Trends in Cognitive Sciences

CellPress
REVIEWS

Opinion

How We Know What Not To Think

Jonathan Phillips,[1,*] Adam Morris,[2] and Fiery Cushman[2]

Humans often represent and reason about unrealized possible actions – the vast infinity of things that were not (or have not yet been) chosen. This capacity is central to the most impressive of human abilities: causal reasoning, planning, linguistic communication, moral judgment, etc. Nevertheless, how do we select possible actions that are worth considering from the infinity of unrealized actions that are better left ignored? We review research across the cognitive sciences, and find that the possible actions considered by default are those that are both likely to occur and generally valuable. We then offer a unified theory of why. We propose that (i) across diverse cognitive tasks, the possible actions we consider are biased towards those of general practical utility, and (ii) a plausible primary function for this mechanism resides in decision making.

The Promise of, and Problem with, Possibilities

Humans can only experience the world as it is, but our most impressive powers of thought require us to imagine how it could be. To choose a city to move to, we imagine different things we might do there. When judging a person's actions, we ask if they could have done otherwise. To assign responsibility for a tragedy, we consider what might have been done to prevent it. When reading that a job candidate is 'punctual', we consider what his recommender might have emphasized instead. In each case, our thoughts are shaped by the alternative possibilities we consider.

This ability is as mysterious as it is pervasive. In theory, there are infinite alternative possibilities we can imagine, and we could therefore spend infinite time generating and considering them. In practice, however, we evaluate alternative possibilities so quickly and effortlessly that the evaluation often goes unnoticed. We must have some rapid and unconscious ability to generate a small set of nonactual possibilities that merit consideration, while excluding the infinity of useless others. How do we know what not to think?

We answer this question at the computational level, as Marr [1] defined the term. Synthesizing recent insights from diverse fields of study, we offer a model for how humans identify one type of possibility: possible actions. These are a crucial part of **modal cognition** (see Glossary) – our general ability to reason about all types of alternative possibilities, including actions as well as events, states of affairs, etc. When humans engage in tasks that require modal cognition (e.g., causal reasoning or responsibility attribution), they sample only a small number of specific possible actions. Moreover, across diverse tasks, people tend to sample possible actions sharing two properties: they are probable (i.e., they occur often) and valuable (i.e., they are usually good). We propose that people's sampling strategy has an adaptive origin: it helps people to efficiently make effective decisions in their own life. We call this the shared adaptive sampling model of modal cognition – it is shared across domains, adaptively organized, and allows us to sample effectively from the vast space of possible actions.

We conclude by exploring the proposed origin of this shared mechanism: decision making. Humans consistently face decisions involving enormous numbers of options (every potential job, next sentence, afternoon activity, etc.). Instead of exhaustively considering all options, people sample a small subset of possible actions for evaluation. Crucially, the possible actions they consider by default are skewed towards actions which are valuable and likely to be chosen. This feature makes perfect sense when trying to efficiently identify effective actions. Hence, we suggest that the shared sampling mechanism essential for modal cognition – one biased towards valuable and probable actions – was principally designed for decision making. In short, our most basic sense of which actions are possible reflects a first draft of what we would choose. Before laying out our proposal in more detail, we review previous work on reasoning about possibilities.

Highlights

Representations of possible actions pervade human high-level cognition, and shape how we plan, attribute causal responsibility, comprehend language, and make moral judgments.

There are too many 'possible actions' for us to consider them all. Recent studies offer a strikingly convergent picture of how we call to mind a limited, useful set of possible actions to consider.

This process of 'sampling' alternative possibilities has a distinctive fingerprint: it focuses on possible actions that are valuable and probable. This fingerprint arises across many diverse tasks that rely on the representation of alternative possibilities.

We provide a novel theoretical proposal that helps to explain this convergence: by default, the possibilities that come to mind are those worth considering during decision making.

[1]Program in Cognitive Science, Dartmouth College, 201 Reed Hall, Hanover, NH 03755, USA

[2]Department of Psychology, Harvard University, 33 Kirkland Street, Cambridge, MA 02138, USA

*Correspondence:
jonathan.s.phillips@dartmouth.edu

Traditional Approaches to Reasoning about Possibilities

There is a strong tradition of research on possibilities across the cognitive sciences. Within philosophy and linguistics, it is common to model possibilities by defining a space of **possible worlds** and then partitioning it in various ways. For instance, we might partition possible worlds by whether they are consistent with what is known (defining what 'might be') or by whether they involve actions an agent is physically capable of performing (defining what she 'could do'). Once the relevant part of the partition is selected (e.g., the part corresponding to what 'might be'), traditional approaches often induce **orderings** of the possibilities within the relevant part of the partition (e.g., from most to least probable). The relevant dimension of ordering will depend on the task at hand [2–9]. It could involve ordering possibilities according to their probability, morality, accessibility, etc. The most highly ranked possibilities are typically those most relevant to the modal task faced. Because this tradition has typically not sought to model the psychological processes involved in modal cognition, these models have had relatively little to say, however, about which possibilities come consciously to mind from within the potentially vast space delimited by the partition.

Within psychology, modal thought has been studied as an essential part of many cognitive domains. Past research includes the role of modal cognition in 'mental models' that are used in reasoning [10–12], explicit counterfactual thought (reviewed in [13]), reflective assessments of which events are 'possible' (e.g., [14,15]), and, recently, the brain networks that support episodic simulation of nonactual future or past events (e.g., [16,17]). Like philosophers, psychologists assume we have some way to define the relevant types of possibilities we are interested in, which can again be modeled as inducing a partition followed by ordering of possibilities within the relevant part of the partition. By partitioning possibilities at the outset, we might, for example, typically not consider possibilities where the premises of an argument are violated when reasoning with mental models [10] or engaging in counterfactual reasoning [18], or follow a given syntactic structure when generating new names for pubs [19]. Further, by ordering possibilities and then focusing on the 'best' of that ordering, we might, for example, focus on morally better counterfactual alternatives when making moral judgments [20], or consider relatively probable alternative possibilities when making causal judgments [21].

Abstracting away from the details of any one particular approach, we can think of all of these models as proposing that we generate the 'alternative possibilities' fundamental to modal cognition by (i) delimiting a task-relevant partition within the vast space of conceivable possibilities, (ii) considering a smaller subset of particular possibilities within the relevant part of that partition, and then (ii) ordering or evaluating them in task-relevant ways to inform our final modal judgments (Figure 1A, Key Figure). Notably, there has been much research indicating how we may construct task-specific partitions and orderings, but there is relatively little work addressing how, before explicit evaluation, specific options are sampled from within the relevant part of the partition. This understudied part of the architecture is our focus.

Possible Actions: A Shared Adaptive Sampling Proposal

Unlike much of this previous work, our proposal focuses specifically on representations of possible actions, a particularly important type of possibility representation that plays a crucial role in high-level cognition. We build on the groundwork laid by traditional approaches in two key ways. First, our proposal gives crucial new structure to 'step 2' – the process by which particular possibilities within the relevant part of the partition are generated. Such spaces often include vast numbers of potentially relevant actions (e.g., 'possible ways to get to an airport', 'gifts one could buy for Christmas', 'ways to spend an afternoon', etc.), but ordinarily we only have the time and resources to explicitly evaluate a few. We propose that our minds adaptively sample a small set of possible actions from within the relevant space: those of probable practical utility are prioritized for consideration.

Second, although the traditional model assumes that the important work (partitioning and ordering) is task-specific (Figure 1B), our model proposes that a key intermediate step (sampling) is not. In principle, different adaptive sampling procedures might be used in different tasks (i.e., one for deciding what people should do, another for predicting their actions, and so on; Figure 1C); however, we

Glossary

Modal auxiliaries: in natural language, terms such as 'must', 'can', 'may', 'ought', etc. are called modal auxiliaries. They are used to make statements about non-actual possibilities and vary in primarily two ways: (i) the 'force' of the modal – whether they suggest that something is necessary (e.g., 'must') or merely possible (e.g., 'can'), and (ii) the 'flavor' of the modal – whether they concern what is good or right, and thus have a 'deontic flavor' (e.g., 'ought'), or what is known, and thus have an 'epistemic flavor' (e.g., 'might'), or some other flavor instead [6].

Modal cognition/modality: modal concepts allow us to represent and reason about sets of non-actual possibilities, for instance by grouping various 'possible worlds' in different ways. Modal statements often have to do with necessity or possibility: which things must be the case and which things could be, which are typically taken to correspond to the possibility and necessity operators in modal logic [115].

Ordering: possible worlds can be 'ordered'. For instance, some are considered to be 'closer' to the real world (e.g., worlds with no newspapers) and others are 'farther' (e.g., worlds with no Earth). Ordering is useful because it allows us to define sets based on points in the resulting order. For instance, some threshold on a probability ordering might separate 'relevant' worlds from irrelevant ones. Possibilities can be ordered in many different ways. We may, for example, order on 'ethicality': worlds with murder are worse than worlds without it, and various points along this ordering might define the set of worlds judged to be 'impermissible', 'permissible', or 'best'.

