Narratives, Probabilities, and the Currency of Thought

Samuel G. B. Johnson [1]
Avri Bilovich [2]
David Tuckett [2,4]
d.tuckett@ucl.ac.uk
http://www.ucl.ac.uk/centre-decision-making-uncertainty

[1] University of Warwick, Department of Psychology
University Road, Coventry CV4 7AL, United Kingdom

[2] University College London, Centre for the Study of Decision-Making Uncertainty
Gower Street, London W1CE 6BT, United Kingdom

[3] University of Bath School of Management
The Avenue, Claverton Down, Bath BA2 7AY, United Kingdom

[4] University of Oxford, Blavatnik School of Government
120 Walton St, Oxford OX2 6GG

Corresponding Author: Sam Johnson

Email for Correspondence: sgbjohnson@gmail.com

Word Count: 8,997 (main text)

Last Revised: 19 November 2022 (submitted to journal)

Abstract

Whereas most commentators agree about the centrality of narratives in decision-making, the commentaries revealed little consensus about the nature of radical uncertainty. Here we consider 13 objections to our views, including our characterization of the uncertain decision environment and associated cognitive, affective, and social processes. We conclude that under radical uncertainty, narratives rather than probabilities are the currency of thought.

Narratives, Probabilities, and the Currency of Thought

R0. Three Theorists in Search of a Brawl

Midnight. A full moon illuminates the town square, as a desperate trio vaults onto the scene. They pause, huffing and puffing as their academic lungs reach beyond their intended capacity, their sweat glistening. Torches. Pitchforks. Crazed villagers scurry from all directions: Your heroes are surrounded. Fifty-four manic eyes stare them down…

A scene much like this was on our minds when we published this paper. Many of the proposed ideas confront not only neoclassical economics, but also most behavioral approaches, including much of JDM research. Formal decision theory, AI research, and standard economics are all deeply uncomfortable with the idea that there *could* often be situations requiring action in which there is no right answer, where probabilities are a poor guide. While behavioral economics rightly rejects the idea that we are optimal probabilistic decision-makers, it retains optimal probabistic reasoning as a normative standard. We challenge this perspective too. Under radical uncertainty, we are not "good" or "bad" probabilistic actors. Probability simply has nothing to do with it.

We were looking for trouble.

Yet, in the town square that is this journal, we were confronted not by thugs, but by 27 sets of experts in a dazzling array of fields, including cognitive psychology, neuroscience, philosophy, psychiatry, management studies, kinesiology, social psychology, affective science, Earth science, experimental humanities, semiotics, literature, organizational behavior, religious studies, and—the dismal science itself—economics. We were delighted to see commentaries from practitioners and applications of CNT to issues in clinical practice (**Linkovski and Eitan**; **Siegel**), education (**Sheskin, Bogucki, Perry, and McAllister**), and society (**Campbell and Fonagy**). Pleasingly, not one of the commentators proposed immolation, defenestration, drawing-and-quartering, or even a good whipping. (Well, maybe a lash or two.)

As far as we can tell, *every* commentator agrees with our basic premise that narrative thinking is crucial to human decision-making. But the commentaries reveal a deep divide in the field over the relevance and even existence of radical uncertainty. Some agree that we need credible theories of how to take action under radical uncertainty. Others are more inclined to reduce radical uncertainty to risk—a position which we believe has limited applicability to the kind of real-life decisions that really matter as opposed to "small worlds" contexts. We doubt we can do much to shake these theoretical commitments—as **Grossmann, Meyers, and Eibach** suggest, people may have strong priors on the existence of radical uncertainty—but will do what we can to taxonomize the disagreements.

R1. The Problem

Real-world decisions differ in many ways from those studied in laboratories and textbooks. We must often make massively consequential choices when data is incomplete, the options are ambiguous, the future may not resemble the past, and the axioms of standard decision theory (see **Ford and Kay**) are not satisfied. Such decisions often require both commitment and flexibility over time. Above all, such decision-making is about taking action in the face of potentially paralyzing uncertainty. Choices like monetary gambles and perhaps even that facing our widow—who has been put through such an ordeal between our target article and interlocutors—may be amenable to standard analyses. But it is far less clear how career choices, climate change mitigation, pandemic preparedness, and economic development policy can be effectively understood this way.

While many of the problems associated with such everyday decisions differ from those of textbook "small world" choices, two central issues must, as a matter of logic, be solved by any theory of decision-making. A central idea in cognitive science is that we do not act on the world directly, but through our mental *representations* of that external world (Fodor & Pylyshyn, 1981). These representations or beliefs need to be stored in a format that can act as a "currency of thought" that mediates between the information we pick up from the external world and our decisions to act on that world (the "mediation problem"; target article Figure 1). In principle, many types of representations could fit this bill—anything sufficiently isomorphic to the world. But our decisions are determined not only by our beliefs, but by our values. We need some "driver of action" that combines beliefs and values into choices (the "combination problem"). In our view, the core logical challenge faced by behavioral decision theory is identifying cognitive architecture that can solve both problems simultaneously.

One important approach is to assign probabilities to various possibilities (our beliefs), utilities to each of those possibilities (our values), and choose the option that yields the greatest expected utility (our choice). In our terminology, the currency of thought is probability (our beliefs are represented in a way that both summarizes the world and feeds into our choices) and the driver of action is utility maximization (our choices reflect the weighted combination of beliefs and values). Determining those probabilities and utilities may be complicated, but the central idea is remarkably simple.

This style of theorizing has often been attributed to economics. But **Ford and Kay** note that formalizations of decision theory in terms of exected utility maximization (Savage, 1954) do not actually make *any* assumptions about psychology, so (i) such formal theories do not strictly treat probabilities as a "currency of thought" or utility maximization as a "driver of action," and (ii) CNT is not inherently in conflict with these theories. Yet this has not stopped economists, psychologists, and many policymakers from treating these constructs in exactly the way we describe—tacitly taking expected utility-maximization as the psychological basis of "rational" economic behavior. As one example among many, Becker (1968) writes of his economic theory of crime that "The approach taken here follows the economists' usual analysis of choice and assumes that a person commits an offense if the expected utility to him exceeds the utility he could get by using his time and other resources at other activities" (pg. 176).

Behavioral economists have criticized this approach at length because people often assess probabilities wrongly. Fair enough. But we are pointing to a different and deeper problem: Often, decisions have to be taken—for instance whether to invest large sums in a new vaccine or to take costly steps to mitigate climate change—where *there are no probabilities to be wrong about*. In a world of risk, where outcomes are enumerable and probabilities calculable, expected utility maximization may make sense as a normative or even descriptive theory. In a world of uncertainty, where there is no way to assign probabilities even in principle, expected utility maximization not only is the wrong descriptive account, but is unintelligible as a normative theory. For these situations, behavioral economics inherits this unintelligibility by accepting the normative premise that uncertainty can be reduced to risk even while denying the descriptive claim that people do so optimally. The problem with using expected utility theory to explain decision under radical uncertainty is not that people are *bad* at probabilistic reasoning. The problem is that, under radical uncertainty, people *couldn't* be good at it.

