

Memory-based Decision-making with Heuristics: Evidence for a Controlled Activation of Memory Representations

Patrick H. Khader¹, Thorsten Pachur², Stefanie Meier¹,
Siegfried Bien¹, Kerstin Jost³, and Frank Rösler⁴

Abstract

■ Many of our daily decisions are memory based, that is, the attribute information about the decision alternatives has to be recalled. Behavioral studies suggest that for such decisions we often use simple strategies (heuristics) that rely on controlled and limited information search. It is assumed that these heuristics simplify decision-making by activating long-term memory representations of only those attributes that are necessary for the decision. However, from behavioral studies alone, it is unclear whether using heuristics is indeed associated with limited memory search. The present study tested this assumption by monitoring the activation of specific long-term-memory representations with fMRI while participants made memory-based decisions using the “take-the-best” heuristic. For different decision trials, different numbers and types of information had to be retrieved and

processed. The attributes consisted of visual information known to be represented in different parts of the posterior cortex. We found that the amount of information required for a decision was mirrored by a parametric activation of the dorsolateral PFC. Such a parametric pattern was also observed in all posterior areas, suggesting that activation was not limited to those attributes required for a decision. However, the posterior increases were systematically modulated by the relative importance of the information for making a decision. These findings suggest that memory-based decision-making is mediated by the dorsolateral PFC, which selectively controls posterior storage areas. In addition, the systematic modulations of the posterior activations indicate a selective boosting of activation of decision-relevant attributes. ■

INTRODUCTION

Imagine you are a contestant on the TV show *Who Wants to Be a Millionaire*. As your final \$1 million question, the host asks you: “Which of the following Spanish cities has more inhabitants: Barcelona or Madrid?” You have heard of both cities but do not know their sizes exactly. Given that time is ticking away and that you cannot search the Internet, you are bound to retrieve relevant information about the cities from long-term memory (LTM). For instance, you may recall that Madrid is the capital of Spain or that Barcelona is located by the sea and use this information to make a decision. Behavioral studies have shown that to make memory-based decisions people often rely on simple strategies (e.g., Bröder & Schiffer, 2003, 2006). Gigerenzer, Todd, and the ABC Research Group (1999; see also Gigerenzer, Hertwig, & Pachur, 2011) proposed a set of simple strategies (“heuristics”) to model decision-making. A key heuristic is take-the-best (TTB; Gigerenzer & Goldstein, 1996), in which attributes are processed sequentially according to their importance. In TTB, the most relevant attribute is inspected first, and if this attribute discriminates between the alternatives,

search is stopped and a decision is made; that is, no further attribute is inspected. If the attribute does not discriminate, then the second-most important attribute is inspected, and so on. Assuming that people rely on such a heuristic when making memory-based decisions, the memory representation of an attribute (such as the fact that Madrid is the capital of Spain) should, strictly speaking, be activated only if it is necessary for a decision; all other representations should not be activated. Moreover, the order of the activation in the relevant representational areas should be guided by the hierarchy of attribute importance.

However, in contrast to this assumed controlled activation, memory retrieval is usually assumed to lead to an automatic spread of activation to all information associated with a particular memory entry (e.g., Anderson, 1983). Some researchers have, therefore, proposed an automatic activation of all attributes (in our example, the capital, the location, etc.) during decision-making, upon which, in a second step, controlled processes work to focus on one or the other activated piece of information (e.g., Glöckner & Betsch, 2008; Juslin & Persson, 2002; Beach & Mitchell, 1987). Behavioral studies alone, however, are unable to show whether memory search under the use of TTB is automatic or selective (and thus modulated by

¹Philipps-University of Marburg, ²University of Basel, ³RWTH Aachen University, ⁴University of Potsdam

the relevance of the respective memory contents for the decision). Brain activation measures might help to address this question.

LTM representations are most likely located in cell assemblies of posterior brain regions (see, e.g., Martin, 2007, for a review). Activation of these cell assemblies can be monitored by brain imaging methods, hence specific activation patterns can be used to indicate which type of LTM representation is retrieved for a given decision. Automatic activation should be reflected by simultaneous and uniform neural responses in all areas representing information that is associated with the decision alternatives, whereas controlled and restricted activation should be reflected by activation patterns that match the sequential retrieval demands according to TTB.

Posterior activation during memory retrieval is assumed to be controlled by the PFC (e.g., Badre & Wagner, 2007; Buckner & Wheeler, 2001). If the same principle holds for memory-based decision-making, then the PFC should be involved here, too. In particular, we expect increasing frontal activation with an increasing number of attributes that have to be processed.

METHODS

Overview of the Experimental Design

In two fMRI experiments, participants were instructed to use the TTB heuristic to decide which of two companies will be more successful in the future. Decisions had to be made on the basis of memorized attribute information about the companies. For different trials, different numbers and types of attributes had to be retrieved. For each company, participants learned information on four attributes, consisting of four types of visual stimuli that are assumed to be represented in different parts of the posterior cortex (e.g., Ishai, Ungerleider, Martin, & Haxby, 2000; Mishkin, Ungerleider, & Macko, 1983). In a learning phase, each company name was associated with one of two spatial locations (“location of the company”), faces (“manager of the company”), visual objects (“product the company is producing”), and colors (“color of the product”; see Figure 1A). According to TTB, the attributes should be inspected sequentially according to their importance, and information search is stopped and a decision is made as soon as one attribute discriminates between the companies. That is, the two companies are first compared on the top-ranked (most predictive) attribute (e.g., the location of the company). If the companies differ on this attribute, no further attribute will be inspected, and the company with a positive value (i.e., the one that indicates higher success) on that attribute is selected as the one that is more successful in the future (in our experiments, participants learned which attribute values indicate success before the decision task). If, however, the top-ranked attribute does not discriminate, then the attribute with the next-highest rank (e.g., the manager of the company) will be inspected, and

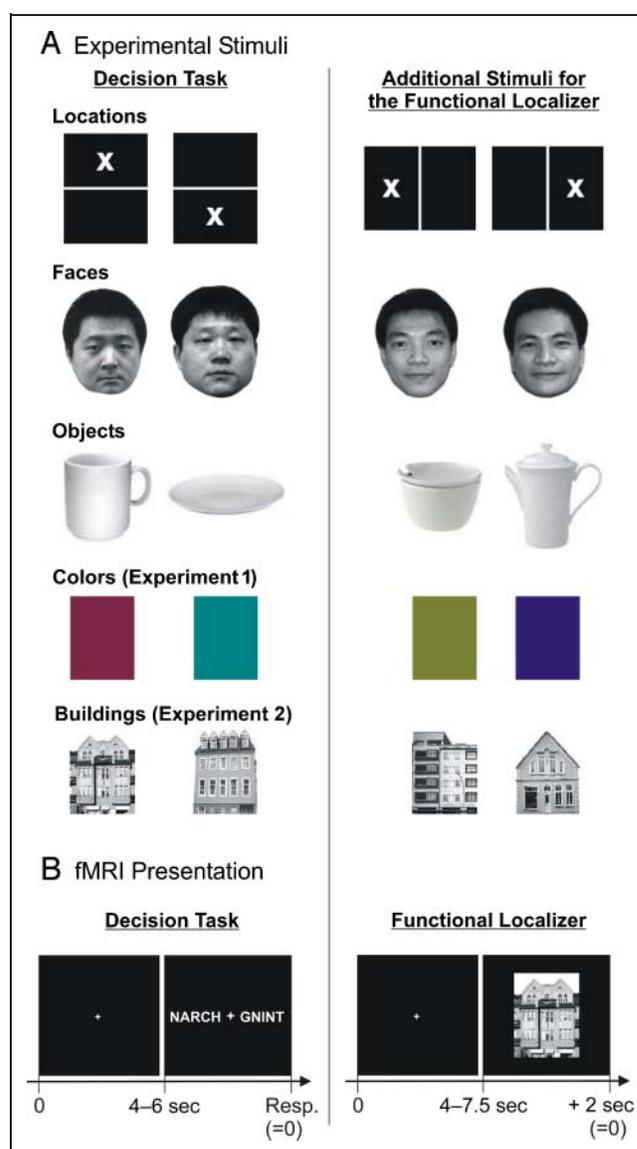


Figure 1. (A) Overview of stimuli used as attributes in the decision task and as stimuli in the localizer scans. The attributes for the decision task consisted of one of two locations (the location of the company), faces (the manager of the company), visual objects (the item the company produces), and colors (the color of the product). As the color attribute in Experiment 1 did not yield a reliable activation in the localizer scans, we replaced this attribute with buildings (i.e., front views of houses) in Experiment 2, which yielded a reliable activation in the parahippocampal gyrus (see Results). For the localizer scans, two new stimuli from each attribute category were added and participants were instructed to make an old/new distinction on each trial (see Procedure for details). (B) In the decision task, participants were shown pairs of company names and decided, using the TTB heuristic, which company would be more successful in the next year. Importantly, throughout the decision task only the names of the companies were shown, so participants had to retrieve the attribute information from memory.

so on. Therefore, depending on which companies are compared, TTB requires the retrieval of one, two, three, or all the four attributes from LTM to make a decision.

To identify the brain areas involved in the representation of the attributes, we ran so-called functional localizer

scans in which fMRI was recorded while participants processed the visual stimuli representing the attributes. Brain areas identified as most responsive to the respective stimuli were defined as ROIs for the analysis of brain activation during decision-making. These ROIs are assumed to be involved in the storage of the attributes' LTM representations. This assumption is based on theories concerning the role of the neocortex in storing LTM representations. According to these theories, the memory representation and the perceptual processing of a stimulus are mediated by the same cortical areas (e.g., O'Reilly & Rudy, 2001; McClelland, McNaughton, & O'Reilly, 1995; Squire & Alvarez, 1995; for recent reviews, see Danker & Anderson, 2010; Khader & Rösler, 2009).