Possible worlds: philosophers and linguists often describe alternative possibilities in terms of 'possible worlds'. In the actual world – the one we are in – everything is a particular way. However, we can consider other ways the world might be. Each of these is a 'possible world' – for instance, the world in which everything is the same, except that you are reading this paper propose that the adaptive sampling process is shared across diverse tasks (Figure 1D). This leads us to predict patterns of convergence among diverse forms of modal cognition.

This three-step process – task-specific partitioning, shared adaptive sampling to construct an initial set of possible actions, and finally, flexible and task-specific reasoning about the possibilities in that set – forms the heart of our proposal.

Our proposal generates its most unique predictions when people are engaged in quick thinking. Naturally, the number of possible actions we can sample increases with time (Box 1). More uniquely, when time or effort is highly limited, our proposal predicts that we will 'default' to reasoning about systematically constrained types of possible actions. Specifically, because the set of possible actions we consider will be heavily influenced by the shared sampling process (rather than by the task-specific ordering process), features of the sampling process will be inherited by any downstream cognitive processes that depend on representations of possible actions. As time increases, however, larger numbers of possible actions can be sampled. This larger, less-restricted set allows increased downstream differentiation in the subset of possible actions evaluated as being relevant based on task-specific ordering. At the limit of infinite decision time, our model converges to traditional models on which all possible actions in the relevant part of the partition are reasoned over, because the adaptive sampling process will no longer influence which particular possible actions are judged to be most relevant for a given task. In short, the shared adaptive sampling process makes unique predictions about default modal cognition, while also accommodating existing evidence about deliberative modal cognition. We next review evidence suggesting that default modal cognition – the type strongly influenced by shared sampling – predominates across diverse kinds of thought and behavior.

Default Representations of Possible Actions

Our model is inspired by recent studies in diverse areas of cognitive science, and which all point to a remarkably consistent structure for default representations of possible actions: (i) judgments of possibility and spontaneous ideation, (ii) causal attribution, (iii) formal semantics and psycholinguistics, (iv) moral judgment and free will, and others. In each case, two factors have an influence on the default representations of possible actions: probability and **value** [22–32]. This fingerprint of shared adaptive sampling allows us to track its promiscuous role in cognition, and provides vital clues about its ultimate function.

The simplest way to discover which possible actions are initially sampled is to ask, for example – what is the amount of television watching per day that first comes to mind? Across 40 different ordinary actions, the possibilities that first come to our minds (the 'defaults') reflect two inputs: beliefs about which amounts are most probable, and beliefs about which amounts are ideal (i.e., valuable) [22]. For example, the average amount of TV watching that came to mind was 2.87 h/day, far below the actual amount [33], but between the amount participants thought was most probable (3.38 h/day) and the amount they thought was ideal (1.63 h/day).

According to the shared adaptive sampling model, the contours of the sampling process should most strongly constrain what actions we consider to be possible when decision time is short, and therefore a limited number of possibilities can be sampled (Box 1). In a recent study, people judged whether various actions are 'possible' under time pressure versus delay [26]. Participants read about people in difficult situations (e.g., Josh's car breaks down on his way to the airport). As different potential actions were proposed, participants judged whether each was 'possible' or 'impossible'. Some solutions were physically impossible (transporting to the airport), others were very typical (calling a friend), and, crucially, others had a low moral value (stealing a car) or were improbable (getting a ride from a stranger). Under time pressure, participants were especially likely to judge it 'impossible' to perform low-value actions, especially low moral value actions.

The same signature properties of default modal representation also arise early in development. Children say, for example, that it is impossible to paint polka-dots on an airplane, and that this would

tomorrow instead of today. There are infinitely many such possible worlds [116].

Reinforcement learning: a computational framework for learning and decision making that is widely used in computer science, neuroscience, and psychology. Reinforcement learning methods formalize how an agent can learn to maximize reward in their environment by acting within it. Reinforcement learning methods vary in how value is assigned to actions. One common method is to assign values based on an explicit model of the environment: the value of an action is determined by the reward it will achieve in the model of the environment. By contrast, model-free methods assign values to actions from the direct experience of rewards from previously doing those actions, and do not involve explicitly modeling the environment.

Value: many human behaviors are performed 'instrumentally'; that is, to maximize some reward or reinforcement signal, such as pleasure. The 'value' of an action is the expected total quantity of reward that it will obtain in the long term. Much research suggests that humans often compute the value of different actions and then tend to choose the actions that are associated with the highest value during decision making [107,117].

CellPress REVIEWS

Key Figure

Schematic Models of Modal Cognition

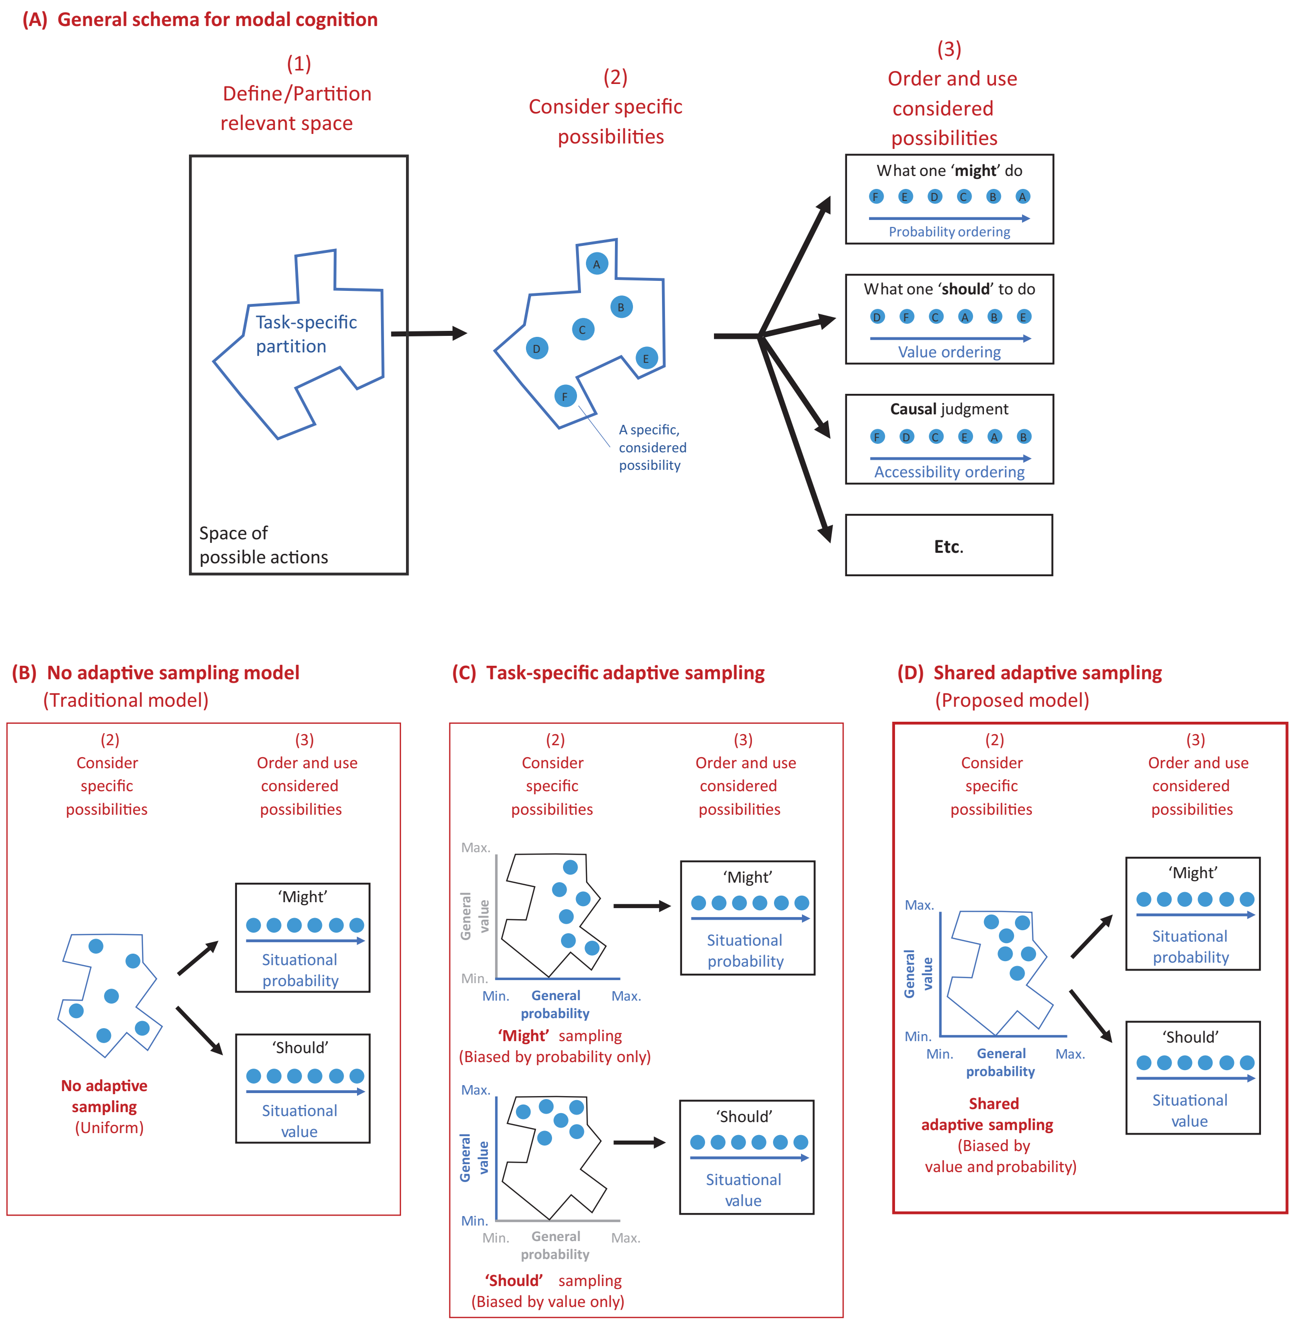

Trends in Cognitive Sciences

(Figure legend at the bottom of the next page.)

