5-18-2011

Ventromedial Frontal Lobe Damage Disrupts Value Maximization in Humans

Nathalie Camille
Montreal Neurological Institute, McGill

Cathryn A. Griffiths
Montreal Neurological Institute, McGill

Khoi Vo
University of Pennsylvania

Lesley K. Fellows
Montreal Neurological Institute, McGill

Joseph W. Kable
University of Pennsylvania, kable@psych.upenn.edu

Ventromedial Frontal Lobe Damage Disrupts Value Maximization in Humans

Abstract

Recent work in neuroeconomics has shown that regions in orbitofrontal and medial prefrontal cortex encode the subjective value of different options during choice. However, these electrophysiological and neuroimaging studies cannot demonstrate whether such signals are necessary for value-maximizing choices. Here we used a paradigm developed in experimental economics to empirically measure and quantify violations of utility theory in humans with damage to the ventromedial frontal lobe (VMF). We show that people with such damage are more likely to make choices that violate the generalized axiom of revealed preference, which is the one necessary and sufficient condition for choices to be consistent with value maximization. These results demonstrate that the VMF plays a critical role in value-maximizing choice.

Disciplines

Cognition and Perception | Cognitive Neuroscience | Neuroscience and Neurobiology | Neurosciences | Other Psychiatry and Psychology

Behavioral/Systems/Cognitive

Ventromedial Frontal Lobe Damage Disrupts Value Maximization in Humans

Nathalie Camille,[1] **Cathryn A. Griffiths,**[1] **Khoi Vo,**[2] **Lesley K. Fellows,**[1] **and Joseph W. Kable**[2]

[1]Department of Neurology and Neurosurgery, McGill University, Montreal Neurological Institute, Montreal, Quebec H3A 2B4, Canada, and [2]Department of Psychology, University of Pennsylvania, Philadelphia, Pennsylvania 19104

Recent work in neuroeconomics has shown that regions in orbitofrontal and medial prefrontal cortex encode the subjective value of different options during choice. However, these electrophysiological and neuroimaging studies cannot demonstrate whether such signals are necessary for value-maximizing choices. Here we used a paradigm developed in experimental economics to empirically measure and quantify violations of utility theory in humans with damage to the ventromedial frontal lobe (VMF). We show that people with such damage are more likely to make choices that violate the generalized axiom of revealed preference, which is the one necessary and sufficient condition for choices to be consistent with value maximization. These results demonstrate that the VMF plays a critical role in value-maximizing choice.

Introduction

Value is a central concept in economic decision making. To make optimal choices, individuals must estimate and compare the subjective value of the available options. The brain mechanisms supporting value-driven choice have been a focus of research in the emerging field of neuroeconomics, but important questions about these mechanisms remain.

Converging evidence suggests that the ventromedial frontal lobe (VMF) plays a role in the coding of the subjective value of options. Electrophysiological (Murray et al., 2007; Kepecs et al., 2008; Padoa-Schioppa and Assad, 2008) and neuroimaging (McClure et al., 2004; Kable and Glimcher, 2007; Plassmann et al., 2007; Tom et al., 2007; Hare et al., 2008; Chib et al., 2009) studies have found activity in VMF related to value assessment in the context of choice, supporting the theory that VMF plays a key role in using value to guide choice.

However, economic decision making often includes complex considerations of risk, delay, and effort. The multidimensionality of such choices can make it challenging to establish the specific role played by VMF. Furthermore, most work in humans has relied on functional neuroimaging, leaving open whether value-related activations in VMF are, in fact, critical for decision mak-

ing. Lesion studies in animals do suggest a critical role for VMF (Murray et al., 2007; Schoenbaum et al., 2009; Walton et al., 2010), but in these studies it is difficult to disentangle the role of VMF in value-maximizing choices from its role in value learning (Fellows and Farah, 2003; Tsuchida et al., 2010).

If VMF plays a central role in representing subjective value, then damage to this area should disrupt even very simple value-based choices. We previously found support for this hypothesis in a single study of simple preference judgments in patients with focal frontal lobe damage (Fellows and Farah, 2007). This question is so central to understanding the role of VMF in decision making that it seemed important to further pursue it in a new sample, and with different measures.

Here we took advantage of an experimental economic paradigm to empirically measure and quantify violations of utility theory. This paradigm allows us to identify disruptions in value-based choice. Specifically, we tested for violations of the General Axiom of Revealed Preference (GARP) (Samuelson, 1938; Houthakker, 1950; Afriat, 1967), the one necessary and sufficient condition for utility maximization. A person's choices are consistent with utility maximization if and only if they satisfy GARP. The basic notion embodied in GARP is that a person's choices should be internally consistent and transitive. For example, if someone picks X over Y, then they should not at another point select Y when X was also available. If someone chooses X over Y, and Y over Z, then they should also choose X over Z.

