Anterior cingulate engagement in a foraging context reflects choice difficulty, not foraging value

Amitai Shenhav[1,2], Mark A Straccia[1,2], Jonathan D Cohen[1,2] & Matthew M Botvinick[1,2]

Previous theories predict that human dorsal anterior cingulate (dACC) should respond to decision difficulty. An alternative theory has been recently advanced that proposes that dACC evolved to represent the value of 'non-default', foraging behavior, calling into question its role in choice difficulty. However, this new theory does not take into account that choosing whether or not to pursue foraging-like behavior can also be more difficult than simply resorting to a default. The results of two neuroimaging experiments show that dACC is only associated with foraging value when foraging value is confounded with choice difficulty; when the two are dissociated, dACC engagement is only explained by choice difficulty, and not the value of foraging. In addition to refuting this new theory, our studies help to formalize a fundamental connection between choice difficulty and foraging-like decisions, while also prescribing a solution for a common pitfall in studies of reward-based decision making.

The dACC stands out as one of the most extensively studied regions of the brain, and yet its basic functions are still a matter of intensive debate[1–7]. Historically, functions attributed to this region have included the encoding of pain[5], surprise[3,8], value[9], and level of cognitive demand[5,10,11], including the difficulty posed by conflict between competing choices[12–16].

A recent study by Kolling, Behrens, Mars and Rushworth[17] (KBMR) presented a significant and intriguing detour from previous approaches to understanding the function of this region. The authors echo previous proposals that the dACC is responsive to the value of choice options. However, they propose that this is restricted to the value of diverging from one's default behavior in a given context. They suggest that this function has its evolutionary roots in the encoding of the overall value of foraging for food in a new patch rather than continuing to engage the current food patch. To test this, they designed a task to model foraging decisions. KBMR showed that, in this task, dACC activity was positively associated with the value of foraging for better rewards and negatively associated with the value of engaging currently available ones (the default behavior). Their conclusion, that dACC is involved in foraging decisions *per se*, has already had a significant impact on perspectives regarding the function of this region[1,18] and has generated several high-profile follow-up studies that reach similar conclusions[19–21].

In the present work, we challenge KBMR's interpretation of their findings and present strong evidence that dACC's role in foraging-like decisions is instead connected with decision difficulty. Although it has not been widely remarked, standard theories of foraging imply an intimate relationship between foraging and decision difficulty. There are two key quantities in the dynamics of foraging: the rate of intake within a patch and the value of foraging—that is, leaving for another patch (**Fig. 1**). According to optimal foraging theory[22–24], the best moment to leave a patch is when these two values coincide. If, like KBMR, we

view foraging choices as involving comparisons between two values or utilities, the optimal moment to forage is precisely the moment at which the value-based decision becomes most difficult: that is, the moment at which the values to be compared are most similar.

A more careful look at the pattern described above (**Fig. 1**) forces a reconsideration of KBMR's findings. Note in particular the relationship between foraging value and decision difficulty. These two are closely linked: as foraging value rises, so does decision difficulty. This correlation raises a serious concern in connection with the theory advanced by KBMR and related work. Specifically, it suggests that foraging value and decision difficulty might be confounded in the experiments that motivate the theory.

We show here, on the basis of two functional magnetic resonance imaging (fMRI) experiments, not only that these two factors have been confounded in previous studies, but that when they are adequately dissociated, dACC activity is found clearly to track choice difficulty rather than foraging value. We begin with a replication of the KBMR study, revealing a fundamental problem with the measure of choice difficulty used in that study, and introducing a more principled measure based on observed choice behavior. In a second experiment, we take advantage of features of the KBMR task that distinguish it from standard foraging tasks in order to deconfound choice difficulty from foraging value. When we do this, we obtain precisely the opposite result from the one they reported: dACC activity in foraging tasks tracks choice difficulty rather than foraging value. The set of results we report highlight (i) a fundamental point about the structure of foraging-like tasks, and the potential role of dACC in the performance of such tasks, and (ii) a key methodological consideration in studies of value-based choice.

RESULTS

In experiment 1 we scanned 15 human participants while they performed KBMR's original foraging task (**Fig. 2**). Each trial involved

Figure 1 The role of choice difficulty in a standard foraging setting. In a typical patch-leaving scenario, an animal faces the recurring decision of whether to continue harvesting a patch with decreasing marginal returns (for example, fewer ripe fruits; blue line) or leave for a new patch. The expected value of switching to a new patch (green line) accounts for the expected reward in the new patch, as well as the travel time between patches (gray region). The best moment to depart the current patch for a new patch is the point at which the value of switching exceeds the value of staying. Therefore, an optimally foraging animal will exit the patch at the point where blue and green lines meet (dashed horizontal black line), which is also the indifference point between these options. Dashed blue and green lines indicate theoretical values of staying and switching that an optimally foraging animal typically would not encounter (having already departed the patch) but that could theoretically be examined with a task like KBMR's that examines cross-sections through a foraging-like context (see **Fig. 2**). For simplicity, we assume here a situation where the value of each new patch, the (exponential) reward decay rate, and the travel time to a new patch remain constant across patches.
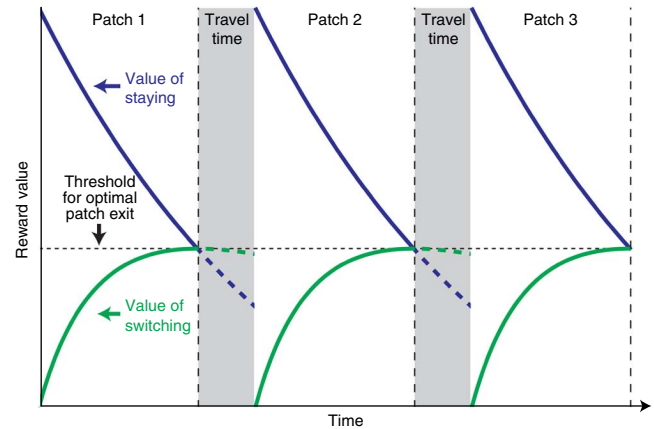
two stages of decision-making (**Fig. 2a**). In stage 1, the participant was offered a pair of potential rewards (the "engage" set) and a set of six alternative possible rewards (the "forage" set), all presented as abstract symbols designating numerical points that could be earned. The participant could choose to proceed directly to stage 2 (engage option) or first swap the current engage pair for a new pair randomly selected from the forage set as many times as they wished (forage option). Each swap generated a new forage set but incurred a cost (a designated number of lost points, as well as a time delay). Choosing to engage advanced the trial to stage 2, at which time a probability was paired with each of the two options in the engage set. The participant then chose between the reward-probability pairs on the left versus right side of the screen, and received the chosen reward based on the outcome of a random draw with the corresponding probability.

KBMR reported three key findings from this task. First, in stage 1 they found a strong bias to engage. Second, they found that dACC tracked the degree to which the points offered for each option favored the foraging option—that is, increased its relative value (RV_{forage}) as compared to the engage option. Third, in stage 2, they found no prepotent bias between the two options (left versus right), and found that dACC activity increased as the expected value of the two engage options became more similar. We replicated all of these findings (**Fig. 3a,b** and **Supplementary Fig. 1**). KBMR point out that, although the third finding is consistent with accounts that predict that dACC activity should track the difficulty of a given choice[4,14], the second finding cannot be explained in these terms. Crucially, they argue that a difficulty account of dACC should predict greatest activity when choices to forage and engage were equivalent in point value (that is, at $RV_{forage} = 0$). But this was not so; rather, dACC activity continued to increase past this point, as forage value continued to increase. KBMR interpreted this as evidence that dACC activity encoded the value of the foraging option and not choice difficulty. Furthermore, noting the strong overall bias to engage, they argued that the default action was to engage and that the association of dACC activity with the value of foraging was consistent with a role in encoding the value of the non-default action. For stage 2, they postulated that the unchosen option was prima facie the non-default, and thus dACC activity increased as its value increased (that is, it became more similar to the chosen one).

