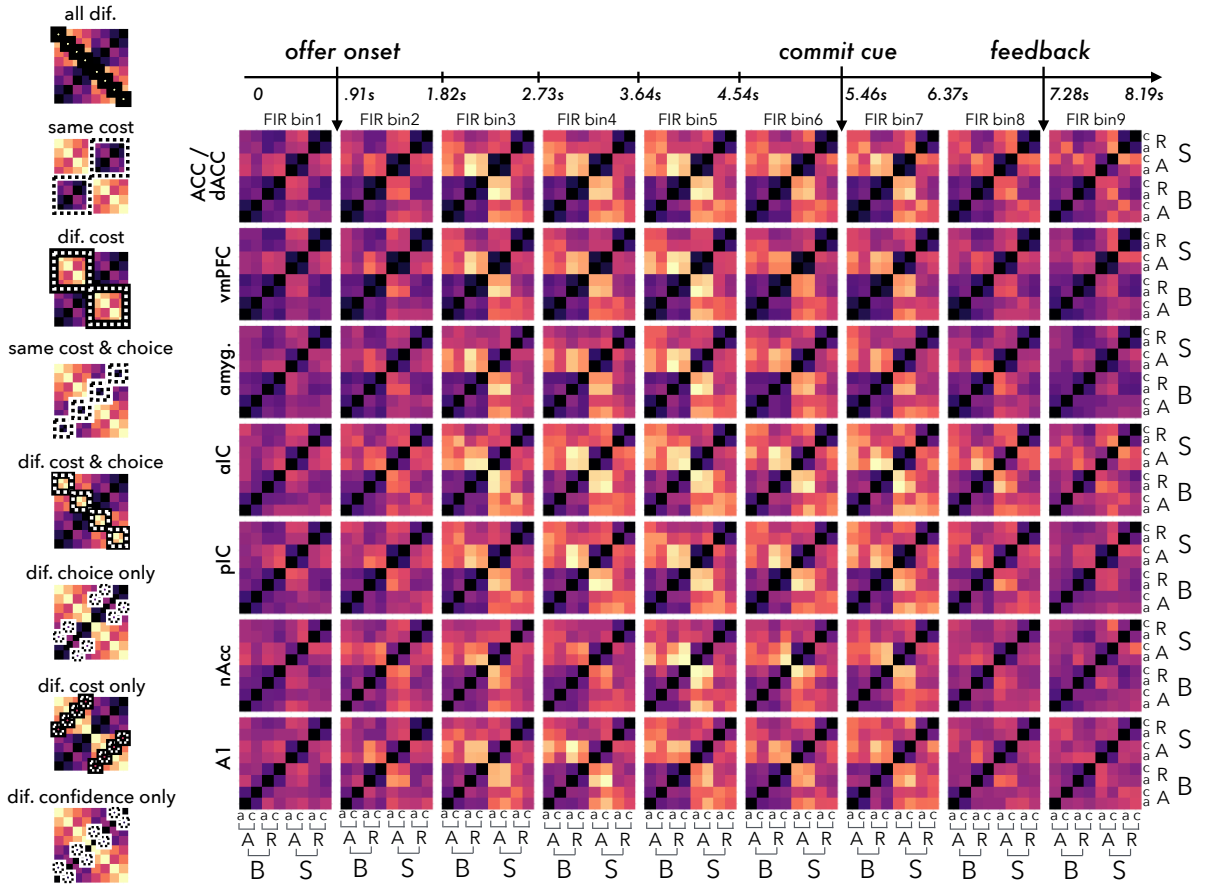

(Left): Comparisons of potential interest over basic predictions (not data). (Right) Plot displays representational dissimilarity results, shown in normalized distances (a.u.) with darker colors indicating smaller distances and brighter colors indicating larger distances. Rows are individual ROIs (control ROI, A1, in bottom row), columns are FIR bins 1-9 (FIR bin 1 is the baseline time point). Within each grid square are pairwise comparisons of all trial conditions observed within the corresponding ROI during the corresponding FIR bin. Trial conditions are noted on bottom and right with B = bike, S = shock, A = accept, R = reject, a = ambivalent, c = confident. Timeline displayed across top of plot indicates trial events. Note that at the beginning of the trial (leftmost columns) distances between trial conditions are relatively small, indicating that there are no meaningful differences in voxel response patterns between trial conditions. Gradually over the course of the trial (moving rightward across columns), differences between trial conditions begin to emerge (brighter shades) across all ROIs, and then diminish back to baseline values. A detailed description is provided in the main body of the text.

Figure 18: Representational dissimilarity of decision costs across all conditions

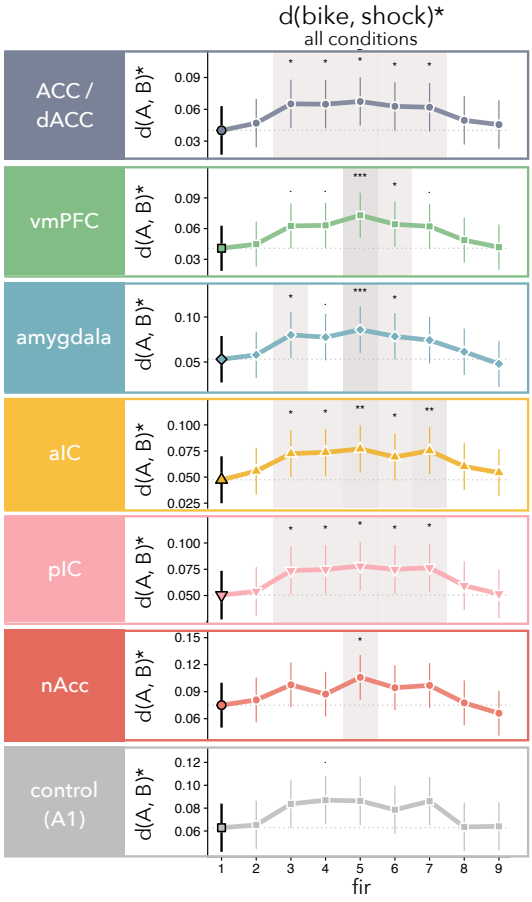

The mean normalized distance across all trial conditions that differ only by cost are plotted with respect to FIR time bin for each ROI and the control region. A significant difference from baseline indicates representational dissimilarity of decision costs within that ROI at that time point. Error bars are 95% confidence intervals around model fitted mean. Shaded regions mark significant differences from baseline (FIR bin 1) when the fixation cue was on the display (*** $p < .005$; ** $p < .01$; * $p < .05$; . $p < .1$; Dunnett's test corrected). All ROIs differ from baseline during at least one FIR bin, whereas the control region does not.

Figure 19: Representational dissimilarity of decision costs within each condition

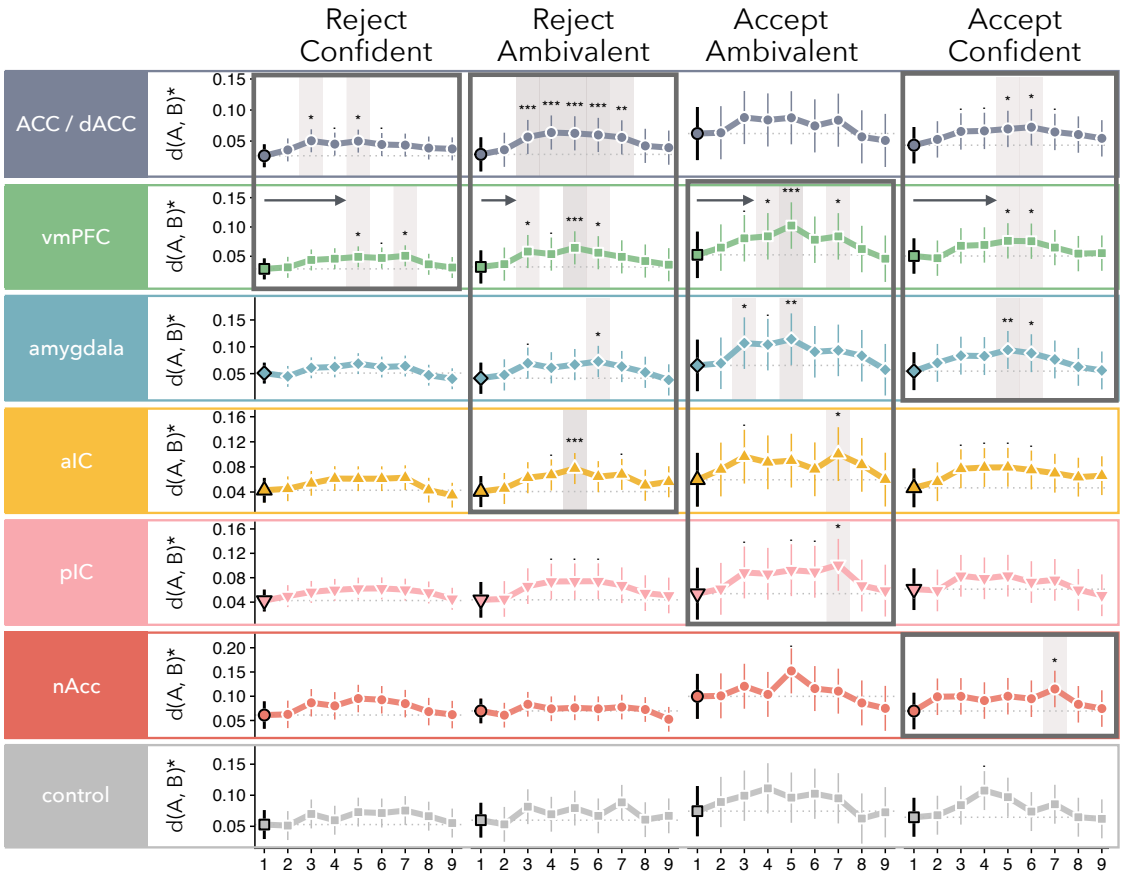

Mean normalized distances for each trial condition are plotted with respect to FIR time bin for each ROI and the control region. A significant difference from baseline indicates representational dissimilarity of decision costs within that ROI, at that time point, for that condition. Error bars are 95% confidence intervals around model fitted mean. Shaded regions mark significant differences from baseline, FIR bin 1 (*** $p < .005$; **$p < .01$; *$p < .05$; . $p < .1$; Dunnett's test corrected). Black outlines indicate the trial conditions where there is a significant effect. Arrows overlaid on vmPFC plot illustrate relative onsets of cost differences in vmPFC – differences in voxel response patterns corresponding to pain and effort costs emerge relatively earlier during ambivalent choices than during confident choices.