Subjects repeatedly chose between different bundles of real food items. We then determined whether their choices respected GARP and the number and severity of any violations. We hypothesized that intact VMF is necessary for optimal representation of value in this choice setting and predicted that damage to this region would therefore lead to more GARP violations than are seen in demographically similar healthy subjects. This design isolates subjective value, captured as simply as possible, from other aspects of decision making. As such, it provides a strong test of the
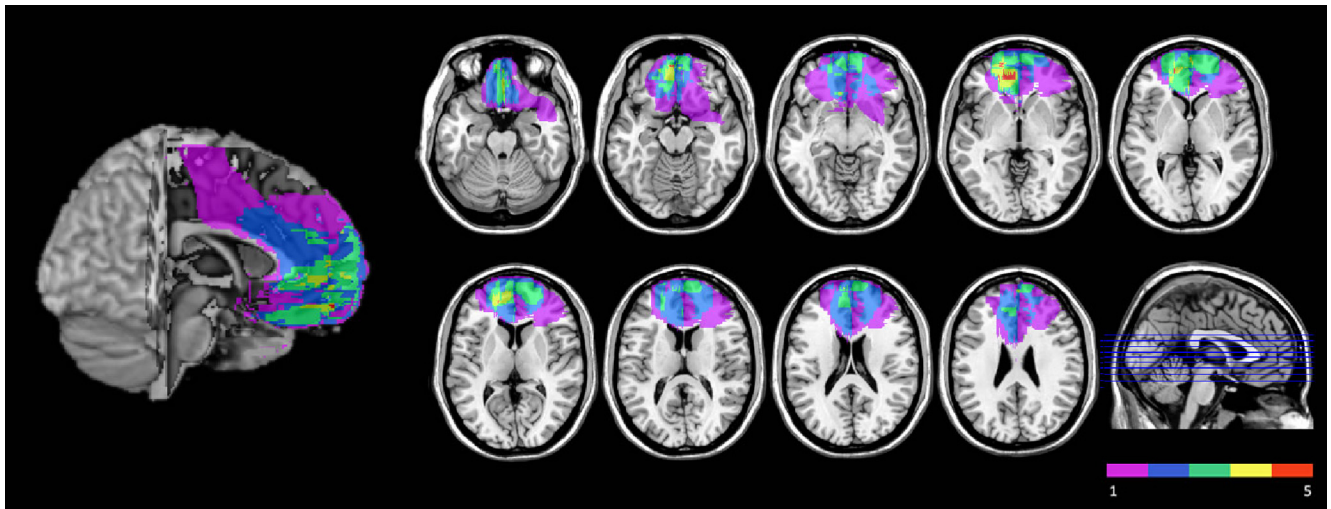
Figure 1. Location and overlap of brain lesions of the nine subjects with VMF damage, projected on axial slices of the MNI brain. Different colors indicate the number of subjects who had damage involving a particular area in common, as indicated in the color key.

Table 1. Background information [mean (SD)]

Group	N	Sex ratio (F:M)	Age (years)	Education (years)	MoCA	BDI	ANART IQ estimate	Sentence comprehension accuracy	Verbal fluency FAS	Verbal fluency Animals	Digit span Forward	Digit span Backward
Control	22	13:9	60.09 (10.8)	14.95 (3.5)	28.4 (1.2)	5.23 (6.1)	124.06 (6.4)	0.95 (0.1)	20.93 (4.0)	12.39 (4.7)	6.7 (1.7)	5.27 (1.6)
VMF	9	7:2	58.11 (15.9)	14.11 (3.2)	25* (4.4)	10.33 (8.3)	116.06* (10.5)	0.98 (0.04)	15.89 (5.9)	9.33* (3.1)	6.33 (1.1)	4.50 (1.7)

F, Female; M, male; MoCA, Montreal Cognitive Assessment; BDI, Beck Depression Inventory; ANART, American National Adult Reading Test. *Values that differ significantly between the two groups, t test, p < 0.05.

claim that VMF encodes subjective value and that this representation contributes to decision making in humans.

Materials and Methods

Subjects. Originally, 11 VMF patients and 26 controls subjects were tested. Two participants from the patient group and four participants from the control group were not included in the analysis (see Data analysis, below). Subjects with VMF damage (N = 9; 2 males, 7 females; mean age: 58.11 years; SD: 15.92 years; age range: 32–79 years) were identified through the McGill Cognitive Neuroscience Research Registry and the Center for Cognitive Neuroscience database at the University of Pennsylvania. This group, defined a priori, included patients with medial orbitofrontal and/or ventral medial prefrontal cortex damage, following the boundaries laid out in Stuss and Levine (2002). Figure 1 shows the lesion overlap in this sample, which covered most of the area within VMF, including both orbital and ventral medial surfaces.

Lesions were secondary to anterior communicating artery aneurysm rupture (N = 2), ischemic stroke (N = 2), hemorrhagic stroke (N = 1), or resections of low-grade brain tumors (N = 4) and had occurred at least 6 months before testing. Five patients had lesions limited to the right hemisphere, one to the left hemisphere, and three patients had bilateral lesions. Individual lesions were traced from the most recent clinical magnetic resonance imaging onto the standard Montreal Neurological Institute brain by a neurologist with experience in image analysis and blind to the behavioral data, using MRIcro software (Rorden and Brett, 2000).

The age- and education-matched healthy control group was comprised of 22 healthy volunteers recruited from the community by advertisement (N = 22, 9 males, 13 females; mean age: 60.09 years; SD: 10.80 years; age range: 37–79 years). Volunteers were free from current or past neurologic or psychiatric disease, closed head injury, and substance abuse; none were taking psychoactive medication. All scored >25/30 on the Montreal Cognitive Assessment, a screening test for mild cognitive impairment (Nasreddine et al., 2005). Groups did not differ significantly in age or education (ANOVA, ps > 0.83). Subjects with brain damage were administered a short neuropsychological battery for screening purposes. Demographic information and performance on the neuropsychological screening tests are summarized in Table 1.