We need a theory that accounts for action *with no probabilities at all.* This is where Conviction Narrative Theory comes in.

Objection 0: Such a theory may not be possible.

Some commentators took issue, less with CNT's specific claims, but with its scope. Not to put too fine a point on it, **Newell and Szollosi** muse whether CNT is "a theory of everything or nothing." They see a tension between our aspiration to explicate the "currency of thought" and our claim to be developing a theory not *of* causation, analogy, emotion, etc., but of how these component processes *fit together* to account for decision-making under uncertainty. We do not see the contradiction. CNT says that, for decisions under radical uncertainty (though not necessarily in other contexts), narratives are the currency of thought. We go into considerable detail about what a narrative is and what processes it participates in. But our account sits at a higher level of analysis than theories of causation, analogy, and so on. There are multiple competing perspectives on these lower-level representations and processes, and we mostly avoid these debates because they are not part of CNT's theoretical commitments.

As inventions go, CNT is more like a car than a carburetor. The automobile's innovation was the particular way that the components are connected to one another, not the components themselves, such as wheels, brakes, or carburetors. A car, like a theory, is a tool for solving a particular problem. It is not a strike against its inventor that it includes parts invented by others, nor does the inventor deserve credit for their inclusion. But the fact remains: They take us further than a horse.

King asks whether a general theory of decision-making under radical uncertainty is even *possible*. The answer is probably 'no' if King seeks a model, like expected utility theory, that can generate choices from some tractable set of inputs (e.g., subjective probabilities and utility functions). We doubt that even a descriptive, much less a normative, theory of this sort is possible. Descriptively, narratives are likely to depend on unique features of the situation and the decision-maker's idiosyncratic beliefs and values, which may be forbiddingly difficult to systematize. And normative theories of decision-making under radical uncertainty in general are likely to fall short for similar reasons to probabilistic approaches—who is to say what the right action is in a non-stationary world where the underlying model is not only unknown but changing?

But less ambitiously, we do believe an *explanatory* rather than predictive or normative theory is possible. We believe we can characterize the *kind* of mental representations and processes that are typically used to make such decisions, even if the *content* of those decisions is going to depend on too many unspecifiable, context-sensitive factors for a complete predictive theory. The more precisely the field can characterize those lower-level components (e.g., causation, analogy, affect), the more explanatory the resulting product (CNT + the best theories of causation, analogy, etc.) will be. We also believe that by focusing decision research on the conditions in which conviction (rather than a "correct" choice) is achieved, an agenda is created that is of much greater practical use.

R2. Context

To restate CNT briefly: For many everyday decisions, people adopt *narratives*—mental representations coordinating causal, analogical, temporal, and valence structure—to explain decision-relevant information and link proposed actions to intended outcomes. Narratives are the joint product of individual cognition and the social environment: People pick up on "narrative fragments" from others which feel right in context and are incorporated into individual narratives to support action. Once a narrative is adopted that makes sense of the evidence, people use the causal structure in that narrative to simulate or imagine the future that would unfold if they chose a particular course of action. They then deploy the same emotional processes they use to appraise actually-present situations to appraise the imagined future. If the appraisal yields approach emotions, then they choose that course of action; if it yields avoidance emotions, they choose a

different course. Since decisions must manage doubts and eventualities over time, people can do so either adaptively in an "integrated state" that maintains commitment while being open to evidence and potentially course corrections, or maladaptively in a "divided state" that repels doubts by blindly clinging to a narrative until reality intervenes. People may chronically differ in their tendency toward either stance, but situational forces such as organizational structure and normative practices can shift people toward one or the other (Fenton O'Creevy & Tuckett, 2022).

We view this process not as infallible, but as by-and-large adaptive given the constraints posed by the environment under which many everyday decisions are taken. We describe four such constraints at the beginning of the target article: Such decisions are often take under radical uncertainty (meaningful probabilities do not exist) and fuzzy evaluation (a precise and commensurable utility function cannot be characterized), the decision often must be sustained over time (they require a long-term commitment), and they are often socially embedded (information is processed not only through individual cognition, but through our interactions with others). Thus, we agree with **Marewski** that both blades of Simon's (1990) scissors are needed to understand decision-making—the external environment and internal processes.

Our point is that narratives can help to cope with all of these features of the decision environment while implementing the internal psychological processes (sense-making and imagination) that so often seem to characterize the phenomenology of decision-making. Narrative evaluation relies on heuristics rather than probabilities. Narrative simulation relies on affect rather than utilities. Narratives can resist change, permitting commitment (though also risking ossification, as **Breithaupt, Hicks, Hiskes, and Lagrange** and **Campbell and Fonagy** rightly note). And narratives can incorporate information from the social environment, permitting the forces of cultural evolution to apply selection pressures to our beliefs alongside individual cognition. This socially distributed aspect of narrative decision-making is an important part of how people often make reasonably good decisions despite environmental and cognitive constraints.

Although we received minimal pushback on our points about fuzzy evaluation, commitment, or social embeddedness, we encountered several objections to our characterization of radical uncertainty.

Objection 1: Radical uncertainty does not exist.

We sense a deep divide among our commentators about the very existence of radical uncertainty, which we believe to represent a disagreement across many fields more broadly. For example, **Jinich-Diamont and Christov-Moore** suggest that uncertainty falls on a continuum, where our probability distributions in more uncertain situations depend more on prior probabilities (generated from a higher-order probability distribution) than on new evidence. In fact, this is an eliminative account of radical uncertainty—it's probabilities all the way down. **Friston**'s commentary can also be read this way. Conversely, many others seem to view the issue as decided in the opposite direction, including **Gigerenzer**, **Kay**, and **King**. **Marewski** also notes Simon and Newell's distinction between well-structured and ill-structured problems, which we see as closely related, since one way a situation can be uncertain is when the outcomes cannot be enumerated. Although most commentators left their view on this topic implicit—and we are reluctant to take words written between the lines and place them firmly in people's mouths—we suspect many commentators and readers find themselves attracted more to one or the other of these positions.

What we are *not* saying is that, under radical uncertainty, every possibility is equally plausible and anything goes, but rather that probabilities are incalculable. For example, here is Keynes' (1937) famous passage:

> By "uncertain" knowledge, let me explain, I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty....The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention, or the position of private wealth-owners in the social system in 1970. About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know.