require magic [31,34]. Only later in development do they reflectively judge that, although improbable, such actions are possible. Crucially, recent work also shows that children's judgments of possibility are constrained by value, and especially by moral value [28,29]. For instance, young children often judge that it is impossible (or would require magic) for someone to steal a candy bar, or to lie. At 3–5 years of age, these moral violations are considered to be equally impossible as violations of physical laws (e.g., turning your hat into a candy bar) [29].

In line with the adaptive sampling model, this work demonstrates that default representations of possible actions clearly reflect the contours of the adaptive sampling process: what is possible reflects what is likely to have practical utility, even when practical utility is not relevant for the task at hand. We next review evidence that this same 'fingerprint' of default modal cognition arises across diverse cognitive domains.

Causal Reasoning

Humans have a basic motive to describe, explain, and predict the world using causal relations. When a house burns down, it is not possible to avoid searching for the cause. This ability has been argued to rely on modal cognition: we identify whether something was a cause in part by considering the alternative possibilities in which that action did not occur (e.g., [21,35–38]; but see [39] for an opposing perspective). For example, if one wants to know whether the fire was caused by a stove being left on, one asks what would have happened in the alternative possibility where the stove was not left on. Several recent studies identify which types of counterfactual possibilities about actions come to mind during causal reasoning. When there is a house fire, we tend not to consider the possibility that a person might never have purchased a stove, or the possibility that they could have built their house entirely out of fireproof material. Instead, we focus on counterfactuals where not doing the action was relatively probable: not leaving the stove on. Put differently, we tend to say that leaving the stove on caused the fire because the relevant alternative action – not doing that – was relatively probable.

This feature is present both in how we think about events and in agents' actions [7,36,40,41]. In a similar way, people tend to focus on counterfactual possibilities involving actions that have a relatively high value. Consequently, they also tend to select lower-value events as causes [32,36,40,42–45]. Box 2 provides a detailed overview of this evidence. In brief, causal judgments depend on reasoning about alternative possibilities, and the alternative possible actions that come to mind tend to be probable and valuable [7,36,38,40,42–45]. This is, of course, the signature of default modal cognition.

Modality in Natural Language

We often use language to specify alternative actions, either directly or indirectly. According to standard linguistic theory, a statement such as 'Angelika might be lecturing in London' means, roughly, that in all the possibilities that are relevant given what we know about Angelika, they involve her lecturing in London [6,46]. Linguists call words like might 'modals' (or **modal auxiliaries**) precisely because they make claims about what is possible or necessary. However, what governs the possible actions that we consider when communicating?

Figure 1. Much of the previous work on modal cognition can be fit into the model of information processing depicted in (A). This three-step process involves (1) delimiting a task-relevant partition within the space of conceivable possibilities, (2) considering a small number of particular possibilities within the relevant part of the partitioned space, and (3) ordering or evaluating these possibilities in task-relevant ways. Our proposal, which focuses specifically on reasoning about possible actions, expands on the second step of this process. Unlike traditional models, we propose that people sample a small set of possibilities in an adaptive way. In particular, we propose that the sampling mechanism exhibits a bias towards possible actions of high practical utility: those that are both likely to be done and typically valuable when done (contrast B and D). In addition, our model proposes that this adaptive sampling mechanism is shared across a diversity of modal tasks. Instead of exhibiting a bias towards possibilities that are likely to be relevant given the specific task faced (e.g., only the probability that an action will be done when judging what someone might do, or only the value of an action when judging what someone should do), the biased sampling mechanism is relatively insensitive to which features are relevant for the modal task being performed (contrast C and D). Accordingly, people exhibit a bias towards sampling actions of high practical utility regardless of the task being faced, and thus these possibilities disproportionately inform a wide range of modal judgments.

CellPress REVIEWS

Box 1. From Default to Reflective Modal Cognition

In contrast to traditional models, the shared adaptive sampling hypothesis makes unique predictions about the effect of time on judgments that involve modal cognition about possible actions. When people are given little time to consider possibilities before responding, the set of possible actions they consider should be relatively small and heavily influenced by the adaptive sampling process used to generate them. Because this representation is then shared across different types of modal judgments, the less time one has, the more these different modal judgments will look similar to one another. By contrast, when people are given more time to respond, various modal judgments should look comparatively dissimilar to one another because these judgments will now largely reflect task-specific constraints on how the possible actions are 'ordered' or reasoned about (Figure I).

By analogy, consider how consumer choice is shaped by a website that 'samples' products for new customers to view. Suppose that many first-time customers with highly different preferences are choosing which of 100 sweaters to buy. The website presents a 'sample' of the available sweaters to each customer, hoping that the customer will buy one of these sweaters; each customer then chooses one sweater from the set presented. If the website only shows each customer a similar set of three or four sweaters, then all customers will end up choosing similar sweaters regardless of their different individual preferences. However, as the number of sweaters 'sampled' by the website increases, the choices of the customers increasingly reflect their individual preferences, rather than the sampling process used by the website.

A recent study supports this prediction by demonstrating an effect of time pressure on modal judgments (e.g., judgments of what someone 'should' or 'could' do) [26]. When participants were given time to reflect before responding, their judgments were clearly distinct – what one 'could' do bears little resemblance to what one 'should' do. However, when participants were made to respond quickly, all the different modal judgments began to look increasingly similar to one another (Figure II). This pattern is particularly striking given that a simple increase in noise under time pressure predicts the opposite pattern (less similarity when answering quickly).

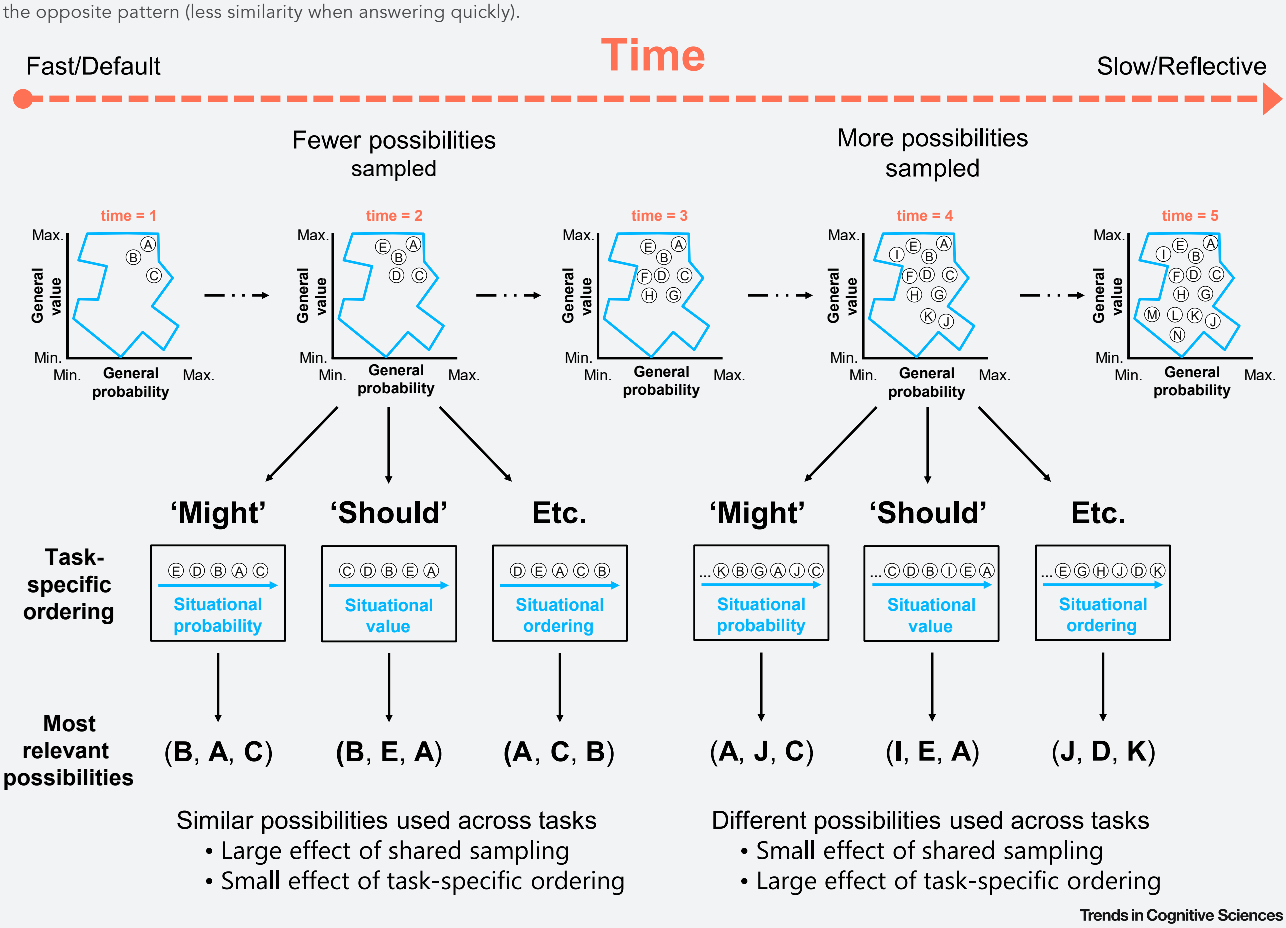

Figure I. Effect of Time on Shared Adaptive Sampling.