In the decision task, in which only the company names were presented and participants had to retrieve the relevant attribute information from LTM, we examined activation in the ROIs as a function of the number of the to-be-retrieved attributes. If this activation is systematically modulated by the relative importance of the information for making a decision, this would be evidence for a controlled retrieval of decision-relevant attributes. To test the generality of our results, we repeated the experiment with a modified attribute hierarchy. For example, whereas the location of the company was the most important attribute for predicting a company's success in Experiment 1, it was the least important attribute in Experiment 2.

Participants

Eighteen (14 women; mean age = 22.7 years, $SD = 1.4$ years) and 16 students (11 women; mean age = 20.9 years, $SD = 1.6$ years) participated in Experiments 1 and 2, respectively. In both experiments, one participant had to be excluded because of excessive head movement and fMRI signal inhomogeneity. The final samples, thus, consisted of 17 students (13 women; mean age = 22.6 years, $SD = 1.4$ years) in Experiment 1 and 15 students (10 women; mean age = 21 years, $SD = 1.6$ years) in Experiment 2. Because of technical problems, the functional localizer scan of one participant in Experiment 2 could not be analyzed. Thus, these data were based on 14 participants. All participants were students of the University of Marburg, right-handed, healthy with normal or corrected-to-normal vision, and native speakers of German. They had no history of neurological illness, were naive with respect to the objective of the experiment, gave informed consent to participate, and received either money or course credit for their participation.

Material

In both experiments, the same 16 company names were used, which consisted of pronounceable five- to six-letter pseudowords: GNINT, NARCH, CLEEF, KUSQUE, SLARB,

KLILK, COLKS, TIRCH, CRYPSE, WRELKS, KNARFS, BLOOR, TINKS, SMAUDS, BLAUB, and SNILM (cf. Newell & Shanks, 2004). These items were taken from the MRC nonword database (www.psy.uwa.edu.au/MRCDataBase/uwa_mrc.htm; Rastle, Harrington, & Coltheart, 2002). Figure 1A shows the visual stimuli that the participants learned to associate with these company names. The stimuli for the "objects" attribute consisted of grayscale photographs of either a cup or a plate. Colors consisted of squares that were filled with one of two specific colors. To avoid that the participants recoded color information as verbal labels, we used colors that are difficult to name (red green blue [RGB] values: 128/0/64 and 0/128/128). The spatial locations, which indicated whether a company is located in the northern or in the southern part of Germany, were represented as crosses located either in the upper or the lower compartment of a rectangle. Face stimuli consisted of grayscale pictures of one of two male faces (courtesy of S. Schweinberger, University of Jena, Germany). Like the colors, the faces could not easily be tagged with a verbal label. This was confirmed by the majority of our participants in postexperimental debriefings. Size, brightness, and contrast were kept constant for both stimuli of each type of attribute.

As mentioned above, the selection of the stimuli types for the four attributes was based on previous work showing that they differ in terms of the brain regions involved in both their encoding and retrieval. Specifically, recalling faces activates the fusiform gyrus (Khader, Burke, Bien, Ranganath, & Rösler, 2005; Ishai, Ungerleider, Martin, et al., 2000), recalling locations activates the posterior parietal cortex (Khader et al., 2005, 2007; Moscovitch, Kapur, Köhler, & Houle, 1995), recalling objects activates various areas in the visual association cortex (Khader et al., 2007; Ishai, Ungerleider, Martin, et al., 2000; Moscovitch et al., 1995; Mishkin et al., 1983), and recalling colors activates the fusiform gyrus (Chao & Martin, 1999; Martin, Haxby, Lalonde, Wiggs, & Ungerleider, 1995) anterior to regions associated with color perception (e.g., Zeki et al., 1991). As, however, the color attribute in Experiment 1 did not yield a reliable ROI in the functional localizer scan, this attribute was replaced by buildings (i.e., front views of houses; see Figure 1A) in Experiment 2. Views of buildings are known to activate the parahippocampal cortex (Ishai, Ungerleider, Martin, et al., 2000; O'Craven & Kanwisher, 2000).

For the localizer scans, both old (i.e., those used in the decision task) and new stimuli from the four attribute categories were presented (see Figure 1A) and participants were instructed to make an old/new decision on each trial (see Procedure for details). The new objects were grayscale photographs of a jug and a can; the new colors had RGB values of 32/0/128 and 118/128/0; the new faces were two other Japanese male faces; the new spatial locations were crosses located in either the left or right compartment of a box; and the new buildings were front views of two other houses.

Procedure

All parts of the experiment were presented using the Presentation software (Neurobehavioral Systems; Version 11.0).

Learning of the Attribute Information

Participants learned by trial and error to associate four stimuli with each of the 16 company names. Each company was associated with a specific stimulus pattern. In each trial, participants were presented with a company name and the two picture stimuli of a specific attribute (e.g., the two faces) and had to indicate (by pressing the left or right Alt key on a computer keyboard) the correct stimulus representing the attribute value of a company. After giving a response, participants received feedback (“correct” or “incorrect”) for 1000 msec, followed by the presentation of the correct stimulus, which remained on the screen until the participants pressed the space bar. Participants were instructed not to verbalize the stimuli, and to carefully inspect and encode them by forming an integrated picture of the display. After all attributes of a single company had been learned (correct responses twice in a row), the next company was presented, etc. The order in which the four stimuli were presented for a company name was determined randomly for each learning cycle. The learning phase was completed as soon as correct responses were given for all 16 company names twice in a row. This took about 2 hr 30 min (ranges: 1 hr 30 min to 3 hr 15 min in Experiment 1 and 1 hr to 3 hr 30 min in Experiment 2). Participants returned to the laboratory on the following day. They freshened up the learned attribute knowledge until again reaching perfect performance. This took between 10 and 30 min.

Strategy Training

Subsequently, participants were trained to make decisions using the TTB heuristic. To avoid participants making inferences about the companies during the strategy training, we used a fictitious job-selection scenario. Participants were instructed to decide, on the basis of four attributes (e.g., programming experience) and using TTB, which of two applicants is more suitable for a job. Pairs of applicants were presented along with their values (“+” or “–”) on the four attributes (e.g., the applicant either had programming experience or not). Moreover, for each attribute its relative importance was indicated by the numbers 7, 8, 9, or 10, with a higher number indicating higher importance. Participants indicated the decision according to TTB by pressing the left or right Alt key. After each trial, feedback concerning the accuracy of the decision (i.e., whether it was indeed the decision mandated by TTB) was provided for 500 msec. For the strategy training only those nine pairs of applicants were used in which TTB makes a different prediction than a strategy

that integrates across all attributes (e.g., a weighted-additive strategy). This allowed us to ensure that participants had indeed understood and used TTB. Participants were presented with the 9 pairs until they had given the correct response to each of them four times in a row. The hierarchy of attribute importance (for deciding between the applicants) was varied randomly across participants.

Learning of the Attribute Hierarchy

Next, participants learned by trial and error the importance of the different attributes for predicting which of two companies would be more successful. Each attribute category was shown separately and participants had to indicate its importance (7, 8, 9, or 10, with higher numbers for more important attributes) by pressing the F7, F8, F9, or F10 key on a computer keyboard. After each response, they received feedback (“correct” or “incorrect”) for 500 msec, followed by the presentation of the correct answer, which remained on the screen until the participants pressed the space bar. This procedure was repeated until correct responses were given to each of the four attribute categories three times in a row, which took about 5–10 min. Participants learned the attribute hierarchy faces > locations > colors > objects in Experiment 1 and objects > houses > locations > faces in Experiment 2.

Learning of the Attribute Direction

Next, participants learned for each attribute category which stimulus was predictive of higher success (i.e., the attribute direction). Participants saw the two stimuli of each attribute category (e.g., a cup and a plate for the “objects” attribute) and had to indicate, by pressing the left or right Alt key, which stimulus is associated with higher success. After each response, participants received feedback (“correct” or “incorrect”) for 500 msec, followed by the presentation of the correct stimulus (which remained on the screen until the participants pressed the space bar). The complete procedure was repeated until participants gave the correct response three times in a row, which took between 5 and 10 min. The attribute direction was varied randomly across participants.

To ensure that participants were able to apply their knowledge about the correct attribute hierarchy and attribute directions quickly during the subsequent decision task, the respective learning phases were repeated under time pressure (i.e., participants had only 2 sec to respond).

Decision Task

The decision task, during which fMRI was recorded, took place immediately after the last learning task. Participants saw pairs of company names and had to decide, using TTB, which company will be more successful in the next year. In each trial, the company names appeared on the left and right sides of a fixation cross until the participant’s

response (see Figure 1B). Participants indicated whether the company presented on the left or right side would be more successful by pressing the left or right response button. They were instructed to respond as quickly and as accurately as possible. The interval between a decision and the next trial varied randomly between 2, 4, or 6 sec. No feedback was provided in the MRI scanner.

From the 16 company names, $16 \times 15/2 = 120$ pairs can be constructed. Sixty-four of these pairs require the retrieval of the most important attribute only (because it already discriminates between the companies), for 32 pairs TTB requires the retrieval of the most and second-most important attributes, for 16 pairs the retrieval of three attributes, and for 8 pairs the retrieval of all four attributes. The number of attributes that TTB requires to be retrieved was our main experimental factor. We randomly drew 37 of the 64 possible pairs for the first level (comparing one attribute) and merged the 16 + 8 pairs of the third and fourth levels into one level. Furthermore, there were seven trials included in the first level and two trials in the second level for which TTB as a noncompensatory strategy leads to a different response than a compensatory strategy (e.g., a weighted-additive strategy). Therefore, we decided to present these “critical” trials twice. In total, there were, thus, three factor levels: retrieving one (44 trials), two (34 trials), and three or four attributes (24 trials). Furthermore, the decision task also included 30 control trials (interspersed among the other trials), in which no attributes had to be retrieved. Here, the same company name was shown both on the left and the right side, and participants were instructed to press both response buttons simultaneously. Note that in these trials participants were exposed to the same visual stimulation as in the experimental trials (i.e., company names) and were also required to give a response but did not need to retrieve any information from memory to respond correctly.