The local ethics committee approved the study and all participants gave written informed consent before beginning the experiment, in accordance with the Declaration of Helsinki. As compensation for their time, all participants received a nominal fee and the food items from one of their choices, drawn at random at the end of the experiment.

Experimental design. An adaptation of Harbaugh et al.'s (2001) food preference task, originally used to study GARP violations in children, was used. Participants made selections from combinations of different quantities of chocolate bars and juice boxes. Participants decided whether the task would involve milk or dark chocolate and apple or orange juice. This was done to ensure that participants liked the food items they were subsequently asked to make choices about.

The task was administered as a series of questionnaires. On each sheet of paper, the participants saw one printed list of food bundles, known as a choice set, from which they had to select their preferred bundle. Each line of the list described a single food bundle, which was made up of a combination of a certain number of chocolate bars (zero to nine) and a certain number of juice boxes (zero to nine). For example, one bundle may have been four bars of chocolate and one box of juice. Each list began with bundles offering more juice boxes than chocolate bars and ended with bundles offering more chocolate bars than juice boxes. As such, the number of chocolate bars increased across successive bundles while the number of juice boxes decreased.

Following the procedure used by Harbaugh et al. (2001), participants saw a total of 11 choice sets, each consisting of between three and seven bundles. The order of presentation of the 11 choice sets was randomized across participants. Implicit in each choice set was a situation where bars of chocolate and boxes of juice were each a fixed price, and the chooser had a fixed income. For example, the choice set "0 bars of chocolate and 6 boxes of juice; 1 bar of chocolate and 3 boxes of juice; or 2 bars of chocolate and 0 boxes of juice" assumes that bars of chocolate are three times as expensive as boxes of juice and that the chooser can only afford two bars of chocolate, or the equivalent. Furthermore, the choice sets only included integer combinations of goods and assumed that choosers selected the maximum amounts possible (i.e., the choice set above did not include "0.5 bars of chocolate and 4.5 boxes of juice" nor did it include "1 bar of chocolate and 2 boxes of juice"). Given this constraint,
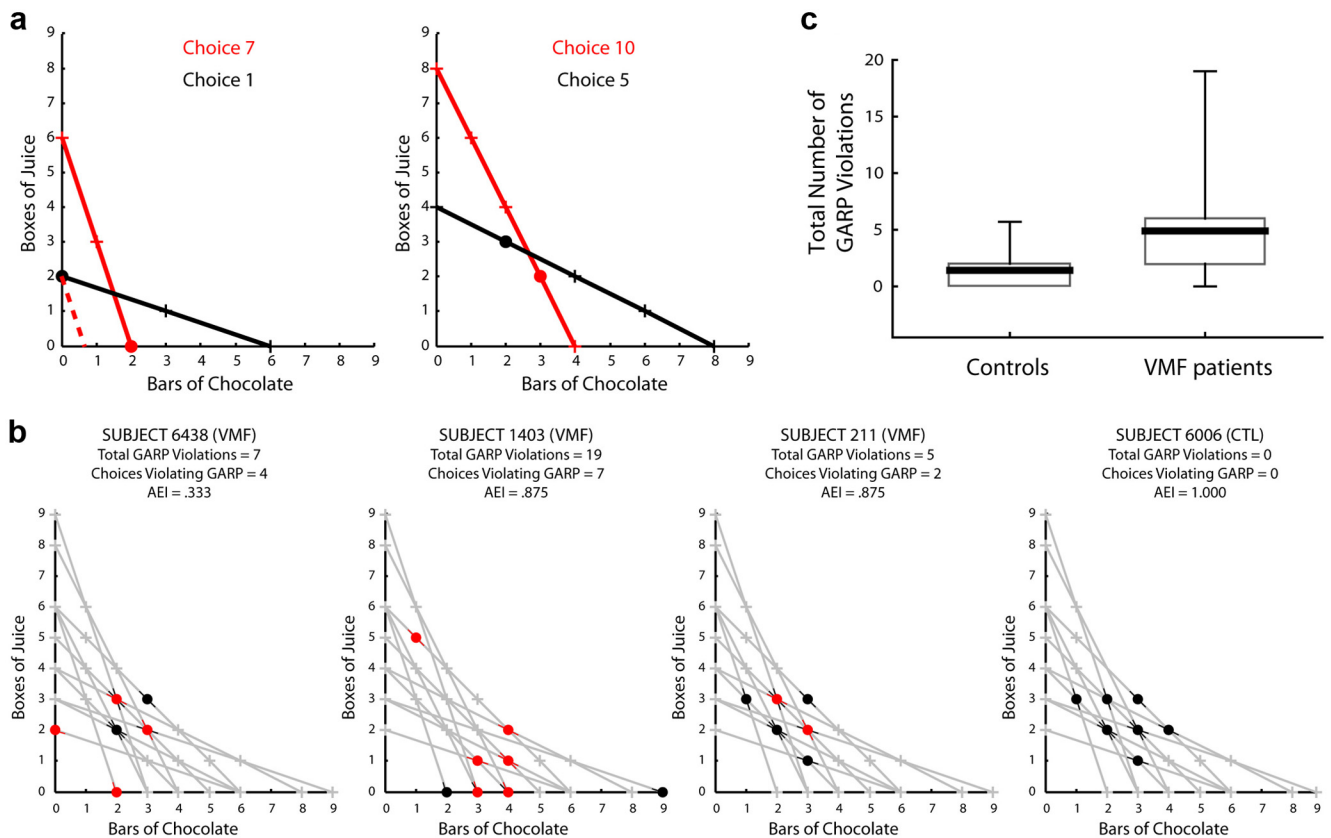
Figure 2. *a*, Two examples of choices that violate GARP. Both were made by subject 6438. Left, choice 7 (red) violates GARP with respect to choice 1 (black). Right, choice 10 (red) violates GARP with respect to choice 5 (black). For each choice, the selected bundle is marked with a circle and all other bundles in the choice set are marked with crosses. A solid line connects all of the bundles in each choice set. The dotted red line illustrates how AEI is calculated in this case. *b*, The full set of choices made by four example subjects. The gray lines illustrate all 11 choice sets that each subject faced. Bundles that were selected are marked with circles, and those that were not selected are marked with crosses. Selections that violated GARP are shown in red, those that do not violate GARP are shown in black. The colored hashes through each circle clarify for which choice set(s) the bundle was selected. A control subject (6006) who made no GARP violations, the VMF subject (211) closest to the average number of GARP violations and AEI, the VMF subject (1403) with the most total violations, and the VMF subject (6438) with the smallest AEI are shown. *c*, Box-and-whisker plots of total number of GARP violations by group. The solid line shows the group mean, the surrounding boxes delimit the 25th to 75th percentiles, and the whiskers the 10th to 90th percentiles.