Figure 20: Comparison of behavioral choices in bike and shock tasks for an example participant

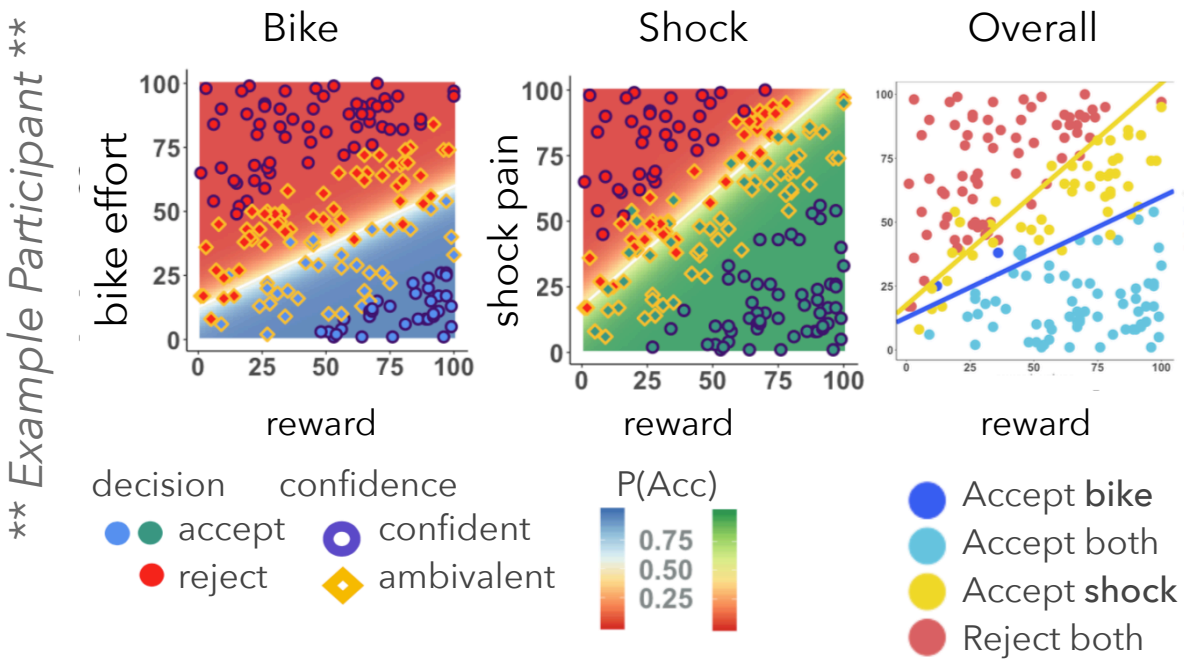

Behavioral choice data from an individual subject in the bike task (left), shock task (middle), and a comparison plot (right). Interaction plots (left and middle) are as described in Figure 2. Note that now, blue corresponds to bike data and green/yellow corresponds to shock data. The comparison plot overlays the decision boundaries from both tasks (yellow = shock, blue = bike) with a summary of decisions from the two tasks (light blue = offer was accepted in both tasks, dark blue = offer was only accepted in bike task, yellow = offer was only accepted in shock task, red = offer was rejected in both tasks). This participant's choices suggest that the bike costs were perceived as more aversive than the shock costs as several trials that were accepted in the shock task were rejected in the bike task. Furthermore, the slope of the decision boundary is less steep in the bike task than the shock task. This indicates that in order to cancel out the value of a reward (at the bound, the relative values of costs and reward are equal), the magnitude of a shock cost needs to be much greater than the magnitude of a bike cost that would cancel out the same reward.

Figure 21: Comparison of behavioral choices in bike and shock tasks – all participants

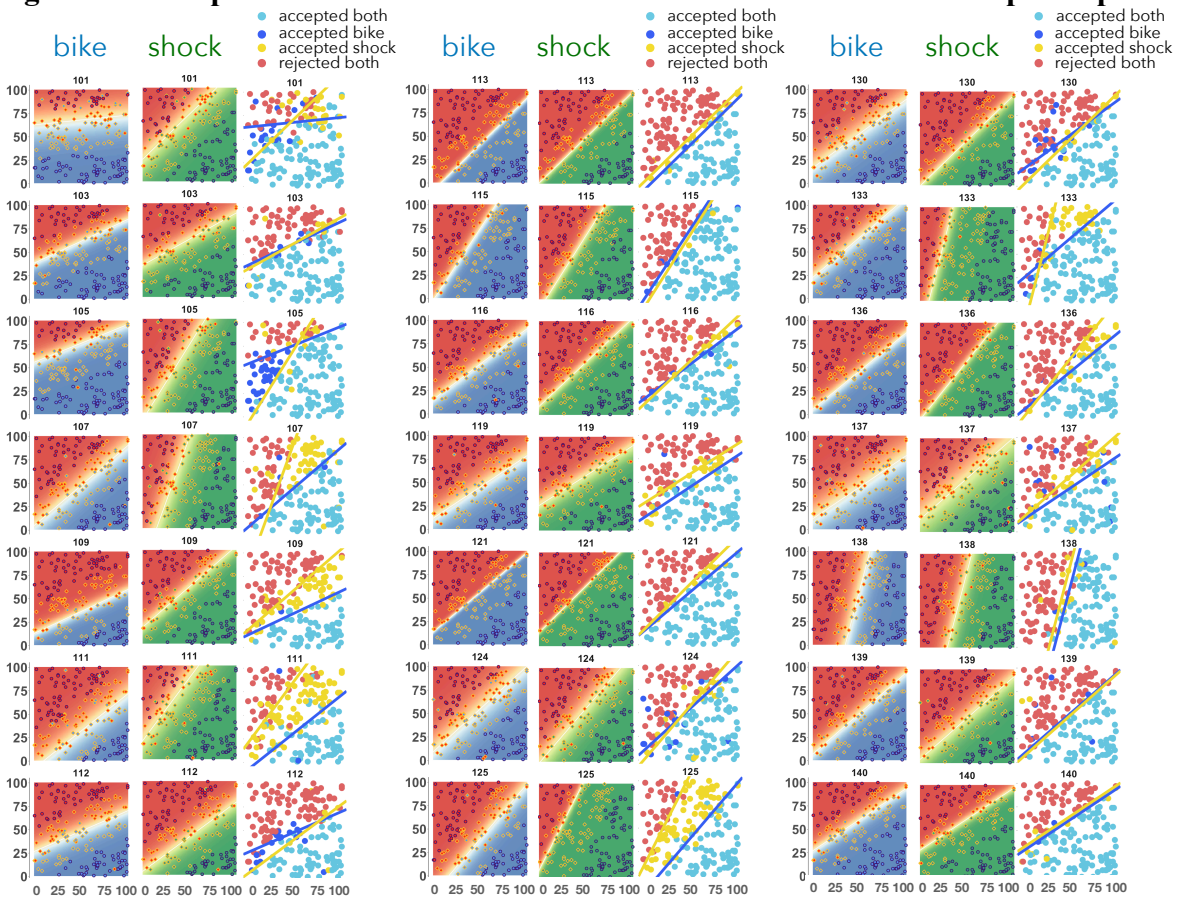

Figure 22: Individual differences in choice behavior for bike and shock tasks

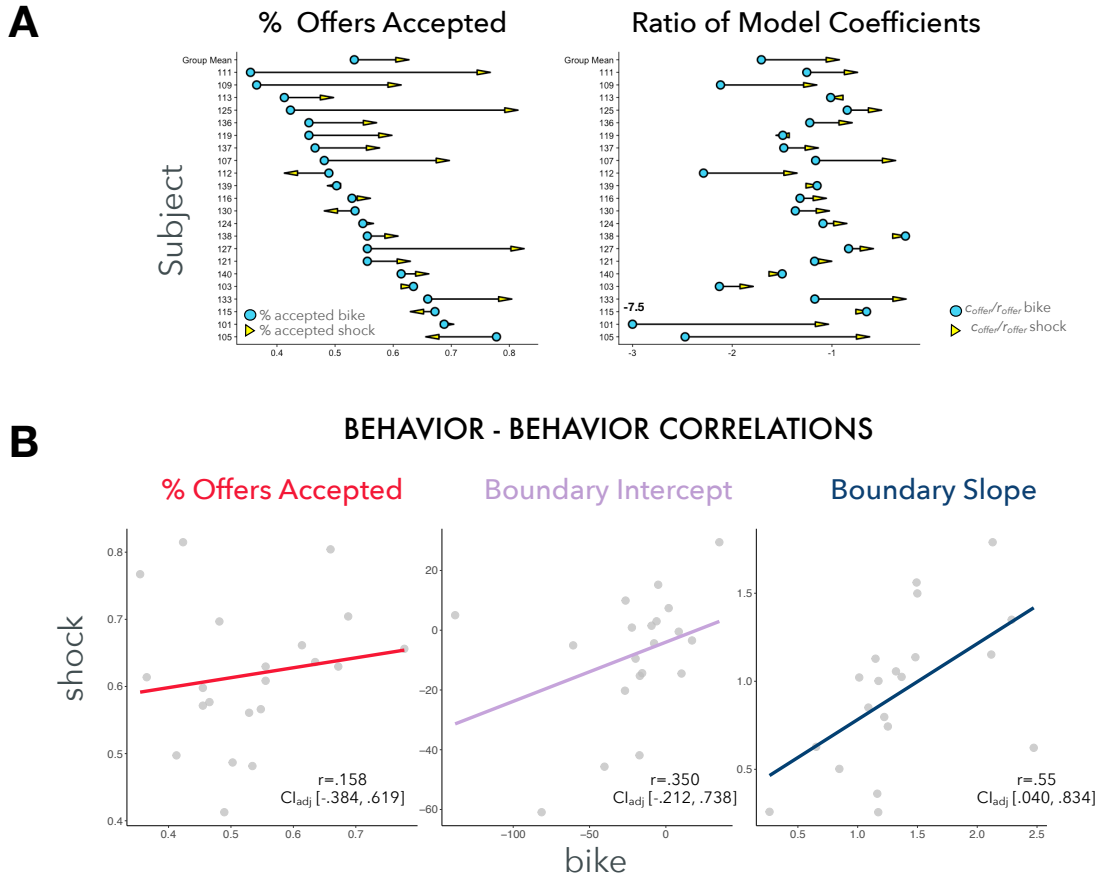

(A) Changes in behavior between the bike (blue circles) and shock (yellow arrows; arrow tip denotes data value) for individual participants, which are row-aligned in both plots for ease of comparison, and group means in top row. Left: Change in proportion of offers accepted in the bike and shock tasks. Right: Change in ratio of model weights (cost weight/reward weight) between bike and shock tasks. Note that there is no obvious relationship between the two plots – the change in the proportion of offers accepted between the two tasks does not meaningfully correspond to the change in the ratio of model weights. **(B)** Correlations between behavior on bike and shock tasks with respect to proportion of offers accepted (Left), decision boundary intercept (Middle), and decision boundary slope (Right). There was a correlation between individuals' decision boundary slopes in the bike and shock tasks, but not between the proportion of offers accepted in the two tasks, nor the decision boundary intercept.