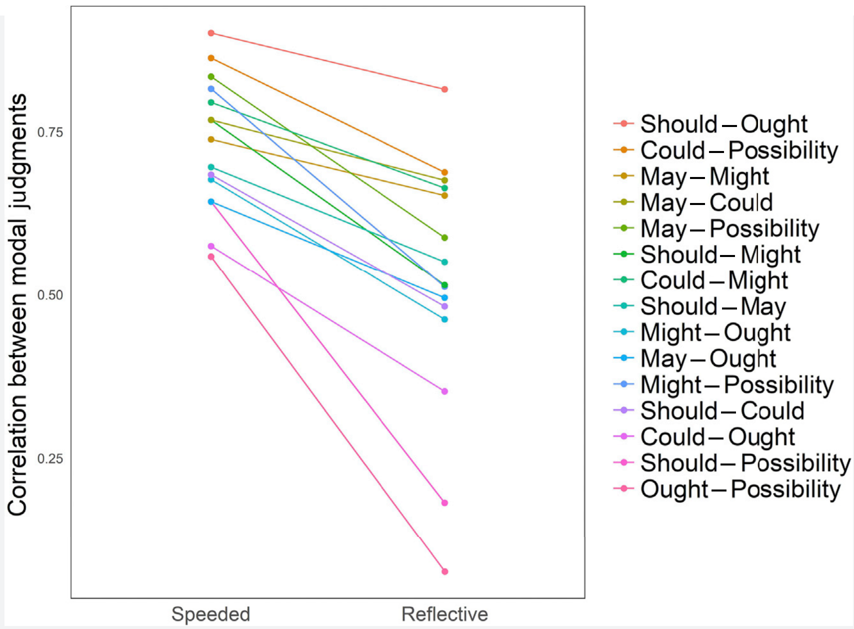

Figure II. Effect of Time on the Similarity of Modal Judgments.

Recent psycholinguistic research shows that modal auxiliaries default to picking out possibilities involving actions that are probable and valuable [26,27,47]. In one study, participants read about a student who forgot her homework. They then made truth-value judgments of sentences such as 'Mary ought to forge an excuse note', or 'Mary might forge an excuse note', or 'Mary could forge an excuse note', and so on. Some participants were instructed to do this task very quickly, increasing their reliance on default modal representations, whereas others were allowed to reflect before responding. When participants were forced to answer quickly, all their different modal judgments began to look increasingly similar to one another (see Figure II in Box 1). This suggests that participants were relying on some common default set of available possibilities across diverse modal judgments. Crucially, this default representation was highly constrained by value: the speeded judgments of all modal auxiliaries resemble our typical judgments of what a person 'ought' to do (e.g., that you 'ought' not to forge an excuse note) [26].

This same feature characterizes young children's linguistic development. Children first learn modals such as 'can' and 'could' [48–50] that can be used to make claims both about the set of possible actions that are likely to occur (e.g., 'John could be eating pizza for dinner') and the set of possible actions that do not have a low value (e.g., 'you can drive at 50 mph here'), and only later come to learn the meaning of modal terms with only one of these constraints, such as 'might' [51–53]. Moreover, children often interpret all modal questions as being about a combination of what would be good to do and which things are likely to be actually done, even when this interpretation is inappropriate to adult speakers. For example, children will say others 'can't' do actions that are unlikely to work, would hurt others, or would violate their own preferences [54–56].

Finally, roughly half of languages worldwide have been estimated to contain modal forms that are sensitive to both probability and value [57,58]. The widespread nature of this pattern suggests that this feature may arise from a universally shared default [59–61].

Box 2. The Role of Value in Causal Judgment

Value strongly influences causal judgment. Consider this scenario: the receptionist in the philosophy department keeps her desk stocked with pens. The administrative assistants are allowed to take pens, but faculty members must buy their own. The administrative assistants typically take the pens. Unfortunately, so do the faculty members. The receptionist has repeatedly emailed them reminders that only administrators are allowed to take the pens.

On Monday morning, both Professor Smith and an administrative assistant take pens at the same time. Later that day, the receptionist needs to take an important message, but she has a problem. There are no pens left.

Who caused the secretary's problem? Most people select Professor Smith, not the administrator [44]. This effect – that actions with low moral value are judged more causal – is highly general and has been documented across a wide range of stimuli (from vignette studies [36,42–44] to visual displays [21,68]), methods (from survey questions to eye-tracking [21,41]), and populations (from adults to children [21,36,41–45,68]).

Evidence suggests that this pattern arises from the structure of default modal cognition. People judge an action (taking a pen) to have caused an effect (the secretary's problem) by considering whether the effect would not have occurred if the potential cause had been different [35–38]. To do this, participants must call to mind counterfactual possibilities ('what else could have happened?'). Crucially, because we naturally call to mind higher-value possible actions, our minds turn more readily to the correction of an immoral act than to the alteration of a permissible one; people find themselves thinking: 'if only Professor Smith had followed the rules and had not taken a pen ...'. This, in turn, makes the necessity of the immoral action for the outcome more salient [27,32].

Three further findings support this account. First, the pattern is eliminated if the professor's action does not have a low value [42]. Second, the pattern is eliminated if the low-value action was not necessary for the effect to occur, implicating an assessment of counterfactual dependence [36,42–44]. Third, the effect is mediated by judgments of how relevant each counterfactual is: people judge the counterfactual where the professor did not take the pen to be more relevant than the one where the administrator did not take the pen [32,43]. These results suggest that our reasoning about possible actions is sensitive to value, and provide a case study of default modal reasoning in high-level cognition.

Freedom and Responsibility

There is a great difference between handing a robber your bank codes unprovoked and doing so with a gun to your head. The difference, of course, resides in the nature of the alternative possible actions [62,63]. In one case you have a much better alternative (say nothing!), whereas in the second you only have a much worse alternative (be shot). Accordingly, we tend to say that only the person with a gun to their head was 'forced' to give their bank codes, that they 'had no choice', and that they were not responsible for doing so. Empirical work on judgments of freedom and responsibility has confirmed this general relationship: people tend to say that a person was not fully responsible for their action if the only alternative actions were of low value [20,26,32,64–68], consistent with the possibility that the alternatives were never represented as 'possibilities' at all. These judgments thus bear part of the fingerprint of default modal cognition. In contrast to value, the role of probability in freedom/responsibility has not been directly tested and remains an important direction for future research.

Summary

Across the examples of modal cognition that we have reviewed, the possible actions that people consider by default are probable and valuable. This fingerprint of adaptive sampling also emerges throughout a diversity of other cognitive domains: in judgments of intentional action [69–71], in how people engage in counterfactual reasoning [13,20,72–75], in judgments of normality [76], in people's predictions of the future actions of others [77–79], and more [32].

Traditional models of modal cognition have difficulty in explaining this consistency. First, they do not differentiate between the structure of default and reflective modal cognition; second, they cannot explain why a common set of factors would reliably shape default modal cognition across diverse tasks. The shared adaptive sampling model explains these data by positing that default modal cognition is dominated by a restricted sampling process, and that this process is shared across diverse tasks. What remains to be explained is why representations of possibility exhibit a sensitivity to value and probability.

A Functional Proposal for Default Modal Cognition

When considering possible actions, why do humans think of valuable and probable actions by default? Although these constraints are puzzling in many of the tasks we have reviewed, they are perfectly natural in at least one context: decision making. We want to make good decisions, and hence we should direct our limited capacity for careful thought towards actions with the highest practical utility. Pursuing this basic idea, we argue that a promising possibility is that default modal cognition about actions (representing 'what is possible') derives its distinctive structure from its role in decision making (choosing 'what to do').

The Role of Value and Probability in Generating Consideration Sets

Similarly to other domains, decision-making depends crucially on which actions are represented as possible. Reasoning about actions takes effort: we cannot exhaustively consider all the actions that, on reflection, we would judge to be technically possible. Consider the problem of choosing which movie to watch, and begin enumerating the possibilities alphabetically: A-Team, Aaron Loves Angela, ABBA, Ace Ventura, etc. There are too many movies to consider, and most are bad. The same problem arises in artificial intelligence (AI) approaches to games such as chess and Go: not all legal moves can be fully evaluated in a practical amount of time.