the bundles in each choice set all fall on a line in the two-dimensional space defined by the numbers of bars of chocolate and boxes of juice (Fig. 2). The slope of this line depends on the ratio of the price of chocolate to the price of juice, with the steeper slopes occurring when chocolate is more expensive, and the shallower slopes when juice is more expensive. The intercept depends on the income constraint, with larger intercepts occurring when income is greater. Different choice sets were constructed by varying the implied prices and income. The relative prices of chocolate and juice varied from 3:1 to 1:3. The 11 choice sets used are the most informative of those that can be constructed when bundles involve a whole number of goods <10 and choice sets involve >2 bundles. The possible choice sets that were omitted intersect rarely, if at all, with the ones used, so they would provide little information about the presence or absence of GARP violations (to see how violations can only arise when choice sets intersect, see Fig. 2*a*).

During the task, participants were instructed to circle their preferred bundle in each choice set. They were not directed to be consistent in their choices between sets. To increase participants' motivation for choosing their preferred bundles, participants were informed that, at the end of the task, they would receive their chosen bundle from a choice set randomly selected by the experimenter. To ensure that participants had sufficient time to read each list thoroughly and carefully consider all the bundles, we allowed 30 s for each choice set. If subjects made their choices before 30 s had elapsed, they were required to wait before turning the page to the next list.

Following the method of Harbaugh et al. (2001), after participants had made their first choices, each list was administered a second time. On this second opportunity, participants examined each choice set again, this

time for 15 s. The sets showed their previous choices, and they were given the opportunity to change their minds, marking any desired changes by crossing out the old and circling the new decisions. Before this second administration of the questionnaire, participants were informed that their second choices would be the ones awarded if the choice set in question was randomly picked at the end of the task. As in the Harbaugh et al. (2001) study, this was done to provide subjects the opportunity to be more consistent, by allowing them to review all their choices.

Data analysis. Our analysis tested whether the 11 selections made by each subject were consistent with GARP. We provide a formal technical definition of GARP below, but the key concept is that choices should be internally consistent and transitive. A violation of GARP occurs, for example, when a person selects bundle of goods *a* over a bundle that has the same or more of each good than bundle *b*, and then selects bundle *b* over a bundle that has the same or more of each good than bundle *a*.

Figure 2*a*, left, provides concrete examples of GARP violations. It shows the seventh choice of subject 6438. This choice violates GARP with respect to choice 1. In choice 7, the subject selected two bars of chocolate and no boxes of juice instead of six boxes of juice and no bars of chocolate. However, in choice 1, the subject selected two boxes of juice and no bars of chocolate instead of six bars of chocolate and no boxes of juice. Choice 7 implies that the subject strongly prefers chocolate to juice whereas choice 1 implies the opposite. Figure 2*a*, right, shows a second example of a GARP violation made by the same subject. In this case, the tenth choice of this subject violates GARP with respect to the fifth choice.

For each selection, we determined: (1) whether that selection constitutes a GARP violation or not, with respect to any of the other 10 selections (scored 0 or 1); and (2) with respect to how many of the other 10

selections that selection violates GARP (scored 0–10). We then summed these measures over all 11 selections to obtain two aggregate measures of GARP violations. We refer to the sum of (1) as the number of choices that violated GARP and to the sum of (2) as the total number of GARP violations.