During the decision task, participants lay in the MRI scanner in supine position with their head immobilized by a soft foam pillow to minimize involuntary head movements. Additionally, headphones were used to dampen scanner noise. Participants gave responses by pressing one or both of two buttons on an MRI-compatible (fiber-optic) response device attached to their thighs. Stimuli were projected on a canvas that the participants could see via mirrors mounted on the MRI head coil. The experiment consisted of three runs with 44 trials each. Between the second and third run, an anatomical reference volume was recorded, which took about 10 min. Before the first run, several practice trials were presented to familiarize the participants with the testing procedure.

The three runs were followed by the localizer scan. Here, both old (i.e., the experimental) and new stimuli from the four attribute categories were presented in random order and participants had to make an old/new distinction by pressing the left or right button on the response device (counterbalanced across participants).

Each stimulus was presented in the center of the screen for 2 sec, followed by a fixation cross for 4–7.5 sec (varied randomly in steps of 500 msec; see Figure 1B). Each of the two old and the two new stimuli from each of the four attribute categories was presented six times, resulting in 96 trials. The localizer scan took about 10 min. In total, participants were in the MRI scanner for approximately 1 hr 15 min.

fMRI Data Acquisition, Preprocessing, and Statistical Analysis

Anatomical and functional imaging was performed with a 1.5-T MR scanner (Sigma, GE Medical Systems). Functional BOLD images with 19 oblique slices covering the whole brain were acquired with a T2*-weighted EPI sequence (repetition time = 2 sec, echo time = 60 msec, flip angle = 80°, field of view = 240/240 mm, matrix = 64 × 64, ascending slice acquisition, slice thickness = 5 mm, interslice gap = 1 mm, in-plane resolution = 3.75 × 3.75 mm) using a standard quadrature head coil. Anatomical whole-head images were acquired from 124 axial slices (1.4-mm thick) using a spoiled gradient-echo recalled acquisition sequence (field of view = 240 × 180 mm, TE/TR = 6.0 msec/33.0 msec, flip angle = 40°, acquisition matrix = 256 × 192, in-plane resolution = 0.9375 × 0.9375 mm).

Preprocessing and statistical analysis were performed with the BrainVoyager2000/QX software package (www.brainvoyager.com). The first four volumes of each run were discarded to allow for signal equilibration. After motion and slice scan time correction, temporal filtering (0.01 Hz highpass) and linear trend removal, the functional data were aligned with the anatomical reference from the same session, transformed into Talairach space (Talairach & Tournoux, 1988), spatially smoothed with a Gaussian kernel (FWHM = 8 mm) and *z*-standardized for each run.

To delineate brain areas that are involved in controlling the retrieval of the attributes, we tried to isolate voxels that exhibit uniform signal increases with an increasing number of to-be-retrieved attributes. To this end, we performed voxelwise multisubject multiple regression analysis using a general linear model with separate predictors for trials in which either zero (control condition), one, two, or three/four attributes had to be retrieved. To account for the fact that RTs varied substantially across trials, the regressor functions were adjusted to the specific RTs associated with each trial by convolving the model hemodynamic response function (HRF) with a boxcar function with length = RT (“RT-convolved HRF analysis”; see Christoff et al., 2001). With these experimental regressors, a parametric contrast ($0 < 1 < 2 < 3/4$ attributes) was computed with contrast coefficients $-3, -1, 1, 3$. To get only positive parametric signal modulations, we added a contrast with coefficients 1, 1, 1, 1 and computed a conjunction analysis. This analysis, which was thresholded at $p < .005$ (random effects, uncorrected; revealing voxels

for which both individual contrasts surpass this significance level), has to be regarded as exploratory and served to find those voxels that responded most parametrically to the experimental manipulation. In a second step, we elevated the statistical threshold to $p < .001$. This value was chosen a posteriori, because here for both experiments a highly circumscribed cluster of voxels remained that resembled the locus of maximum parametric variation. In a further step, we extracted event-related time courses from these areas to evaluate the parametric increases with t tests for dependent samples. Therefore, we did not intend to adhere to a specific a priori defined significance level but just increased the threshold to detect the place of maximum parametric activation.

For the localizer scan, which served to isolate attribute-specific representation areas, t contrasts were computed with $p < .0001$, in which each attribute category was contrasted with the three other categories (with contrast coefficients 3, -1 , -1 , -1). The activated areas served as ROIs from which event-related BOLD signals were extracted and averaged for the different experimental conditions of the decision task. Differences between these averaged BOLD responses were evaluated by t tests for dependent samples to determine possible deviations from a systematic signal increase with the number of retrieved attributes.

RESULTS

Behavioral Data

The RTs and error rates in the decision task are shown in Figure 2, separately for Experiments 1 and 2. Incorrect responses and outliers ($RT > 2.5 \cdot SD$, computed separately for each participant and experimental condition) were excluded from the RT analysis (there were, on average, 3.05 [$SD = 1.43$] and 3.53 [$SD = 1.30$] outliers per participant in Experiments 1 and 2, respectively). Assuming that the attribute retrieval is sequential, TTB predicts that RTs increase as a function of the number of attributes to be retrieved (whereas a compensatory strategy would predict no such increase; cf. Bergert & Nosofsky, 2007; Bröder & Gaissmaier, 2007). As can be seen, an increase was clearly visible in both experiments, indicating that participants used the TTB strategy as instructed.

Repeated-measurements ANOVAs showed significant main effects of the number of to-be-retrieved attributes (excluding the control condition), $F(2, 32) = 133.98$, $p < .001$; ϵ (after Huynh & Feldt, 1976) = .624 for Experiment 1 and $F(2, 28) = 69.70$, $p < .001$; $\epsilon = .682$ for Experiment 2. Planned t tests for dependent samples showed that RTs were significantly higher for trials requiring the retrieval of two attributes, as compared with one attribute, as well as for trials requiring the retrieval of three or four, as compared with two, attributes. Furthermore, RTs were longer for trials requiring the retrieval of one attribute compared with trials in the control condition ($p < .001$ for all t tests).

Although participants were instructed to implement TTB at every decision trial, it is in principle possible that

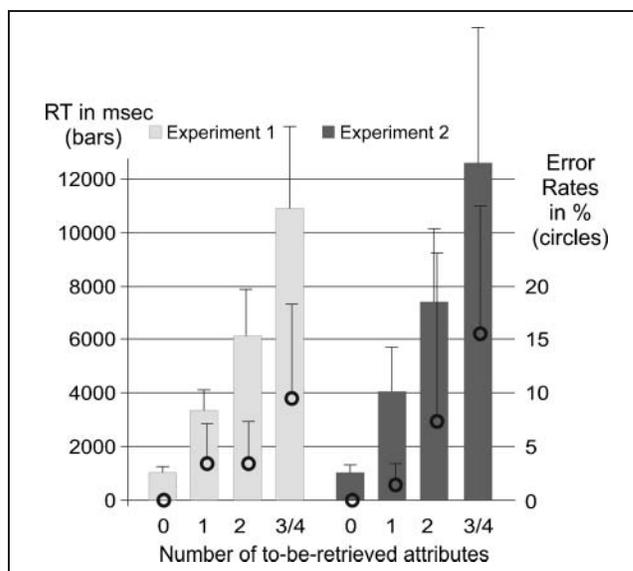


Figure 2. RTs and error rates (i.e., responses deviating from the correct responses according to the TTB heuristic) in the decision task, separately for Experiments 1 and 2 (the error bars are the standard deviations across participants). Consistent with the predictions of the TTB heuristic, RT increased as a function of the number of to-be-retrieved attributes. Moreover, error rates followed the same pattern, indicating that more incorrect responses were made, the more attributes had to be retrieved and compared.

they successively learned the relative success of each company and increasingly used this knowledge to make a decision. If this was the case, however, the magnitude of the uniform RT increase with increasing comparisons should diminish across the fMRI runs of the retrieval phase. Neither visual inspection of the data nor postexperimental debriefings provided any evidence for this possibility (in either experiment). In other words, we have no reason to assume that participants did not adhere to the TTB strategy throughout the decision task.

Figure 2 shows that the error rates were, on average, rather low. Visual inspection suggests a pattern that parallels the pattern obtained for the RTs. ANOVAs showed main effects of the number of to-be-retrieved attributes, $F(2, 28) = 7.27$, $p = .003$; $\epsilon = .929$ for Experiment 1 (on the basis of data from 15 participants, as two participants had to be excluded because of a coding error in the log-files) and $F(2, 28) = 9.97$, $p = .001$; $\epsilon = 1$ for Experiment 2. Planned t tests showed that error rates were lower for trials requiring the retrieval of one attribute, as compared with two attributes, in Experiment 2 and for trials requiring two attributes, as compared with three or four attributes, in Experiment 1 ($p < .01$ for both t tests). Furthermore, error rates were higher for trials requiring the retrieval of one attribute compared with trials in the control condition ($p < .01$ for the first and $p < .05$ in the second experiment). These results show that the greater the number of attributes that have to be retrieved and processed, the more errors occurred in the use of TTB.

fMRI Data

Prefrontal Control Processes: Material-unspecific Activation that Reflects the Number of To-be-retrieved Attributes

On the basis of models of memory control (Badre & Wagner, 2007; Buckner & Wheeler, 2001), we expected that control processes in the decision task would be reflected by parametric activation in the PFC that increases systematically with the number of retrieved attributes. As can be seen in the top part of Figure 3, the left dorsolateral PFC (DLPFC; middle frontal gyrus, BA 9/46) responded maximally to increasing processing demands (i.e., to the number of to-be-retrieved attributes). The close agreement of this maximum across both experiments (which had different attribute hierarchies; cf. Table 1) suggests that the activation of the PFC is independent of the specific type of attribute that has to be retrieved and processed. Another area (also the DLPFC) showing a material-unspecific uniform signal increase was found in the superior parietal cortex (see Figure 3). The bottom two panels of Figure 3

show the event-related hemodynamic signals (time-locked to the presentation of the pair of companies) extracted from the areas that showed the maximum activation in the two experiments. As can be seen, there was a uniform signal increase with the number of to-be-retrieved attributes in both experiments. In addition, it can be seen that in the control condition—in which no attribute information had to be retrieved—there was no response at all (in fact, Figure 3 suggests that there was even a small negative response in this condition). It might appear puzzling that the event-related signal for the control condition shows a strong response beginning at about 9 sec. Note, however, that this activation reflects the (average) response to the subsequent trial.