In research on decision making, people have been found to begin by sampling a small set of options worth considering [80–84], often called a 'consideration set' [85–90]. Sampling a consideration set is a crucial step in the choice process; in one analysis of consumer purchasing decisions, the content of people's consideration sets explained nearly 80% of the variation in their final choice [85,91]. Nevertheless, despite their importance, the sampling processes that generate consideration sets have received relatively little attention [82,83].

Existing studies do, however, consistently find that the first options which come to mind tend to be higher in value, and are chosen more often (i.e., are more probable) [80–83,88–90]. In [92], chess experts were shown board positions and were asked to name the first moves that came to mind. The initial moves that came to mind were much better, and much more commonly chosen, than chance moves. Indeed, the first move that chess grandmasters consider is often the best move [92]. Similar results have been found with handball players [81,93], firefighters [94], children [95], and consumer purchasing decisions [85–90]. In real-world decisions, the initial sampling process is so skewed towards high-value, high-probability actions that people often end up simply choosing the first possibility that comes to mind [80,81]. In addition, when faced with decisions, patients with damage to the ventromedial prefrontal cortex – a region implicated in storing and computing option values [96] – produce fewer options, suggesting that value computations are important for sampling possible actions [97]. Finally, these results accord with work showing that sampling in other parts of the decision process – for example, sampling of possible outcomes – is influenced by both probability and value [98–100].

Notably, cutting-edge AI systems have converged on a similar architecture. For AI, as for humans, complex decision problems often involve too many choices for exhaustive evaluation. In games such as chess and Go, for instance, it is computationally prohibitive to evaluate every legal sequence of moves. Instead, planning algorithms must select a small subset of sequences to evaluate. When done probabilistically, this is often called a 'Monte Carlo tree search' (MCTS), which refers to the process of selecting branches of a decision tree for evaluation [101]. Notably, dominant MCTS methods focus evaluation on actions that are associated with higher heuristic value estimates. For instance, the popular Upper Confidence bounds applied to Trees (UCT) algorithm selects actions for evaluation based on (i) the current value estimate of an action, and (ii) the amount of evaluation already devoted to that action [101]. The newer application of 'deep Reinforcement Learning' to Go, as demonstrated by AlphaGo [102] and AlphaZero [103], depended on extensions of UCT. Crucially, in both cases, the selection of actions for rigorous evaluation was guided broadly by estimates of value and probability.

Why Two Stages?

If people can bias consideration sets towards higher-value options, and if the first thing that comes to mind is often of such a high value that they choose it, then why bother ever thinking of more than one

Box 3. Adaptive Consideration Sets in Decision Making

We present a simple illustration of how value-guided consideration sets could be adaptive in decision making (Figure I). This proposal is closely related to that in [113], which provides an algorithmic implementation of our computational proposal for a specific decision problem. When facing a decision with N options, people first generate a consideration set, C, of $K << N$ options, where the probability of inclusion, i, of each option in C is proportional to the 'cached value' of that option [denoted $CV(i)$], which correlates imperfectly with the expected value of the option in the specific decision context. This sampling process produces higher-value, higher-probability options for consideration, and can be accomplished efficiently using, for instance, the algorithm described in [114]. Second, people compute the true expected value of each option in the consideration set [denoted $EV(j)$], and choose the option with the highest EV. We assume that the computation of EV incurs a significant computational cost that scales linearly with the number of evaluations.

Building on this approach, an optimality analysis was conducted [113], asking when it is optimal to sequence heuristic and deliberative processes (a 'hybrid' model) or instead to act either purely heuristically or purely deliberatively (two 'pure' models). Exploring a wide range of parameter settings, they find that the hybrid approach is often favored, pure habit is sometimes favored, and pure deliberation is rarely favored. The main advantage of the hybrid approach is to screen off costly deliberation over the large majority of actions that are almost certainly poor choices by a computationally frugal method, while investing deliberation in improving the final choice among a few good candidates.

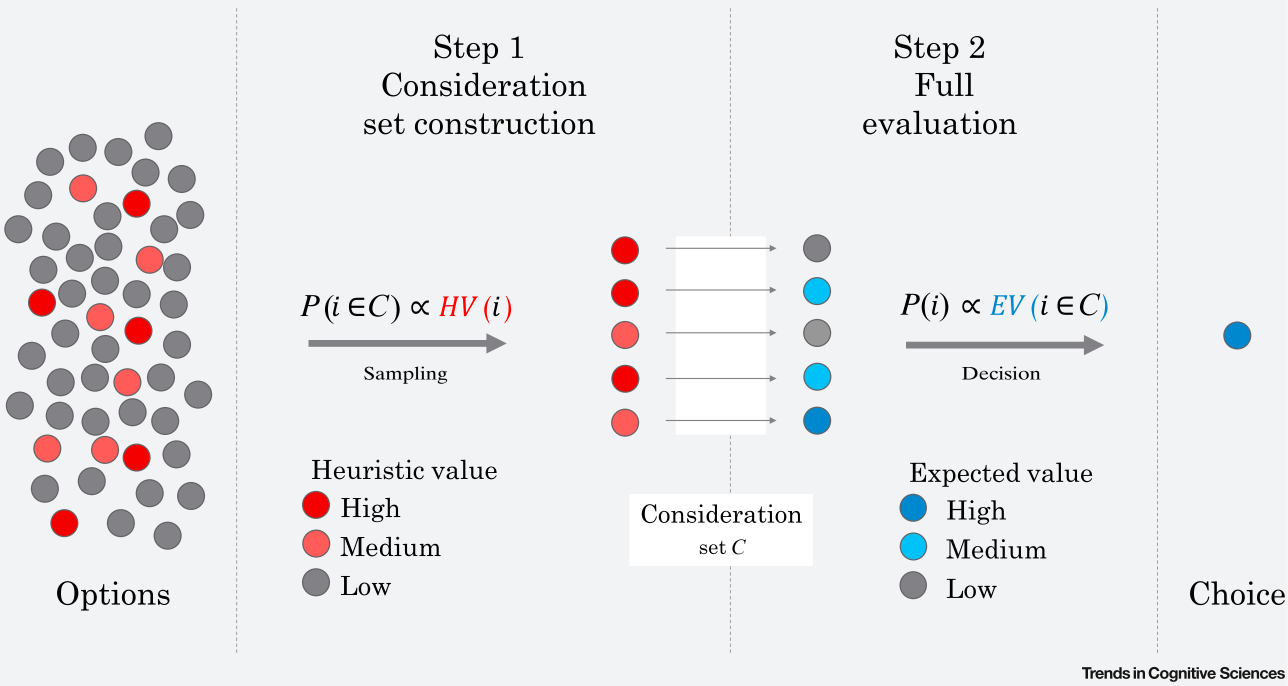

Figure I. Schematic Depiction of the Two-Stage Choice Architecture.

thing during decision making? Why would our decision-making process involve two stages – sampling a consideration set, and then choosing a decision from it?

This two-stage architecture may reflect a basic distinction between two different ways in which we can assign value: one that is computationally cheap but less accurate, and another that is computationally expensive but more accurate [104–106]. Building on prior suggestions [83,89], we propose that people rely on the 'cheap/inaccurate' system to generate a consideration set of reasonable candidates, and then switch to the 'expensive/accurate' system to choose the very best option from within the set generated. Box 3 provides a discussion of simulations that demonstrate the adaptiveness of such a mechanism.

More specifically, we propose that the type of value representation used for sampling possible actions is not computed at decision time; instead, it is cached (i.e., 'precompiled') from previous experience [83,89,104–106]. According to contemporary models of decision making (in particular, model-free **reinforcement learning** [107]), people cache the values of actions they experience; these cached values are then available automatically in future decisions, and influence choices above and beyond people's explicit calculations of the expected value of an action in the current context [106,108]. Moreover, people also appear to learn cached values through observation [109,110] and instruction [111,112], allowing them to acquire cached values for actions that they have not personally experienced. These cached values could plausibly be used to quickly sample higher-value possibilities for further consideration. Later, more precise value estimates can be obtained by subjecting each member of the consideration set to model-based reinforcement learning – a computationally expensive method that uses more information to estimate the precise, situation-specific value to candidate actions.

In sum, consideration sets could plausibly be constructed according to a readily available, computationally cheap representation of the value of an action 'in general', whereas actions are chosen from the consideration by computing online a representation of the value of an action 'specifically' for the present circumstance [83,89]. Initial evidence indeed suggests both a crucial role for cached values in the generation of the consideration set, and a separate role for task-specific values in the choice of a final option from the consideration set [113]. For instance, when choosing dinner for a guest with unusual tastes, people tend to generate potential dinners by sampling options that they themselves typically enjoy (relying on cached value representations), and then choose from within the options in this consideration set the particular dinner that best satisfy their guest's tastes (a task-specific value representation) [113].