Figure 23: Value transformations in ROIs

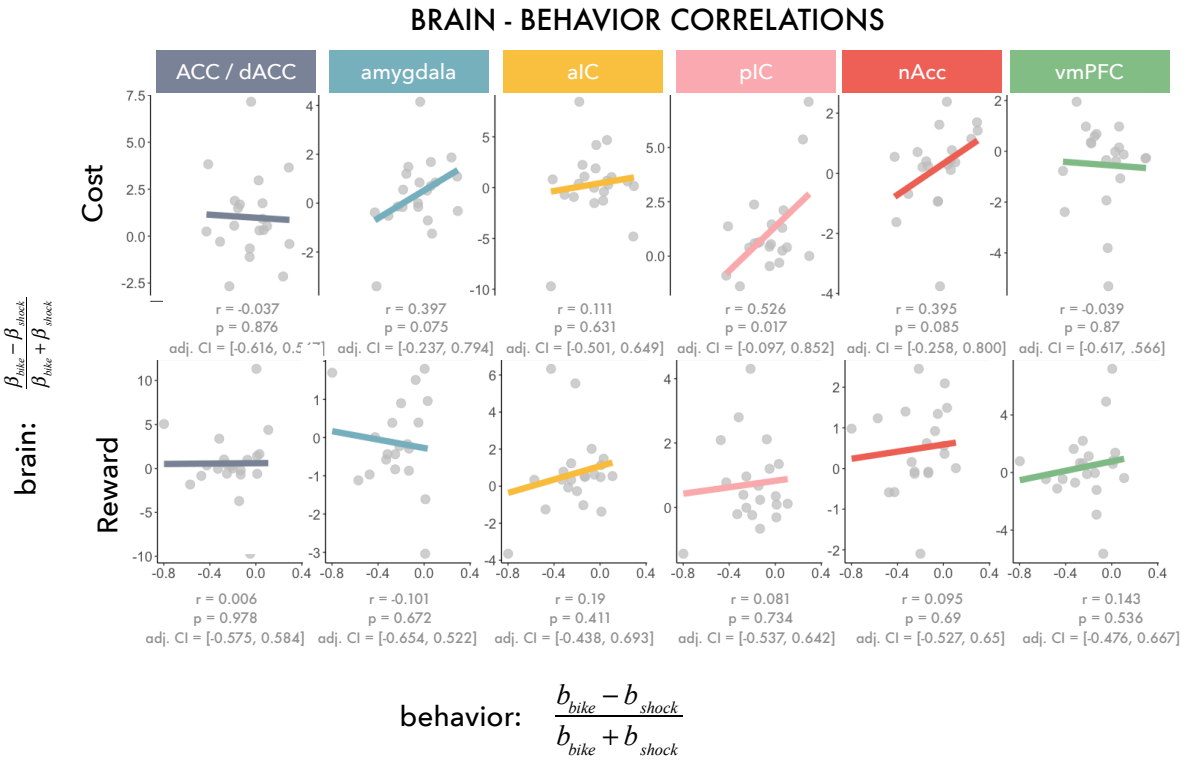

Correlations between change in behavioral model weights and change in ROI parameter estimates for objective costs (top) and objective rewards (bottom). Text on figure provides person's R and p values, along with a 95% confidence interval that has been Bonferroni corrected for multiple comparisons. A strong correlation between the change in the behavioral cost (or reward) weights between the two tasks and the change in parameter estimates in an ROI would suggest that the extent to which an ROI responds to a given stimulus modulates the extent to which behavior is influenced by that stimulus. There were no significant correlations after corrections for multiple comparisons.

Table 1: Significant cluster activation from GLM1 and GLM2

Full details of significant task-related activation. Clusters <10 voxels excluded. Sub-headings give the analysis and decision phaseThe first three columns identify the analysis, decision phase when BOLD responses were measured, and the variable measured in the statistic (Con. = confiden(t)ce, Amb.=ambivalen(t)ce. The next column specifies each cluster and indicates the total number of voxels in that cluster. The remaining columns are details about local maxima within that cluster, including t-value (* signifies cluster peaks illustrated in figures), MNI coordinates, region from AAL of MNI space, and Brodmann area (- signifies peak voxel falls outside of labelled cortex, region provided is the closest).

.

GLM1: SV, Offer Phase							
		MNI					
Voxels	***t***	**x**	**y**	**z**		**Region**	
19370	7.13	42	-66	32	R	Angular gyrus	BA39
	5.99	-16	-52	42	L	Precuneus -	BA31
	5.97	-38	-66	30	L	Middle occipital gyrus	BA39
	5.79	66	-28	-4	R	Middle temporal gyrus	BA21
	5.26	28	-14	10	R	Lenticular nucleus, putamen	
	5.23	32	-16	10	R	Insula	BA13 -
6613	7.28	32	20	52	R	Middle frontal gyrus	BA8
	6.69	**8**	**44**	**-8**	† **R**	**Superior frontal gyrus, medial orbital**	**BA10**
	6.13	-12	4	-22	L	Parahippocampal gyrus	
	5.79	-8	44	0	L	Anterior cingulate and paracingulate gyri	BA32
	5.09	-4	38	-16	L	Gyrus rectus	BA11
	4.95	10	16	-8	R	Caudate nucleus	
1106	5.04	-60	-34	-26	L	Inferior temporal gyrus	BA20
	5.02	-42	-22	10	L	Heschl gyrus	BA41
	4.92	-66	-30	-16	L	Inferior temporal gyrus	BA21
	4.51	-62	-56	-18	L	Inferior temporal gyrus	BA37
	4.38	-62	-44	-12	L	Middle temporal gyrus	BA37
	4.26	-64	-50	-14	L	Inferior temporal gyrus	BA37
480	6.51	-24	18	44	L	Middle frontal gyrus	BA8
	4.53	-26	28	36	L	Middle frontal gyrus	BA9
	4.32	-26	24	48	L	Middle frontal gyrus	BA8
	4.02	-20	24	32	L	Middle frontal gyrus -	BA8 -
	3.87	-24	24	26	L	Middle frontal gyrus -	BA9 -
	3.79	-28	28	42	L	Middle frontal gyrus	BA9
249	4.97	0	-6	34	L	Median cingulate and paracingulate gyri	BA24
	4.36	-2	-4	38	L	Median cingulate and paracingulate gyri	BA24

	2.95	12	0	42	R	Median cingulate and paracingulate gyri	BA32 -
	2.81	10	2	38	R	Median cingulate and paracingulate gyri	BA32
	2.71	6	2	40	R	Median cingulate and paracingulate gyri	BA32
75	3.15	-36	28	-22	L	Inferior frontal gyrus, orbital part	BA47
	3.05	-32	28	-18	L	Inferior frontal gyrus, orbital part	BA47
	2.78	-36	38	-18	L	Inferior frontal gyrus, orbital part	BA47
	2.74	-38	38	-12	L	Inferior frontal gyrus, orbital part	BA47
	2.63	-34	40	-12	L	Middle frontal gyrus, orbital part	BA47
72	3.48	-38	4	-16	L	Temporal pole: superior temporal gyrus -	BA13 -
	3.02	-36	-4	-10	L	Insula	BA13 -
	2.87	-30	6	-16	L	Amygdala -	BA34
58	4.66	-38	-68	-44	L	Crus II of cerebellar hemisphere	
	4.03	-32	-66	-48	L	Lobule VIIB of cerebellar hemisphere	
44	4.37	-42	-54	10	L	Middle temporal gyrus	BA37
	3.72	-44	-52	0	L	Middle temporal gyrus	BA37 -
	3.59	-42	-50	16	L	Middle temporal gyrus	BA39 -
	3.38	-42	-52	20	L	Middle temporal gyrus	BA39 -
	3.31	-44	-46	0	L	Middle temporal gyrus -	BA21 -
22	2.78	-32	-10	-14	L	Hippocampus	
	2.64	-30	-10	-4	L	Lenticular nucleus, putamen	
	2.60	-30	-6	-8	L	Lenticular nucleus, putamen -	
15	3.03	12	-22	48	R	Median cingulate and paracingulate gyri	BA6
	2.94	12	-24	42	R	Median cingulate and paracingulate gyri	BA31
6	3.31	8	-54	-12	R	Lobule IV, V of cerebellar hemisphere	
6	3.38	-38	52	2	L	Middle frontal gyrus	BA10
6	2.99	28	12	32	R	Inferior frontal gyrus, opercular part -	BA8
4	2.81	30	10	-12	R	Olfactory cortex -	BA13 -
3	4.08	-42	-16	28	L	Postcentral gyrus -	BA1 -
3	3.03	-24	-26	76	L	Postcentral gyrus	BA1 -
1	2.85	38	-68	-34	R	Crus I of cerebellar hemisphere	

GLM1: Confidence (negative), Offer Phase

		MNI					
Voxels	***t***	**x**	**y**	**z**		**Region**	**BA**
2043	-7.14	16	-90	0	R	Calcarine fissure and surrounding cortex	BA17
	-6.63	-10	-96	-6	L	Calcarine fissure and surrounding cortex	BA18
	-6.50	-14	-92	2	L	Superior occipital gyrus	BA18
	-6.35	16	-88	-8	R	Lingual gyrus	BA18

	-6.10	0	-80	8		L	Calcarine fissure and surrounding cortex	BA17
	-5.46	-8	-82	4		L	Calcarine fissure and surrounding cortex	BA17
2035	-6.65	46	-32	36		R	Supramarginal gyrus -	BA40
	-6.35	50	-28	46		R	Postcentral gyrus	BA1
	-6.33	38	-42	38		R	Inferior parietal -	BA7 -
	-6.25	42	-42	42		R	Inferior parietal	BA40
	-6.02	46	-32	46		R	Supramarginal gyrus	BA40
	-5.94	44	-46	48		R	Inferior parietal	BA40
2030	-6.42	22	8	44		R	Middle frontal gyrus -	BA8
	-6.13	**2**	**20**	**36**	†	R	**Median cingulate and paracingulate gyri**	**BA8**
	-6.04	2	16	40		R	Median cingulate and paracingulate gyri	BA8
	-5.54	-10	28	28		L	Anterior cingulate and paracingulate gyri	BA32
	-5.44	-4	28	28		L	Anterior cingulate and paracingulate gyri	BA32
	-5.32	24	-2	42		R	Superior frontal gyrus, dorsolateral -	BA6 -
131	**-6.41**	**40**	**42**	**6**	†	**R**	**Middle frontal gyrus**	**BA46**
	-4.38	44	44	16		R	Middle frontal gyrus	BA10
14	-3.89	10	-74	42		R	Precuneus	BA7
11	-4.02	38	52	-6		R	Middle frontal gyrus, orbital part	BA10
6	-3.79	12	-66	46		R	Precuneus	BA7
5	-4.21	10	-62	54		R	Precuneus	BA7
3	-5.57	16	-54	-38		R	Lobule VIII of cerebellar hemisphere -	