Is Activation in Material-specific Posterior Areas Modulated by the Sequential Retrieval Demands of TTB?

In a first step, we identified material-specific processing areas in the posterior cortex using the functional localizer

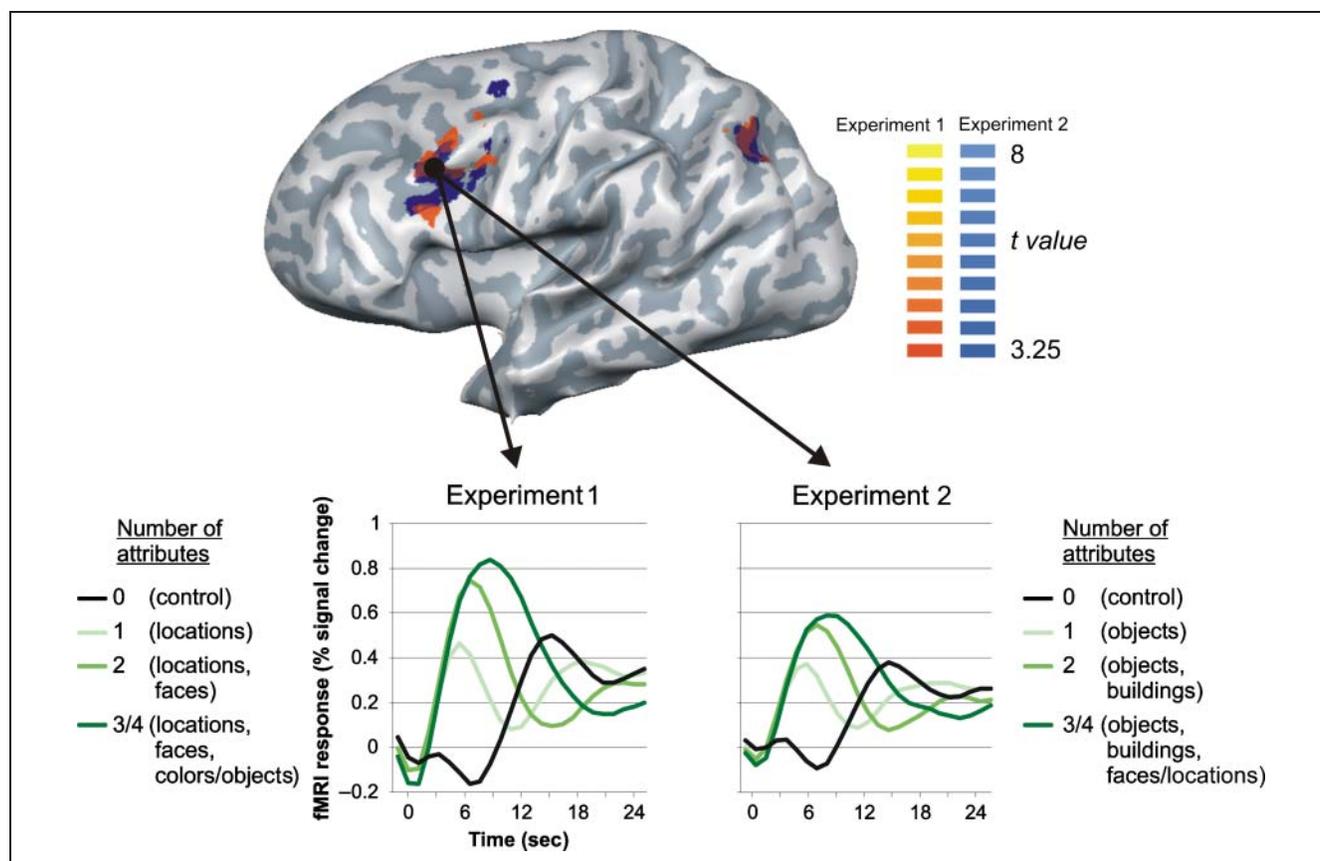


Figure 3. Brain areas showing a systematic activation increase with the number of to-be-retrieved attributes. The parametric contrasts of the number of to-be-retrieved attributes in both Experiments (thresholded at $p < .005$, RFX uncorrected) revealed the strongest parametric activation in the left DLPFC (middle frontal gyrus, BA 9), showing that this area reflects general processing demands that are independent of the specific attribute hierarchy, that is, the order in which the attributes have to be compared. Also the left DLPFC, the superior parietal cortex also exhibited this effect, albeit to a weaker extent. Statistical maps were projected onto a slightly inflated cortex reconstruction of one participant, on which concave curvature (i.e., sulci) appears in dark and convex curvature (i.e., gyri) in light gray. Below are shown the plots of event-related hemodynamic signals (averaged to the presentation of the pair of companies) from the area of maximum activation in each experiment. This area was derived by elevating the statistical threshold to $p < .001$, yielding closely matching locations across the two experiments located in the left DLPFC (cf. Table 1).

Table 1. Prefrontal Control Processes Involved in Memory-based Decision-making: Locations (Anatomical Label and Brodmann's Area), Peak *t* Values, and Numbers of Significantly Activated Voxels of Brain Areas that Reflect the Number of To-be-retrieved Attributes

<i>Region</i>	<i>BA</i>	<i>X</i>	<i>Y</i>	<i>Z</i>	<i>t</i>	<i>Voxels</i>
<i>Experiment 1</i>						
<i>p</i> < .005						
L middle frontal gyrus	6/9/46	-50	19	34	4.34	1984
L middle frontal gyrus	6	-43	6	39	5.02	1837
L superior frontal gyrus	6	-48	2	55	3.70	81
L inferior parietal lobe	7/19	-33	68	43	4.18	1055
<i>p</i> < .001						
L middle frontal gyrus	9	-48	22	35	4.34	85
<i>Experiment 2</i>						
<i>p</i> < .005						
L middle frontal gyrus	6/9/46	-47	17	34	4.92	3135
L middle frontal gyrus	6	-41	1	38	4.03	38
L middle frontal gyrus	6	-31	-1	55	4.11	325
L medial frontal gyrus	6	-8	11	52	3.97	346
L inferior parietal lobe	7/19	-33	-66	42	4.34	1532
<i>p</i> < .001						
L middle frontal gyrus	9	-47	17	36	4.92	598

The activation at the higher significance level was used to extract the event-related hemodynamic signal curves shown in Figure 3. L = left.

scans. These areas were defined as ROIs for the activation of material-specific stimuli representing the attributes that were required in the decision task. As can be seen in Figure 4 and Table 2, we were able to identify material-specific ROIs in Experiment 1 in the left middle occipital gyrus for locations, in the right fusiform gyrus for faces, and in the left lingual gyrus for objects and in Experiment 2 in the left superior-to-inferior parietal lobe for locations, in the right fusiform gyrus for faces, and in the left cuneus for objects. Moreover, buildings (replacing colors in Experiment 2) yielded a reliable ROI in the left parahippocampal cortex (part of the fusiform gyrus).

In a second step, we extracted from the ROIs the averaged hemodynamic responses during the decision task, separately for the different levels of the experimental factor “number of to-be-retrieved attributes.” As can be seen in Figures 5 and 6, there is a monotonically increasing response in all ROIs; that is, the maximum of the activation increases with the number of attributes required for a decision. At a first glance, the fact that such an increase is found for basically all ROIs seems to speak against a controlled activation of LTM representations. Instead, it sug-

gests that there is an automatic activation of all attributes associated with a company—even if an attribute is not relevant for the decision. However, we argue that such a conclusion is not warranted. For example, that the face-specific ROI is activated even when positions have to be retrieved could simply be due to the fact that the face area is, to a certain degree, generally responsive to positions. As Figure 4 shows, a coding overlap was apparent in all ROIs. It seems physiologically implausible to assume that areas respond in an all-or-none fashion only to one specific kind of visual stimulus, but not to others. Therefore, we refrain from inferring that there is an automatic retrieval process on the basis of the observation of a uniform increase across all ROIs. Instead, we ask, is there evidence for systematic deviations from a uniform signal increase? Such deviations would indicate the existence of a controlled modulation of representations, even if part of the activation is automatic.

A closer look at the BOLD signals extracted from the ROIs indeed reveals such systematic deviations. As argued below, these deviations reflect the importance of the attributes for decision-making. In the following, we focus on those experimental conditions in which more than one attribute has to be retrieved; we do not consider the condition where only one attribute has to be retrieved, because it is possible that, in comparison with the control condition, activation in this condition may just indicate the retrieval of visual stimuli in general, rather than the retrieval of a specific attribute. If TTB leads to a controlled retrieval, there should be a selective boosting of activation specifically in those ROIs that represent the attributes that are relevant for a comparison.

Let us first turn to the results of Experiment 1. The top left plot of Figure 5 shows the responses of the face-specific ROI (which represents the second-most important attribute), which has to be retrieved only when locations (the most important attribute) do not discriminate between the companies. As can be seen, there is some activation in the face-specific ROI even when TTB requires only the retrieval of locations (light green line; the maximum of this activation is indicated by a blue arrow). However, when faces do become relevant for a decision (i.e., when two attributes—locations and faces—have to be retrieved because locations do not discriminate; medium green line), there is an additional increase in activation in the face-specific ROI, reflecting the retrieval of this attribute (indicated by another blue arrow). Importantly, however, when three or four attributes have to be retrieved (i.e., colors, and colors and objects, respectively; dark green line), the additional increase in activation in the face-specific ROI is rather small. In other words, there is a strong activation increase in the face-specific ROI only in those trials in which faces are relevant for the decision.