From Choice to Possibility

We propose that modal cognition about human action – that is employed during causal reasoning, language interpretation, moral judgment, etc. – relies on a similar mechanism for consideration set construction (Figure 2). In other words, that the same value- and probability-based sampling process that identifies practical candidates for choice is at the heart of shared adaptive sampling when we reason about possible actions that we ourselves are not doing.

Of course, showing that the tendency to sample higher-value, higher-probability possible actions could be adaptive for decision-makers does not prove that this tendency actually originated as a decision-making adaptation. One relevant consideration is that sensitivity to value and probability is not obviously beneficial in many of the tasks where modal cognition is involved (e.g., judgments of magic, or predictions of the actions of others), which suggests that the observed similarity is not merely a coincidence or the product of convergent evolution. Nonetheless, future research should test this claim more directly: are these processes linked developmentally, neurally, and across individual differences?

Importantly, our model predicts that the shared adaptive sampling process used in modal cognition will show a signature effect of employing value and probability 'in general', rather than the values and probabilities specific to the particular situation in question. For instance, when making causal judgments by considering counterfactuals, the counterfactual actions sampled will be those that are 'generally' valuable and probable, whereas the subsequent ordering of counterfactual relevance will instead be situation-specific (we may consider a counterfactual that is typically good and likely, but not allow it to inform our judgments because it is not actually possible in the specific situation at hand). These remain important topics for further research.

Concluding Remarks

Three elements of our proposal stand out for further development (see also Outstanding Questions). First, although our proposal was framed at the computational level, it makes clear predictions that are so far untested. For example, although most of the literature on modal cognition has focused on the role of moral value in particular, our model suggests that (i) this is only a special case of effects of value more generally, and thus other types of value (e.g., prudential value) will give rise to similar effects, (ii)

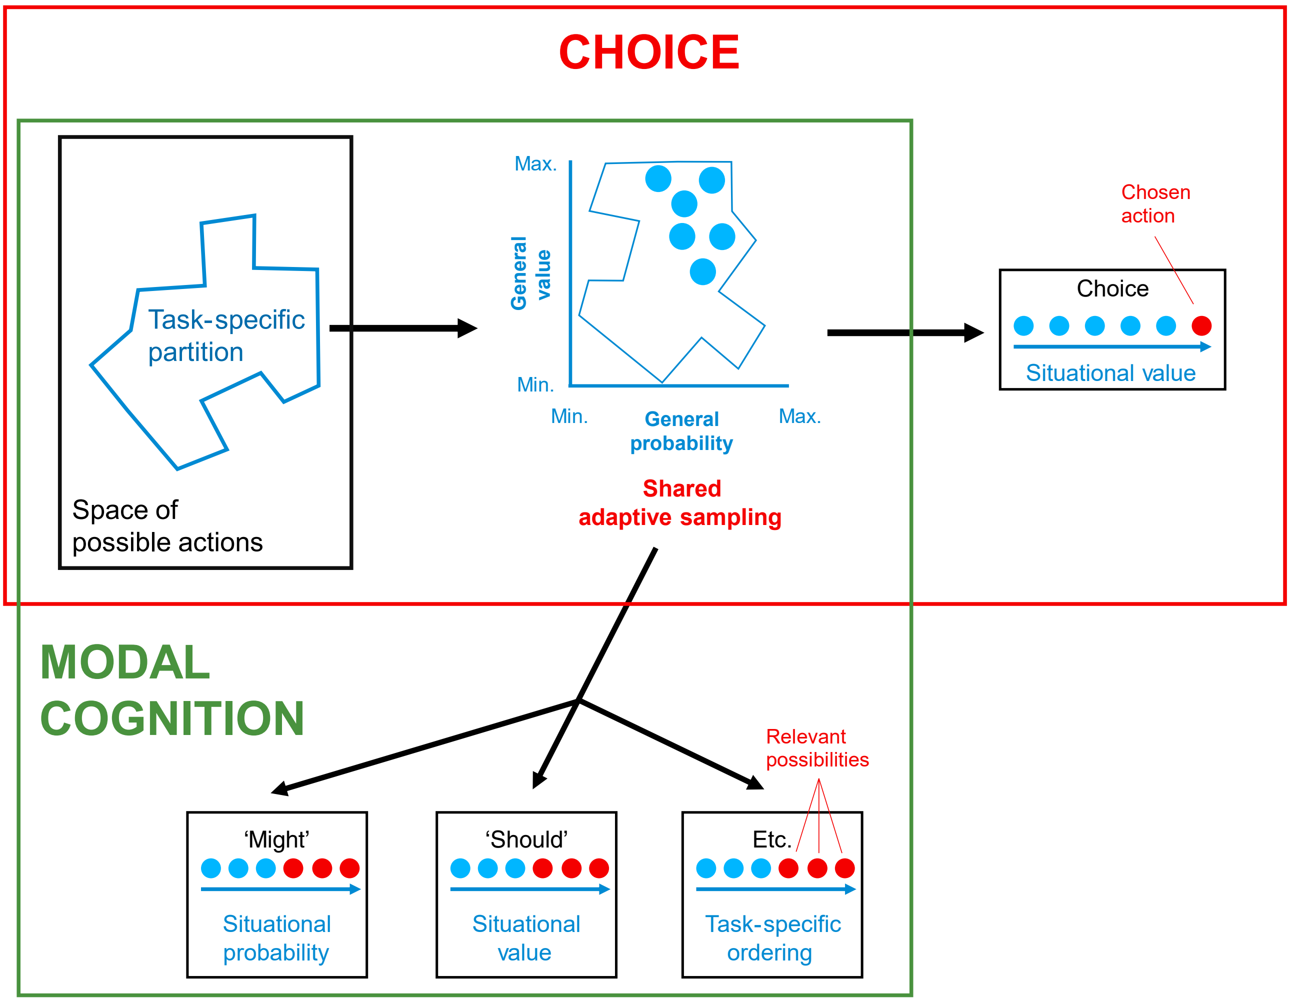

Trends in Cognitive Sciences

Figure 2. Relationship between Choice and Modal Cognition
Schematic illustration of the proposed shared adaptive sampling hypothesis in choice contexts (red box) and its shared recruitment across diverse cognitive tasks involving modal thought (green box).

this value is encoded in a heuristic, context-free way rather than being computed online in a task-specific way, (iii) the effect of this value will be strikingly similar across different types of modal cognition (causal reasoning, psycholinguistics, judgments of responsibility, etc.), and (iv) the effect of such value will increase when the number of possibilities that can be considered is limited, whether through time pressure, cognitive load, or individual differences.

Second, although our proposal focuses on the representation of possible actions, actions are only one of the many types of possibilities we consider. We also consider possible events, states of the world, objects, and so on. Here again, we often face intractably large spaces of possibilities. It remains an important question whether the processes of sampling possibilities in these other types of modal cognition are structured similarly to the shared adaptive sampling mechanism we have proposed for actions. These may differ in important respects. For instance, when choosing actions, it is clearly beneficial to consider only the good ones, but when anticipating events beyond our control, it may make more sense to focus on their 'importance', irrespective of whether they are good or bad [98].

Third, although we have focused on how possible actions are generated for consideration, there are presumably similar mechanisms that serve to reject possibilities from consideration. Indeed, data from experiments that involve judging the possibility of proposed actions may be most naturally accounted for by mechanisms of heuristic rejection rather than generation. Further research should establish whether probability and value structure the rejection of possibilities from a consideration set in the same way as they structure their introduction.

In sum, the cognitive sciences have long recognized that many of human's most impressive forms of thinking require us to reason about 'possible worlds' in addition to the actual one. Nonetheless, it is obvious that humans cannot and do not explicitly consider all of them. We have argued that an essential form of modal representation is the set of actions, both likely and good, that arise for consideration by default. Our general capacity to generate possible actions may be intimately linked to our capacity for decision making. Thus, a core component of ordinary modal cognition may not be possible worlds, but practical worlds – those involving actions that are typically worth considering.

Outstanding Questions

We have considered how people focus on only a few of the many possible actions they may take. However, each action comprises many different features. Do sampling processes focus our attention on only the most relevant features?

How do individual differences in subjective value affect the representation of possibilities? For example, if individuals with psychopathy do not consider harmful behavior 'low value', are they more likely to consider it possible by default?

Humans value some things for personal reasons and others for social reasons; we can also determine how good something is relative to its class (e.g., a good robber). How do we integrate these various types of value, especially when they conflict?

In AI applications of reinforcement learning, the values of options are often learned by extensive simulation, or 'self-play' in games. Do humans learn the value of possibilities in a similar way? What other forms of value learning – such as cultural transmission – are uniquely available to humans?

Acknowledgments

We would like to thank participants in the exploratory seminar 'What is good and what is possible: Searching for an interdisciplinary language' funded by the Radcliffe Institute for Advanced Study. This research was supported by grant N00014-19-1-2025 from the Office of Naval Research and grant 61061 from the John Templeton Foundation.