GLM1: Confidence, Commit Phase

		MNI					
Voxels	***t***	**x**	**y**	**z**		**Region**	**BA**
3894	4.91	-52	-38	14	L	Superior temporal gyrus	BA22
	4.73	-64	-18	-6	L	Middle temporal gyrus	BA21
	4.57	-64	-24	0	L	Middle temporal gyrus	BA22
	4.49	-40	-56	14	L	Middle temporal gyrus	BA39
	4.40	-60	6	-22	L	Middle temporal gyrus	BA38 -
	4.38	-48	-2	-26	L	Middle temporal gyrus	BA38
2160	4.73	62	-20	8	R	Superior temporal gyrus	BA41
	4.49	64	-20	2	R	Superior temporal gyrus	BA22
	4.43	50	0	-6	R	Superior temporal gyrus	BA22 -
	4.36	48	-34	18	R	Superior temporal gyrus	BA22
	4.34	38	4	-26	R	Amygdala -	BA38 -
	4.33	48	12	-34	R	Temporal pole: middle temporal gyrus	BA38
585	5.29	-10	-72	-8	L	Lingual gyrus	BA18

	4.91	-4	-86	18	L	Cuneus	BA18
	4.00	-8	-92	24	L	Cuneus	BA18
	3.90	-14	-88	16	L	Superior occipital gyrus	BA18
	3.79	-4	-86	26	L	Cuneus	BA19
	3.73	-8	-88	22	L	Cuneus	BA18
349	4.38	10	-66	-2	R	Lingual gyrus	BA18
	4.34	14	-68	-8	R	Lingual gyrus	BA19
	3.80	20	-84	0	R	Lingual gyrus	BA18
	3.66	10	-62	-12	R	Lobule VI of cerebellar hemisphere	
	3.42	12	-74	0	R	Lingual gyrus	BA18
	3.41	8	-66	4	R	Lingual gyrus	BA18
237	5.30	-30	-18	-30	L	Fusiform gyrus	BA36
	5.25	-26	-18	-28	L	Parahippocampal gyrus	BA36
	3.92	-20	-12	-28	L	Parahippocampal gyrus	BA36
	3.90	-34	-18	-24	L	Fusiform gyrus	BA36 -
	3.30	-24	-6	-26	L	Parahippocampal gyrus	
183	3.84	-38	-16	42	L	Postcentral gyrus	BA4
	3.73	-38	-14	28	L	Insula -	BA4 -
	3.72	-48	-18	42	L	Postcentral gyrus	BA1
	3.56	-44	-22	30	L	Postcentral gyrus -	BA1
	3.36	-44	-16	46	L	Postcentral gyrus	BA4
	3.30	-48	-16	32	L	Postcentral gyrus	BA1
137	**5.05**	**2**	**38**	**-20** †	**R**	**Gyrus rectus**	**BA11**
	5.00	2	34	-18	R	Gyrus rectus	BA11
	4.97	2	38	-24	R	Gyrus rectus	BA11
	4.58	-2	34	-12	L	Superior frontal gyrus, medial orbital	BA11
	4.39	-6	34	-12	L	Superior frontal gyrus, medial orbital	BA11
	4.35	4	30	-10	R	Superior frontal gyrus, medial orbital	BA11
136	5.47	2	-24	54	R	Supplementary motor area	BA6
	4.45	-4	-32	56	L	Paracentral lobule	BA1
	3.26	4	-20	62	R	Supplementary motor area	BA6
51	3.15	-34	2	0	L	Lenticular nucleus, putamen -	
	3.10	-28	0	6	L	Lenticular nucleus, putamen	
	3.02	-38	6	-2	L	Insula	BA13
	3.00	-30	2	2	L	Lenticular nucleus, putamen	
45	3.93	-4	-60	20	L	Precuneus	BA31
	3.89	-10	-58	10	L	Calcarine fissure and surrounding cortex	BA23
	3.83	-6	-60	14	L	Calcarine fissure and surrounding cortex	BA31

45	4.18	18	-74	26	R	Cuneus	BA19
	3.77	24	-68	22	R	Superior occipital gyrus	BA18 -
42	3.56	36	-30	46	R	Postcentral gyrus	BA1
	3.46	28	-28	46	R	Postcentral gyrus -	BA4 -
	3.42	30	-32	44	R	Postcentral gyrus	BA1
39	4.05	48	-10	52	R	Precental gyrus	BA6
	3.19	54	-10	44	R	Precental gyrus	BA6
39	3.95	48	-24	54	R	Postcentral gyrus	BA1
37	4.28	28	-22	-20	R	Parahippocampal gyrus	BA36
	3.90	22	-18	-24	R	Parahippocampal gyrus	BA36
	3.80	18	-10	-24	R	Parahippocampal gyrus	BA36 -
	3.79	18	-6	-26	R	Parahippocampal gyrus	BA36 -
32	3.52	60	8	22	R	Precental gyrus	BA6
	3.48	60	2	24	R	Postcentral gyrus	BA6
	3.35	60	-2	24	R	Postcentral gyrus	BA6
26	3.75	62	-10	32	R	Postcentral gyrus	BA1
23	4.23	8	-58	12	R	Calcarine fissure and surrounding cortex	BA18
10	3.84	20	-58	14	R	Calcarine fissure and surrounding cortex	BA18
	3.76	18	-56	8	R	Calcarine fissure and surrounding cortex	BA23
6	3.33	44	22	-34	R	Temporal pole: middle temporal gyrus	BA38
	3.09	40	24	-36	R	Temporal pole: middle temporal gyrus	BA38
3	4.26	16	-52	4	R	Lingual gyrus	BA18
2	3.54	-14	-44	40	L	Precuneus	BA31
1	3.52	-16	-96	26	L	Superior occipital gyrus	BA18

GLM2: Acc > Rej, Offer Phase

		MNI					
Voxels	***t***	**x**	**y**	**z**		**Region**	**BA**
22910	7.08	-38	-72	34	L	Middle occipital gyrus	BA39
	6.38	-38	-74	42	L	Inferior parietal	BA39
	6.10	-34	-76	42	L	Inferior parietal	BA39
	6.02	68	-28	-4	R	Middle temporal gyrus	BA21
	5.89	28	-46	70	R	Postcentral gyrus	BA5 -
	5.87	-24	28	36	L	Middle frontal gyrus	BA9
4829	6.93	-10	62	20	L	Superior frontal gyrus, medial	BA10
	5.96	**8**	**44**	**-12**	† **R**	**Superior frontal gyrus, medial orbital**	**BA11**
	5.84	12	46	-16	R	Gyrus rectus	BA11
	5.83	-4	40	-12	L	Superior frontal gyrus, medial orbital	BA11

	5.81	16	34	-12	R	Superior frontal gyrus, orbital part -	BA11 -
	5.70	30	24	52	R	Middle frontal gyrus	BA8
2416	4.97	-6	-50	38	L	Precuneus	BA31
	4.76	-8	-42	38	L	Median cingulate and paracingulate gyri	BA31
	4.75	-4	-66	22	L	Calcarine fissure and surrounding cortex	BA18
	4.62	-12	-50	40	L	Precuneus	BA31
	4.54	10	-60	18	R	Calcarine fissure and surrounding cortex	BA18
	4.50	10	-58	22	R	Precuneus	BA23
205	3.85	10	-92	12	R	Calcarine fissure and surrounding cortex	BA17
	3.40	-6	-100	14	L	Cuneus	BA18
	3.24	-2	-96	8	L	Calcarine fissure and surrounding cortex	BA17
	3.11	-10	-102	4	L	Middle occipital gyrus	BA18
27	3.59	-10	46	48	L	Superior frontal gyrus, medial	BA8
	3.11	-14	42	52	L	Superior frontal gyrus, dorsolateral	BA8
7	3.42	20	-86	2	R	Calcarine fissure and surrounding cortex	BA18
4	3.42	34	-86	34	R	Middle occipital gyrus	BA39
3	2.73	-58	16	12	L	Inferior frontal gyrus, opercular part	BA44
1	3.30	20	-102	10	R	Cuneus -	BA18
1	2.64	-30	6	62	L	Middle frontal gyrus	BA6 -

GLM2: Confident < Ambivalent, Offer Phase

		MNI					
Voxels	***t***	**x**	**y**	**z**		**Region**	**BA**
20543	-10.50	28	-64	46	R	Angular gyrus	BA39
	-10.40	34	-74	32	R	Middle occipital gyrus	BA39
	-9.89	42	-42	44	R	Inferior parietal	BA40
	-9.15	44	-34	40	R	Supramarginal gyrus	BA40
	-8.93	44	-34	36	R	Supramarginal gyrus -	BA40
	-8.78	38	-52	36	R	Angular gyrus	BA39 -
12133	**-10.50**	**4**	**16**	**40**	† **R**	**Median cingulate and paracingulate gyri**	**BA8**
	-8.49	-8	12	44	L	Median cingulate and paracingulate gyri	BA32
	-8.01	-4	30	18	L	Anterior cingulate and paracingulate gyri	BA24
	-7.77	26	4	44	R	Superior frontal gyrus, dorsolateral -	BA6
	-7.76	**38**	**48**	**4**	† **R**	**Middle frontal gyrus**	**BA10**
	-7.50	28	4	48	R	Middle frontal gyrus	BA6
118	-6.83	-42	14	-8	L	Insula	BA13 -
	-6.74	-28	22	0	L	Insula -	BA13