Formally, we can denote each bundle as a vector of consumption goods. For example, the vector [1;3] would represent one bar of chocolate and three boxes of juice. Furthermore, each choice set implies a set of prices, and these prices can also be expressed as a vector. For example, the price vector [3,1] denotes that a bar of chocolate costs three times as much as a box of juice (since only relative prices matter in our analysis, we adopt the convention of setting the price of a box of juice to 1). Following Varian (1993), we then say that bundle x^a is directly revealed preferred (denoted R^D) to bundle x^b if x^a is chosen over x^b, or over a bundle with at least as much of every good as in x^b and more of at least one good. That is, $x^a R^D x^b$ if and only if $p^a x^a \geq p^a x^b$, where p^a is the vector of prices at which the vector of goods x^a is chosen. We also say that bundle x^a is revealed preferred (denoted R) to bundle x^b if bundle x^a is directly revealed preferred to bundle x^b, or if there is a sequence of directly revealed preferred relations linking x^a to x^b (e.g., $x^a R^D x^i$ and $x^i R^D x^b$). Thus, R is the transitive closure of the preference relation R^D. Finally, we can define GARP: If $x^a R x^b$, then it is not the case that $p^b x^b > p^b x^a$. In other words, GARP states that if bundle of goods a is revealed preferred to bundle b, then b cannot be chosen when a is affordable.

To check for GARP violations, we first constructed an 11×11 matrix (M) that expressed the direct revealed preference relations among the 11 selected bundles. $M(i,j)$ was set equal to 1 if the bundle selected from the ith choice set was directly revealed preferred to the bundle selected from the jth choice set (if $p^j x^i \geq p^j x^j$), and set to 0 otherwise. Following the procedure by Varian (1993), we found the transitive closure of M by raising it to the 11th power (because we had 11 observations) and taking the sign of the result. In the resulting matrix (MT), $MT(i,j)$ equals 1 if the bundle selected from the ith choice set is revealed preferred (directly or indirectly) to the bundle selected from the jth choice set. We then checked for GARP violations by testing whether $p^j x^j \geq p^j x^i$ when $MT(i,j) = 1$, which if true would mean that the selection from the jth choice was a violation of GARP with respect to the selection from the ith choice set. As described above, since each choice could have violated GARP with respect to multiple other selections, we calculated the number of choices that violated GARP and the total number of GARP violations for each subject.

Satisfying GARP is binary. Either a participant's choices are consistent with GARP or they are not. Although the number of GARP violations provides one measure of the inconsistency in someone's choices, a participant could also have a large number of relatively minor violations. Afriat (1972) proposed a measure, called Afriat's efficiency index (AEI), to quantify the severity of GARP violations. AEI ranges from 0 to 1, with 1 indicating choices that satisfy GARP and lower numbers indicating more serious violations. Conceptually, $1 - $ AEI is the degree of inefficiency brought about by a set of selections that violates GARP in terms of the proportion of income wasted on account of the violation. For example, an AEI of 0.33 means that the least expensive bundle that would have satisfied GARP was 33% as expensive as the one actually chosen, which violated GARP.

Figure 2a illustrates how AEI is calculated. Notice that the GARP violation in the left panel occurs because the selected item from choice 1 is within the budget set of choice 7. That is, the black circle is below (on the 0,0 side of) the red line connecting the offered bundles in choice 7. One measure of the seriousness of a violation is how far inside the budget set it is. In this case, this would be the fraction you would need to multiply every point on the red line by in order for the black circle to no longer be below the line. This is shown in Figure 2a, left, by the dotted red line. The dotted red line has the same slope as the solid red one, but every point is multiplied by 0.33; the black circle is just above this dotted line. This same calculation can be made for each violation independently, and AEI is the minimum fraction across all violations. For subject 6438, the violation shown in Figure 2a, left, was the most severe, so the AEI for this subject was 0.33.

Formally, again following Varian (1993), we can say that x^a is directly revealed preferred to x^b at an efficiency of e if $ep^a x^a > p^a x^b$, define R_e as the

transitive closure of this analog of the direct revealed preference measure, and define GARP$_e$ as the analog of GARP: if $x^a R_e x^b$, then it is not the case that $ep^b x^b > p^b x^a$.

AEI is simply the efficiency e at which this condition is satisfied. We calculated AEI by substituting R_e and GARP$_e$ into the process described above for finding GARP violations, and iteratively searching from $e = 1$ downwards until there were no GARP$_e$ violations.

There were several advantages of using discrete choice sets, but this aspect of the paradigm also necessitated some important assumptions. First, our analysis of GARP violations assumes that preferences are strictly monotonic with the number of items in each bundle. Put simply, we assume that people always prefer more of one item to less of that item. Three boxes of juice and three bars of chocolate are preferred to three boxes of juice and two bars of chocolate, and so on. Although monotonicity is not an unreasonable assumption, our analyses could misestimate GARP violations if our subjects' preferences violated monotonicity (for example, a situation that violates monotonicity is one where there is a strong complementarity of goods. A subject might prefer having two matched pairs of right and left shoes over the same plus three unmatched right shoes). Second, discrete choice sets could lead to a misestimation of AEI. Subjects might have preferred a fractional bundle of juice and chocolate to the integer bundle they selected, and if allowed to make this fractional choice they might have exhibited an increased or decreased AEI. Importantly for our purposes, both of these known limitations of the current paradigm (Harbaugh et al., 2001) would apply equally to patients and controls.

Note that a person's choices would be consistent with GARP if none of their selections contained any of one of the items—for example, if they chose bundles of all chocolate and no juice each time. Such a pattern could be caused by a strong preference for one item over the other, or by following a decision strategy to simplify the task. Since such a behavioral pattern is automatically and simplistically consistent with GARP, the analysis excluded all subjects who chose in this manner. This criterion eliminated two of 11 participants from the original patient group and four of 26 participants from the original control group ($\chi^2(1) = 0.67$, $p = 0.8$).