References

1. Marr, D. (1982) *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*, MIT Press
2. von Wright, G.H. (1953) An essay in modal logic. *Philos. Q.* 3, 287
3. Kripke, S.A. (1963) Semantical considerations on modal logic. *Acta Philosophica Fennica* 16, 83–94
4. Lewis, D. (1981) Ordering semantics and premise semantics for counterfactuals. *J. Philos. Logic* 10, 217–234
5. Stalnaker, R.C. (1968) A theory of conditionals. In *Ifs: Conditionals, Belief, Decision, Chance, and Time* (W.L. Harper *et al.* eds), pp. 41–55, Springer
6. Kratzer, A. (2012) *Modals and Conditionals: New and Revised Perspectives*, Oxford University Press
7. Halpern, J.Y., and Hitchcock, C. (2014) Graded causation and defaults. *Brit. J. Philos. Sci.* 66, 413–457
8. Lassiter, D. (2011) *Measurement and Modality: The Scalar Basis of Modal Semantics* (PhD Dissertation), New York University. https://semanticsarchive.net/Archive/WMzOWU2O/Lassiter-diss-Measurement-Modality.pdf
9. McCoy, J. *et al.* (2019) Modal prospection. In *Metaphysics and Cognitive Science* (A. Goldman and B. McLaughlin eds), pp. 235–267, Oxford University Press
10. Johnson-Laird, P.N. (1980) Mental models in cognitive science. *Cogn. Sci.* 4, 71–115
11. Hinterecker, T. *et al.* (2016) Modality, probability, and mental models. *J. Exp. Psychol. Learn. Mem. Cogn.* 42, 1606
12. Khemlani, S.S. *et al.* (2018) Facts and possibilities: a model-based theory of sentential reasoning. *Cogn. Sci.* 42, 1887–1924
13. Byrne, R.M. (2016) Counterfactual thought: from conditional reasoning to moral judgment. *Annu. Rev. Psychol.* 67, 133–157
14. Stanley, M.L. *et al.* (2017) Counterfactual plausibility and comparative similarity. *Cogn. Sci.* 41, 1216–1228
15. Shtulman, A., and Tong, L. (2013) Cognitive parallels between moral judgment and modal judgment. *Psychon. Bull. Rev.* 20, 1327–1335
16. De Brigard, F. *et al.* (2013) Remembering what could have happened: neural correlates of episodic counterfactual thinking. *Neuropsychologia* 51, 2401–2414
17. Schacter, D.L. *et al.* (2015) Episodic future thinking and episodic counterfactual thinking: intersections between memory and decisions. *Neurobiol. Learn. Mem.* 117, 14–21
18. Rafetseder, E. *et al.* (2010) Counterfactual reasoning: developing a sense of 'nearest possible world'. *Child Dev.* 81, 376–389
19. Ullman, T. *et al.* (2016) Coalescing the vapors of human experience into a viable and meaningful comprehension. In *Proceedings of the 38th Annual Meeting of the Cognitive Science Society* (A. Papafragou *et al.* eds), pp. 1493–1498, Cognitive Science Society
20. Byrne, R.M., and Timmons, S. (2018) Moral hindsight for good actions and the effects of imagined alternatives to reality. *Cognition* 178, 82–91
21. Gerstenberg, T. *et al.* (2017) Eye-tracking causality. *Psychol, Sci.* 28, 1731–1744
22. Bear, A. *et al.* (2020) What comes to mind? *Cognition* 194, 104057
23. Cesana-Arlotti, N. *et al.* (2018) Precursors of logical reasoning in preverbal human infants. *Science* 359, 1263–1266
24. Cesana-Arlotti, N. *et al.* (2012) The probable and the possible at 12 months: intuitive reasoning about the uncertain future. *Adv. Child Dev. Behav.* 43, 1–25
25. Téglás, E. *et al.* (2007) Intuitions of probabilities shape expectations about the future at 12 months and beyond. *Proc. Natl. Acad. Sci. U S A* 104, 19156–19159
26. Phillips, J., and Cushman, F. (2017) Morality constrains the default representation of what is possible. *Proc. Natl. Acad. Sci. U S A* 114, 4649–4654
27. Phillips, J., and Knobe, J. (2018) The psychological representation of modality. *Mind Lang.* 33, 65–94
28. Shtulman, A., and Phillips, J. (2018) Differentiating 'could' from 'should': developmental changes in modal cognition. *J. Exp. Child Psychol.* 165, 161–182

29. Phillips, J., and Bloom, P. (2017) Do children believe immoral events are possible? *OSFHOME*. Published online December 18, 2017. https://osf.io/en7ut/
30. Smith, K.A., and Vul, E. (2015) Prospective uncertainty: the range of possible futures in physical prediction. In *37th Annual Meeting of the Cognitive Science Society: Mind, Technology and Society* (D.C. Noelle *et al.* eds), pp. 2230–2235, Cognitive Science Society
31. Shtulman, A., and Carey, S. (2007) Improbable or impossible? how children reason about the possibility of extraordinary events. *Child Dev.* 78, 1015–1032
32. Phillips, J. *et al.* (2015) Unifying morality's influence on non-moral judgments: the relevance of alternative possibilities. *Cognition* 145, 30–42
33. Nielsen. (2016) *The Total Audience Report: Q1 2016*, Nielsen
34. Shtulman, A. (2009) The development of possibility judgment within and across domains. *Cogn. Dev.* 24, 293–309
35. Pearl, J. (2009) *Causality*, Cambridge University Press
36. Icard, T.F. *et al.* (2017) Normality and actual causal strength. *Cognition* 161, 80–93
37. Lewis, D. (1974) Causation. *J. Philos.* 70, 556–567
38. Morris, A. *et al.* (2019) Quantitative causal selection patterns in token causation. *PLoS One* 14, e0219704
39. Wolff, P. (2007) Representing causation. *J. Exp. Psychol. Gen.* 136, 82
40. Kominsky, J.F. *et al.* (2015) Causal superseding. *Cognition* 137, 196–209
41. Gerstenberg; T.; Icard; T. (2019) Expectations affect physical causation judgments. *J. Exp. Psychol. Gen.* Published online September 12, 2019. https://doi.org/10.1037/xge0000670
42. Hitchcock, C., and Knobe, J. (2009) Cause and norm. *J. Philos.* 106, 587–612
43. Kominsky, J., and Phillips, J. (2019) Immoral professors and malfunctioning tools: counterfactual relevance accounts explain the effect of norm violations on causal selection. *Cogn. Sci.* e12792
44. Knobe, J., and Fraser, B. (2008) Causal judgment and moral judgment: two experiments. *Moral Psychology* 2, 441–448
45. Samland, J. *et al.* (2016) The role of prescriptive norms and knowledge in children's and adults' causal selection. *J. Exp. Psychol. Gen.* 145, 125–130
46. Portner, P. (2009) *Modality*, Oxford University Press
47. Knobe, J., and Szabó, Z.G. (2013) Modals with a taste of the deontic. *Semant. Pragmat.* 6, 1–42
48. Shatz, M., and Wilcox, S. (1991) Constraints on the acquisition of English modals. In *Perspectives on Language and Thought: Interrelations in Development* (S. Gelman and J. Byrnes eds), pp. 319–353, Cambridge University Press
49. Papafragou, A. (1998) The acquisition of modality. *Mind Lang.* 13, 370–399
50. van Dooren, A. *et al.* (2017) Learning what must and can must and can mean. In *Proceedings of the 21st Amsterdam Colloquium* (A. Cremers *et al.* eds), pp. 225–234, Institute for Logic, Language and Computation
51. Shepherd, S.C. (1982) From deontic to epistemic: an analysis of modals in the history of English, creoles, and language acquisition. In *Papers from the Fifth International Conference on Historical Linguistics* (A. Ahlqvist ed), pp. 316–323, Benjamins
52. Cournane, A., and Pérez-Leroux, A.T. (2018) Leaving obligations behind: epistemic incrementation in preschool English. In *Linguistics: Conference on Language Development*, p. P6395, Boston University
53. Wells, G. (1979) Learning and using the auxiliary verb in English. In *Language Development* (V. Lee ed), pp. 250–270, Helm
54. Chernyak, N. *et al.* (2013) A comparison of American and Nepalese children's concepts of freedom of choice and social constraint. *Cogn. Sci.* 37, 1343–1355
55. Kalish, C. (1998) Reasons and causes: children's understanding of conformity to social rules and physical laws. *Child Dev.* 69, 706–720
56. Browne, C.A., and Woolley, J.D. (2004) Preschoolers' magical explanations for violations of physical, social, and mental laws. *J. Cogn. Dev.* 5, 239–260
57. Hacquard, V., and Cournane, A. (2015) Themes and variations in the expression of modality. In *NELS 46: Proceedings of the Forty-Sixth Annual Meeting of the North East Linguistics Society* (Vol. 2) (C. Hammerly and B. Prickett eds), pp. 21–42, Concordia University
58. van der Auwera, J., and Ammann, A. (2008) Overlap between situational and epistemic modal marking. In *The World Atlas of Language Structures Online* (M.S. Dryer and M. Haspelmath eds) Max Planck Institute for Evolutionary Anthropology
59. Strickland, B. *et al.* (2015) Event representations constrain the structure of language: sign language as a window into universally accessible linguistic biases. *Proc. Natl. Acad. Sci. U S A* 112, 5968–5973
60. Christiansen, M.H., and Chater, N. (2008) Language as shaped by the brain. *Behav. Brain Sci.* 31, 489–509
61. Culbertson, J. *et al.* (2013) Cognitive biases, linguistic universals, and constraint-based grammar learning. *Top. Cogn. Sci.* 5, 392–424
62. (R. Kane ed) (2011). The Oxford Handbook of Free Will Oxford University Press
63. Hart, H.L.A., and Honoré, T. (1985) *Causation in the Law*, Oxford University Press
64. Mandelkern, M., and Phillips, J. (2018) Sticky situations: 'force' and quantifier domains. In *Proceedings of the 28th Semantics and Linguistic Theory Conference* (S. Maspong *et al.* eds), pp. 474–492, Linguistic Society of America
65. Young, L., and Phillips, J. (2011) The paradox of moral focus. *Cognition* 119, 166–178
66. Phillips, J., and Knobe, J. (2009) Moral judgments and intuitions about freedom. *Psychol. Inq.* 20, 30–36
67. Chakroff, A., and Young, L. (2015) Harmful situations, impure people: an attribution asymmetry across moral domains. *Cognition* 136, 30–37
68. Gerstenberg, T. *et al.* (2018) Lucky or clever? From expectations to responsibility judgments. *Cognition* 177, 122–141
69. Halpern, J.Y., and Kleiman-Weiner, M. (2018) Towards formal definitions of blameworthiness, intention, and moral responsibility. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, 1853–1860* AAAI
70. Knobe, J. (2010) Person as scientist, person as moralist. *Behav. Brain Sci.* 33, 315–329
71. Uttich, K., and Lombrozo, T. (2010) Norms inform mental state ascriptions: a rational explanation for the side-effect effect. *Cognition* 116, 87–100
72. Byrne, R.M. (2017) Counterfactual thinking: from logic to morality. *Curr. Dir. Psychol. Sci.* 26, 314–322
73. Kahneman, D., and Miller, D. (1986) Norm theory: comparing reality to its alternatives. *Psychol. Rev.* 93, 136–153
74. Wells, G. *et al.* (1987) The undoing of scenarios. *J. Pers. Soc. Psychol.* 53, 421–430
75. N'gbala, A., and Branscombe, N. (1995) Mental simulation and causal attribution: when simulating