KBMR's task confounds foraging value and choice difficulty

On the surface, KBMR's evidence for a foraging account of dACC may seem persuasive. However it depends on a problematic assumption, namely that difficulty is greatest when the points offered for the two

options are most similar (that is, when RV_{forage} is closest to zero). The validity of this assumption can be checked against the empirically observed choice behavior, by plotting RV_{forage} against the observed frequency of choices to forage. If KBMR's assumption is correct, then the point on the curve at which RV_{forage} equals zero should coincide with the point at which the forage and engage options each had a 50% likelihood of being selected; that is, the empirical indifference point. It should also coincide with the point at which decisions take the longest, reflecting a maximum in choice difficulty. For the unbiased choices made in stage 2, both conditions held (**Fig. 2c**). However, in stage 1, neither was the case: KBMR assumed that choice difficulty would reach its maximum at $RV_{forage} = 0$; our results showed that both the empirical indifference point (**Fig. 2b**, left) and the point of longest response times (RTs) (**Fig. 2b**, right) aligned at a similar RV_{forage} value, but that this point was substantially to the right of $RV_{forage} = 0$.

Given these observations, a more appropriate measure of choice difficulty would adjust the inferred value of the options such that they are shifted to the right and their point of equal relative value is centered on each participant's empirical indifference point. To identify these empirical indifference points, we fit a standard model of decision making[4,25–29] to each participant's choice and RT data (**Fig. 2b,c**; see Online Methods and **Supplementary Fig. 2**). The best fit of the model to these behavioral data required setting the indifference point, on average, at $RV_{forage} = 1.65$, significantly to the right of 0 (s.e.m. = 0.26, $t_{14} = 6.5$, $P = 1.4 \times 10^{-5}$; **Fig. 2b**), and consistent with the observed bias to engage. We used each participant's estimated shift from this zero point to generate estimates of relative foraging value that better match their empirical choice behavior. We refer to this corrected estimate as $RV_{forage-C}$.

While these analyses reveal a problem with KBMR's index of difficulty, they also uncover a more serious concern: when a more appropriate index of difficulty is applied (on the basis of $RV_{forage-C}$), it becomes evident that KBMR's experimental design confounds foraging value and choice difficulty. Consider the range of foraging choices used in their study (and in our replication). This range was reasonably evenly distributed around the zero point for RV_{forage}, but not for $RV_{forage-C}$ (**Fig. 2b**). Critically, the vast majority of choices fell to the left of or near participants' actual choice indifference points (at $RV_{forage-C} = 0$). This resulted in a high correlation between foraging value and an empirically derived estimate of value similarity (and therefore choice difficulty) based on the absolute magnitude of $RV_{forage-C}$ ($-|RV_{forage-C}|$; average Spearman's $\rho = 0.80$, $t_{14} = 14.1$, $P = 1.2 \times 10^{-9}$). Accordingly, when we regressed fMRI activity on this empirical measure of difficulty, we identified the same regions of

Figure 2 KBMR's estimates of value and choice difficulty do not align with behavioral data. (**a**) Schematic of an example trial in KBMR's task. In stage 1 (top), participants are offered a pair of potential rewards (center numbers). They can choose to forage for a better pair of rewards from the set shown at the top of the screen (red box), in which case a random pair from that set is swapped with the current offer and they incur a forage cost (shown at left below red box) and a delay until the new choice is shown. They can forage as many times as they prefer (or not at all) before opting to proceed to stage 2 (lower panel) and engage in the selected choice. At that point, a probability is randomly assigned to each reward (height of violet bar beside each number), and they choose which reward-probability pair to attempt. They receive the outcome of this gamble as points that accumulate at the bottom of the screen (not shown). While potential rewards were indicated numerically in experiment 2 (as shown here), abstract symbols with learned reward associations were used in the original task and experiment 1. (**b**,**c**) Choice (left) and RT (right) data from the two stages of experiment 1 (black curves). Gray bars show histograms of trial frequencies. (**b**) In stage 1, both the indifference point in the choice curve and the peak in response times exhibit clear

(and comparable) shifts to the right of $RV_{forage} = 0$. (**c**) In contrast, in stage 2, both the indifference point and RT peak coincide with the point on the x axis where the relative value of choosing right versus left (RV_{right}) = 0. Red curves show the predicted RTs and choice probabilities, and corresponding indifference points (vertical dashed lines) based on fits of the decision-making model (see Online Methods). We corrected RV_{forage} so that it was centered on this empirical indifference point ($RV_{forage-C} = 0$) and defined choice difficulty as value similarity with respect to this corrected measure ($-|RV_{forage-C}|$; blue-red shading). These data further show that in experiment 1, as in KBMR's study, the indifference point occurs toward the higher end of forage values tested, confounding forage value and choice difficulty. For display purposes, **b**,**c** (and **Fig. 4b**) show the result of a fixed-effects model across all participants (error bars reflect s.e.m.), but all analyses reported in the main text were based on individual participant fits. Note also that continuous data were used in all fits but are shown here binned. We have also truncated the x axis to only show RV_{forage} bins with an average of five or more trials per participant, but we show complete fits in **Supplementary Figure 2**.
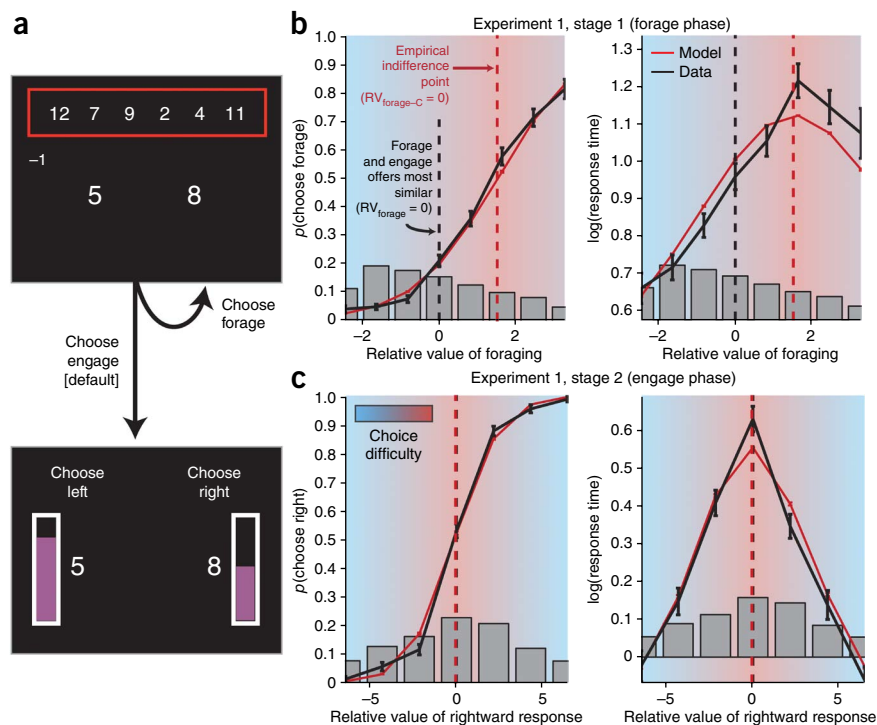
dACC as KBMR found to track foraging value in stage 1 and the value of the unchosen option in stage 2 (**Fig. 3c** and see below).

Difficulty explains dACC activity better than forage value

KBMR's original analyses thus appear both to incorrectly model choice difficulty and to confound it with foraging value. As suggested above, the latter is a consequence of the range of choices tested in their study. Critically, in this range both interpretations of the findings—in terms of foraging or difficulty—make the same prediction: activity of dACC should increase as the value of foraging increases and the pair of options approaches the indifference point (**Fig. 4a**). However, the two theories make different predictions as the value of foraging increases further and begins to strongly favor selection of that option. The foraging theory predicts that dACC activity should continue to increase (or perhaps approach an asymptote) in this range as foraging continues to increase in value. In contrast, a difficulty-based account predicts that dACC should decrease past the indifference point, as the value of foraging more decisively exceeds that of engagement and thus choices become easier. In other words, the foraging account predicts a monotonic relationship between foraging value and dACC activity, whereas the difficulty account predicts a non-monotonic relationship, with activity maximal at the point of indifference and dropping off as one option or the other becomes more clearly preferred, and the choice becomes easier.