We can read **Friston** and **Jinich-Diamont and Christov-Moore** in two ways. On the one hand, coming from the relatvely stable world of biological functioning (e.g., vision or motor control), we can suppose they are mainly thinking about contexts which change very slowly and are strongly path-dependent. It means they can usually discount the significance of the kind of rapid and unexpected innovation that chacterizes the social and economic phenomena with which Keynes was preoccupied. In this case, our differences with them are ones of focus and framing. On the other hand, if not and they really have in mind decisions about pandemics, climate change, economic policy, etc., then their approach, in which decision-makers rely on prior probabilities even under radical uncertainty, is hard to support. We doubt that people have prior probability distributions over such events as those Keynes is describing. Perhaps, when pressed, people might provide such numbers. But such estimates could only be improvised in the moment and have little relationship to behavior (here, we suspect we are aligned with **Chater**).

We accept we will not convince everyone with our negative case against probabilistic approaches. Preexisting philosophical and disciplinary commitments preclude this. Yet, in the target article, we also marshal many sources of positive evidence for narrative influences in decision-making. Some of this evidence is potentially reconcilable with probabilistic approaches to narrative, while some would be very difficult to reconcile (e.g., the "digitization" results in Johnson, Merchant, & Keil, 2020, as **Greco** notes). For **Friston**, we suspect his proposed probabilization of narrative decision-making would be a satisfying theory that could formalize narratives as a tool for decision-making under risk (we take it that for Friston, all decisions are risky in the sense that they can be formalized using probability theory). We are much less bullish on this approach because we reject the rejection of radical uncertainty. This leads naturally to the next objection.

Objection 2: CNT [does] [does not] explain decisions under risk.

CNT has little to say about decision-making under risk. Probabilistic theories, such as classical decision theory or Bayesian models or heuristics-and-biases models that take probabilities as inputs, are sensible contenders and we do not attempt to adjudicate this debate. **Beach and Wise** believe this is an unnecessary limitation. We see it as an empirical question, as either answer could be justified theoretically.

If so many decisions are taken under radical uncertainty that our cognitive architecture is primarily attuned to such contexts, then plausibly narrative cognition applies to risky decisions as well. One way that this could be true is if we think of some narratives as *licensing* probabilistic reasoning—a point closely related to **Grossmann et al.**'s suggestion that radical uncertainty is itself a narrative and to **Gigerenzer**'s observation that the idea of a "small world" is a compelling narrative. To use Beckert and Bronk (2018)'s term, "calculative devices" can be embedded in a narrative, as when a gambler decides to adopt a particular mathematical procedure as an input to their decision-making strategy. Less sanguinely, policymakers may reify a particular economic model as literally true, applying its mathematical formalism to a situation in which it is inappropriate. Orthogonal to whether the calculative device is *appropriate* to a situation, its mathematical application may itself

be *deployed* correctly (as when seasoned weather forecasters apply their models) or incorrectly (as in the many mistakes identified in judgment and decision-making research).

Alternatively, since probabilistic reasoning is indeed a valuable way to decide under risk, we do not rule out the possibility that we have separate mechanisms for doing so. Indeed, CNT is not trying to explain *all* decisions. Some commentaries point to very low-level decisions that we agree may not be mediated by narratives. **Cañal-Bruland and Raab** note that many motor decisions, such as where and when to hit a baseball, do not seem to be especially cognitively mediated. **Andreetta, Spalla, and Treves** look to decisions such as selecting the correct answer in a recognition memory task, questioning whether such decisions are affectively mediated. We don't know; we simply see this as a non-overlapping magisterium. Problems such as when and where to hit a baseball are far more constrained and structured than the sorts of problems CNT is intended to address.

If we do have separate mechanisms for dealing with risk versus uncertainty, then as **Newell and Szollosi** note, we would need to have some way to *identify* such situations. Alternatively, if even risky decision-making is mediated through narratives, we still need principles for identifying when to apply narratives that include calculative devices. **Grossmann et al.** make an impressive array of intriguing suggestions in this connection, including the role of folk theories, individual differences, moral norms, motivation, and culture.

R3. Representations

We characterize narratives as "higher-order" representations; they do not necessarily confine themselves to causal structure, but can coordinate causal structure with analogical, temporal, and valence information. We believe that this kind of representational flexibility is integral to the sophistication of human thought. Despite its importance, there is little theoretical or empirical work examining how these different kinds of lower-level representations are coordinated. We sketched how these representations might look, but we readily acknowledge that much more needs to be done. Here, we consider three challenges to our sketch.

Objection 3: CNT's representational framework is too complicated.

Newell and Szollosi recommend that, in lieu of narratives, we focus on simpler representations. They give as examples the study of how people represent integers and event frequency.

We certainly have no objection to research that takes a more minimalist or even reductionist approach to understanding mental representations or processes. Such strategies have been massively successful in cognitive science for understanding countless issues. But this does not entail that they are the *only* legitimate scientific approach or that they are well-suited to solving *every problem*. Indeed, Newell and Szollosi unfavorably juxtapose our view that qualitative methods are superior to experiments for understanding the *prevalence* of narrative thinking with our view that experiments are superior to qualitative methods for understanding the *processes* of narrative thinking. But this is not a contradiction! Different tools are suitable for different purposes. Lab experiments risk creating thoughts or behaviors in artificial settings, where they would not exist in a naturalistic setting—so they are not ideal for understanding how frequent a behavior is in the real world. Interviews can suffer from post-hoc rationalization and the opacity of introspection—so they may not be ideal for understanding some psychological processes. Methodological imperialism is equally myopic whether it is staked by psychologists on behalf of lab experiments, behavioral economists on behalf of field experiments, economic theorists on behalf of formal models, or social scientists on behalf of interviews and ethnographies. Complex problems require all the help we can muster if we are ever to triangulate the truth.

As for the specific suggestion to build up from simple to complex representations, we believe success is far more likely if we approach the problem from both angles. After all, there are *many* investigations of simpler representations and processes, including much of our own work and arguably *most* of cognitive psychology. We would channel here (a different) Newell (1973), who argues that "you can't play 20 questions with nature and win." The elder Newell means that we need to conjoin experimental evidence that dissects component processes ("playing 20 questions") with integrative models that explain how these processes work in concert to yield thought and behavior. In the case of CNT, this is not (yet) developed to the level of a mathematical model—and, indeed, part of the reason is that cognitive science does not yet have complete models of the component processes. Yet, we believe that theorizing the *relationships* between cognitive processes must take an equal role to examining how those component processes work. We need both cars and carburetors.

Objection 4: CNT's representational framework is too simplistic.