GLM2: Confident > Ambivalent, Commit Phase

		MNI					
Voxels	***t***	**x**	**y**	**z**		**Region**	**BA**
4192	5.65	2	-18	46	R	Median cingulate and paracingulate gyri	BA24
	5.45	6	-34	56	R	Paracentral lobule	BA1
	5.30	-58	-12	-2	L	Superior temporal gyrus	BA22
	5.26	-52	-12	12	L	Heschl gyrus	BA1
	5.15	6	-34	52	R	Paracentral lobule	BA5
	5.12	-44	-22	30	L	Postcentral gyrus -	BA1
936	4.93	64	-18	0	R	Superior temporal gyrus	BA22
	4.69	50	0	-4	R	Superior temporal gyrus	BA22 -
	4.68	64	2	-6	R	Superior temporal gyrus	BA22
	4.59	54	-6	-2	R	Superior temporal gyrus	BA22
	4.46	38	-28	18	R	Rolandic operculum	BA40
	4.36	56	-32	14	R	Superior temporal gyrus	BA41
767	5.59	14	-70	18	R	Calcarine fissure and surrounding cortex	BA18
	4.98	-4	-60	16	L	Precuneus	BA31
	4.94	-12	-58	8	L	Calcarine fissure and surrounding cortex	BA23
	4.92	18	-56	12	R	Calcarine fissure and surrounding cortex	BA23
	4.72	18	-72	24	R	Cuneus	BA19
	4.72	-12	-62	10	L	Calcarine fissure and surrounding cortex	BA23
416	**5.80**	**2**	**40**	**-20**	† **R**	**Gyrus rectus**	**BA11**
	5.71	4	40	-26	R	Gyrus rectus -	BA11
	5.65	8	26	-10	R	Olfactory cortex -	BA11 -
	5.62	10	30	-10	R	Superior frontal gyrus, medial orbital -	BA11 -
	5.58	-4	44	-18	L	Gyrus rectus	BA11
	5.45	6	28	-18	R	Gyrus rectus	BA11
394	5.90	38	-28	46	R	Postcentral gyrus	BA1
	5.07	32	-24	42	R	Postcentral gyrus -	BA1 -
	4.28	46	-22	56	R	Postcentral gyrus	BA1
	3.80	42	-12	44	R	Precental gyrus	BA6
	3.74	44	-10	58	R	Precental gyrus	BA6
	3.71	52	-10	46	R	Precental gyrus	BA6
192	5.07	-4	-86	20	L	Cuneus	BA18
	5.01	-8	-90	18	L	Cuneus	BA18
	5.01	-8	-90	24	L	Cuneus	BA18
	4.64	-8	-84	20	L	Cuneus	BA18
79	5.44	-14	-66	-12	L	Lobule VI of cerebellar hemisphere	

	5.30	-10	-72	-10	L	Lobule VI of cerebellar hemisphere	
	4.80	-18	-72	-12	L	Lobule VI of cerebellar hemisphere	
53	4.30	-16	-10	68	L	Precental gyrus	BA6
	4.04	-16	-14	78	L	Precental gyrus	BA6
	3.42	-20	-10	60	L	Superior frontal gyrus, dorsolateral -	BA6 -
41	6.72	0	12	-18	R	Olfactory cortex -	BA25 -
	5.88	0	6	-14	L	Olfactory cortex -	
31	4.15	42	-4	-24	R	Inferior temporal gyrus -	BA21 -
	4.00	42	4	-24	R	Temporal pole: superior temporal gyrus -	BA38 -
31	3.57	36	-16	12	R	Insula	BA13
	3.46	32	-14	10	R	Insula -	BA13 -
	3.26	42	-20	10	R	Heschl gyrus	BA41
22	5.49	60	6	22	R	Precental gyrus	BA6
15	3.76	50	12	-32	R	Temporal pole: middle temporal gyrus	BA38

GLM2: Confident < Ambivalent, Commit Phase

		MNI					
Voxels	***t***	**x**	**y**	**z**		**Region**	**BA**
10993	**-7.95**	**6**	**36**	**24**	† **R**	**Anterior cingulate and paracingulate gyri**	**BA9**
	-7.32	**36**	**48**	**-14**	† **R**	**Middle frontal gyrus, orbital part**	**BA10**
	-6.79	32	50	12	R	Middle frontal gyrus	BA10
	-6.62	**36**	**22**	**-12**	† **R**	**Inferior frontal gyrus, orbital part**	**BA47**
	-6.51	42	20	-12	R	Inferior frontal gyrus, orbital part	BA47
	-6.34	10	40	14	R	Anterior cingulate and paracingulate gyri	BA32
2403	-7.16	48	-54	40	R	Angular gyrus	BA39
	-6.55	46	-48	36	R	Angular gyrus	BA39
	-5.90	42	-54	44	R	Inferior parietal	BA39
	-5.76	16	-78	56	R	Superior parietal gyrus	BA7
	-5.68	44	-52	48	R	Inferior parietal	BA39
	-5.60	26	-74	48	R	Superior occipital gyrus	BA7
1895	-7.78	-30	-70	-34	L	Crus I of cerebellar hemisphere	
	-7.12	-26	-72	-32	L	Crus I of cerebellar hemisphere	
	-6.63	-34	-68	-30	L	Crus I of cerebellar hemisphere	
	-6.32	-42	-70	-30	L	Crus I of cerebellar hemisphere	
	-6.30	-40	-78	-30	L	Crus I of cerebellar hemisphere	
	-5.87	-38	-64	-48	L	Lobule VIIB of cerebellar hemisphere	
1008	-5.32	40	-78	-26	R	Crus I of cerebellar hemisphere	
	-5.28	50	-54	-24	R	Inferior temporal gyrus	BA37

	-5.25	50	-56	-36	R	Crus I of cerebellar hemisphere	
	-5.19	50	-56	-32	R	Crus I of cerebellar hemisphere	
	-5.15	38	-84	-32	R	Crus I of cerebellar hemisphere	
	-5.13	62	-60	-12	R	Inferior temporal gyrus -	BA37

GLM2 vmPFC ROI: Accept Confident > Accept Ambivalent, Commit Phase

		MNI					
Voxels	***t***	**x**	**y**	**z**		**Region**	**BA**
678	**4.54**	**-10**	**44**	**-10**	† **L**	**Superior frontal gyrus, medial orbital -**	**BA10**
	4.29	0	42	-16	R	Gyrus rectus	BA11
	4.17	2	42	-24	R	Gyrus rectus	BA11
	4.11	6	22	-22	R	Gyrus rectus	BA11
	4.09	0	38	-30	L	Gyrus rectus -	BA11 -
	3.94	-6	46	-12	L	Superior frontal gyrus, medial orbital	BA10
49	4.18	0	12	-18	† L	Olfactory cortex -	BA25 -

GLM2 vmPFC ROI: Reject Confident > Reject Ambivalent, Commit Phase

		MNI					
Voxels	***t***	**x**	**y**	**z**		**Region**	**BA**
698	**4.98**	**2**	**12**	**-18**	† **R**	**Olfactory cortex -**	**BA25 -**
	4.39	6	28	-22	R	Gyrus rectus	BA11
	4.14	2	36	-16	R	Gyrus rectus	BA11
	3.87	2	38	-20	R	Gyrus rectus	BA11
	3.84	8	28	-10	R	Superior frontal gyrus, medial orbital -	BA11
	3.77	-2	38	-20	L	Gyrus rectus	BA11
5	3.27	4	30	-2	R	Anterior cingulate and paracingulate gyri	BA24

Table 2: Comparison of linear mixed models predicting vmPFC responses from value

The first column identifies the model and variables specified as fixed effects from that model. Columns 2-7 give the model estimate, 95% confidence interval (95CI), standard error (SE), and t statistic and p value estimated from Satterthwaite approximation. Column 8 is the standard deviation (SD) of the random effect (Intercept | Subject). Columns 9-12 are overall model statistics: variance explained by fixed effects (Marg. = marginal R^2), variance explained by fixed and random effects together (Cond. = conditional R^2), and overall model AIC, and BIC. The final two columns are with $\chi 2$ statistics and p values from likelihood ratio tests (after refitting models with maximum likelihood estimates), describing improvement of fit from adding quadratic term to the original linear model.

	Fixed Effects						**Rand. Effect**	**Model Fit**		**Model Comparison**	
	Est.	95CI lower	95CI upper	SE	*t*	*p*	*SD*	Marg. R^2 Cond.R^2	AIC BIC	$\chi 2$	*p*
Offer Phase											
Lin. Model							14.58	0.05 0.7	901.1 911.98		
Intercept	-19.08	-24.74	-13.28	2.91	-6.56	<0.001					
P(accept)-P(reject)	5.1	2.73	7.41	1.17	4.38	<0.001					
Quad. Model							14.57	0.05 0.7	903.1 916.69	0.002	0.966
Intercept	-19	-25.78	-12.2	3.45	-5.50	<0.001					
P(accept)-P(reject)	5.11	2.76	7.4	1.18	4.33	<0.001					
$(P(accept)-P(reject))^2$	-0.12	-5.76	5.44	2.86	-0.04	0.967					
Commit Phase											
Lin. Model							19.17	0.01 0.55	1014.5 1025.3		
Intercept	-9.07	-17.08	-1.02	3.98	-2.28	0.031					
P(accept)-P(reject)	3.63	-0.48	7.78	2.06	1.76	0.082					
Quad. Model							19.18	0.07 0.61	1002 1015.6	14.43	< 0.001
Intercept	-20.86	-30.57	-10.96	4.96	-4.2	<0.001					
P(accept)-P(reject)	2.87	-0.96	6.63	1.92	1.49	0.14					
$(P(accept)-P(reject))^2$	18.16	9.21	27.27	4.65	3.91	<0.001					

Table 3: Cost distances varying from baseline (all conditions together)

Results of linear mixed analysis predicting cost distances from categorical time bin. P values were adjusted with Dunnett's correction for multiple comparisons.