Experiment 2 ($N = 14$) tested these predictions. We used a modified version of KBMR's task that maintained an engage bias (mean

indifference point = 5.0, s.e.m. = 1.3, $t_{13} = 3.9$, $P = 0.0018$; see also **Supplementary Fig. 1**) but included choice sets that spanned a wider range, from ones that strongly favored the engage option to others that strongly favored the forage option (**Fig. 4a,b**; see Online Methods). When testing a range of foraging values comparable to that in experiment 1, we again found a linear relationship between foraging value and dACC activity (**Supplementary Fig. 3**). However, as we tested value ranges that increasingly favored the forage option (using a sliding-windowed analysis), dACC's relationship with foraging value became less positive and then reversed, such that, at the upper end of foraging values, dACC had a significant negative correlation with foraging value (**Supplementary Fig. 3b**; $t_{13} = -3.4$, $P = 0.0046$; see also **Supplementary Fig. 8**). The results of experiment 2 thus clearly demonstrated a non-monotonic relationship between foraging value and dACC activity, whereby activity was least when the choice options strongly favored either engaging or foraging, and greatest when they were equiprobable (**Fig. 5c**; compare **Fig. 5a**). Furthermore, as in experiment 1, we found that choice difficulty explained dACC activity at both stages of the task (**Fig. 5d**; compare **Fig. 5b**). Finally, because choice difficulty and foraging value were orthogonal in this task ($M_\rho = 0.07$, $t_{13} = 0.8$, $P = 0.45$), we were also able to directly compare the ability of each to predict dACC activity. We did so by entering them into the same general linear model (GLM) and found a significant effect of difficulty ($t_{13} = 2.8$, $P = 0.014$) but not foraging value ($t_{13} = 0.47$, $P = 0.64$) (**Supplementary Fig. 4a**). In two more tests, we confirmed that dACC activity in this study was better accounted for by choice
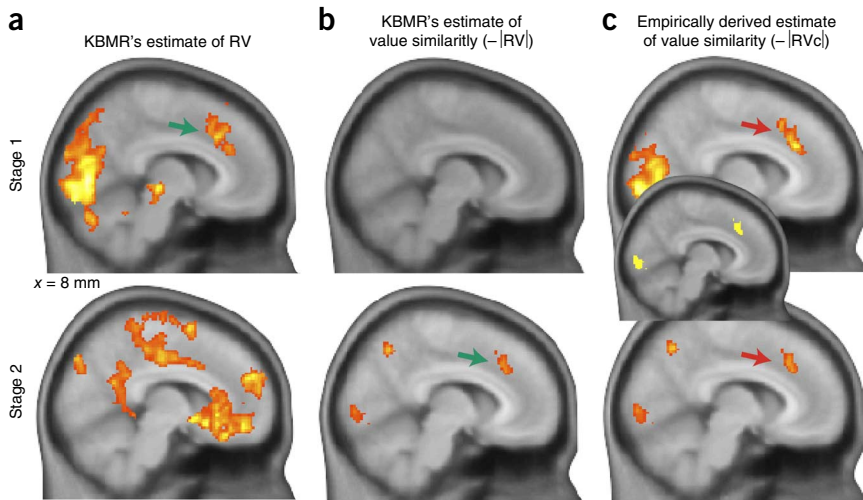
Figure 3 Experiment 1: choice difficulty accounts for dACC activation in both stages of KBMR's original task. (**a**) Whole-brain contrast for brain regions tracking KBMR's estimate of the relative value of forage versus engage options in stage 1 (RV_{forage}) (top) and the chosen versus unchosen option in stage 2 (bottom). (**b**) Regions tracking the similarity of option values ($-|RV_{forage}|$) for the same options represented in **a**. We replicate the finding of significant dACC activity in the contrasts shown in **a**, top, and **b**, bottom (green arrows), consistent with the foraging value account. (**c**) However, using an estimate of value similarity corrected to align with the behavioral data ($-|RV_{forage-C}|$ in stage 1), the same region of dACC is found to be associated with choice difficulty in both stage 1 and stage 2 (red arrows). A conjunction of these two contrasts (inset) indicates a large degree of overlap in dACC. Statistical maps in **a–c** are thresholded at voxelwise $P < 0.01$, extent threshold of 200 voxels. Color maps represent t statistics, ranging from 2.6 to 8.0.

activations are found when simply regressing dACC activity on RT, which is assumed to provide a 'model-free' estimate of difficulty for two-alternative forced-choice tasks like this one. However, because RT is a noisy estimate of difficulty, we tested for and found a significant contribution of our earlier model-based estimate of difficulty ($-|RV_{forage-C}|$) to dACC activity even after removing variance accounted for by RT (experiment 1 $t_{14} = 4.0$, $P = 0.0013$; experiment 2 $t_{13} = 2.4$, $P = 0.034$). Thus, despite the close relationship between choice difficulty and RT on this task (as in many), we were still able to rule out a simple 'time-on-task' (that is, purely RT-based[30,31]) account of the choice difficulty effects we observed in dACC.

Consistent with the inherent relationship between foraging value and choice difficulty in standard foraging settings (**Fig. 1**), the potential for confounding these two variables lurks in any study that attempts to link dACC to foraging/non-default valuation. This appears to be the case for three prominent neuroimaging studies that have been argued to support KBMR's foraging account. Mobbs et al.[21] had participants perform an analog to a dynamic patch-leaving task and showed that dACC activity increased as conditions favored leaving the patch (more competition and/or less reward for current resources) and decreased as conditions favored staying. To the extent that participant behavior in this study approximated optimal foraging, dACC activity was simultaneously indexing choice difficulty (for precisely the reasons described in **Fig. 1**). Two more studies—one by Boorman et al.[19], the other a more recent study from Kolling et al.[20]—employed tasks that share properties with KBMR's. In both studies, participants made risky choices partly based on explicit values manipulated by the experimenter and partly based on an inherent bias

difficulty than foraging value. First, a direct contrast of the regressors in the aforementioned GLM showed a significantly greater average parameter estimate for difficulty than foraging value (paired $t_{13} = 2.2$, $P = 0.046$; see also **Supplementary Fig. 4b**). Second, we performed a Bayesian model comparison over separate GLMs that accounted only for foraging value or only for choice difficulty, and found that the difficulty model was favored across dACC (**Supplementary Fig. 4c**), including within our a priori region of interest (**Supplementary Fig. 4d**). These findings weigh heavily in favor of choice difficulty and against foraging value as an account of responses in dACC.

In keeping with analyses performed in KBMR's study, the analyses above assume that choice difficulty is a simple linear function of value similarity. Further analyses showed that our findings hold for alternate formulations of difficulty that instead focus on the relative likelihood of choosing one option or another, based on either the decision model described above (the drift diffusion model[25]; **Supplementary Fig. 5b**) or a simpler model of the decision process (a logistic regression; **Supplementary Fig. 5c**). These measures of difficulty have the benefit of accounting for nonlinearities in choice behavior (**Figs. 2b,c** and **4b**), including the tendency for choice probabilities to become asymptotic beyond a certain point on the relative value scale. Similar