Coming from the totally opposite end, **Caldwell** suggests that our representational model of narratives is too simple. We do not intend for our sketch of a representational theory of narratives to be the final word on the subject, and we agree with Caldwell that it can be developed further. Some of Caldwell's extensions are plausible and should be considered in future work. Yet, extensions of this framework should be (i) mindful of the empirical evidence, and (ii) careful to preserve the features that make it suitable for dealing with action taken under radical uncertainty.

As an example of the first issue, Caldwell suggests that people represent some causal relationships as stronger than others. This may be true, but the point is controversial in the causal cognition literature (e.g., Yin & Sun, 2021) and we are ourselves unsure.

As an example of the second issue, we suggest caution in installing "a scalar affective value" into each node. We are unsure that this is an advance over representing utilities. As Walasek and Brown (in press) forcefully demonstrate, the problem with utilities is shoving preferences into a one-dimensional common currency, as a scalar affective value would. CNT abandons this approach in favor of affect because affect is multidimensional, motivates action, and is an evolved system for adaptively dealing with real situations. As no small bonus, emotions are known to exist, whereas no one has ever actually seen a utility function in the wild (though we understand that there is a large bounty for one caught alive).

Objection 5: Narratives are too incoherent to be the building blocks of knowledge.

Chater asks whether narratives are stored in long-term memory and applied off-the-shelf to situations as they arise, or instead improvised in the moment and constructed on the fly. On one extreme, Chater suspects that narratives are constructed as post-hoc justifications. On the other extreme, **Caldwell** suggests that all of our knowledge is constructed in one grand narrative from which we extract a subset at any given time—a point **Beach and Wise** echo in their invocation of a "prime narrative" in the Theory of Narrative Thought.

While we agree with **Chater** that narratives can be inconsistent, this is not a sufficient basis to reject the idea that we store narratives that support action. Narratives are beliefs, and although Chater (2018) believes that "no one, at any point in human history, has ever been guided by inner beliefs or desires," we do not share this belief. (Unless Chater is right, in which case neither of us have beliefs to disagree about!) Most beliefs are much more pedestrian than unconscious desires or repressed memories: Bees can sting, fertilizer makes crops grow, my boss is grumpy on

Mondays, J-walking is taboo. Such beliefs are not consciously accessible most of the time, but plainly guide our decisions; they are not post hoc confabulations for avoiding the beehive, purchasing fertilizer, complimenting the boss, or waiting for the traffic signal. The most extreme version of Chater's view appears untenable: The building blocks of narratives—including causal and analogical knowledge, social norms, and probably simple intuitive theories—seem clearly to be stored as beliefs in long-term memory.

Yet, we would not go as far as **Caldwell** or **Beach and Wise** in the opposite direction. Caldwell suggests that narratives are excerpted as subgraphs from a "causal graph…representing the agent's whole mental model of the world." One needn't embrace a Chaterian view of mental flatness to find this implausible. It is unlikely that people could maintain a consistent model of this size in their minds, both for theoretical reasons (e.g., the "frame problem" mentioned by commentators) and empirical reasons mentioned by Chater. Indeed, our own prior work suggests that causal representations are remarkably discrete (Johnson & Ahn, 2015, 2017). People believe that there is a causal relationship, for instance, between wine (A) and sleep (B) and between sleep (B) and dreaming (C), but not between wine (A) and dreaming (C). People have separate, schematized causal mechanisms for the soporifics of wine and for the phenomenology of dreams. Causation is mentally represented in "islands" rather than "networks."

Rather than impaling ourselves on the horns of this dilemma, we tentatively offer the following resolution. We agree with Chater that narratives are not the "building blocks" of knowledge. That distinction goes to the lower-level representations that are *coordinated* by narratives (e.g., causal and analogical schemata). But those lower-level schemata must often be fleshed out in more complex narratives to support sense-making, prediction, and action. Sometimes this fleshing out happens slowly, through cultural evolution and social development, whereas other times it may indeed occur on-the-fly as Chater suggests, to fit a particular situation.

Unlike Chater, however, we believe that narratives can also be stored, retrived, and elaborated over time. This is central to CNT's account of why decision-making can be reasonably adaptive, even under radical uncertainty. Sustained representations of narratives enable two feedback loops. First, they allow feedback from our actions to impact our narratives on an individual level. Decision-makers potentially revise narratives in light of whether actions taken on their behalf work out well or not. Second, sustained individual representations of narratives underlie the shared narratives that are shaped by cultural evolution. Entire societies and subcultures can fall under the spell of a shared narrative, for good or for ill, as illustrated by examples from **Campbell and Fonagy**, **Gigerenzer**, and **Lightner**. This depends on shared representations, such as transactive or collective memory (Hirst et al., 2018), which conjoins memories stored in individual minds with social processes.

Thus, we reject the claim that narratives are merely epiphenomenal constructions after the fact, made to justify our decisions. It is certainly true that narratives do often play a justifactory role (Cushman, 2020; Mills, 1940), as we note in the target article. But we think it clear that they can also cause decisions. For one, the anticipated need to justify a behavior can *itself* change our behavior; as Mills (1940) puts it, our justifications must make sense to collaborators, requiring that they fit a "vocabulary of motive." But even beyond their justificatory role, causal narratives seem to cause behavior. Surely causal cognition functions to allow successful interventions on the world (Woodward, 2003) and analogical reasoning functions to apply lessons from one situation to another, supporting action (Holyoak, 1985). Many experiments cited in the target article show how altering one aspect of a mental representation (causal, analogical, temporal, or valence information) impacts downstream decision-making. Narratives can serve a justificatory function, and these communicated narratives may even differ from the narratives that drove behavior. But recognizing

that narratives can be constructed *post hoc* does not mean that a decision was not also driven by a narrative *ex ante.*

R4. Explanation

Turning from representations to processes, we first highlight the role of narratives in making sense of decision-relevant information by imposing structure. The target article made two claims about this process. First, it relies on heuristics; second, explanatory fit is experienced affectively (some narratives 'feel right'). Many of the more critical commentaries highlighted concerns about this process.

Objection 6: CNT does not provide an account of how narratives are generated.

Several commentaries point out that we do not provide an account of how narratives are generated, only how they are evaluated once they are generated (**Caldwell**, **Greco**, **Newell and Szollosi**). This is not an issue that CNT considers in detail, although it is indeed important. As Greco notes, solving this issue is an essential part of the transition from a 'grand world' problem to a more tractable 'small world' problem. For instance, **Friston** seems to equate the problem of Bayesian model selection with the problem of hypothesis generation, but such a process requires candidate models to be considered. Some process that sits outside of probabilistic reasoning would be necessary to generate the structures to be selected, even if the selection process operates along Bayesian lines.