Model Fixed Effects						Model Fit			Contrast with FIR1	
				Satterthwaite's			95CI			
Model Term	Estimate	SE	t	df	Pr(>|t|)	fit	lower	upper	P (adj.)	
ACC / dACC										
(Intcpt)	0.040	0.012	3.443	34.54	0.002	0.040	0.017	0.063		
FIR2	0.007	0.009	0.802	727	0.423	0.047	0.024	0.070	0.845	
FIR3	0.025	0.009	2.921	727	0.004	0.065	0.042	0.088	0.024	*
FIR4	0.025	0.009	2.888	727	0.004	0.065	0.042	0.088	0.024	*
FIR5	0.027	0.009	3.186	727	0.002	0.067	0.045	0.090	0.012	*
FIR6	0.023	0.009	2.65	727	0.008	0.063	0.040	0.086	0.040	*
FIR7	0.022	0.009	2.551	727	0.011	0.062	0.039	0.085	0.043	*
FIR8	0.010	0.009	1.11	727	0.267	0.050	0.027	0.072	0.801	
FIR9	0.005	0.009	0.64	727	0.523	0.046	0.023	0.068	0.845	
Amygdala										
(Intcpt)	0.053	0.013	4.047	34.23	2.80E-04	0.053	0.027	0.079		
FIR2	0.005	0.010	0.507	727	0.612	0.058	0.032	0.084	1	
FIR3	0.027	0.010	2.827	727	0.005	0.080	0.054	0.106	0.033	*
FIR4	0.024	0.010	2.554	727	0.011	0.078	0.052	0.103	0.053	.
FIR5	0.033	0.010	3.433	727	0.001	0.086	0.060	0.112	0.005	**
FIR6	0.025	0.010	2.649	727	0.008	0.078	0.053	0.104	0.049	*
FIR7	0.021	0.010	2.213	727	0.027	0.074	0.049	0.100	0.108	
FIR8	0.008	0.010	0.867	727	0.386	0.061	0.036	0.087	1	
FIR9	-0.005	0.010	-0.531	727	0.596	0.048	0.022	0.074	1	
aIC										
(Intcpt)	0.047	0.011	4.144	36.24	1.96E-04	0.047	0.025	0.070		
FIR2	0.008	0.009	0.955	727	0.340	0.056	0.033	0.078	0.680	
FIR3	0.025	0.009	2.873	727	0.004	0.072	0.050	0.095	0.020	*
FIR4	0.026	0.009	3.014	727	0.003	0.074	0.051	0.096	0.015	*
FIR5	0.030	0.009	3.39	727	0.001	0.077	0.055	0.099	0.006	**
FIR6	0.022	0.009	2.5	727	0.013	0.069	0.047	0.092	0.050	*
FIR7	0.028	0.009	3.196	727	0.001	0.075	0.053	0.098	0.010	**
FIR8	0.013	0.009	1.461	727	0.144	0.060	0.038	0.083	0.432	
FIR9	0.007	0.009	0.786	727	0.432	0.054	0.032	0.077	0.680	
pIC										
(Intcpt)	0.050	0.012	4.284	37.58	1.23E-04	0.050	0.027	0.073		
FIR2	0.004	0.009	0.395	727	0.693	0.054	0.031	0.077	1	
FIR3	0.023	0.009	2.534	727	0.012	0.074	0.051	0.097	0.047	*

FIR4	0.024	0.009	2.657	727	0.008	0.075	0.052	0.098	0.047	*
FIR5	0.028	0.009	3.002	727	0.003	0.078	0.055	0.101	0.021	*
FIR6	0.024	0.009	2.663	727	0.008	0.075	0.052	0.098	0.047	*
FIR7	0.026	0.009	2.833	727	0.005	0.076	0.053	0.100	0.032	*
FIR8	0.009	0.009	0.99	727	0.323	0.059	0.036	0.083	0.967	
FIR9	0.001	0.009	0.138	727	0.890	0.052	0.029	0.075	1	
nAcc										
(Intcpt)	0.075	0.013	5.927	41.64	0.000	0.075	0.050	0.100		
FIR2	0.006	0.011	0.55	727	0.582	0.081	0.056	0.106	1.000	
FIR3	0.023	0.011	2.138	727	0.033	0.098	0.073	0.122	0.227	
FIR4	0.012	0.011	1.163	727	0.245	0.087	0.062	0.112	0.979	
FIR5	0.031	0.011	2.927	727	0.004	0.106	0.081	0.131	0.027	*
FIR6	0.019	0.011	1.837	727	0.067	0.094	0.070	0.119	0.331	
FIR7	0.022	0.011	2.083	727	0.038	0.097	0.072	0.122	0.227	
FIR8	0.003	0.011	0.247	727	0.805	0.078	0.053	0.102	1.000	
FIR9	-0.009	0.011	-0.851	727	0.395	0.066	0.041	0.091	1.000	
vmPFC										
(Intcpt)	0.041	0.011	3.638	36.03	0.001	0.041	0.019	0.063		
FIR2	0.004	0.009	0.472	727	0.637	0.045	0.023	0.067	1	
FIR3	0.022	0.009	2.557	727	0.011	0.063	0.041	0.085	0.053	.
FIR4	0.022	0.009	2.633	727	0.009	0.063	0.041	0.085	0.051	.
FIR5	0.032	0.009	3.783	727	1.67E-04	0.073	0.051	0.095	0.001	**
FIR6	0.023	0.009	2.756	727	0.006	0.064	0.042	0.086	0.041	*
FIR7	0.021	0.009	2.506	727	0.012	0.062	0.040	0.084	0.053	.
FIR8	0.008	0.009	0.935	727	0.350	0.049	0.027	0.071	1	
FIR9	0.001	0.009	0.128	727	0.898	0.042	0.020	0.064	1	
control (A1)										
(Intcpt)	0.063	0.011	5.867	48.17	3.96E-07	0.063	0.042	0.084		
FIR2	0.002	0.010	0.253	727	0.800	0.065	0.044	0.086	1.000	
FIR3	0.021	0.010	2.161	727	0.031	0.084	0.063	0.105	0.154	
FIR4	0.024	0.010	2.501	727	0.013	0.087	0.066	0.108	0.099	.
FIR5	0.023	0.010	2.433	727	0.015	0.086	0.065	0.107	0.105	
FIR6	0.016	0.010	1.63	727	0.104	0.079	0.058	0.100	0.412	
FIR7	0.023	0.010	2.415	727	0.016	0.086	0.065	0.107	0.105	
FIR8	0.001	0.010	0.068	727	0.946	0.064	0.043	0.085	1.000	
FIR9	0.001	0.010	0.132	727	0.895	0.064	0.043	0.085	1.000	

Table 4: Cost distances varying from baseline (separate conditions)

Results of linear mixed analysis predicting cost distances from categorical time bin. P values were adjusted with Dunnett's correction for multiple comparisons.

	Model Fixed Effects						Model Fit			Contrast with FIR1	
					Satterthwaite			95CI			
cond	Model Term	Est.	SE	t	df	Pr(>|t|)	fit	lower	upper	P (adj.)	
ACC / dACC											
Acc Amb	(Intcpt)	0.062	0.02	2.865	25.23	0.008	0.062	0.019	0.105		
	FIR2	0.001	0.01	0.123	160	0.902	0.063	0.021	0.106	1.000	
	FIR3	0.026	0.01	2.384	160	0.018	0.088	0.045	0.131	0.137	
	FIR4	0.022	0.01	2.026	160	0.044	0.084	0.041	0.127	0.256	
	FIR5	0.025	0.01	2.354	160	0.020	0.088	0.045	0.130	0.137	
	FIR6	0.012	0.01	1.143	160	0.255	0.074	0.032	0.117	1.000	
	FIR7	0.022	0.01	1.992	160	0.048	0.084	0.041	0.126	0.256	
	FIR8	-0.005	0.01	-0.490	160	0.625	0.057	0.014	0.100	1.000	
	FIR9	-0.011	0.01	-1.026	160	0.307	0.051	0.008	0.094	1.000	
Acc Con	(Intcpt)	0.043	0.01	2.921	28.89	0.007	0.043	0.014	0.073		
	FIR2	0.009	0.01	1.018	160	0.310	0.053	0.023	0.082	0.458	
	FIR3	0.022	0.01	2.413	160	0.017	0.066	0.036	0.095	0.079	.
	FIR4	0.023	0.01	2.507	160	0.013	0.066	0.037	0.096	0.073	.
	FIR5	0.026	0.01	2.849	160	0.005	0.070	0.040	0.099	0.031	*
	FIR6	0.029	0.01	3.155	160	0.002	0.072	0.043	0.102	0.013	*
	FIR7	0.021	0.01	2.321	160	0.022	0.065	0.035	0.094	0.081	.
	FIR8	0.017	0.01	1.868	160	0.064	0.061	0.031	0.090	0.185	
	FIR9	0.011	0.01	1.204	160	0.231	0.054	0.025	0.084	0.458	
Rej Amb	(Intcpt)	0.029	0.01	2.049	28.55	0.050	0.029	0.001	0.056		
	FIR2	0.007	0.01	0.870	160	0.385	0.036	0.008	0.064	0.404	
	FIR3	0.028	0.01	3.292	160	0.001	0.057	0.029	0.084	0.005	***
	FIR4	0.035	0.01	4.138	160	5.64E-05	0.064	0.036	0.091	2.80E-04	***
	FIR5	0.034	0.01	3.988	160	1.01E-04	0.062	0.035	0.090	4.66E-04	***
	FIR6	0.031	0.01	3.680	160	3.18E-04	0.060	0.032	0.087	0.001	**
	FIR7	0.027	0.01	3.219	160	0.002	0.056	0.028	0.083	0.005	***
	FIR8	0.014	0.01	1.612	160	0.109	0.042	0.015	0.070	0.321	
	FIR9	0.011	0.01	1.276	160	0.204	0.039	0.012	0.067	0.404	
Rej Con	(Intcpt)	0.026	0.01	2.807	40.48	0.008	0.026	0.008	0.045		
	FIR2	0.009	0.01	1.208	160	0.229	0.036	0.017	0.054	0.328	
	FIR3	0.024	0.01	3.101	160	0.002	0.050	0.032	0.069	0.015	*
	FIR4	0.019	0.01	2.415	160	0.017	0.045	0.027	0.064	0.094	.
	FIR5	0.024	0.01	3.035	160	0.003	0.050	0.031	0.068	0.017	*
	FIR6	0.018	0.01	2.332	160	0.021	0.044	0.026	0.063	0.099	.
	FIR7	0.017	0.01	2.207	160	0.029	0.044	0.025	0.062	0.109	