Figure 4 Experiment 2 de-confounds foraging value and choice difficulty. (**a**) Because of the original confound (**Fig. 2b**), forage value (green) and choice difficulty (red) accounts make similar predictions with respect to dACC activity in experiment (Exp) 1 (**Fig. 3**). However, experiment 2 dissociates the two accounts: foraging value predicts that dACC should increase monotonically, whereas choice difficulty predicts that dACC should decrease as foraging value increases beyond the indifference point (**Fig. 5**). (**b**) As in experiment 1, behavior and model fits from experiment 2 exhibited a shift in the indifference point in stage 1 (compare **Fig. 2b**), but a wider range of choices allowed us to de-confound difficulty and RV_{forage}.
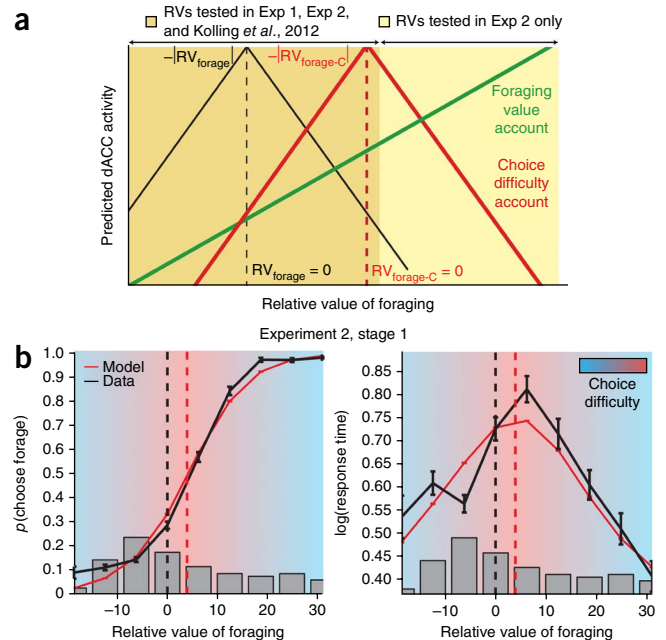
**Figure 5** Experiment 2: choice difficulty but not foraging value accounts for dACC activation when a wider range of foraging values is used. (**a**) Average fMRI activity for each of six RV_{forage} quantiles in experiment 1, taken from a dACC region of interest around peak coordinates from the contrast shown in **Figure 3a**, top. This plot is provided for visual reference, but note that, unlike those in the remaining panels, this analysis is circular because it is intentionally biased toward the dACC region that is maximally sensitive to RV_{forage}. We therefore omit error bars. (**b**) Given the high correlation between foraging value and choice difficulty in experiment 1, the same pattern of activity is observed when dACC activity is binned by choice difficulty. (**c,d**) Results from experiment 2, using a wider range of foraging values that orthogonalized this with respect to choice difficulty (**Fig. 4b**). (**c**) Whole-brain contrasts during stage 1, showing that dACC activity exhibits a quadratic but not linear relationship to RV_{forage}; plot in bottom panel confirms that, over the fuller range of RV_{forage} values used, dACC exhibits the non-monotonic pattern of activity predicted by the choice difficulty account (compare **Fig. 4a**; also see **Supplementary Figs. 3** and **8**). For ease of comparison with contrasts in **d**, the color map for the quadratic contrast is inverted so that negative coefficients (suggesting an inverted U-shape) appear in red–yellow. (**d**) A whole-brain contrast for a linear relationship with choice difficulty again identifies dACC, plotted in bottom panel; inset shows conjunction with the same contrast for stage 2. *Linear contrast in **c** is shown at a liberal voxelwise threshold of $P < 0.05$, no cluster extent threshold. All other statistical maps are shown at voxelwise $P < 0.01$, extent threshold of 200 voxels. Error bars reflect between-subjects s.e.m.

toward the safest (most probable) of the outcomes available. Again, their analyses show that dACC activity increases with the value of the non-default option (or, conversely, decreases with the value of the default). We simulated choices in both of these contexts, with approximations to all relevant decision-making parameters (including bias), and found that choice difficulty increased under the relevant conditions in which these experiments found increasing dACC activity (Online Methods; **Supplementary Figs. 6** and **7**).

Taken together with our neuroimaging findings, these analyses make clear the importance of properly estimating choice difficulty in studies of reward-based choice. Of course, difficulty is not the only domain-general decision parameter that needs to be accounted for when attempting to relate dACC activity to valuation in the context of decision-making. We describe additional factors that may have contributed to the foraging value findings in KBMR and some subsequent work (**Supplementary Figs. 7–9**). For instance, in KBMR's data we found foraging value to be correlated with the degree of surprise participants might have experienced when encountering a given option set (on the basis of past experience) and found that this surprise signal correlated with activity in a rostral region of dACC in both of our experiments (Online Methods; **Supplementary Fig. 9**), consistent with previous context-general accounts[2,3,8,32–36]. We also highlight potential concerns arising from region of interest selection[37], as well as some fundamental inconsistencies in the predictions made by foraging accounts across the relevant publications (**Supplementary Figs. 7** and **8**). All in all, when considered alongside our findings on decision difficulty, these additional considerations further undermine the theoretical conclusions of the papers in question.

DISCUSSION
KBMR's foraging account of dACC predicts that the region tracks the value of the current non-default option (for example, switching to a new patch or performing a different task), irrespective of and potentially negating any role the dACC has in tracking choice difficulty. We have provided evidence that refutes this prediction, showing instead that dACC responds to foraging value only to the extent that it offers a proxy for choice difficulty. This conclusion resonates
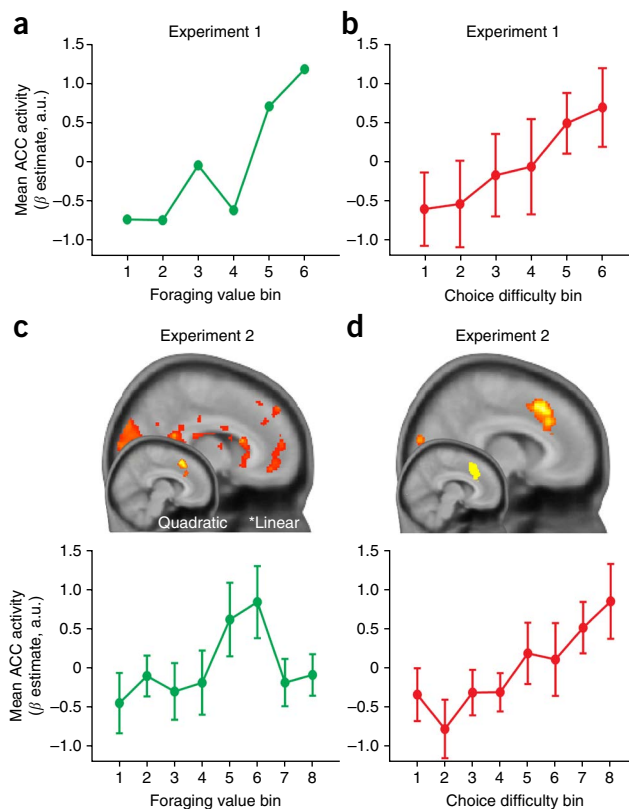
with two closely related theories of dACC function, which focus respectively on conflict monitoring and value comparison. Conflict monitoring accounts[14,15] predict that dACC should track one's level of indifference in a decision-making task because higher indifference trials require the allocation of greater control, and may be aversive for the same reason. The value comparator account[4] similarly predicts greater dACC involvement on more difficult choice trials, but instead because this region is assumed to be directly involved in the process of comparing the values of one's options (using an accumulator model similar to the one used to model the present data). We do not purport to adjudicate between these two accounts with the current data, and we note that doing so may in general be difficult, as the two make a number of overlapping predictions (see refs. 4,14,38).

Both conflict and value comparator theories share in common the assumption that dACC's role in decision-making and/or cognitive control is general rather than specific to a particular decision-making context. Our findings therefore argue against a specific role for dACC in foraging-like decisions. Rather, to the extent that dACC is responsive to a non-default option, our results are consistent with previous theories proposing that this reflects its role in engaging the control processes needed to override the (typically more automatic) default option. We recently described an integrative theory of dACC function, which proposed that the dACC is responsible for estimating the expected value of control-demanding behaviors (EVC) and selecting which to execute[2]. Like the KBMR theory, this theory predicts that dACC activity should track the expected reward for engaging in non-default behavior, inasmuch as this can be considered to be control-demanding. However, the EVC theory specifies that determining the expected value of a controlled behavior also requires estimating the demands for control as well as the costs of control itself. Both of these quantities should correlate with choice difficulty, whereas the value of the outcome (that is, the expected reward) can in principle be dissociated from difficulty (as in experiment 2). As we have shown,

in the particular case of KBMR's task, their findings provide evidence only for dACC's role in encoding the demand for or costs of control, and not the value of the forage option *per se*.

The present work also illustrates how a widely used quantitative model of decision making can provide an estimate of choice difficulty in the presence of strong biases toward one option (or one attribute[19]), a situation that otherwise risks confounding difficulty and value. In such cases, there is guaranteed to be a range of values for which the relative value of the less favored (that is, non-default) option appears to be positive and increasing simultaneously with the difficulty of the decision. (As discussed in the introduction, optimal behavior in a classic patch-leaving foraging task represents a particularly clear case of this.) The degree to which this range of options dominates in a given study can be evaluated qualitatively by comparing relative value and RT distributions (in KBMR's study these showed a strong positive relationship; **Fig. 2b**). Extracting this information from RTs further requires that the subject respond immediately on reaching a decision (non-default valuation studies, as with many other neuroimaging studies, typically forbid responses within an initial window of approximately 2–4 s; refs. 17,19,20). More generally, our findings highlight the risks of estimating choice difficulty in value-based decision tasks on the basis of objective value (for example, points) without taking account of behavioral data (that is, choice and RT distributions). Here as well, the use of a formal model quantitatively fit to the behavioral data can be helpful.