Results

The choices of four exemplar subjects are shown in Figure 2b, with aggregate performance shown in Figure 2c. Overall, compared with healthy subjects, VMF patients were less consistent in their choices, according to both the total number of GARP violations and the number of choices that violated GARP at all. VMF patients had a significantly greater total number of GARP violations (Mann–Whitney U test, Z = -2.4, p = 0.01). On average, control subjects had 1.41 GARP violations (SD = 2.54), with a majority making no GARP violations (15 of 22 subjects; 68%). VMF patients had a mean of 4.89 GARP violations (SD = 5.71), with a large majority making two or more GARP violations (8 of 9 subjects; 89%). VMF patients also made a significantly greater number of choices that violated GARP (Z = -2.5, p = 0.01). VMF patients made 2.45 (SD = 2.07) choices that violated GARP on average; in the control group, the mean number of choices that violated GARP was only 0.73 (SD = 1.28).

Not only were VMF subjects more inconsistent in their choices compared with controls, they also more severely violated GARP, as indexed by a lower AEI (Z = -2.1, p = 0.02). The mean AEI in controls was 0.95 (SD = 0.15). In contrast, the mean AEI in VMF patients was 0.90 (SD = 0.22).

On the second administration of the questionnaire, approximately half the controls and one of the patients made one or more changes to their initial choices (12 of 22 controls and 1 of 9 VMF patients, Z = -1.7, p = 0.05). Of those who made changes, the mean number of changes was low (controls: 1.5 changes, SD = 0.8; patient: 2 changes). The pattern of results was identical whether first or final choices were used in the analysis (statistical

tests for final choices: total number of GARP violations, $Z = -2.4$, $p = 0.01$; number of choices that violated GARP, $Z = -2.5$, $p = 0.01$; AEI, $Z = -2.1$, $p = 0.02$).

Discussion

This study applied a well established experimental economics framework to test whether VMF plays a necessary role in value-based choice. If, as is widely supposed, a subjective value signal in this region contributes critically to decision making, then damage to this region should lead to more violations of GARP compared with healthy subjects. We observed exactly this pattern using a simple paradigm involving repeated choices between bundles of consumable foods: patients with VMF damage were more inconsistent in their choices, with significantly more violations of GARP, compared with demographically similar healthy participants. Further, these violations were significantly more severe.

These findings confirm and extend the conclusions of a previous study that showed that VMF damage disrupted the transitivity of even very simple pairwise preference judgments (Fellows and Farah, 2007). The present study also involved a very simple decision task, but one that is amenable to a more formal, utility function-based analysis. This task had the additional advantages of involving real (as opposed to hypothetical) choices, of minimizing any reliance on explicit memory, and of avoiding complex instructions. The design thus focuses tightly on value-based choice, using a preference judgment paradigm simple enough to have been successfully used with children, and that avoids the more complex considerations of delay, risk, or ambiguity that complicate the interpretation of much of the existing neuropsychological work on this topic.

These results are thus strong evidence that VMF plays a critical role in value-based choice in humans. This provides important converging support for a growing body of neuroimaging work in healthy subjects that has reported activations within ventromedial and orbitofrontal prefrontal cortex in relation to preference judgment (Arana et al., 2003; Paulus and Frank, 2003; McClure et al., 2004; Kim et al., 2007; Chaudhry et al., 2009; Cunningham et al., 2009; Lebreton et al., 2009) and in proportion to the subjective value of stimuli in other tasks (Kringelbach et al., 2003; Plassmann et al., 2007; Kable and Glimcher, 2009; Hare et al., 2010; Peters and Büchel, 2010).

Our results also complement a large body of work in animal models showing that the VMF is critical for learning stimulus values or stimulus–outcome associations (Murray et al., 2007; Walton et al., 2010). One explanation for these findings is that the VMF encodes the reward predicted by a stimulus. A very similar type of quantity could be used to make value-maximizing choices, and animals with VMF lesions also exhibit disrupted preferences regarding food (Baylis and Gaffan, 1991), at least in some contexts (Izquierdo et al., 2004; Machado and Bachevalier, 2007).

Despite the simplicity of the task used here, we cannot establish whether the VMF is involved in representing the value of each option, in making a comparison between these values, or perhaps both. All of these possibilities have some support from other work (Tremblay and Schultz, 1999; Elliott et al., 2008; Fujiwara et al., 2008; Padoa-Schioppa and Assad, 2008; Grabenhorst and Rolls, 2009; Simmons et al., 2010), and any of these would result in the deficit we observed.

Though there were many benefits of the simplicity of the task we used, there were also some important assumptions. Because participants chose from discrete bundles, our analysis of GARP violations had to assume that preferences were strictly mono-

tonic (that people would always prefer to get more of a good). Using discrete choices could also lead to misestimates of AEI. These are both recognized limitations of this paradigm (Harbaugh et al., 2001) that would apply equally to all groups in the current experiment. Further work could use more complex continuous choice sets (Choi et al., 2007) or examine newly developed measures of the severity of GARP violations (Dean and Martin, 2009).