To substantiate this visual impression statistically, we extracted the maximum amplitudes of the BOLD signals in the face-specific ROI for each participant, separately for the different experimental conditions. Consistent

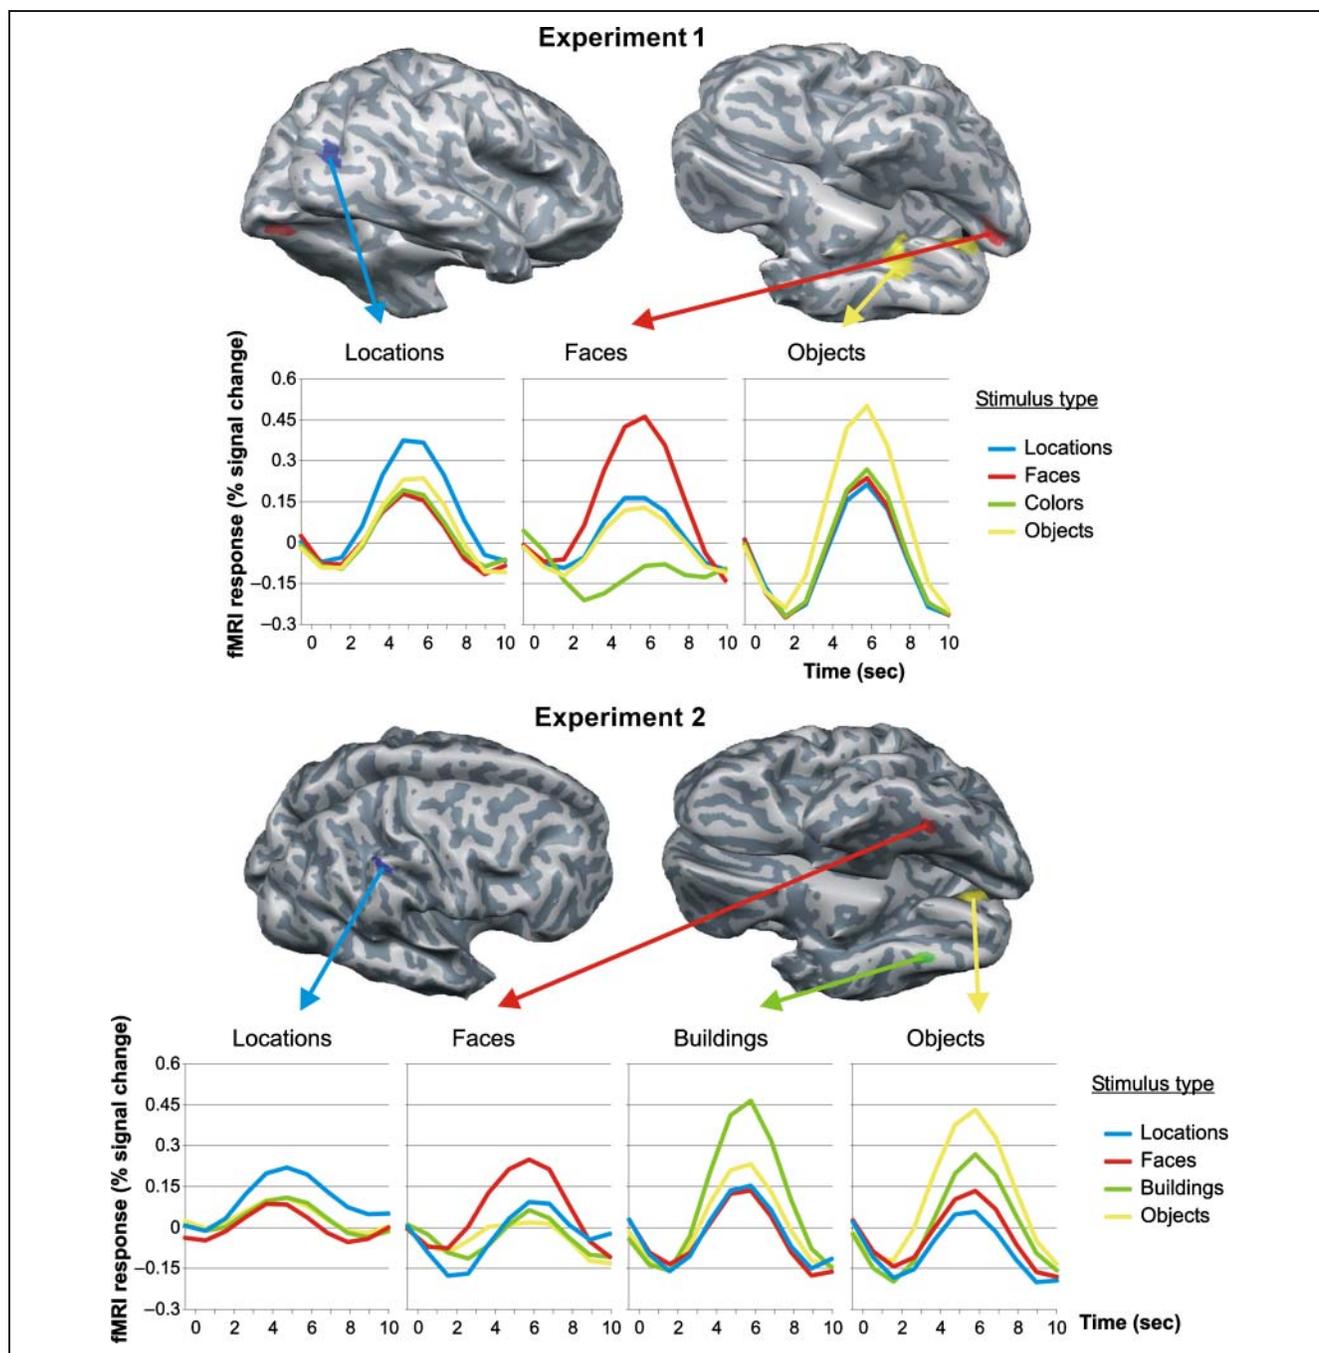


Figure 4. Material-specific ROIs in the posterior cortex identified by the functional localizer scans. On the basis of the assumption (detailed in the overview of the experimental design) that stimulus representations are located in the same cortical areas that are also involved in the perceptual processing of the information, these areas were defined as target areas for the representation of the attributes.

with the visual impression, dependent-sample t tests indicated a significant increase in activation when two attributes, as compared with one attribute, have to be retrieved, $t(16) = 6.89$; $p < .001$. However, when three or four attributes, as compared with two, have to be retrieved, the increase was rather small, $t(16) = 1.97$; $p = .067$. Another t test confirmed that the increase from one to two attributes was significantly larger than the increase from two to three or four attributes, $t(16) = 4.63$; $p < .001$.

To validate this finding we next looked for a corresponding pattern in Experiment 2. Importantly, note that in Experiment 2 the buildings rather than the faces were the second-most important attribute (the different hierarchies of attribute importance are depicted in the center panels of Figures 5 and 6). In close correspondence with the findings for Experiment 1, the top right plot in Figure 5 shows that there is a pronounced signal increase in the building-specific ROI in those trials in which two attributes (objects and buildings) have to be retrieved,

Table 2. Material-specific Processing Areas in the Posterior Cortex: Locations, Peak t Values, and Numbers of Significantly Activated Voxels ($p < .0001$, Uncorrected) of Material-specific Posterior Areas Identified by the Functional Localizer Scans

Region	BA	X	Y	Z	t	Voxels
<i>Experiment 1</i>						
Locations: L middle occipital gyrus	18/19	-28	-79	10	7.26	1408
Faces: R fusiform gyrus	19	22	-76	-8	10.84	2901
Objects: L lingual gyrus	18/19	-19	-54	0	10.44	3951
<i>Experiment 2</i>						
Locations: L superior/inferior parietal lobe	7/40	-37	-45	40	7.02	823
Faces: R fusiform gyrus	37	39	-43	-16	7.93	116
Objects: L cuneus	18	-3	-77	14	10.27	6846
Buildings: L fusiform gyrus	37	-28	-58	-10	11.53	4472

L = left; R = right.