In summary, our study corroborates the importance of dACC in foraging decisions, but not for the reasons reported by KBMR. Rather than reflecting the value of the foraging option itself, their findings and ours suggest that dACC activity can be most parsimoniously and accurately interpreted as reflecting choice difficulty (or a correlate of this, such as total evidence accumulated), as has been observed in many other contexts. Importantly, in a changing environment, indifference may mark the optimal point of transition between default and foraging behavior and thus explain its engagement in the context of the task introduced by KBMR. These observations should help unite the long-standing literature on dACC function with the more recent one that has begun to emerge concerning its involvement in foraging decisions.

METHODS

Methods and any associated references are available in the online version of the paper.

Note: Any Supplementary Information and Source Data files are available in the online version of the paper.

AUTHOR CONTRIBUTIONS
A.S., J.D.C. and M.M.B. designed the experiments; A.S. and M.A.S. performed the experiments; A.S. analyzed the data; and A.S., J.D.C. and M.M.B. wrote the paper.

COMPETING FINANCIAL INTERESTS
The authors declare no competing financial interests.

1. Rushworth, M.F.S., Kolling, N., Sallet, J. & Mars, R.B. Valuation and decision-making in frontal cortex: one or many serial or parallel systems? *Curr. Opin. Neurobiol.* **22**, 946–955 (2012).
2. Shenhav, A., Botvinick, M.M. & Cohen, J.D. The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron* **79**, 217–240 (2013).
3. Alexander, W.H. & Brown, J.W. Medial prefrontal cortex as an action-outcome predictor. *Nat. Neurosci.* **14**, 1338–1344 (2011).
4. Hare, T.A., Schultz, W., Camerer, C.F., O'Doherty, J.P. & Rangel, A. Transformation of stimulus value signals into motor commands during simple choice. *Proc. Natl. Acad. Sci. USA* **108**, 18120–18125 (2011).
5. Shackman, A.J. *et al.* The integration of negative affect, pain and cognitive control in the cingulate cortex. *Nat. Rev. Neurosci.* **12**, 154–167 (2011).
6. Venkatraman, V. & Huettel, S.A. Strategic control in decision-making under uncertainty. *Eur. J. Neurosci.* **35**, 1075–1082 (2012).
7. Holroyd, C.B. & Yeung, N. Motivation of extended behaviors by anterior cingulate cortex. *Trends Cogn. Sci.* **16**, 122–128 (2012).
8. Behrens, T.E.J., Woolrich, M.W., Walton, M.E. & Rushworth, M.F.S. Learning the value of information in an uncertain world. *Nat. Neurosci.* **10**, 1214–1221 (2007).
9. Rushworth, M.F.S., Noonan, MaryAnn P., Boorman, E.D., Walton, M.E. & Behrens, T.E. Frontal cortex and reward-guided learning and decision-making. *Neuron* **70**, 1054–1069 (2011).
10. Duncan, J. & Owen, A.M. Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends Neurosci.* **23**, 475–483 (2000).
11. Paus, T., Koski, L., Caramanos, Z. & Westbury, C. Regional differences in the effects of task difficulty and motor output on blood flow response in the human anterior cingulate cortex: a review of 107 PET activation studies. *Neuroreport* **9**, R37–R47 (1998).
12. Pochon, J.-B., Riis, J., Sanfey, A.G., Nystrom, L.E. & Cohen, J.D. Functional imaging of decision conflict. *J. Neurosci.* **28**, 3468–3473 (2008).
13. FitzGerald, T.H.B., Seymour, B. & Dolan, R.J. The role of human orbitofrontal cortex in value comparison for incommensurable objects. *J. Neurosci.* **29**, 8388–8395 (2009).
14. Botvinick, M.M., Braver, T.S., Barch, D.M., Carter, C.S. & Cohen, J.D. Conflict monitoring and cognitive control. *Psychol. Rev.* **108**, 624–652 (2001).
15. Botvinick, M.M. Conflict monitoring and decision making: reconciling two perspectives on anterior cingulate function. *Cogn. Affect. Behav. Neurosci.* **7**, 356–366 (2007).
16. Shenhav, A. & Buckner, R.L. Neural correlates of dueling affective reactions to win-win choices. *Proc. Natl. Acad. Sci. USA* doi:10.1073/pnas.1405725111 (14 July 2014).
17. Kolling, N., Behrens, T.E.J., Mars, R.B. & Rushworth, M.F.S. Neural mechanisms of foraging. *Science* **336**, 95–98 (2012).
18. Pearson, J.M., Watson, K.K. & Platt, M.L. Decision making: the neuroethological turn. *Neuron* **82**, 950–965 (2014).
19. Boorman, E.D., Rushworth, M.F. & Behrens, T.E. Ventromedial prefrontal and anterior cingulate cortex adopt choice and default reference frames during sequential multi-alternative choice. *J. Neurosci.* **33**, 2242–2253 (2013).
20. Kolling, N., Wittmann, M. & Rushworth, M.F.S. Multiple neural mechanisms of decision making and their competition under changing risk pressure. *Neuron* **81**, 1190–1202 (2014).
21. Mobbs, D. *et al.* Foraging under competition: the neural basis of input-matching in humans. *J. Neurosci.* **33**, 9866–9872 (2013).
22. Hayden, B.Y., Pearson, J.M. & Platt, M.L. Neuronal basis of sequential foraging decisions in a patchy environment. *Nat. Neurosci.* **14**, 933–939 (2011).
23. Charnov, E.L. Optimal foraging, the marginal value theorem. *Theor. Popul. Biol.* **9**, 129–136 (1976).
24. Stephens, D.W. & Krebs, J.R. *Foraging Theory* (Princeton Univ. Press, 1986).
25. Ratcliff, R. A theory of memory retrieval. *Psychol. Rev.* **85**, 59–108 (1978).
26. Bogacz, R., Brown, E., Moehlis, J., Holmes, P. & Cohen, J.D. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* **113**, 700–765 (2006).
27. Milosavljevic, M., Malmaud, J., Huth, A., Koch, C. & Rangel, A. The Drift Diffusion Model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgm. Decis. Mak.* **5**, 437–449 (2010).
28. Krajbich, I., Armel, C. & Rangel, A. Visual fixations and the computation and comparison of value in simple choice. *Nat. Neurosci.* **13**, 1292–1298 (2010).
29. Basten, U., Biele, G., Heekeren, H.R. & Fiebach, C.J. How the brain integrates costs and benefits during decision making. *Proc. Natl. Acad. Sci. USA* **107**, 21767–21772 (2010).
30. Grinband, J. *et al.* The dorsal medial frontal cortex is sensitive to time on task, not response conflict or error likelihood. *Neuroimage* **57**, 303–311 (2011).
31. Weissman, D.H. & Carp, J. The congruency effect in the posterior medial frontal cortex is more consistent with time on task than with response conflict. *PLoS ONE* **8**, e62405 (2013).
32. Wessel, J.R., Danielmeier, C., Morton, J.B. & Ullsperger, M. Surprise and error: common neuronal architecture for the processing of errors and novelty. *J. Neurosci.* **32**, 7528–7537 (2012).
33. Garrison, J., Erdeniz, B. & Done, J. Prediction error in reinforcement learning: a meta-analysis of neuroimaging studies. *Neurosci. Biobehav. Rev.* **37**, 1297–1310 (2013).
34. Hayden, B.Y., Heilbronner, S.R., Pearson, J.M. & Platt, M.L. Surprise signals in anterior cingulate cortex: neuronal encoding of unsigned reward prediction errors driving adjustment in behavior. *J. Neurosci.* **31**, 4178–4187 (2011).
35. Bryden, D.W., Johnson, E.E., Tobia, S.C., Kashtelyan, V. & Roesch, M.R. Attention for learning signals in anterior cingulate cortex. *J. Neurosci.* **31**, 18266–18274 (2011).
36. Cavanagh, J.F. & Frank, M.J. Frontal theta as a mechanism for cognitive control. *Trends Cogn. Sci.* doi:10.1016/j.tics.2014.04.012 (15 May 2014).
37. Kriegeskorte, N., Simmons, W.K., Bellgowan, P.S.F. & Baker, C.I. Circular analysis in systems neuroscience: the dangers of double dipping. *Nat. Neurosci.* **12**, 535–540 (2009).
38. Yeung, N., Botvinick, M.M. & Cohen, J.D. The neural basis of error detection: conflict monitoring and the error-related negativity. *Psychol. Rev.* **111**, 931–959 (2004).