	FIR8	0.012	0.01	1.601	160	0.111	0.039	0.020	0.057	0.328	
	FIR9	0.011	0.01	1.427	160	0.156	0.037	0.019	0.056	0.328	
Amygdala											
Acc Amb	(Intcpt)	0.065	0.02	2.684	28.48	0.012	0.065	0.017	0.113		
	FIR2	0.004	0.01	0.269	160	0.788	0.069	0.021	0.117	1.000	
	FIR3	0.041	0.01	2.806	160	0.006	0.107	0.059	0.155	0.035	*
	FIR4	0.039	0.01	2.618	160	0.010	0.104	0.056	0.152	0.053	.
	FIR5	0.049	0.01	3.308	160	0.001	0.114	0.066	0.162	0.008	**
	FIR6	0.025	0.01	1.718	160	0.088	0.091	0.043	0.139	0.343	
	FIR7	0.028	0.01	1.894	160	0.060	0.093	0.045	0.141	0.291	
	FIR8	0.018	0.01	1.216	160	0.226	0.083	0.035	0.131	0.672	
	FIR9	-0.008	0.01	-0.543	160	0.588	0.057	0.009	0.105	1.000	
Acc Con	(Intcpt)	0.055	0.02	3.052	31.52	0.005	0.055	0.019	0.090		
	FIR2	0.015	0.01	1.262	160	0.209	0.070	0.035	0.105	0.621	
	FIR3	0.029	0.01	2.367	160	0.019	0.084	0.048	0.119	0.107	
	FIR4	0.028	0.01	2.303	160	0.023	0.083	0.047	0.118	0.107	
	FIR5	0.040	0.01	3.237	160	0.001	0.094	0.059	0.129	0.010	**
	FIR6	0.033	0.01	2.740	160	0.007	0.088	0.053	0.123	0.043	*
	FIR7	0.022	0.01	1.784	160	0.076	0.076	0.041	0.112	0.297	
	FIR8	0.008	0.01	0.661	160	0.509	0.063	0.027	0.098	1.000	
	FIR9	0.001	0.01	0.110	160	0.912	0.056	0.021	0.091	1.000	
Rej Amb	(Intcpt)	0.041	0.01	2.812	35.07	0.008	0.041	0.012	0.070		
	FIR2	0.006	0.01	0.562	160	0.575	0.048	0.019	0.077	1.000	
	FIR3	0.028	0.01	2.537	160	0.012	0.069	0.040	0.099	0.078	.
	FIR4	0.019	0.01	1.763	160	0.080	0.061	0.032	0.090	0.312	
	FIR5	0.025	0.01	2.301	160	0.023	0.067	0.038	0.096	0.128	
	FIR6	0.031	0.01	2.831	160	0.005	0.073	0.044	0.102	0.037	*
	FIR7	0.022	0.01	1.973	160	0.050	0.063	0.034	0.092	0.243	
	FIR8	0.011	0.01	1.000	160	0.319	0.052	0.023	0.082	0.951	
	FIR9	-0.003	0.01	-0.284	160	0.777	0.038	0.009	0.067	1.000	
Rej Con	(Intcpt)	0.051	0.01	5.181	53.98	3.37E-06	0.051	0.032	0.070		
	FIR2	-0.006	0.01	-0.652	160	0.515	0.045	0.025	0.064	1.000	
	FIR3	0.010	0.01	1.041	160	0.299	0.061	0.041	0.080	1.000	
	FIR4	0.011	0.01	1.218	160	0.225	0.063	0.043	0.082	1.000	
	FIR5	0.018	0.01	1.866	160	0.064	0.069	0.049	0.088	0.497	
	FIR6	0.011	0.01	1.195	160	0.234	0.062	0.043	0.082	1.000	
	FIR7	0.013	0.01	1.392	160	0.166	0.064	0.045	0.084	1.000	
	FIR8	-0.004	0.01	-0.416	160	0.678	0.047	0.028	0.067	1.000	
	FIR9	-0.011	0.01	-1.113	160	0.268	0.040	0.021	0.060	1.000	
aIC											
	(Intcpt)	0.060	0.02	2.740	29.73	0.010	0.060	0.017	0.102		

Acc Amb	FIR2	0.016	0.01	1.177	160	0.241	0.076	0.033	0.119	0.694	
	FIR3	0.037	0.01	2.645	160	0.009	0.096	0.053	0.139	0.057	.
	FIR4	0.028	0.01	1.995	160	0.048	0.087	0.044	0.130	0.230	
	FIR5	0.030	0.01	2.191	160	0.030	0.090	0.047	0.133	0.171	
	FIR6	0.017	0.01	1.197	160	0.233	0.076	0.033	0.119	0.694	
	FIR7	0.041	0.01	2.937	160	0.004	0.100	0.058	0.143	0.027	*
	FIR8	0.024	0.01	1.708	160	0.090	0.083	0.040	0.126	0.351	
	FIR9	0.000	0.01	0.018	160	0.986	0.060	0.017	0.103	0.986	
Acc Con	(Intcpt)	0.047	0.02	2.976	36.59	0.005	0.047	0.016	0.078		
	FIR2	0.010	0.01	0.802	160	0.424	0.056	0.025	0.087	0.423	
	FIR3	0.030	0.01	2.504	160	0.013	0.077	0.046	0.108	0.074	.
	FIR4	0.032	0.01	2.667	160	0.008	0.079	0.048	0.110	0.056	.
	FIR5	0.033	0.01	2.698	160	0.008	0.079	0.048	0.110	0.056	.
	FIR6	0.028	0.01	2.344	160	0.020	0.075	0.044	0.106	0.095	.
	FIR7	0.023	0.01	1.923	160	0.056	0.070	0.039	0.101	0.218	
	FIR8	0.017	0.01	1.405	160	0.162	0.064	0.033	0.095	0.329	
	FIR9	0.019	0.01	1.599	160	0.112	0.066	0.035	0.097	0.329	
Rej Amb	(Intcpt)	0.041	0.01	3.251	40.93	0.002	0.041	0.016	0.065		
	FIR2	0.005	0.01	0.465	160	0.643	0.046	0.021	0.070	0.674	
	FIR3	0.022	0.01	2.110	160	0.036	0.063	0.038	0.088	0.140	
	FIR4	0.026	0.01	2.530	160	0.012	0.067	0.042	0.092	0.069	.
	FIR5	0.037	0.01	3.507	160	0.001	0.077	0.053	0.102	0.004	***
	FIR6	0.024	0.01	2.268	160	0.025	0.064	0.040	0.089	0.117	
	FIR7	0.027	0.01	2.612	160	0.010	0.068	0.043	0.093	0.063	.
	FIR8	0.010	0.01	0.960	160	0.338	0.051	0.026	0.076	0.674	
	FIR9	0.016	0.01	1.493	160	0.137	0.056	0.032	0.081	0.407	
Rej Con	(Intcpt)	0.043	0.01	4.280	51.16	8.22E-05	0.043	0.023	0.063		
	FIR2	0.002	0.01	0.250	160	0.803	0.045	0.025	0.065	1.000	
	FIR3	0.011	0.01	1.173	160	0.243	0.054	0.034	0.074	0.964	
	FIR4	0.019	0.01	1.985	160	0.049	0.061	0.042	0.081	0.330	
	FIR5	0.018	0.01	1.964	160	0.051	0.061	0.041	0.081	0.330	
	FIR6	0.018	0.01	1.965	160	0.051	0.061	0.041	0.081	0.330	
	FIR7	0.020	0.01	2.134	160	0.034	0.063	0.043	0.083	0.263	
	FIR8	0.000	0.01	0.014	160	0.989	0.043	0.023	0.063	1.000	
	FIR9	-0.008	0.01	-0.836	160	0.404	0.035	0.015	0.055	1.000	
pIC											
Acc Amb	(Intcpt)	0.054	0.02	2.501	31.46	0.018	0.054	0.011	0.096		
	FIR2	0.008	0.01	0.542	160	0.588	0.062	0.019	0.104	1.000	
	FIR3	0.035	0.01	2.370	160	0.019	0.089	0.046	0.131	0.089	.
	FIR4	0.033	0.01	2.221	160	0.028	0.086	0.044	0.129	0.105	
	FIR5	0.038	0.01	2.624	160	0.010	0.092	0.050	0.135	0.061	.
	FIR6	0.036	0.01	2.447	160	0.015	0.090	0.047	0.132	0.086	.

	FIR7	0.047	0.01	3.224	160	0.002	0.101	0.059	0.144	0.010	*
	FIR8	0.013	0.01	0.909	160	0.365	0.067	0.025	0.110	1.000	
	FIR9	0.005	0.01	0.344	160	0.732	0.059	0.016	0.101	1.000	
Acc Con	(Intcpt)	0.061	0.02	3.535	36.01	0.001	0.061	0.027	0.095		
	FIR2	-0.002	0.01	-0.155	160	0.877	0.059	0.025	0.093	1.000	
	FIR3	0.022	0.01	1.654	160	0.100	0.083	0.049	0.117	0.765	
	FIR4	0.018	0.01	1.338	160	0.183	0.079	0.045	0.113	1.000	
	FIR5	0.022	0.01	1.666	160	0.098	0.083	0.049	0.118	0.765	
	FIR6	0.012	0.01	0.878	160	0.381	0.073	0.039	0.107	1.000	
	FIR7	0.015	0.01	1.149	160	0.252	0.076	0.042	0.111	1.000	
	FIR8	-0.001	0.01	-0.080	160	0.936	0.060	0.026	0.094	1.000	
	FIR9	-0.010	0.01	-0.753	160	0.452	0.051	0.017	0.085	1.000	
Rej Amb	(Intcpt)	0.044	0.01	2.953	38.36	0.005	0.044	0.015	0.073		
	FIR2	0.001	0.01	0.118	160	0.906	0.045	0.016	0.074	1.000	
	FIR3	0.022	0.01	1.889	160	0.061	0.066	0.037	0.095	0.236	
	FIR4	0.030	0.01	2.526	160	0.013	0.074	0.045	0.103	0.081	.
	FIR5	0.030	0.01	2.555	160	0.012	0.074	0.045	0.103	0.081	.
	FIR6	0.031	0.01	2.571	160	0.011	0.074	0.045	0.104	0.081	.
	FIR7	0.024	0.01	2.014	160	0.046	0.068	0.038	0.097	0.220	
	FIR8	0.011	0.01	0.923	160	0.358	0.055	0.025	0.084	1.000	
	FIR9	0.007	0.01	0.601	160	0.548	0.051	0.022	0.080	1.000	
Rej Con	(Intcpt)	0.043	0.01	4.671	61.60	1.67E-05	0.043	0.025	0.061		
	FIR2	0.007	0.01	0.786	160	0.433	0.050	0.032	0.068	0.864	
	FIR3	0.014	0.01	1.532	160	0.128	0.057	0.039	0.075	0.503	
	FIR4	0.017	0.01	1.893	160	0.060	0.060	0.042	0.078	0.319	
	FIR5	0.020	0.01	2.124	160	0.035	0.062	0.044	0.080	0.244	
	FIR6	0.020	0.01	2.164	160	0.032	0.063	0.045	0.081	0.244	
	FIR7	0.018	0.01	1.934	160	0.055	0.060	0.042	0.078	0.319	
	FIR8	0.013	0.01	1.432	160	0.154	0.056	0.038	0.074	0.503	
	FIR9	0.003	0.01	0.314	160	0.754	0.046	0.028	0.064	0.864	
nAcc											
Acc Amb	(Intcpt)	0.074	0.02	3.626	36.47	0.001	0.100	0.053	0.146		
	FIR2	0.015	0.02	0.931	160	0.353	0.101	0.054	0.147	1.000	
	FIR3	0.025	0.02	1.571	160	0.118	0.121	0.074	0.167	1.000	
	FIR4	0.037	0.02	2.320	160	0.022	0.104	0.057	0.150	1.000	
	FIR5	0.022	0.02	1.373	160	0.172	0.152	0.106	0.199	0.083	.
	FIR6	0.028	0.02	1.773	160	0.078	0.116	0.069	0.162	1.000	
	FIR7	0.021	0.02	1.306	160	0.193	0.111	0.064	0.157	1.000	
	FIR8	-0.012	0.02	-0.748	160	0.456	0.086	0.040	0.133	1.000	
	FIR9	-0.002	0.02	-0.112	160	0.911	0.075	0.029	0.121	1.000	
Acc Con	(Intcpt)	0.070	0.02	3.639	42.73	0.001	0.070	0.032	0.107		
	FIR2	0.029	0.02	1.798	160	0.074	0.099	0.061	0.137	0.434	