Future studies might also further examine the range of preferences for which VMF is critical. Recent work has argued that the brain may use a common currency to evaluate all kinds of goods, with VMF a key region in this process (Chib et al., 2009; FitzGerald et al., 2009). The present study supports a key role for VMF in simple preference judgments of consumable goods, but leaves open the question of the generality of this value representation. FMRI studies suggest that activation within VMF can reflect values in tasks ranging from economics to esthetics to ethics (Greene et al., 2001; Moll et al., 2002; Knutson et al., 2003; O'Doherty et al., 2003; Kawabata and Zeki, 2004; Ursu and Carter, 2005; Hare et al., 2010; Shenhav and Greene, 2010). Further work in lesion models will be needed to fully delineate the range of choices that rely critically on such signals. However, studies to date suggest a role for this region in preference judgments that would seem to capture economic, moral, social, or esthetic value at varying levels of complexity (Camille et al., 2004; Fellows and Farah, 2007; Young et al., 2010).

These findings in neuroscience are clearly relevant for the decision sciences. Economists (Samuelson, 1938; von Neumann and Morgenstern, 1944), psychologists (Kahneman and Tversky, 1979), and computer scientists (Sutton and Barto, 1998) have often started decision models with the assumption that agents choose as if they maximize a value or utility function. Though many violations of value maximization have been noted (Tversky, 1969, 1972; Payne et al., 1993), decision scientists have debated the frequency, severity, and importance of these violations (Choi et al., 2007). Previous work has shown that under the conditions we studied, most people choose in accordance with value maximization (Harbaugh et al., 2001). Our findings add to a growing body of work that suggests that under conditions like these, the principle of value maximization can be extended to the mechanistic level. A promising avenue for future research, which is already being explored, will be to explicitly link decision models to the underlying neural mechanisms. An understanding of the neural processes that occur during choice might provide a unified explanation for both decisions that are consistent with value maximization and those that deviate from it.

References

Afriat SN (1967) The construction of utility functions from expenditure data. Int Econ Rev 8:67–77.

Afriat SN (1972) Efficiency estimation of production functions. Int Econ Rev 13:568–598.

Arana FS, Parkinson JA, Hinton E, Holland AJ, Owen AM, Roberts AC (2003) Dissociable contributions of the human amygdala and orbitofrontal cortex to incentive motivation and goal selection. J Neurosci 23:9632–9638.

Baylis LL, Gaffan D (1991) Amygdalectomy and ventromedial prefrontal ablation produce similar deficits in food choice and in simple object discrimination learning for an unseen reward. Exp Brain Res 86:617–622.

Camille N, Coricelli G, Sallet J, Pradat-Diehl P, Duhamel JR, Sirigu A (2004) The involvement of the orbitofrontal cortex in the experience of regret. Science 304:1167–1170.

Chaudhry AM, Parkinson JA, Hinton EC, Owen AM, Roberts AC (2009) Preference judgements involve a network of structures within frontal, cingulate and insula cortices. Eur J Neurosci 29:1047–1055.

Chib VS, Rangel A, Shimojo S, O'Doherty JP (2009) Evidence for a common representation of decision values for dissimilar goods in human ventromedial prefrontal cortex. J Neurosci 29:12315–12320.

Choi S, Fisman R, Gale D, Kariv S (2007) Consistency and heterogeneity of individual behavior under uncertainty. Am Econ Rev 97:1921–1938.

Cunningham WA, Kesek A, Mowrer SM (2009) Distinct orbitofrontal regions encode stimulus and choice valuation. J Cogn Neurosci 21:1956–1966.

Dean M, Martin D (2009) How consistent are your choice data? Working paper, Economics Department, New York University.

Elliott R, Agnew Z, Deakin JF (2008) Medial orbitofrontal cortex codes relative rather than absolute value of financial rewards in humans. Eur J Neurosci 27:2213–2218.

Fellows LK, Farah MJ (2003) Ventromedial frontal cortex mediates affective shifting in humans: evidence from a reversal learning paradigm. Brain 126:1830–1837.

Fellows LK, Farah MJ (2007) The role of ventromedial prefrontal cortex in decision making: judgment under uncertainty or judgment per se? Cereb Cortex 17:2669–2674.

FitzGerald TH, Seymour B, Dolan RJ (2009) The role of human orbitofrontal cortex in value comparison for incommensurable objects. J Neurosci 29:8388–8395.

Fujiwara J, Tobler PN, Taira M, Iijima T, Tsutsui K (2008) Personality-dependent dissociation of absolute and relative loss processing in orbitofrontal cortex. Eur J Neurosci 27:1547–1552.

Grabenhorst F, Rolls ET (2009) Different representations of relative and absolute subjective value in the human brain. Neuroimage 48:258–268.

Greene JD, Sommerville RB, Nystrom LE, Darley JM, Cohen JD (2001) An fMRI investigation of emotional engagement in moral judgment. Science 293:2105–2108.

Harbaugh WT, Krause K, Berry TR (2001) GARP for kids: on the development of rational choice behavior. Am Econ Rev 91:1539–1545.

Hare TA, O'Doherty J, Camerer CF, Schultz W, Rangel A (2008) Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. J Neurosci 28:5623–5630.

Hare TA, Camerer CF, Knoepfle DT, Rangel A (2010) Value computations in ventral medial prefrontal cortex during charitable decision making incorporate input from regions involved in social cognition. J Neurosci 30:583–590.

Houthakker HS (1950) Revealed preference and the utility function. Economica 17:159–174.

Izquierdo A, Suda RK, Murray EA (2004) Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. J Neurosci 24:7540–7548.