ONLINE METHODS

Participants. Healthy right-handed individuals were recruited to participate in a neuroimaging study involving choices and rewards. 15 individuals completed experiment 1 (9 female; M_{age} = 23.5, SD_{age} = 4.1) and 14 independent individuals completed experiment 2 (8 female; M_{age} = 20.6, SD_{age} = 2.4). Additional participants were excluded a priori for excessive head movement (1), incomplete sessions (4), misunderstanding instructions as assessed during a structured post-experiment interview (2) or not meeting KBMR's criterion of foraging on at least eight trials (6). Of the 14 included participants in experiment 2, we excluded one of three trial blocks for one participant who reported falling asleep during that block. No statistical tests were used to predetermine sample sizes, but our sample sizes are within the standard range in the field. Participants provided informed consent in accordance with policies of the Princeton University institutional review board.

Procedure. Experiment 1 followed the procedure described by KBMR[17] (**Fig. 2a**). Briefly, participants learned fixed reward values associated with 12 abstract symbols. Participants then performed a decision task that proceeded in two stages. In stage 1, the participant chose whether to engage a pair of symbols offered or to forage for a better pair from a set of six other symbols also shown. Each time they chose to forage, the current engage pair was swapped out for a random pair from the forage set. Foraging came with an explicit search cost that was indicated at the start of stage 2 by a colored box surrounding the search set, and feedback regarding this cost was provided during the delay period that was also incurred when choosing to forage. Search costs were associated with point losses of varying (but known) magnitudes and were incurred on 70% of forage choices; the remaining forage choices were associated with no point loss. Once the participant chose to engage the offered pair (and no longer forage, or not to forage in the first place), they proceeded to stage 2. In stage 2, each of the two symbols was independently and randomly assigned an explicit probability of success (range: 20–90%) and the participant chose which of the two symbol-probability pairs (gambles) they would like to play. If the gamble chosen was successful, the participant received the full reward value associated with that symbol. Participants were given feedback both after each forage choice (regarding whether they incurred the search cost) and after making their stage 2 choice (regarding the success of both the chosen and unchosen options). Rewards received were tallied by a progress bar shown at the bottom of the screen; each time this bar filled to a goal line, the participant received $1.00 and the bar restarted.

Jittered Poisson-distributed intertrial intervals (ITIs) were added between choosing to forage and receiving feedback regarding whether the search cost was incurred on that trial (range, 2–6 s; M = 3.0 s), between foraging feedback (1–2 s duration) and the next stage 1 choice set for that trial (range, 2–4.5 s; M = 2.7 s), between choosing to engage and being shown the engage probabilities (range, 3–8 s; M = 4.5 s), between making a stage 2 choice and being given feedback (range, 3–8 s; M = 4.5 s), and between stage 2 feedback and beginning stage 1 of the next trial (range, 2–4.5 s; M = 2.7 s).

Relative to KBMR's experiment, our procedure differed only in that participants in our study were able to submit a response as soon as they were provided with the relevant information at a given stage. KBMR, by contrast, included a jittered 'monitor phase' at the start of each stage that prevented participants from responding for 2–4 s at the start of stage 1 and 1–4 s at the start of stage 2, resulting in truncated RT distributions that preclude modeling with the drift diffusion model (DDM; see "Behavioral analysis" below). To maintain similar overall timing, we buffered ITIs immediately after stage 1 and stage 2 responses by an additional 2.7 s or 1.8 s (the average length of the respective monitor phases) minus the RT on that trial (that is, any extra time that would have been captured by the monitor phase during the response period was instead added to the ITI, without the participants' knowledge).

As in the original experiment, Exp 1 included 135 trials, broken up into 2 blocks, and the 6 forage and 2 engage values presented at the start of each trial were drawn at random, without replacement, from a uniform distribution of the 12 total symbol values. This necessarily resulted in an over-representation of trials in which the average forage and engage values were similar, relative to those in which forage values were much higher or much lower than engage values. To explore the distribution of relative forage values more fully, particularly those trials in which the relative value of foraging exceeded the individual's subjective indifference point, Exp 2 made a number of modifications.

First, we used a much wider distribution of potential reward values (69 rather than 12). To accomplish this without imposing excess memory load on our participants, we used explicit numeric reward values (2–70) during the task rather than abstract symbols, and all explicit reward values were scaled down by a factor of 10 relative to Exp 1 (where the range shown was 20–130 points). These scaled-down numeric values mapped onto similar monetary reward values as those in Exp 1 (for example, the point value '100' in Exp 1 was associated with similar monetary reward as the point value '10' in Exp 2); we therefore applied this constant scaling factor to Exp 1 values during modeling to allow more direct comparison between the studies. Search costs were also indicated numerically rather than with the color of the search box. Furthermore, rather than uniform random sampling of forage and engage values on each trial, forage and engage values were preselected across trials so as to uniformly sample a range of relative forage values. Specifically, for each session we randomly sampled sets of forage and engage values until we found ten within each of twenty relative forage value ranges (where relative foraging value is defined as the difference between the average of the initial foraging set and the average of the pair of initial engage options) evenly spaced between approximately −29 and 56. This resulted in a total of 200 trials, performed across 3 task blocks. An asymmetric value range was used to account for the strong baseline bias to engage observed in the original task. Trial order was randomized and, aside from this initial arrangement of stage 1 values, all other sampling of values (for example, forage/engage options after choosing to forage, probabilities in stage 2) proceeded as in Exp 1.

To slightly shorten and reduce the complexity of the individual trials, Exp 2 also used deterministic search costs rather than including a possibility that the search cost would not be incurred. As a result, the feedback portion of the forage delay was removed and the ITI between choosing to forage and receiving the next forage/engage options was set to a range of 2.3–7 s (M = 3.0). The range of search costs used was otherwise the same as in Exp 1 (appropriately scaled, as described above). In spite of these differences, behavior was qualitatively similar in the two studies (see **Supplementary Fig. 1**). While average stage 1 RTs were faster in Exp 2 (two-sample t_{27} = −2.3, two-tailed P = 0.030), possibly owing to the larger number of trials and/or numeric rather than symbolic reward (reducing the need for episodic memory retrieval), the average range of RTs was similar (t_{27} = 0.71, P = 0.48), as was the pattern of regression estimates for relevant decision variables for both task stages. While Exp 2 participants also exhibited a smaller overall bias, individual difference analyses performed by KBMR (see also ref. 20) indicate that their account predicts that these participants should, if anything, encode foraging value more strongly (that is, exhibit a stronger readiness to forage) than participants in Exp 1 (see legends to **Supplementary Figs. 1** and **7**).

Stimulus presentation and response acquisition was performed using Matlab (MathWorks) with the Psychophysics Toolbox[39]. Participants responded with MR-compatible response keypads.

Image acquisition. Scanning was performed on a Siemens Allegra 3-T MR system. Following KBMR, we used the following sequence parameters for volumes acquired during task performance: 3-mm³ isotropic voxels, repetition time (TR) = 3.0 s, echo time (TE) = 30 ms, flip angle (FA) = 87°, 43 slices, with slice orientation tilted 15° relative to the AC/PC plane. At the start of the imaging session, a high-resolution structural volume (MPRAGE) was also collected, with the following sequence parameters: 2-mm × 1-mm × 1-mm voxels, TR = 2.5 s, TE = 4.38 ms, FA = 8°.