	FIR3	0.030	0.02	1.849	160	0.066	0.100	0.062	0.137	0.434	
	FIR4	0.022	0.02	1.318	160	0.189	0.091	0.053	0.129	0.563	
	FIR5	0.031	0.02	1.866	160	0.064	0.100	0.062	0.138	0.434	
	FIR6	0.025	0.02	1.550	160	0.123	0.095	0.057	0.133	0.484	
	FIR7	0.045	0.02	2.779	160	0.006	0.115	0.077	0.153	0.044	*
	FIR8	0.014	0.02	0.848	160	0.398	0.083	0.046	0.121	0.793	
	FIR9	0.005	0.02	0.309	160	0.758	0.075	0.037	0.112	0.793	
Rej Amb	(Intcpt)	0.070	0.01	5.391	67.81	9.55E-07	0.070	0.044	0.095		
	FIR2	-0.009	0.01	-0.646	160	0.519	0.061	0.035	0.086	1.000	
	FIR3	0.014	0.01	1.020	160	0.309	0.083	0.058	0.109	1.000	
	FIR4	0.004	0.01	0.330	160	0.741	0.074	0.049	0.100	1.000	
	FIR5	0.006	0.01	0.476	160	0.635	0.076	0.051	0.102	1.000	
	FIR6	0.005	0.01	0.336	160	0.737	0.074	0.049	0.100	1.000	
	FIR7	0.008	0.01	0.595	160	0.553	0.078	0.052	0.103	1.000	
	FIR8	0.003	0.01	0.218	160	0.828	0.073	0.047	0.098	1.000	
	FIR9	-0.017	0.01	-1.273	160	0.205	0.052	0.027	0.078	1.000	
Rej Con	(Intcpt)	0.061	0.01	4.262	91.60	4.92E-05	0.061	0.033	0.090		
	FIR2	0.001	0.02	0.087	160	0.931	0.063	0.034	0.091	1.000	
	FIR3	0.025	0.02	1.546	160	0.124	0.087	0.058	0.115	0.733	
	FIR4	0.019	0.02	1.159	160	0.248	0.080	0.052	0.109	0.986	
	FIR5	0.034	0.02	2.078	160	0.039	0.095	0.067	0.124	0.302	
	FIR6	0.032	0.02	1.922	160	0.056	0.093	0.064	0.121	0.382	
	FIR7	0.024	0.02	1.440	160	0.152	0.085	0.056	0.113	0.750	
	FIR8	0.007	0.02	0.434	160	0.665	0.068	0.040	0.097	1.000	
	FIR9	0.001	0.02	0.055	160	0.956	0.062	0.034	0.090	1.000	
vmPFC											
Acc Amb	(Intcpt)	0.052	0.02	2.612	27.58	0.014	0.052	0.013	0.092		
	FIR2	0.012	0.01	1.057	160	0.292	0.065	0.025	0.104	0.871	
	FIR3	0.028	0.01	2.421	160	0.017	0.081	0.041	0.120	0.077	.
	FIR4	0.032	0.01	2.709	160	0.007	0.084	0.044	0.124	0.047	*
	FIR5	0.050	0.01	4.282	160	0.000	0.102	0.063	0.142	1.48E-04	***
	FIR6	0.026	0.01	2.208	160	0.029	0.078	0.039	0.118	0.109	
	FIR7	0.031	0.01	2.698	160	0.008	0.084	0.044	0.124	0.047	*
	FIR8	0.010	0.01	0.848	160	0.398	0.062	0.023	0.102	0.871	
	FIR9	-0.007	0.01	-0.569	160	0.570	0.046	0.006	0.085	0.871	
Acc Con	(Intcpt)	0.050	0.02	3.302	28.66	0.003	0.050	0.020	0.081		
	FIR2	-0.004	0.01	-0.393	160	0.695	0.047	0.017	0.077	1.000	
	FIR3	0.017	0.01	1.859	160	0.065	0.068	0.038	0.098	0.315	
	FIR4	0.019	0.01	2.016	160	0.045	0.069	0.039	0.099	0.263	
	FIR5	0.026	0.01	2.761	160	0.006	0.076	0.046	0.106	0.046	*
	FIR6	0.025	0.01	2.700	160	0.008	0.076	0.046	0.106	0.049	*
	FIR7	0.014	0.01	1.540	160	0.126	0.065	0.035	0.095	0.494	

	FIR8	0.004	0.01	0.437	160	0.662	0.055	0.024	0.085	1.000	
	FIR9	0.005	0.01	0.527	160	0.599	0.055	0.025	0.086	1.000	
Rej Amb	(Intcpt)	0.032	0.01	2.223	28.74	0.034	0.032	0.004	0.060		
	FIR2	0.005	0.01	0.545	160	0.587	0.037	0.008	0.065	1.000	
	FIR3	0.026	0.01	2.984	160	0.003	0.058	0.030	0.086	0.020	*
	FIR4	0.022	0.01	2.489	160	0.014	0.054	0.025	0.082	0.064	.
	FIR5	0.033	0.01	3.701	160	2.96E-04	0.064	0.036	0.093	0.001	***
	FIR6	0.024	0.01	2.756	160	0.007	0.056	0.028	0.084	0.035	*
	FIR7	0.017	0.01	1.952	160	0.053	0.049	0.021	0.077	0.204	
	FIR8	0.010	0.01	1.142	160	0.255	0.042	0.014	0.070	0.760	
	FIR9	0.004	0.01	0.428	160	0.669	0.036	0.007	0.064	1.000	
Rej Con	(Intcpt)	0.029	0.01	3.140	41.13	0.003	0.029	0.011	0.046		
	FIR2	0.003	0.01	0.346	160	0.730	0.031	0.013	0.049	1.000	
	FIR3	0.015	0.01	2.000	160	0.047	0.044	0.026	0.062	0.182	
	FIR4	0.017	0.01	2.279	160	0.024	0.046	0.028	0.064	0.113	
	FIR5	0.021	0.01	2.697	160	0.008	0.049	0.031	0.067	0.049	*
	FIR6	0.019	0.01	2.446	160	0.016	0.047	0.029	0.065	0.087	.
	FIR7	0.022	0.01	2.935	160	0.004	0.051	0.033	0.069	0.027	*
	FIR8	0.008	0.01	1.027	160	0.306	0.036	0.018	0.054	0.913	
	FIR9	0.002	0.01	0.304	160	0.762	0.031	0.013	0.049	1.000	
control (A1)											
Acc Amb	(Intcpt)	0.074	0.02	3.626	36.47	0.001	0.074	0.034	0.115		
	FIR2	0.015	0.02	0.931	160	0.353	0.089	0.049	0.130	1.000	
	FIR3	0.025	0.02	1.571	160	0.118	0.099	0.059	0.140	0.697	
	FIR4	0.037	0.02	2.320	160	0.022	0.111	0.071	0.152	0.163	
	FIR5	0.022	0.02	1.373	160	0.172	0.096	0.056	0.137	0.849	
	FIR6	0.028	0.02	1.773	160	0.078	0.102	0.062	0.143	0.534	
	FIR7	0.021	0.02	1.306	160	0.193	0.095	0.055	0.136	0.849	
	FIR8	-0.012	0.02	-0.748	160	0.456	0.063	0.022	0.103	1.000	
	FIR9	-0.002	0.02	-0.112	160	0.911	0.073	0.032	0.113	1.000	
Acc Con	(Intcpt)	0.065	0.02	4.063	60.08	1.42E-04	0.065	0.033	0.096		
	FIR2	0.003	0.02	0.210	160	0.834	0.068	0.037	0.099	1.000	
	FIR3	0.020	0.02	1.230	160	0.220	0.084	0.053	0.115	1.000	
	FIR4	0.043	0.02	2.706	160	0.008	0.108	0.076	0.139	0.055	.
	FIR5	0.032	0.02	2.040	160	0.04	0.097	0.066	0.128	0.290	
	FIR6	0.009	0.02	0.578	160	0.564	0.074	0.042	0.105	1.000	
	FIR7	0.021	0.02	1.320	160	0.189	0.086	0.054	0.117	1.000	
	FIR8	0.000	0.02	0.007	160	0.994	0.065	0.033	0.096	1.000	
	FIR9	-0.003	0.02	-0.161	160	0.872	0.062	0.031	0.093	1.000	
Rej Amb	(Intcpt)	0.060	0.01	4.212	53.39	9.77E-05	0.060	0.032	0.088		
	FIR2	-0.007	0.01	-0.495	160	0.621	0.053	0.025	0.081	1.000	
	FIR3	0.022	0.01	1.599	160	0.112	0.082	0.054	0.110	0.769	

	FIR4	0.010	0.01	0.706	160	0.481	0.069	0.041	0.097	1.000
	FIR5	0.019	0.01	1.437	160	0.15	0.079	0.051	0.107	0.905
	FIR6	0.007	0.01	0.510	160	0.611	0.067	0.039	0.095	1.000
	FIR7	0.029	0.01	2.121	160	0.036	0.089	0.061	0.117	0.271
	FIR8	0.001	0.01	0.061	160	0.951	0.061	0.033	0.089	1.000
	FIR9	0.007	0.01	0.514	160	0.608	0.067	0.039	0.095	1.000
Rej Con	(Intcpt)	0.053	0.01	4.473	82.35	2.45E-05	0.053	0.030	0.076	
	FIR2	-0.002	0.01	-0.125	160	0.900	0.051	0.028	0.074	1.000
	FIR3	0.017	0.01	1.303	160	0.195	0.070	0.047	0.093	0.963
	FIR4	0.007	0.01	0.532	160	0.595	0.060	0.036	0.083	1.000
	FIR5	0.020	0.01	1.531	160	0.128	0.073	0.050	0.096	0.880
	FIR6	0.019	0.01	1.418	160	0.158	0.071	0.048	0.095	0.937
	FIR7	0.023	0.01	1.721	160	0.087	0.075	0.052	0.099	0.682
	FIR8	0.014	0.01	1.033	160	0.303	0.066	0.043	0.090	1.000
	FIR9	0.002	0.01	0.186	160	0.853	0.055	0.032	0.079	1.000