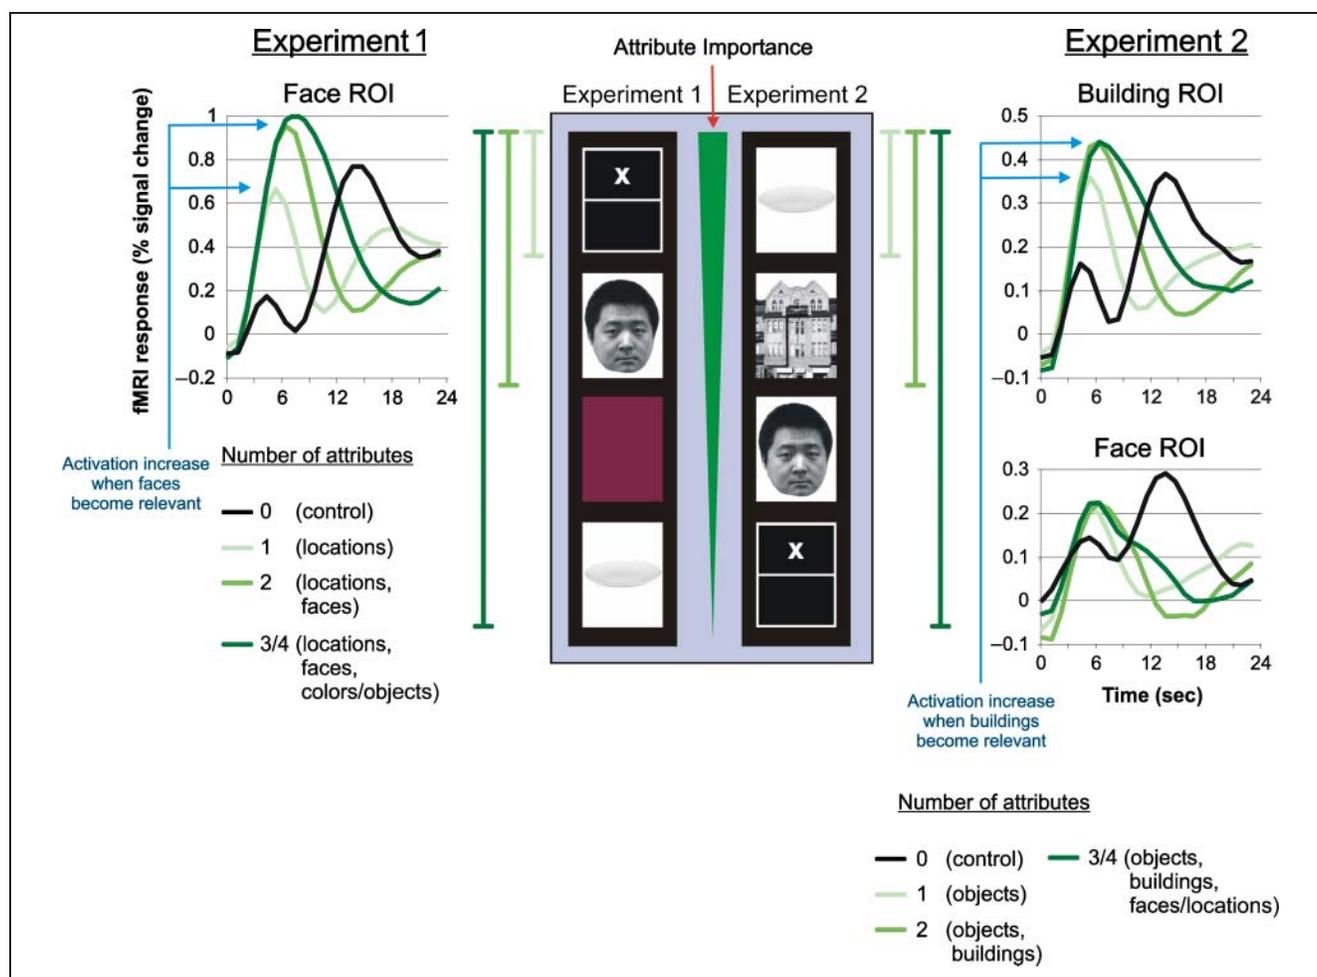


Figure 5. Selective modulation of systematic signal increases in attribute-specific regions. The areas representing the second most important attribute show a substantial increase of the maximum hemodynamic response when not only the most but also the second-most important attributes have to be retrieved. Importantly, however, when three or four attributes have to be retrieved, the additional increase in activation is rather small. In other words, there is a strong activation increase only in those trials in which the respective attribute becomes relevant for the decision.

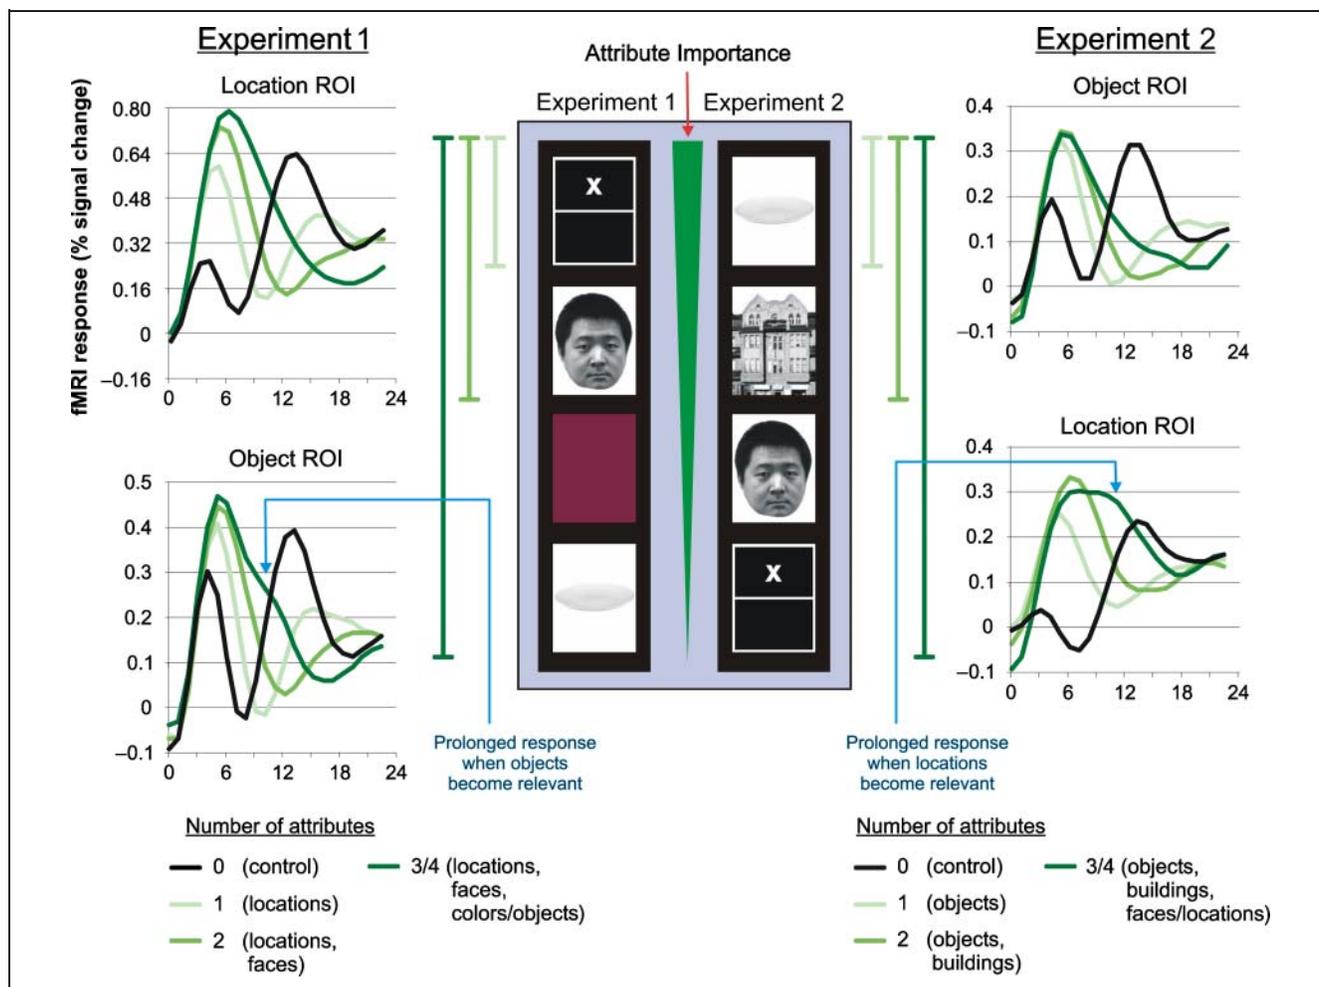


Figure 6. Selective modulation of systematic signal increases in attribute-specific regions. Only for those trials in which up to four attributes have to be retrieved is the descending flank of the hemodynamic response in the areas that represent the least important attribute (i.e., objects in Experiment 1 and positions in Experiment 2) prolonged, rather than descending smoothly to baseline activity after having reached its maximum—as one would expect with a strictly uniform signal modulation (see bottom left and right maps). In other words, attributes that are relevant relatively late during the decision process (i.e., those with low importance) generally show a prolonged response during the descending part of the hemodynamic response. When the attributes are high in importance, by contrast, this is not the case (see top left and right maps).

but no further increase when also faces or faces and locations have to be retrieved (again indicated by two blue arrows). As in Experiment 1, this visual impression was substantiated statistically: In the building-specific ROI, there was a significant increase when two attributes, as compared with one attribute, have to be retrieved, $t(14) = 2.74$; $p = .016$, but not when three or four, rather than two, attributes have to be retrieved, $t(14) = 0.31$; $p = .763$. Furthermore, the increase from one to two attributes was larger than the increase from two to three or four attributes, $t(14) = 1.90$; $p = .078$.

To summarize, areas representing the attributes with the *second-highest* importance (faces in Experiment 1 and buildings in Experiment 2) show an increase of the maximum BOLD response when the *second-most important* attribute becomes relevant as compared with when already the most important attribute leads to a decision. However, when three or four attributes have to be retrieved, the additional increase in the areas representing

the second-most important attribute is much smaller or even completely absent. Importantly, in Experiment 2, in which the faces are no longer the second-most important attribute (Figure 5, bottom right plot), the face-specific ROI does *not* show this effect. This selective modulation of neural activation that follows the retrieval order according to TTB supports the notion of controlled retrieval processes. Overall, the maximum of the BOLD response in attribute-specific ROIs seems to be boosted only when the respective attribute is relevant for a decision.

If the second-most important attribute affects the maximum of the BOLD response, then the third-most and least important attributes should affect signal parts after the maximum. Indeed, a closer look at the bottom right plot in Figure 5 shows an interesting pattern in the time courses of the hemodynamic response in the face-specific ROI in Experiment 2: Only for those trials in which three or four attributes have to be retrieved does the descending flank of the hemodynamic response seem to be prolonged

rather than descending smoothly to baseline activity after having reached its maximum—as one would expect with a strictly uniform signal modulation. In other words, attributes that are relevant relatively late during the decision process (i.e., those low in importance) generally show a prolonged response during the descending part of the hemodynamic response. To illustrate, consider the location and object attributes, whose ranks in the attribute hierarchy are exactly reversed between Experiments 1 and 2 (locations are the most important attribute in Experiment 1—and thus relevant in each trial—but the least important attribute in Experiment 2—and thus relevant only when all four attributes have to be retrieved; for the object attribute, the opposite holds). As can be seen in Figure 6 (indicated by blue arrows in the bottom left and right maps), in both the location-specific and object-specific ROIs, there is a prolonged descending part of the hemodynamic response when the attributes are low in importance. When the attributes are high in importance, by contrast, this is not the case (see top left and right maps in Figure 6).

To corroborate this observation of prolonged responses statistically, we required a measure that could capture deviations from a smooth return to baseline activity. Because any such deviation necessarily produces an additional inflection point that occurs *earlier* than would be expected under a smoothly descending signal, we focused on the time of the first inflection point after the signal maximum for the different curves of the posterior ROIs. As a control analysis, we analyzed the signals in the PFC (see Figure 3), in which such prolonged responses should be absent; as a consequence, in the PFC the inflection points should only depend on the maximum amplitude of the signal and, thus, occur later the more attributes have to be retrieved.