Behavioral analysis. *Relative value.* We first estimated the relative value of the options at each stage (stage 1: forage versus engage; stage 2: left versus right). These values were used as regressors to identify regions that tracked relative value (and/or its absolute value) and, crucially, were also used as proxies for drift rate when modeling choices with the DDM (see the following section).

For stage 2, determining relative value simply involved comparing the expected values (reward magnitude × probability) of the two options, as in the original paper. For stage 1 we used KBMR's primary measure of relative foraging value, which they refer to as search evidence (see their equations S2–S5, reproduced as equation (1) below). Determining this quantity involved comparing the average value of the engage pair (consisting of reward values R_1 and R_2, whose exact probabilities P_1 and P_2, each 55% in expectation, will be revealed in stage 2), weighted by the ratio of the two values (to account for a possible preference for easier stage 2

choices), to the average over all n possible pairs drawn from the current forage set (similarly weighted):

$$RV_{forage} = \left(\frac{1}{n} \sum_{i=1}^{n} V_{Offer_i} \right) - V_{CurrentOffer} \qquad (1)$$

where $V_{offer} = 0.55(w_1 R_1 + w_2 R_2)$, $w_1 = p\left(\frac{R_2}{R_1} < \frac{P_1}{P_2} \right)$ and $w_2 = 1 - w_1$.

Note that RV_{forage} omits the explicit cost of foraging on any given trial. In KBMR's study this factor was instead included as a separate regressor in their imaging analyses, so we have done the same to remain consistent with their approach. However, we note that additional analyses (not shown) confirm that all of our findings are robust to including this cost in RV_{forage}, as estimated above. They are also robust to replacing the RV_{forage} equation above with one that makes no assumptions about offer weighting and simply subtracts the average value of engaging and the forage cost from the average value of foraging.

Following the original paper, we also performed a logistic regression to predict forage choices on the basis of the minimum/maximum of the engage set, the search cost and the minimum/mean/maximum of the forage set, including an intercept term. This intercept term was significantly biased toward engaging in both experiments (Wilcoxon signed-rank test, Exp 1: $P = 6.1 \times 10^{-5}$, Exp 2: $P = 0.04$; **Supplementary Fig. 1a**), consistent both with the previous results and with our finding that these explicit reward values alone fail to fully describe the inputs to the decision process. We also used this logistic regression to generate a more 'model-free' estimate of choice difficulty, relative to the DDM (see **Supplementary Fig. 5c**). We did this by generating the log-odds of choosing to forage versus engage on each trial on the basis of the choice values (the ones entered into the logistic regression) and the best-fit regression weights (including the intercept). Choice difficulty was defined as the negative absolute value of this log-odds term.

Modeling behavior. To identify individual subjective indifference points (that is, the points at which choices were most difficult) at each stage of the task, we fit behavioral data (choices and RTs) with the DDM[25]. The model represents the decision process as a particle drifting toward one of two decision boundaries (for example, forage versus engage) with a drift rate that determines how much faster it moves toward one or the other (before some amount of Gaussian noise is additionally inserted into the process). The DDM, and variants thereof, have been shown to provide a reasonable approximation for value-based decision processes such as those in the foraging task[4,26–29], with the value of one option relative to another typically being treated as a proxy for drift rate. Notably, this model is at the heart of the value comparator account of dACC function[4], which suggests that dACC activity should reflect the accumulated activity in the DDM over time. As KBMR point out when attempting to rule out both accounts, conflict and comparator theories both predict that dACC activity should be greatest when choices are most difficult. Using the DDM to identify when choices were most difficult therefore allows us to respond most directly to KBMR's presumed lack of support for this prediction.

We used relative value at each stage (as defined above) to generate distributions of predicted choice probabilities and RTs, which were then fit to the data. We used the five-parameter version of the DDM; of these, we fixed the starting point of the decision process (0.0) but allowed the coefficient for decision noise, decision threshold, non-decision time and value scaling factor (d below) to vary by participant and decision stage. Crucially, we also included an offset term (c below) that was added to all relative values within a given stage when they were translated into drift rates. The equation for drift rate was therefore as follows:

$$A = (RV + c)d \qquad (2)$$

where the presumed relative value (RV) for a given choice is modified by the offset tem (c; this can also be thought of as a latent cost) and the scaling factor (d).

This allowed the distribution of drift rates to shift relative to the assumed distribution of relative values. In practical terms, if the assumed relative values were missing an additional fixed value (for example, a subjective cost or bonus for foraging that was not expressed in the explicit values shown), this term would absorb that value by shifting the subjective indifference point accordingly

(see also ref. 40). Like all other free DDM parameters described above, this offset term was allowed to vary by participant and trial stage.

We generated predicted choice probabilities and average RTs (threshold-crossing times) for a given set of parameters based on an analytical solution to the DDM (see appendix to ref. 26). We fit the predictions of the DDM to our data to identify the participant and trial phase-specific DDM parameters that minimized a combination of (a) the negative log likelihood of our choice data and (b) the sum of the squared error for our RT data (after log-transforming both the actual and expected RTs). We optimized these model fits using a version of the fminsearch function in Matlab that implements bounded parameter ranges.

$RV_{forage-C}$ was generated by combining the original values of RV_{forage} with the best-fit offset term (c) for that individual and choice stage. Our primary measure of choice difficulty was the negative absolute value of these corrected values of RV_{forage} ($-|RV_{forage-C}|$). Note that this measure implicitly assumes that the hidden cost estimated by our offset parameter was in a similar value currency as the explicit values that went into RV_{forage}. However, we also used the DDM to generate an alternative measure of choice difficulty that was agnostic on this front (**Supplementary Fig. 5b**). Similarly to the approach described above for our logistic regression analysis, we used the best-fit parameters for each participant to determine the predicted likelihood of choosing one option rather than the other on each trial (that is, likelihood of choosing to forage rather than engage in stage 1, likelihood of choosing the option on the right rather than the left in stage 2). We then defined choice difficulty with respect to the absolute distance between the predicted choice likelihood and indifference (that is, 50%) for that decision, with indifference being considered the point where choices were most difficult.

fMRI analysis. Imaging data were analyzed in SPM8 (Wellcome Department of Imaging Neuroscience, Institute of Neurology, London, UK). Functional volumes were motion corrected, normalized to a standardized (MNI) template (including resampling to 2-mm isotropic voxels), spatially smoothed with a Gaussian kernel (5 mm FWHM) and high-pass filtered (0.01 Hz cut-off). Separate regressors were included for the stage 1 and stage 2 decision phases, as well as for the feedback associated with foraging (only in Exp 1) and with stage 2 choices. These regressors were all modeled as stick functions, and the two decision phases were further modulated by parametric regressors (as described below).

We ran the following variations on this whole-brain GLM; unless otherwise noted, similar regressors were included for both stages of the task (for stage 1, we use relative value to refer to the value of forage versus engage; for stage 2, we refer to the value of the chosen versus unchosen option):

1. Parametric regressors for linear effects of relative value (RV) (**Figs. 3a** and **5c**).
2. Parametric regressors for explicit value similarity ($-|RV_{uncorrected}|$) (**Fig. 3b**).
3. Parametric regressors for choice difficulty ($-|RV_{corrected}|$) (**Figs. 3c** and **5d**; see also variants in **Supplementary Fig. 5**).
4. Parametric regressors for RV and choice difficulty (**Supplementary Fig. 4a,b**).
5. Parametric regressors for linear and quadratic factors of RV (**Fig. 5c**, inset).
6. Parametric regressors for RT and choice difficulty.
7. Parametric regressors for RV by choice at stage 1 (forage versus engage; **Supplementary Fig. 8**).
8. Separate event regressors for each of six (Exp 1) or eight (Exp 2) binned quantiles of relative value in stage 1 and (in a separate GLM) the same for choice difficulty (**Fig. 5**). Additional bins were used for Exp 2 because of the larger number of trials and wider range of relative values tested.

For all GLMs above, stage 2 outcomes and (when appropriate) stage 1 outcomes were each modeled with an additional regressor. GLM 8 also included an event regressor for the onset of stage 2 decisions, without additional parametric modulators. Also, to test for relative variance captured in a simultaneous regression, the default serial orthogonalization procedure in SPM was turned off for parametric regressors in GLMs 1–4, and these GLMs also included an additional regressor for the explicit forage cost on each trial (see "Behavioral analysis" above). Conversely, to provide a strong test against a pure time-on-task account, GLM 6 did implement serial orthogonalization (that is, tested for an effect of choice difficulty

after removing variance accounted for by RT). For display purposes, all whole-brain t-statistic maps are shown at a voxelwise uncorrected threshold of $P < 0.01$ (corresponding to the $z > 2.3$ criterion used in KBMR's study).