Much of the work of narrarive generation is in fact outsourced to the social, historical, and family environment. Certainly, new 'grand narratives' of the sort described by **Gigerenzer** and **Lightner**, or those discussed by Shiller (2019), arise only once in a great while, emerging slowly through cultural evolution. (In this sense, the question "how are narratives generated?" is rather like "how are species generated?") On a much smaller scale, though, people likely combine causal and analogical knowledge in new ways in everyday life, such as when we consider a new career or relationship. CNT does not provide a detailed account of this process, but we do believe that a narrative approach helps to focus attention in the right direction.

First, narrative accounts can draw on the broader literature on how people generate causal hypotheses, including through heuristics. The central idea is that effective generation strategies (i) rely on easily accessible information and (ii) cue hypotheses that have a reasonable chance of being relevant and correct. Many of these strategies are related to time (e.g., Lagnado et al., 2017; Rottman & Keil, 2012), which is one reason to think narratives coordinate temporal information. One example, briefly mentioned in the target article, is our own prior work on how event structure is used to narrow the space of potential causes for a specific event (Johnson & Keil, 2014). Research on event perception has found that people segment experience into discrete events (Zacks & Tversky, 2001) at multiple levels of granularity (i.e., higher-order events that subsume lower-level events, such as a trip to the mall that includes episodes for visits to particular stores). Johnson and Keil (2014) found that people use two heuristics derived from event structure for narrowing candidate causes. First, events at one level of granularity are preferentially considered as candidate causes (i.e., matching higher-level effects with higher-level causes). Second, for events at a lower level of the hierarchy, other lower-level events that are part of the same higher-level event are preferentially considered as candidate causes.

Second, hypothesis generation must draw not only on information readily available from the immediate environment, but long-term memory. Thus, a detailed understanding of how the lower-level building blocks of narratives are represented and indexed in memory will be crucial. When

narratives must be generated bottom-up, the generation process presumably begins with the evidence to be explained, which acts as retrieval cues for general knowledge that may be incorporated into the narrative. If the evidence includes observations, A, B, and C, then the generation process will involve memory search for causes and analogies related to A, B, and C. Indeed, we know that causal relationships act as stronger retrieval cues than non-causal relationships with equally strong associative strength (Fenker et al. 2005). Plausibly, the narrative would be built up until it is deemed a sufficiently complete explanation (Korman & Khemlani, 2020).

Third, notice that storing narratives in memory greatly simplifies this process, as does adopting shared narratives from the social environment. Rather than carrying out a potentially arduous process of narrative construction, one could simply pull a narrative off the shelf when triggered by a given situation, perhaps with some tailoring. Or multiple narratives might be retrieved and then evaluated. This is one additional reason to be skeptical of the idea, à la Chater, that narratives are always or usually constructed on the fly.

Objection 7: CNT is not specific enough about what heuristics are used.

Greco mentioned that we were not very specific in the target article about what heuristics are used. While CNT is itself agnostic on this question, there is a growing literature on exactly this issue that we reference in the target article (e.g., Horne et al., 2019; Johnson et al., 2016; Lombrozo, 2016). As one example among many, people use simplicity as a cue to an explanation's prior probability and complexity as a cue to its likelihood or fit to the data (Johnson et al., 2014, 2019; Lombrozo, 2007), circumventing the need for exact probabilistic computations. This literature mostly looks at much simpler kinds of representations (of the sort that **Newell and Szollosi** favor studying), so future research should generalize these heuristics to more complex narrative structures, understand how they interact, and study how they are prioritized when in conflict.

Some of the commentators mentioned still other heuristics. **Gigerenzer** brings up the fascinating point that narratives can themselves select heuristics (his Protestant Work Ethic example). **Schwarz** highlights the importance of metacognitive cues, such as fluency, with a related point made by **Jinich-Diamant and Christov-Moore** about "interoception" (observation of internal states). We agree that such cues are important. As Schwarz notes, fluency is closely related to our point that we select narratives that "feel right," and thus to the construct of explanatory "satisfaction" studied in the explanation literature. The causes of explanatory satisfaction are often, but not always, structural features such as simplicity (Lombrozo, 2007) and explanatory scope (Khemlani et al., 2011). That is, feelings are often the proximate cause of behavior, even as they implement intelligent and context-sensitive heuristics or strategies for selecting actions. We suspect that Schwarz is sympathetic to this point, although he also discusses evidence that incidental emotions can influence fluency for reasons unrelated to structure.

Objection 8: Heuristics could be used to estimate probabilities rather than narratives.

There are three versions of this objection.

First, **Greco** argues that probabilistic theories have an equal right to invoke heuristics as CNT does. Thus, evidence that people use heuristics to evaluate explanations should not count as evidence for CNT. We agree with this in principle, but in practice probabilistic theories of cognition tend to reduce heuristics to *approximation* techniques for performing probabilistic calculations. For example, Griffiths et al. (2012) explicitly contrast their view of heuristics that we "can connect directly to optimal solutions" (pg. 417) against "bag of tricks" views. This is a

plausible move for probabilistic accounts, which link to more recent developments in "resource rationality" (Lieder & Griffiths, 2020), and which may be successful for modeling cognition under risk. We remain skeptical of "optimal" solutions, whether exact or approximate, under radical uncertainty.

Second, **Jinich-Diamant and Christov-Moore** argue that CNT actually is sneaking in probabilities under the cover of darkness when it invokes heuristics such as "simpler explanations are *likelier* than complex ones." This is a mistake. A heuristic that compares two possibilities and concludes that one is likelier does not logically necessitate probabilities, but only a comparative (rather than absolute) judgment of likelihood. Radical uncertainty is not the same as saying that anything goes. Some possibilities (there will be another pandemic or development of cheap nuclear fusion) are more plausible than others (there will be a civilization-ending nuclear war or development of faster-than-light travel) even if we cannot assign probabilities to them. This is true even without radical uncertainty. If you see one bag with many black marbles and another with only a few, your approximate number system (Feigenson et al., 2004) can estimate which bag has more marbles in it. You would conclude it is likelier that a marble drawn from one bag is likelier to be black than one drawn from the other, even if you have little idea how many marbles are in each bag. Crucially, while such behavior could be rationalized probabilistically from the outside, probabilities themselves need not be *mentally represented* in making this judgment. To view it as impossible to deem one event more likely than another without representing probabilities is to trivialize probabilistic approaches.