an event does not affect fault assignment. *J. Exp. Soc. Psychol.* 31, 139–162
76. Bear, A., and Knobe, J. (2017) Normality: part descriptive, part prescriptive. *Cognition* 167, 25–37
77. Critcher, C.R., and Dunning, D. (2013) Predicting persons' versus a person's goodness: behavioral forecasts diverge for individuals versus populations. *J. Pers. Soc. Psychol.* 104, 28
78. Critcher, C.R., and Dunning, D. (2014) Thinking about others versus another: three reasons judgments about collectives and individuals differ. *Soc. Personal. Psychol. Compass* 8, 687–698
79. Bernard, S. *et al.* (2016) Rules trump desires in preschoolers' predictions of group behavior. *Soc. Dev.* 25, 453–467
80. Klein, G.A. (1993) A recognition-primed decision (RPD) model of rapid decision making. In *Decision Making in Action: Models and Methods* (G.A. Klein *et al.* eds), pp. 138–147, Ablex
81. Johnson, J.G., and Raab, M. (2003) Take the first: option-generation and resulting choices. *Organ. Behav. Hum. Decis. Process* 91, 215–229
82. Smaldino, P.E., and Richerson, P.J. (2012) The origins of options. *Front.Neurosci.* 6, 50
83. Kalis, A. *et al.* (2013) Why we should talk about option generation in decision-making research. *Front. Psychol.* 4, 555
84. Kaiser, S. *et al.* (2013) The cognitive and neural basis of option generation and subsequent choice. *Cogn. Affect. Behav. Neurosci.* 13, 814–829
85. Meyer, R.J., and Kahn, B.E. (1991) Probabilistic models of consumer choice behavior. In *Handbook of Consumer Behavior* (T. Robertson and H. Kassarjian eds), pp. 85–123, N.J: Prentice-Hall: Englewood Cliffs
86. Hauser, J.R. (2014) Consideration-set heuristics. *J. Bus. Res.* 67, 1688–1699
87. Howard, J.A., and Sheth, J.N. (1969) *The Theory of Buyer Behavior*, John Wiley & Sons
88. Nedungadi, P., and Hutchinson, J. (1985) The prototypicality of brands: relationships with brand awareness, preference and usage. In *Advances in Consumer Research* (Vol. 12) (E. Hirschman and M. Holbrook eds), pp. 489–503, Association for Consumer Research
89. Hauser, J.R., and Wernerfelt, B. (1990) An evaluation cost model of consideration sets. *J. Consum. Res.* 16, 393–408
90. Roberts, J.H., and Lattin, J.M. (1991) Development and testing of a model of consideration set composition. *J. Market. Res.* 28, 429–440
91. Hauser, J.R. (1978) Testing the accuracy, usefulness, and significance of probabilistic choice models: an information-theoretic approach. *Oper. Res.* 26, 406–421
92. Klein, G. *et al.* (1995) Characteristics of skilled option generation in chess. *Org. Behav. Hum. Decis. Proces.* 62, 63–69
93. Raab, M., and Johnson, J.G. (2007) Expertise-based differences in search and option-generation strategies. *J. Exp. Psychol. Appl.* 13, 158
94. Klein, G.A. *et al.* (1986) Rapid decision making on the fire ground. In *Proceedings of the Human Factors Society Annual Meeting* (Vol. 30) *Proceedings of the Human Factors Society Annual Meeting*, pp. 576–580, Sage Publications
95. Musculus, L. *et al.* (2019) A developmental perspective on option generation and selection. *Dev. Psychol.* 55, 745
96. Levy, D.J., and Glimcher, P.W. (2012) The root of all value: a neural common currency for choice. *Curr. Opin. Neurobiol.* 22, 1027–1038
97. Peters, S.L. *et al.* (2017) The ventromedial frontal lobe contributes to forming effective solutions to real-world problems. *J. Cogn. Neurosci.* 29, 991–1001
98. Lieder, F. *et al.* (2014) The high availability of extreme events serves resource-rational decision-making. In *Proceedings of the 36th Annual Meeting of the Cognitive Science Society* (P. Bello *et al.* eds), pp. 2567–2572, Cognitive Science Society
99. Griffiths, T.L. *et al.* (2015) Rational use of cognitive resources: levels of analysis between the computational and the algorithmic. *Top. Cogn. Sci.* 7, 217–229
100. Lieder, F. *et al.* (2012) Burn-in, bias, and the rationality of anchoring. In *Advances in Neural Information Processing Systems 25* (F. Pereira *et al.* eds), pp. 2690–2798, Neural Information Processing Systems Foundation
101. Browne, C.B. *et al.* (2012) A survey of Monte Carlo tree search methods. *IEEE Trans. Comput. Intell. AI Games* 4, 1–43
102. Silver, D. *et al.* (2017) Mastering the game of Go without human knowledge. *Nature* 550, 354
103. Silver, D. *et al.* (2018) A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science* 362, 1140–1144
104. Daw, N.D. *et al.* (2005) Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* 8, 1704
105. Rangel, A. *et al.* (2008) A framework for studying the neurobiology of value-based decision making. *Nat. Rev. Neurosci.* 9, 545
106. Dolan, R.J., and Dayan, P. (2013) Goals and habits in the brain. *Neuron* 80, 312–325
107. Sutton, R.S., and Barto, A.G. (2018) *Reinforcement Learning: An Introduction*, MIT Press
108. Gläscher, J. *et al.* (2010) States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron* 66, 585–595
109. Bellebaum, C. *et al.* (2012) The neural coding of expected and unexpected monetary performance outcomes: dissociations between active and observational learning. *Behav. Brain Res.* 227, 241–251
110. Cooper, J.C. *et al.* (2012) Human dorsal striatum encodes prediction errors during observational learning of instrumental actions. *J. Cogn. Neurosci.* 24, 106–118
111. Li, J. *et al.* (2011) How instructed knowledge modulates the neural systems of reward learning. *Proc. Natl. Acad. Sci. U S A* 108, 55–60
112. Gershman, S.J. *et al.* (2014) Retrospective revaluation in sequential decision making: a tale of two systems. *J. Exp. Psychol. Gen.* 143, 182
113. Morris, A. *et al.* (2019) Habits of thought generate candidate actions for choice. *PsyArXiv*. Published online September 4, 2019. https://doi.org/10.31234/osf.io/2ayw3
114. Efraimidis, P.S., and Spirakis, P.G. (2006) Weighted random sampling with a reservoir. *Inform. Process. Lett.* 97, 181–185
115. Lewis, C.I., and Langford, C.H. (1932) *Symbolic Logic*, Dover Publications
116. Menzel, C. (2017) Possible worlds. In *The Stanford Encyclopedia of Philosophy (Winter 2017 Edition)* (E.N. Zalta ed) Metaphysics Research Laboratory, Stanford University
117. Franklin, N.T., and Frank, M.J. (2018) Compositional clustering in task structure learning. *PLoS Comput. Biol.* 14, e1006116