As can be seen in Figure 3, there is a uniform signal increase in the PFC, and the first inflection point is successively delayed in time as the number of to-be-retrieved attributes increases. More specifically, for Experiment 1 (left graph in Figure 3), the inflection points occurred, on average, at 4.13 sec in the control condition (i.e., when no attribute has to be retrieved), at 7.87 sec when one attribute has to be retrieved, at 8.93 sec when two attributes have to be retrieved, and at 9.73 sec when three or four attributes have to be retrieved. For Experiment 2 (right graph in Figure 3), the inflection points occurred at 5.40, 8.00, 9.77, and 11.41 sec, respectively. *t* tests, for which the inflection points in each condition were extracted for each participant, showed that in Experiment 1 the inflection point occurred significantly later when one, rather than no, attribute has to be retrieved, $t(16) = 5.83; p < .001$. Moreover, the inflection point occurred significantly later when two attributes, rather than one, have to be retrieved, $t(16) = 5.60; p < .001$, and when three or four attributes, rather than two, have to be retrieved, $t(16) = 2.67; p = .017$. Similarly, in Experiment 2, the inflection point occurred significantly later when one, rather

than no, attribute, $t(14) = 6.42; p < .001$, and when two attributes, rather than one, have to be retrieved, $t(14) = 2.87; p = .012$. The increase from two to more than two attributes was not significant.

In contrast to these uniform signal increases in the PFC, the posterior ROIs show a rather different pattern. Specifically, when locations and objects are the least important attributes (i.e., Experiments 2 and 1, respectively), the inflection points of the hemodynamic responses in the corresponding ROIs (bottom left and right maps in Figure 6 for Experiments 2 and 1, respectively) occur *earlier* (rather than later) than when only two attributes have to be retrieved: 7.41 versus 8.06 sec, $t(16) = 2.02, p = .060$ (Experiment 1, object ROI) and 8.93 versus 10.13 sec, $t(14) = 2.28, p = .039$ (Experiment 2, location ROI). By contrast, when locations and objects are the most important attributes, the inflection point in the corresponding brain areas occurs significantly later when three or four, rather than two, attributes have to be retrieved, 8.71 versus 10.29 sec, $t(16) = 7.85; p < .001$ (Experiment 1, location-specific ROI; cf. top left graph in Figure 6) and 8.73 versus 9.27 sec, $t(14) = 3.74; p = .002$ (Experiment 2, object-specific ROI; cf. top right graph in Figure 6). This is consistent with what one would expect from a uniform signal increase, as observed in the PFC.

To conclude, although in the posterior areas there were increasing activation patterns across all ROIs—which might suggest that using TTB is associated with an automatic spreading activation—we found evidence for systematic modulations of this activation that reflect the retrieval order mandated by the attribute hierarchy. These systematic modulations are consistent with the hypothesis of a controlled retrieval of decision-relevant attribute information.

DISCUSSION

In two fMRI experiments, we examined the neural dynamics of heuristic decision-making when information about alternatives has to be retrieved from LTM. Behavioral studies have shown that the TTB heuristic is a key strategy that people often use to make memory-based decisions (e.g., Bröder & Schiffer, 2003, 2006). According to TTB, information retrieval, is controlled and sequential guided by attribute importance. Crucially, information search is terminated as soon as an attribute discriminates between the alternatives; other attributes associated with the alternatives that are not required for making a decision are not activated. Here, we tested TTB's assumption of controlled and limited retrieval by monitoring the activation of specific LTM representations with fMRI.

Participants first learned attribute information (represented by visual stimuli) about companies and then made decisions on the basis of this information using TTB. The attributes consisted of visual stimuli (e.g., faces, objects) that are known to be processed and represented in different parts of the posterior cortex. By means of functional

localizer scans, we identified brain areas involved in processing and representing these specific types of attributes (cf. Table 2). The delineated material-specific brain areas are all fully consistent with findings of previous studies. Specifically, faces are assumed to be represented in the fusiform gyrus (Khader et al., 2005; Ishai, Ungerleider, Martin, et al., 2000), locations in the parietal cortex (Khader et al., 2005, 2007; Moscovitch et al., 1995), objects in the visual association cortex (Khader et al., 2007; Ishai, Ungerleider, Martin, et al., 2000; Moscovitch et al., 1995; Mishkin et al., 1983), and buildings in the parahippocampal cortex (Ishai, Ungerleider, Martin, et al., 2000; O'Craven & Kanwisher, 2000).

The identification of material-specific representation areas allowed us to investigate whether the activation in these areas during decision-making is modulated by the retrieval demands posed by TTB. Specifically, we used the degree to which there was a deviation from a uniform signal increase as an indicator for the controlled activation of material-specific LTM representations. On the one hand, the signal-increase patterns for the different decision trials were rather similar across all ROIs. This might suggest an automatic activation of all attribute information associated with the alternatives. However, we also observed systematic deviations from a uniform signal increase, as the increase was modulated by the necessity (according to TTB) to inspect the respective attribute for the decision: First, the maximum of the BOLD responses in attribute-specific ROIs increased only when the respective attribute was relevant for the decision. Second, the inflection points of the descending flank of the hemodynamic response, reflecting a prolonged duration of the underlying cognitive processes, were also found to be modulated by attribute relevance. Specifically, we found a prolonged response to an attribute only when it was relevant late in the decision process (i.e., when the attribute was low in importance). Overall, we observed systematic modulations of brain activation that reflect the retrieval order mandated by the attribute hierarchy. These systematic modulations are consistent with the notion of a controlled activation of decision-relevant attribute representations. As outlined in the Results section, the observation of a uniform signal increase across ROIs does not necessarily indicate the existence of automatic retrieval processes. The reason is that activation in an ROI might not be exclusively because of the retrieval of the attribute to which the ROI is most sensitive (in fact, Figure 4 shows that none of the ROIs is completely inactive when the other attributes are retrieved). Therefore, we refrain from inferring automatic retrieval process from the observed patterns. One possibility could be that all attributes become initially activated by means of spreading activation (Anderson, 1983), followed by a subsequent controlled activation of the decision-relevant associations via memory-based attentional selection processes. What exactly causes the selective boosting of decision-relevant memory representations, however, is subject to further research.

We found that the left DLPFC responded systematically to the number of attributes needed for a decision, irrespective of the type of the attribute that had to be retrieved. This finding shows that the DLPFC is important during memory-based decision-making when specific visual information has to be retrieved and subsequently processed. But what function does the DLPFC have during the decision process? The DLPFC is often associated with working memory (WM) load (e.g., Linden et al., 2003). WM load, however, is an unlikely factor underlying the observed parametric effects: As, according to the TTB heuristic, only one pair of attribute values needs to be maintained in WM at any time, an increased number of required attributes does not lead to an increased WM load. A more likely candidate is, by contrast, WM updating: TTB requires the sequential comparison of attributes according to the attribute hierarchy, requiring a constant updating of the content in WM. The larger the number of required comparisons, the higher the updating demands. This presumably leads to successive DLPFC activations that add up to the parametric increase observed in the present study. The assumption that the DLPFC mediates WM updating processes is consistent with studies showing that the DLPFC plays a critical role for WM (e.g., Postle, 2005; Curtis & D'Esposito, 2003; Levy & Goldman-Rakic, 2000). According to Funahashi (2001), the DLPFC mediates selection and monitoring in WM tasks. Similarly, Achim and Lepage (2005) found evidence that the DLPFC is involved in post-retrieval monitoring processes, such as tracking the relevance of the retrieved information with respect to the task at hand.

Our experiments also show that the DLPFC activity is material-unspecific and thus seems to reflect general control demands during the decision task that increase with the number of to-be-retrieved attributes. Furthermore, the additional activation of the parietal cortex (see Figure 3) is consistent with the notion of a "fronto-parietal network" (Corbetta & Shulman, 2002) that is assumed to regulate content-unspecific higher cognitive functions, such as mental imagery, WM control, or LTM retrieval (e.g., Mechelli, Price, Friston, & Ishai, 2004; Sohn, Goode, Stenger, Carter, & Anderson, 2003; Ishai, Ungerleider, & Haxby, 2000). The present experiments show that the fronto-parietal network is recruited during decision-making when attribute information has to be retrieved from LTM. This finding is consistent with results by Kahn, Davachi, and Wagner (2004), who showed that the left DLPFC and parietal cortex mediate control processes of LTM retrieval (such as retrieval attempts, maintenance of retrieval cues, monitoring of retrieved information; see also Iidaka, Matsumoto, Nogawa, Yamamoto, & Sadato, 2006; see Wagner, Shannon, Kahn, & Buckner, 2005, for a review).

Numerous studies have emphasized the role of the ventrolateral PFC (encompassing BA 45 and surrounding areas of the inferior frontal gyrus) for the controlled retrieval of memory representations (for a review, see

Badre & Wagner, 2007; but see, e.g., Wheeler & Buckner, 2003; Henson, Rugg, Shallice, & Dolan, 2000; Henson, Shallice, & Dolan, 1999, for evidence for a role of the DLPFC for this function). One may, therefore, wonder why the memory-based decision task used in our experiments did not evoke ventrolateral PFC activation. One possible reason is that our task involved visual information, whereas the studies reviewed by Badre and Wagner (2007) mainly employed verbal material (semantic representations of conceptual knowledge). Potentially, different kinds of representation draw on control processes in different subsections of the PFC—a possibility that should be tested in detail by future studies.

To conclude, we proposed an experimental framework for investigating retrieval processes during memory-based decision-making that builds on monitoring the activation of specific LTM representations with fMRI. Our findings suggest that memory-based decision-making involves the controlled retrieval of stored attribute information, which is implemented on a neural level as a selective modulation of the activation of content-specific representational areas in the posterior cortex.