Third, a possibility not raised by commentators but which troubles us nonetheless, is that even if probabilistic theories could not be appropriate *normative* accounts under radical uncertainty, people might nonetheless *impose* probabilities where they ought not. On the one hand, much of the evidence reported in the target article contradicts this perspective for simple everyday decisions under radical uncertainty. But of course, inappropriate probability-based decisions can be observed, as when value at risk models made disastrous probabilistic assumptions in the lead-up to the 2008 financial crisis, ignoring the possibility of things never previously encountered (Taleb, 2007). In fact, we suspect that the tendency to conflate risk with uncertainty is primarily a (recent) cultural phenomenon that has occurred as the proliferation of risk management tools and other quantitative models have appeared to offer greater levels of "authoritative" precision. This is one way that "the belief that we live in a small world is itself a powerful narrative," as **Gigerenzer** put it.

Objection 9: CNT can be subsumed into probabilistic theories.

Several commentaries compare CNT to probabilistic approaches, such as Free Energy Theory, including **Friston**, **Jinich-Diamant and Christov-Moore**, **Solms**, and **Tuominen**. Insofar as these approaches are committed to probabilistic *representations*, we view this reduction as implausible under radical uncertainty (see Objection 1). For example, **Jinich-Diamant and Christov-Moore** suggest that we can "take narrative to be the structure of higher-level priors in the nested predictive hierarchy" and feelings as "low dimensional representations of interoceptive predictions about the expected consequences of visceromotor commands on the body's internal milieu." We are unsure how literally this is to be taken, but, needless to say, we are wary of such representational claims.

More broadly, we think it crucial to be clear about what level of analysis a theory is operating at. This can help to clear up confusion. For instance, sometimes the claim is made (e.g., by **Jinich-Diamant and Christov-Moore**) that probabilistic theories do not view our minds as *actually* computing probabilities, but that probability distributions can be constructed that *approximate* what

the mind is doing. This is, roughly, the move from Marr's (1982) algorithmic to computational level. We have no problem with this as a research strategy (and said so in Sec. 4.1 of the target article), with three caveats.

The first is that such theories should make falsifiable predictions which should be reasonably close matches to the empirical evidence. Some of the empirical results described in the target article are flatly inconsistent with probabilistic theories, with work on "digitization" (Johnson et al., 2020; Murphy & Ross, 1994) striking at fundamental claims about probabilistic thought. This is not to say that theoretical reconciliations could not be generated, but that we have not yet seen attempts to do so.

The second is that we must be clear on their theoretical status (i.e., what scientific problem they are solving). Computational-level theories are useful for characterizing why the outputs of a cognitive system are intelligent given the problem and constraints, but they explicitly are *not* mechanistic theories. The reason is obvious: If the theory is not saying that people represent probabilities in their minds but instead are doing something well-approximated by Bayesian inference, a mechanistically explicit account would need to say *what that thing is*. We are not claiming that CNT is a fully mechanistically explicit theory in the sense that in its current state it could be simulated on a computer—but it *is* mechanistically explicit in the sense that we think that people *do* represent narratives in their minds, and *do* use them in processes of explanation, simulation, affective evaluation, and communication. There is much more to be said about all of these things, but CNT does specify what representations and processes underlie decisions under radical uncertainty.

The third is that computational-level accounts often depend on a "rational analysis" of a task (Anderson, 1990; Chater & Oaksford, 1999). We agree that this is valuable. However, we are skeptical that a probabilistic analysis of choices under radical uncertainty *is* a rational analysis. Approaches such as Free Energy Theory seem potentially well-suited for understanding evolved biological functions in stationary environments where probability distributions are learnable because they do not shift; under such conditions, the theory's commitment to optimization (Friston, 2010) can be appropriate. Yet, while **Friston** notes that we can be uncertain about hidden states, model parameters, or models themselves, he does not consider that the underlying model *itself* can often shift in complex decision environments such as the economy—they are non-stationary. The assumptions required to "rationally analyze" radical uncertainty using probabilities in fact *assume away radical uncertainty*.

Although the four commentaries written from an explicitly probabilistic perspective all are couched in terms of Free Energy Theory, other probabilistic perspectives are also relevant. One family recently gaining traction is resource rationality (Lider & Griffiths, 2020), which consider bounded rationality from a Bayesian perspective. These accounts invoke tools such as approximation algorithms and attempt to give Bayesian rationales for specific heuristics given cognitive constraints. We are sympathetic to such approaches, particularly to the sampling approaches that have recently begun to accrue empirical support. Such approaches can yield "Bayesian brains without probabilities" (Sanborn & Chater, 2016), a view which seems far more plausible under radical uncertainty compared to approaches that invoke explicit probabilistic representations. Examining whether sampling and narrative approaches can be unified would be an exciting theoretical task.

R5. Simulation

With a narrative in mind to explain the past, CNT holds that decision-makers use the causal and temporal structure embedded in the narrative to simulate the consequences of potential actions. These simulations occur one a time, with people focusing on one particular imagined future that feels accurate rather than weighting combinations of different futures.

Objection 10: Simulation changes narrative plausibility

Enz and Tamir describe some fascinating experiments showing that merely simulating the future changes its perceived likelihood, in part by making alternatives harder to retrieve. This mechanism is compatible with, but distinct from, those invoked by CNT. Given that a narrative is often compatible with multiple possible futures, a further account is needed of how imagined futures are generated from narratives, unifying with greater detail the evidence discussed in the target article. The path-dependence proposed by Enz and Tamir will likely be an important part of such a theory, especially for the important subset of decisions that are made repeatedly or sustained over time.

Objection 11: People can and should simulate multiple possibilities

Breithaupt, **Hicks**, **Hiskes**, and **Lagrange** balk at our suggestion that people simulate a single future at a time. To be clear, we are not suggesting that people only *consider* a single future. Rather, we are saying that people consider one at a time, cannot consider them all, and adopt one (provisionally, in the case of an integrated state) to guide actions sustained over time. Thus, multiple futures can be considered and compared, with only one (e.g., the most plausible one) adopted as true. What we *do not* believe people can typically do is to contemplate multiple futures *simultaneously*, weighing predictions by the relative probability that each of those futures might happen. This is the point made by the "digitization" experiments discussed in the target article (Johnson et al., 2020; Murphy & Ross, 1994).

Yet, even provisional *adoption* of a single imagined future can get one quite far. First, we need not always adopt the most plausible future, but instead sometimes select futures that highlight threats (an important point made by **Beach and Wise**). Indeed, the boundary conditions on "digitization" tend to feature exactly such situations, such as dangerous categories (Zhu & Murphy, 2013) or immoral people (Johnson, Murphy, Rodrigues, & Keil, 2019). Second, despite *cognitively* adopting a single future as true (in the sense that they do not average over alternatives), people can still experience *metacognitive* uncertainty about that belief if the adopted future is not seen as much more plausible than alternatives. In Johnson et al. (2020), participants ignored less-likely alternatives when making predictions, but when more-likely and less-likely alternatives implied different futures, participants were less confident in their predictions. The adoption of a single future, coupled by metacognitive uncertainty that fuels epistemic humility, can support a committed course of action while maintaining an open stance toward alternatives (an "integrated state"), both in the sense of willingness to revise one's narrative as new information accrues and of "hedging" behavior like that described by **Breithaupt et al**. This strikes us as an important direction for future research.