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

Kable JW, Glimcher PW (2009) The neurobiology of decision: consensus and controversy. Neuron 63:733–745.

Kahneman D, Tversky A (1979) Prospect theory: an analysis of decision under risk. Econometrica 47:263–291.

Kawabata H, Zeki S (2004) Neural correlates of beauty. J Neurophysiol 91:1699–1705.

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

Kim H, Adolphs R, O'Doherty JP, Shimojo S (2007) Temporal isolation of neural processes underlying face preference decisions. Proc Natl Acad Sci U S A 104:18253–18258.

Knutson B, Fong GW, Bennett SM, Adams CM, Hommer D (2003) A region of mesial prefrontal cortex tracks monetarily rewarding outcomes: characterization with rapid event-related fMRI. Neuroimage 18:263–272.

Kringelbach ML, O'Doherty J, Rolls ET, Andrews C (2003) Activation of the human orbitofrontal cortex to a liquid food stimulus is correlated with its subjective pleasantness. Cereb Cortex 13:1064–1071.

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.

Machado CJ, Bachevalier J (2007) Measuring reward assessment in a semi-naturalistic context: the effects of selective amygdala, orbital frontal or hippocampal lesions. Neuroscience 148:599–611.

McClure SM, Li J, Tomlin D, Cypert KS, Montague LM, Montague PR (2004) Neural correlates of behavioral preference for culturally familiar drinks. Neuron 44:379–387.

Moll J, de Oliveira-Souza R, Bramati IE, Grafman J (2002) Functional networks in emotional moral and nonmoral social judgments. Neuroimage 16:696–703.

Murray EA, O'Doherty JP, Schoenbaum G (2007) What we know and do not know about the functions of the orbitofrontal cortex after 20 years of cross-species studies. J Neurosci 27:8166–8169.

Nasreddine ZS, Phillips NA, Bédirian V, Charbonneau S, Whitehead V, Collin I, Cummings JL, Chertkow H (2005) The Montreal Cognitive Assessment, MoCA: a brief screening tool for mild cognitive impairment. J Am Geriatr Soc 53:695–699.

O'Doherty J, Winston J, Critchley H, Perrett D, Burt DM, Dolan RJ (2003) Beauty in a smile: the role of medial orbitofrontal cortex in facial attractiveness. Neuropsychologia 41:147–155.

Padoa-Schioppa C, Assad JA (2008) The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. Nat Neurosci 11:95–102.

Paulus MP, Frank LR (2003) Ventromedial prefrontal cortex activation is critical for preference judgments. Neuroreport 14:1311–1315.

Payne JW, Bettman JR, Johnson EJ (1993) The adaptive decision maker. Cambridge: Cambridge UP.

Peters J, Büchel C (2010) Neural representations of subjective reward value. Behav Brain Res 213:135–141.

Plassmann H, O'Doherty J, Rangel A (2007) Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. J Neurosci 27:9984–9988.

Rorden C, Brett M (2000) Stereotaxic display of brain lesions. Behav Neurol 12:191–200.

Samuelson PA (1938) A note on the pure theory of consumer's behavior. Economica 5:61–71.

Schoenbaum G, Roesch MR, Stalnaker TA, Takahashi YK (2009) A new perspective on the role of the orbitofrontal cortex in adaptive behaviour. Nat Rev Neurosci 10:885–892.

Shenhav A, Greene JD (2010) Moral judgments recruit domain-general valuation mechanisms to integrate representations of probability and magnitude. Neuron 67:667–677.

Simmons JM, Minamimoto T, Murray EA, Richmond BJ (2010) Selective ablations reveal that orbital and lateral prefrontal cortex play different roles in estimating predicted reward value. J Neurosci 30:15878–15887.

Stuss DT, Levine B (2002) Adult clinical neuropsychology: lessons from studies of the frontal lobes. Annu Rev Psychol 53:401–433.

Sutton RS, Barto AG (1998) Reinforcement learning: an introduction. Cambridge: MIT.

Tom SM, Fox CR, Trepel C, Poldrack RA (2007) The neural basis of loss aversion in decision-making under risk. Science 315:515–518.

Tremblay L, Schultz W (1999) Relative reward preference in primate orbitofrontal cortex. Nature 398:704–708.

Tsuchida A, Doll BB, Fellows LK (2010) Beyond reversal: a critical role for human orbitofrontal cortex in flexible learning from probabilistic feedback. J Neurosci 30:16868–16875.

Tversky A (1969) Intransitivity of preferences. Psychol Rev 76:31–48.

Tversky A (1972) Elimination by aspects: a theory of choice. Psychol Rev 79:281–299.

Ursu S, Carter CS (2005) Outcome representations, counterfactual comparisons and the human orbitofrontal cortex: implications for neuroimaging studies of decision-making. Brain Res Cogn Brain Res 23:51–60.

Varian H (1993) Microeconomic analysis. New York: W.W. Norton and Co.

von Neumann J, Morgenstern O (1944) The theory of games and economic behavior. Princeton: Princeton UP.

Walton ME, Behrens TE, Buckley MJ, Rudebeck PH, Rushworth MF (2010) Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. Neuron 65:927–939.

Young L, Bechara A, Tranel D, Damasio H, Hauser M, Damasio A (2010) Damage to ventromedial prefrontal cortex impairs judgment of harmful intent. Neuron 65:845–851.