The present experiments should be regarded as a starting point. Future studies will need to delineate more clearly the neural dynamics of memory-based decision-making, such as possible interactions of the DLPFC and posterior areas in which attributes are represented. Moreover, it would be interesting to compare the present results with brain activations when people use a compensatory strategy (e.g., a weighted-additive strategy), in which all attributes are considered for making a decision. In comparison with TTb, such a compensatory strategy additionally includes integration and, potentially, weighting processes and involves alternative-wise rather than attribute-wise (as with TTb) processing of the attributes. The identification of the specific neural signatures of various decision strategies may then help identify the decision processes in people's spontaneous (in contrast to instructed) strategy use. For instance, neural activation data could be used in combination with further indicators of the cognitive process (overt decisions, confidence judgments, RTs) and contribute to a multi-method approach to identify strategies in decision-making.

Acknowledgments

This work was supported by grant KH235/1-1 of the German Research Foundation (DFG) assigned to P. K. and F. R. and by the University of Basel. We thank Katharina Dobs for programming the stimulus presentation and Lilian Weber for help in data acquisition.

Reprint requests should be sent to Patrick H. Khader, Department of Psychology, Gutenbergstrasse 18, 35037 Marburg, Germany, or via e-mail: Khader@uni-marburg.de, Web: www.uni-marburg.de/fb04/team-schuboe/people/khader-en.

REFERENCES

Achim, A. M., & Lepage, M. (2005). Dorsolateral prefrontal cortex involvement in memory post-retrieval monitoring

- revealed in both item and associative recognition tests. *Neuroimage*, *24*, 1113–1121.
- Anderson, J. R. (1983). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior*, *22*, 261–295.
- Badre, D., & Wagner, A. D. (2007). Left ventrolateral prefrontal cortex and the cognitive control of memory. *Neuropsychologia*, *45*, 2883–2901.
- Beach, L. R., & Mitchell, T. R. (1987). Image theory: Principles, goals, and plans in decision making. *Acta Psychologica*, *66*, 201–220.
- Bergert, F. B., & Nosofsky, R. M. (2007). A response-time approach to comparing generalized rational and take-the-best models of decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*, 107–129.
- Bröder, A., & Gaissmaier, W. (2007). Sequential processing of cues in memory-based multiattribute decisions. *Psychonomic Bulletin and Review*, *14*, 895–900.
- Bröder, A., & Schiffer, S. (2003). Take the best versus simultaneous feature matching: Probabilistic inferences from memory and effects of representation format. *Journal of Experimental Psychology: General*, *135*, 207–236.
- Bröder, A., & Schiffer, S. (2006). Stimulus format and working memory in fast and frugal strategy selection. *Journal of Behavioral Decision Making*, *19*, 361–380.
- Buckner, R. L., & Wheeler, M. E. (2001). The cognitive neuroscience of remembering. *Nature Reviews Neuroscience*, *2*, 624–634.
- Chao, L. L., & Martin, A. (1999). Cortical representation of perception, naming, and knowledge of color. *Journal of Cognitive Neuroscience*, *11*, 25–35.
- Christoff, K., Prabhakaran, V., Dorfman, J., Zhao, Z., Kroger, J. K., Holyoak, K. J., et al. (2001). Rostrolateral prefrontal cortex involvement in relational integration during reasoning. *Neuroimage*, *14*, 1136–1149.
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*, 201–215.
- Curtis, C. E., & D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences*, *7*, 415–423.
- Danker, J. F., & Anderson, J. R. (2010). The ghosts of brain states past: Remembering reactivates the brain regions engaged during encoding. *Psychological Bulletin*, *136*, 87–102.
- Funahashi, S. (2001). Neuronal mechanisms of executive control by the prefrontal cortex. *Neuroscience Research*, *39*, 147–165.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650–669.
- Gigerenzer, G., Hedwig, R., & Pachur, T. (Eds.). (2011). *Heuristics. The foundations of adaptive behavior*. New York: Oxford University Press.
- Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart*. Oxford, England: Oxford University Press.
- Glöckner, A., & Betsch, T. (2008). Modeling option and strategy choices with connectionist networks: Towards an integrative model of automatic and deliberate decision making. *Judgment and Decision Making*, *3*, 215–228.
- Henson, R. N., Rugg, M. D., Shallice, T., & Dolan, R. J. (2000). Confidence in recognition memory for words: Dissociating right prefrontal roles in episodic retrieval. *Journal of Cognitive Neuroscience*, *12*, 913–923.
- Henson, R. N., Shallice, T., & Dolan, R. J. (1999). Right prefrontal cortex and episodic memory retrieval: A functional MRI test of the monitoring hypothesis. *Brain*, *122*, 1367–1381.

- Huynh, H., & Feldt, L. S. (1976). Estimation of the box correction for degrees of freedom from sample data in randomized block and split-plot designs. *Journal of Educational Statistics*, *1*, 69–82.
- Iidaka, T., Matsumoto, A., Nogawa, J., Yamamoto, Y., & Sadato, N. (2006). Frontoparietal network involved in successful retrieval from episodic memory. Spatial and temporal analyses using fMRI and ERP. *Cerebral Cortex*, *16*, 1349–1360.
- Ishai, A., Ungerleider, L. G., & Haxby, J. V. (2000). Distributed neural systems for the generation of visual images. *Neuron*, *28*, 979–990.
- Ishai, A., Ungerleider, L. G., Martin, A., & Haxby, J. V. (2000). The representation of objects in the human occipital and temporal cortex. *Journal of Cognitive Neuroscience*, *12*, 35–51.
- Juslin, P., & Persson, M. (2002). PROBABILITIES from EXemplars (PROBEX): A “lazy” algorithm for probabilistic inference from generic knowledge. *Cognitive Science*, *26*, 563–607.
- Kahn, I., Davachi, L., & Wagner, A. D. (2004). Functional-neuroanatomic correlates of recollection: Implications for models of recognition memory. *Journal of Neuroscience*, *24*, 4172–4180.
- Khader, P., Burke, M., Bien, S., Ranganath, C., & Rösler, F. (2005). Content-specific activation during associative long-term memory retrieval. *Neuroimage*, *27*, 805–816.
- Khader, P., Knoth, K., Burke, M., Bien, S., Ranganath, C., & Rösler, F. (2007). Topography and dynamics of associative long-term memory retrieval in humans. *Journal of Cognitive Neuroscience*, *19*, 492–511.
- Khader, P., & Rösler, F. (2009). Where memories are stored discloses how memories are formed: Content specificity of long-term memory representations. In F. Rösler, C. Ranganath, B. Röder, & R. H. Kluwe (Eds.), *Neuroimaging and psychological theories of human memory*. Oxford: Oxford University Press.
- Levy, R., & Goldman-Rakic, P. S. (2000). Segregation of working memory functions within the dorsolateral prefrontal cortex. *Experimental Brain Research*, *133*, 23–32.
- Linden, D. E. J., Bittner, R. A., Muckli, L., Waltz, J. A., Kriegeskorte, N., Goebel, R., et al. (2003). Cortical capacity constraints for visual working memory: Dissociation of fMRI load effects in a fronto-parietal network. *Neuroimage*, *20*, 1518–1530.
- Martin, A. (2007). The representation of object concepts in the brain. *Annual Review of Psychology*, *58*, 25–45.
- Martin, A., Haxby, J. V., Lalonde, F. M., Wiggs, C. L., & Ungerleider, L. G. (1995). Discrete cortical regions associated with knowledge of color and knowledge of action. *Science*, *270*, 102–105.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*, 419–457.
- Mechelli, A., Price, C. J., Friston, K. J., & Ishai, A. (2004). Where bottom-up meets top-down: Neuronal interactions during perception and imagery. *Cerebral Cortex*, *14*, 1256–1265.
- Mishkin, M., Ungerleider, L. G., & Macko, K. A. (1983). Object vision and spatial vision: Two cortical pathways. *Trends in Neurosciences*, *6*, 414–417.
- Moscovitch, C., Kapur, S., Köhler, S., & Houle, S. (1995). Distinct neural correlates of visual long-term memory for spatial location and object identity: A positron emission tomography study in humans. *Proceedings of the National Academy of Sciences, U.S.A.*, *92*, 3721–3725.
- Newell, B. R., & Shanks, D. R. (2004). On the role of recognition in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 923–935.
- O'Craven, K. M., & Kanwisher, N. (2000). Mental imagery of faces and places activates corresponding stimulus-specific brain regions. *Journal of Cognitive Neuroscience*, *12*, 1013–1023.
- O'Reilly, R. C., & Rudy, J. W. (2001). Conjunctive representations in learning and memory: Principles of cortical and hippocampal function. *Psychological Review*, *108*, 311–345.
- Postle, B. R. (2005). Delay-period activity in the prefrontal cortex: One function is sensory gating. *Journal of Cognitive Neuroscience*, *17*, 1679–1690.
- Rastle, K., Harrington, J., & Coltheart, M. (2002). 358,534 nonwords: The ARC nonword database. *Quarterly Journal of Experimental Psychology A*, *55*, 1339–1362.
- Sohn, M. H., Goode, A., Stenger, V. A., Carter, C. S., & Anderson, J. R. (2003). Competition and representation during memory retrieval: Roles of the prefrontal cortex and the posterior parietal cortex. *Proceedings of the National Academy of Sciences, U.S.A.*, *100*, 7412–7417.
- Squire, L. R., & Alvarez, P. (1995). Retrograde amnesia and memory consolidation: A neurobiological perspective. *Current Opinion in Neurobiology*, *5*, 169–177.
- Talairach, J., & Tournoux, P. (1988). *Co-planar stereotaxic atlas of the human brain*. Stuttgart: Thieme.
- Wagner, A. D., Shannon, B. J., Kahn, I., & Buckner, R. L. (2005). Parietal lobe contributions to episodic memory retrieval. *Trends in Cognitive Sciences*, *9*, 445–453.
- Wheeler, M. E., & Buckner, R. L. (2003). Functional dissociation among components of remembering: Control, perceived oldness, and content. *Journal of Neuroscience*, *23*, 3869–3880.
- Zeki, S., Watson, J. D., Lueck, C. J., Friston, K. J., Kennard, C., & Frackowiak, R. S. (1991). A direct demonstration of functional specialization in human visual cortex. *Journal of Neuroscience*, *11*, 641–649.