We used GLMs 1 and 3 for our Bayesian model comparison of choice difficulty versus foraging value in Exp 2. Since we were comparing the two accounts for dACC's role in foraging choices specifically, these GLMs were modified only to exclude parametric modulators from stage 2. To perform these model comparisons, we first used SPM's Bayesian equivalent of the first-level analyses described above[41], which produced within-participant log-evidence maps for each GLM. These log-evidence maps were aggregated into a formal Bayesian model comparison at the group level with a random-effects analysis[42] that produced voxel-wise estimates of exceedance probability for each of the models being compared. An additional Bayesian model comparison related the difficulty and foraging value models to a baseline model in which all of the relevant task events were modeled with indicator functions (as with all of our GLMs) but no additional parametric modulators were included.

For our sliding window analysis of foraging value (**Supplementary Fig. 3**), we used a variant of GLM 1 wherein only a certain subset of trials was modulated by the foraging value parameter. This subset was determined by rank-ordering trials by foraging value and selecting those that fell within a certain percentile range window (for example, 0th–50th percentiles). All foraging trials outside this window were modeled with a single indicator variable and no parametric modulation. We performed two sets of windowed analyses, one using windows of 70% per GLM, the other using windows of 50% per GLM. Consecutive windows were shifted by 10 percentiles, resulting in 86% and 80% overlap between consecutive windows in the respective analyses.

Our key analyses (particularly in Exp 2) relied on a region-of-interest (ROI) approach. We extracted β estimates from a sphere (9 mm diameter) drawn around peak dACC coordinates from the relative foraging value contrast (from GLM 1 above; **Fig. 3a**; MNI coordinates [x, y, z] = [4, 32, 42]). While this results in a circular analysis for **Figure 5a** (which is shown for visual comparison to the patterns of difficulty-related activity in **Fig. 5b**), we chose these coordinates to provide the strongest bias in our Exp 2 analyses of relative value (**Fig. 5c** and **Supplementary Figs. 3, 4** and **8**) in favor of detecting the same linear pattern of activity as in Exp 1. This ROI is used for all other dACC ROI analyses described in the paper, with the exception of the forage choice analyses shown in **Supplementary Figure 8e** (right). For this analysis, an ROI was drawn around the peak dACC coordinates from Exp 2's sliding analysis of foraging value (**Supplementary Fig. 3b**) focusing on the 20th–70th percentile window, which showed the strongest linear foraging value effect across the two windowed analyses (coordinates: [6, 28, 34]). Unless otherwise indicated, statistical inferences for these ROI analyses were performed with two-tailed, one-sample t-tests.

Simulating the role of choice difficulty in more recent studies of dACC and non-default value. A few studies have attempted to provide supporting evidence for dACC's role in non-default valuation since KBMR's study was published. In particular, Boorman $et\ al.$ (2013)[19] and Kolling $et\ al.$ (2014)[20] engaged participants in different kinds of risky choice settings and exploited their default risk aversion to explore dACC responses to the increasing relative value of a non-default (riskier) option (see figure legends to **Supplementary Figs. 6** and **7**). We examined the possibility that choice difficulty might account for their relevant findings. To this end, we simulated sets of agents to make decisions in each of these choice contexts. These simulated decision-makers simply made choices based on a set of decision parameters (including decision noise) that were passed through a binary logistic (Kolling $et\ al.$[20]) or trinary softmax (Boorman $et\ al.$[19]) function, which simultaneously served to provide us with the relative probabilities of each of the choices. These choice probabilities were used to estimate choice difficulty, based either on the proximity to indifference (50%) for the binary choices or the Shannon's entropy of the choice probabilities for the trinary choices. Depending on the relevant set of dACC results, we either show the average choice difficulty for different conditions (**Supplementary Fig. 7d**, left) or show the β estimates resulting from regressing choice difficulty on a relevant

decision variable (**Supplementary Fig. 6b**, bottom, and **Supplementary Fig. 7d**, middle and right).

Decision parameters and trial values were chosen to produce behavioral patterns approximating those reported for a given study (see **Supplementary Fig. 6**, top, and **Supplementary Fig. 7b**) and in each case included a parameter that served to bias choices toward the currently most probable[19] or safer[20] of the choice options. For simplicity, and without loss of generality, we assumed full knowledge of the outcome probabilities on each trial of Boorman $et\ al.$'s study, which were randomly selected on each trial rather than varying gradually over time. We also note that the qualitative patterns of results shown in **Supplementary Figures 6** and **7** were robust to variations in decision parameters that provided similar qualitative fits to the observed behavior in a given experiment. Collectively, these results comport with the intuition (emphasized in the main text) that choosing against a bias becomes easier as more reward is offered for the already biased option and/or as less reward is offered for the alternative (non-default) option(s). In other words, choice difficulty increases alongside foraging/non-default value (within a certain range; see **Fig. 1**). Additional analyses from ref. 20 are discussed in the legend to **Supplementary Figure 7**.

Foraging value prediction errors. KBMR provide additional support for their foraging value account by showing that dACC activity increases with foraging value both when participants choose to engage and when they choose to forage (intuitively, though by no means necessarily, including trials that strongly favored foraging; see **Fig. 2b**). This result did not replicate in our Exp 1 and, in fact, showed a robust effect in the opposite direction in Exp 2 (**Supplementary Fig. 8**). It is therefore possible that KBMR's finding of foraging value signals for both types of forage choice was a byproduct of the choice difficulty effect we describe in the main text (and/or selection bias; see **Supplementary Fig. 8d**). However, we noted that the skewed distribution of RV_{forage} values in the original experiment (**Fig. 2b**) may have also resulted in a greater proportion of unexpected or surprising (and perhaps therefore more salient) choice sets at the upper end of the RV_{forage} spectrum. Assuming such surprise occurred particularly frequently for trials on which participants chose to forage, it would have made it difficult to disentangle foraging value on those trials from surprise (cf. unsigned prediction error) at the infrequent configuration of values that would encourage foraging (for related arguments, see ref. 43). This fact is particularly relevant given that dACC has been previously associated with prediction errors[8,32–35] and interpreted in more basic terms unrelated to foraging value, including for the relevance of surprise as a signal of control demands[2,3,36].

As a coarse measure of the degree to which a given choice set was likely to have generated such an unsigned prediction error signal, we simply took the absolute difference of each trial's RV_{forage} from the observed mean RV_{forage} up to that point. We found that this was in fact correlated with RV_{forage} in our replication study (Exp 1), conditional on the participant's choice whether or not to forage on a given trial. In particular, we saw a strong positive correlation between RV_{forage} and RV_{forage} surprise on the trials where subjects chose to forage ($M_\rho = 0.89$, $t_{14} = 31.4$, $P = 2.2 \times 10^{-14}$). This correlation reversed on trials where subjects chose to engage ($M_\rho = -0.54$, $t_{14} = -10.6$, $P = 4.5 \times 10^{-8}$). These estimates of foraging value prediction error were used to generate the whole-brain analyses shown in **Supplementary Figure 9**.

A **Supplementary Methods Checklist** is available.

39. Brainard, D.H. The Psychophysics Toolbox. *Spat. Vis.* **10**, 433–436 (1997).
40. Mulder, M.J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W. & Forstmann, B.U. Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *J. Neurosci.* **32**, 2335–2343 (2012).
41. Penny, W., Kiebel, S. & Friston, K. Variational Bayesian inference for fMRI time series. *Neuroimage* **19**, 727–741 (2003).
42. Stephan, K.E., Penny, W.D., Daunizeau, J., Moran, R.J. & Friston, K.J. Bayesian model selection for group studies. *Neuroimage* **46**, 1004–1017 (2009).
43. Hare, T.A., O'Doherty, J.P., Camerer, C., Schultz, W. & Rangel, A. Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. *J. Neurosci.* **28**, 5623–5630 (2008).