R6. Affective Evaluation

With an imagined future in mind, the decision-maker uses the same affective system to appraise that future as they would use for an actually-present situation, which in turn feed into our motivational system to approach or avoid choices that lead to those futures. While emotions start out as adaptive mechanisms to prioritize needs, through development they come to internalize social input and reflect cultural wisdom (Scherer, 2005).

We received little push-back on the specific role of emotions invoked by our theory, including our idea that appraisals could be carried out over "default" dimensions (those invoked in standard appraisal theories of emotion) as well as "ad hoc" dimensions (cued by specific goals) and our ideas about how emotions help to manage decisions that must be maintained over time.

Objection 12: CNT unduly prioritizes affect at the expense of cognition.

Despite relatively little pushback on specifics pertaining to emotion, beyond what we mentioned in earlier responses, we did get the sense that some commentators were uneasy with the extent to which CNT prioritizes affect over cognition. We simply think affect versus cognition is a false dichotomy. Rather, CNT is a theory of how affect and cognition relate, support one another, and jointly support action. Affect guides which narratives we adopt (though not *only* affect, we hasten to add in response to **Schwarz**) and affect appraises potential choices via the imagined futures they invoke (though this need not be the only effect of imagination, as **Enz and Tamir** note). Moreover, in adaptive decision-making characterized by "integrated states," cognition and affect work together to maintain conviction and potentially revise narratives in light of new information. Affect and cognition are essential and inseparable in adaptive decision-making, with neither more important than the other.

R7. Communication

CNT encompasses both cognitive and social processes. Fragments of (individual) narratives are communicated within social groups, giving rise to shared narratives that in turn feed back into individual narratives. This is important both in solving the challenging problem of narrative generation (as **Greco** and several other commentators note) and in harnessing the adaptive power of cultural evolution to explain how people can make choices that mostly work, in a world of uncertainty. Narratives facilitate social coordination (via reputation-tracking and persuasion), drive social learning, spread through social networks, and evolve. **Pelletier, McLaughlin, and Boespflug** highlight the usefulness of a semiotic approach for understanding these functions, while **Jones and Hilde-Jones** extend some of these ideas more formally, identifying reasons why narratives act as "cultural attractors" that might be resistant to change (for better and for worse).

Other commentators highlight both the positive (**Lightner**) and negative (**Campbell and Fonagy**) implications of narratives' privileged status in communication. Together with the commentaries by **Linkovski and Eitan** and **Siegel** applying CNT to clinical practice, **Campbell and Fonagy**'s commentary suggests the intriguing possibility that individual psychopathy is often linked to maladaptive (individual) narratives, while "collective psychopathy" is often linked to maladaptive shared narratives. Given the pressing importance of issues such as polarization and misinformation, studying such issues through a CNT lens may be valuable.

Objection 13: There is no guarantee that cultural evolution will select adaptive narratives.

This objection was actually not raised by any of the commentators, but we think we need to make clear it is an important issue for future research.

Consider **Lightner**'s interesting commentary on religious narratives. Lightner asks why narratives are so often false, concluding that even false narratives can enhance fitness by simplifying problems, solving long-term collective action dilemmas, and harnessing powerful emotions to generate conviction to act and coordinate behavior. Whether such narratives are actually true has little to do with whether they are adaptive. This account is quite consonant with CNT.

Presumably the selection story here is a group selection one—groups that evolve such narratives are likelier to survive and expand, and thus the narrative spreads with the group. But an important point raised by Dawkins (1976) is that "memes" such as narratives, like genes, are selfish. That is, selection favors memes that themselves survive and propagate, not necessarily memes that help their hosts to survive. These *can* overlap, as Lightner argues for religious narratives, but they need not. The examples of misinformation raised by **Campbell and Fonagy** may fall into this category.

Thus, what guarantee (if any) is there that cultural evolution of shared narratives promotes adaptive decision-making? An initial observation is that the answer must depend on the reason why the shared narrative is propagating—whether it is spreading *because* it facilitated a successful outcome. This could be the case if its adherents are likelier to survive to tell the tale or because the positive outcome motivated them to share the underlying narrative. In contrast, narratives can act as "cultural attractors" without bearing much resemblance to reality, as when they are used to signal group affiliation. **Sheskin et al.** mention research on science communication (e.g., Kahan et al., 2017), which shows that narrative-confirming beliefs about issues such as climate change and gun control contribute to polarization. Similarly, Caplan (2007) argues that people are often "rationally irrational" in their economic views, which may be wrong from a scientific perspective but useful for gaining social credit with one's in-group, all the while leading to poor government policy through the ballot box.

We view this as a crucial question for future research, and one where formal modeling along the lines suggested by **Jones and Hilde-Jones** may be useful.

R8. Future Directions

The commentators generously provided many ideas for future research. Some commentators provide ideas about how to unify CNT with other approaches such as the Theory of Narrative Thought (**Beach and Wise**), Free Energy Theory (**Friston**), situation models (**Dominey**), and semiotics (**Pelletier et al.**). **Marewski** observes that both cognitive and environmental processes guide decision-making; we agree and would argue that one of CNT's principal insights is that two features of the decision environment are importantly linked: Radical uncertainty and social embeddedness. There are ontogenetic questions about how narratives are used through development (**Tuominen**); we would add the phylogenetic question of how narrative cognition evolves. Interesting questions were raised about the similarities and differences between narratives and other mental representations, such as scripts (**Siegel**), story grammars (**Andreetta, Spalla, and Treves**), and situation models (**Dominey**). More work needs to be done to fill in, computationally model, and experimentally test our sketch of narrative representation; suggestions such as those provided by **Caldwell** will be useful, as will incorporating research on lower-level representations as **Newell and Szollosi** suggest. Much is still to be learned about how narratives are used to simulate futures, particularly the question of *which* futures are simulated and how this interacts with the processes described by **Enz and Tamir**. And an especially important issue for understanding decision-making across levels of analysis is to model the feedback loops between individual and social processes that contribute to narrative evolution (**Jones and Hilde-Jones**).

Space precludes us from speculating about these fascinating issues in detail; but crucially, it would *be* speculation. We hope that other researchers will continue to take up these important questions so that we need no longer speculate.

In the target article, we closed with our hope that CNT can provide a common vocabulary and motivate a shared set of questions for the decision sciences. The commentaries have fueled our belief that CNT will prove to be, indeed, a useful narrative.

References

Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.

Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, *76*, 169–217.

Beckert, J., & Bronk, R. (Eds.) (2018). *Uncertain futures: Imaginaries, narratives, and calculation in the economy*. Oxford, UK: Oxford University Press.

Caplan, B. (2007). *The myth of the rational voter: Why democracies choose bad policies*. Princeton, NJ: Princeton University Press.

Chater, N., & Oaksford, M. (1999). Ten years of the rational analysis of cognition. *Trends in Cognitive Sciences*, *3*, 57–65.

Cushman, F. (2020). Rationalization is rational. *Behavioral and Brain Sciences*, *43*, e28.

Dawkins, R. (1976). *The selfish gene*. Oxford, UK: Oxford University Press.

Feigenson, L., Dehaene, S., & Spelke, E. (2004). Core systems of number. *Trends in Cognitive Sciences*, *8*, 307–314.

Fenker, D. B., Waldmann, M. R., & Holyoak, K. J. (2005). Accessing causal relations in semantic memory. *Memory & Cognition*, *33*, 1036-1046.

Fenton O'Creevy, M., & Tuckett, D. (2022). Selecting futures: The role of conviction, narratives, ambivalence, and constructive doubt. *Futures & Foresight Science*, *4*.

Fodor, J. A., & Pylyshyn, Z. (1981). How direct is visual perception? Some reflections on Gibson's "ecological approach." *Cognition*, *9*, 139–196.

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, *11*, 127–138.

Griffiths, T. L., Chater, N., Norris, D., & Pouget, A. (2012). How the Bayesians got their beliefs (and what those beliefs actually are): Comment on Bowers and Davis (2012). *Psychological Bulletin*, *138*, 415–422.

Hirst, W., Yamashiro, J. K., & Coman, A. (2018). Collective memory from a psychological perspective. *Trends in Cognitive Sciences*, *22*, 438–451.

Holyoak, K. J. (1985). The pragmatics of analogical transfer. *Psychology of Learning and Motivation*, *19*, 59–87.

Horne, Z., Muradoglu, M., & Cimpian, A. (2019). Explanation as a cognitive process. *Trends in Cognitive Sciences*. Advanced online publication.

Johnson, S. G. B., & Ahn, W. (2015). Causal networks or causal islands? The representation of mechanisms and the transitivity of causal judgment. *Cognitive Science*, *39*, 1468–1503.

Johnson, S. G. B., & Ahn, W. (2017). Causal mechanisms. In M. R. Waldmann (Ed.). *Oxford Handbook of Causal Reasoning* (pp. 127–146). New York, NY: Oxford University Press.

Johnson, S. G. B., Jin, A., & Keil, F. C. (2014). Simplicity and goodness-of-fit in explanation: The case of intuitive curve-fitting. In P. Bello, M. Guarini, M. McShane & B. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 701–706). Austin, TX: Cognitive Science Society.

Johnson, S. G. B., & Keil, F.C. (2014). Causal inference and the hierarchical structure of experience. *Journal of Experimental Psychology: General*, *143*, 2223–2241.

Johnson, S. G. B., Merchant, T., & Keil, F.C. (2020). Belief digitization: Do we treat uncertainty as probabilities or as bits? *Journal of Experimental Psychology: General*, *149*, 1417–1434.

Johnson, S. G. B., Murphy, G. L., Rodrigues, M., & Keil, F. C. (2019). Predictions from uncertain moral character. In A. K. Goel, C. M. Seifert, & C. Freksa (Eds.), *Proceedings of the 41st Annual Conference of the Cognitive Science Society* (pp. 506–512). Austin, TX: Cognitive Science Society.

Johnson, S. G. B., Rajeev-Kumar, G., & Keil, F. C. (2016). Sense-making under ignorance. *Cognitive Psychology*, *89*, 39–70.
Johnson, S. G. B., Valenti, J. J., & Keil, F.C. (2019). Simplicity and complexity preferences in causal explanation: An opponent heuristic account. *Cognitive Psychology*, *113*, 101222.
Kahan, D. M., Peters, E., Dawson, E. C., & Slovic, P. (2017). Motivated numeracy and enlightened self-government. *Behavioural Public Policy*, *1*, 54–86.
Keynes, J. M. (1937). The general theory of employment. *Quarterly Journal of Economics*, *51*, 209–223.
Korman, J., & Khemlani, S. (2020). Explanatory completeness. *Acta Psychologica*, *209*, 103139.
Lagnado, D. A., Waldmann, M. R., Hagmayer, Y., & Sloman, S. A. (2007). Beyond covariation: Cues to causal structure. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 154–172). New York, NY: Oxford University Press.
Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*, e1.
Lombrozo, T. (2007). Simplicity and probability in causal explanation. *Cognitive Psychology*, *55*, 232–257.
Lombrozo, T. (2016). Explanatory preferences shape learning and inference. *Trends in Cognitive Sciences*, *20*, 748–759.
Marr, D. (1982). *Vision*. San Francisco, CA: Freeman.
Mills, C. W. (1940). Situated actions and vocabularies of motive. *American Sociological Review*, *5*, 904–913.
Murphy, G. L., & Ross, B. H. (1994). Predictions from uncertain categorizations. *Cognitive Psychology*, *27*, 148–193.
Newell, A. (1973). You can't play 20 questions with nature and win. In W. G. Chase (Ed.), *Visual information processing* (pp. 287–307). New York, NY: Academic Press.
Rottman, B. M., & Keil, F. C. (2012). Causal structure learning over time: Observations and interventions. *Cognitive Psychology*, *64*, 93–125.
Sanborn, A. N., & Chater, N. (2016). Bayesian brains without probabilities. *Trends in Cognitive Sciences*, *20*, 883–893.
Savage, L. J. (1954). *The foundations of statistics*. New York, NY: Wiley.
Scherer, K. R. (2005). What are emotions? And how can they be measured? *Social Science Information*, *44*, 695–729.
Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, *41*, 1–19.
Taleb, N. N. (2007). *The black swan*. New York, NY: Random House.
Walasek, L., & Brown, G. D. A. (in press). Incomparability and incommensurability in choice: No common currency of value? *Perspectives on Psychological Science*.
Woodward, J. (2003). *Making things happen: A theory of causal explanation*. New York, NY: Oxford University Press.
Yin, P., & Sun, J. (2021). Is causation deterministic or probabilistic? A critique of Frosch and Johnson-Laird (2011). *Journal of Cognitive Psychology*, *33*, 899–918.
Zacks, J. M., & Tversky, B. (2001). Event structure in perception and conception. *Psychological Bulletin*, *127*, 3–21.
Zhu, J., & Murphy, G. L. (2013). The effect of emotionally charged information on category-based induction. *PLoS ONE*, *8*, e54286.