

Psychological Review

Toward Nonprobabilistic Explanations of Learning and Decision-Making

Aba Szollosi, Chris Donkin, and Ben R. Newell

Online First Publication, April 7, 2022. <http://dx.doi.org/10.1037/rev0000355>

CITATION

Szollosi, A., Donkin, C., & Newell, B. R. (2022, April 7). Toward Nonprobabilistic Explanations of Learning and Decision-Making. *Psychological Review*. Advance online publication. <http://dx.doi.org/10.1037/rev0000355>

THEORETICAL NOTE

Toward Nonprobabilistic Explanations of Learning and Decision-Making

Aba Szollosi, Chris Donkin, and Ben R. Newell
School of Psychology, University of New South Wales

Referring to probabilistic concepts (such as randomness, sampling, and probability distributions among others) is commonplace in contemporary explanations of how people learn and make decisions in the face of environmental unknowns. Here, we critically evaluate this practice and argue that such concepts should only play a relatively minor part in psychological explanations. To make this point, we provide a theoretical analysis of what people need to do in order to deal with unknown aspects of a typical decision-making task (a repeated-choice gamble). This analysis reveals that the use of probabilistic concepts in psychological explanations may and often does conceal essential, nonprobabilistic steps that people need to take to attempt to solve the problems that environmental unknowns present. To give these steps a central role, we recast how people solve these problems as a type of hypothesis generation and evaluation, of which using probabilistic concepts to deal with unknowns is one of many possibilities. We also demonstrate some immediate practical consequences of our proposed approach in two experiments. This perspective implies a shift in focus toward nonprobabilistic aspects of psychological explanations.

Keywords: learning, decision-making, unknowns, probability theory

Supplemental materials: <https://doi.org/10.1037/rev0000355.supp>

The application of probability theory as a model of unknowns that emerge in complex physical environments has been immensely successful in scientific enquiries. This success has provided an enduring inspiration for psychological theories of how people generally deal with the problems associated with such unknowns. Indeed, most research traditions concerned with learning and decision-making at least partly adopted concepts from probability theory in their explanations—a tendency prominently recognizable in Bayesian (Chater et al., 2006; Oaksford & Chater, 2007; Tenenbaum et al., 2006) and sampling models of cognition (Busemeyer et al., 2019; Fiedler & Juslin, 2006; Ratcliff et al., 2016; Stewart et al., 2006) to name only a few.

In this article, we critically evaluate the role that probabilistic concepts may play in theories of learning and decision-making. Although criticism of particular implementations of probabilistic theories of cognition already exists (e.g., of Bayesian models; Bowers & Davis, 2012; Jones & Love, 2011), the nature and the focus of our present critique is different: Because probabilistic concepts have seeped through a broad range of theories of learning and decision-making (not only Bayesian), we do not focus on the shortcomings of any particular framework, but rather examine and offer potential solutions to the problems associated with the general reliance on such concepts in psychological explanations.

We start with a brief overview of the history and current influence of probability theory in psychological theorizing. We then evaluate how probability theory and its related concepts can help deal with unknowns. Based on this analysis, we argue that the role that probabilistic concepts may play in psychological theories is necessarily minor, mainly because their application is tied to sophisticated argumentation, background knowledge, and assumptions that laypeople rarely possess. More problematically, the reliance on such concepts may and often does conceal the steps required to deal with unknowns. We illustrate these problems and sketch a potential way forward through a theoretical and experimental analysis of sequential effects in probabilistic experimental designs.

Probabilistic Models: From Tools to Theories

Probability theory is a mathematical theory concerned with assigning probability values to potential outcomes in abstract mathematical spaces (e.g., Kolmogorov, 1950). Its simplest iterations consist of a *sample space* of possible outcomes, the actual occurrence of which are governed by *random variables*. The allure of this theory comes from this latter random variable, which can be

Aba Szollosi  <https://orcid.org/0000-0003-3457-542X>

Chris Donkin  <https://orcid.org/0000-0002-4285-8537>

Ben R. Newell  <https://orcid.org/0000-0003-1898-205X>

Aba Szollosi is now at Department of Psychology, University of Edinburgh.

This work was supported by the Australian Research Council (DP160101186, DP190101675). Portions of the data were presented at the 46th Annual Conference of the Australasian Society for Experimental Psychology (Wellington, New Zealand), at the 27th Subjective Probability, Utility and Decision Making Conference (Amsterdam, Netherlands), and at the 60th Annual Meeting of the Psychonomic Society (Montreal, Canada). All data and the analysis script are available in the online Supplemental Materials and at the following link: <https://osf.io/6thwz/>. The authors are grateful to Jake Embrey for his assistance in data collection.

Correspondence concerning this article should be addressed to Aba Szollosi, Department of Psychology, University of Edinburgh, 7 George Square, Edinburgh EH8 9JZ, United Kingdom. Email: aba.szollosi@gmail.com

used to express expectations about outcome occurrences in terms of *probability* values. Probability theory achieved its popularity due to its potential to model a broad range of environments because, in such models, unknown causes of variability can be replaced with the concept of randomness. More specifically, probability distributions can provide a robust mathematical approximation for the variation in the outcomes produced by many complex physical environments.

Perhaps most prominent of such applications is as models of games of chance (which inspired the development of probability theory; Hacking, 2006). For an example, consider how the roll of a six-sided die would be modeled using probability theory. As a first step, one would define the part of the model that represents the anticipated physical outcomes of the die roll: In this case, a mathematical sample space of outcomes ranging 1–6 could be specified to model the six sides of the die, based on the argument that one of these sides is expected to face up after the roll (and that alternative occurrences, such as the die breaking during the roll, are irrelevant for the purposes of the model). Next, one would define the part of the model that represents the physical process of a die roll. In this case, a discrete random variable uniformly distributed across the six possible outcomes could be specified, based on arguments as to why the die is fair (e.g., its sides are of even size, the weight of the material it is made of is distributed evenly). Once these parts are in place, one may use the model to substitute expectations of the outcome of the physical die roll with the expectations of the random variable of the model (e.g., in the long run, one should expect all sides to be rolled a similar number of times).

In addition to its popularity for modeling games of chance, probability theory has also been influential among scientists. Probabilistic models are routinely used to model the behavior of other physical systems, ranging from relatively simple systems, such as gases or liquids, to more complex ones, such as the weather, the climate, traffic flows, or pandemics. The application of such models is widespread in psychological science as well. Intriguingly, however, psychologists have gone beyond using them as tools to model people's behavior and have made probabilistic models—or at least some of the concepts associated with them—integral parts of theories regarding how humans learn and make decisions.

Psychological theories frequently invoke probabilistic concepts both in the explanation of people's cognitive capacities that supposedly help them deal with unknowns, and when characterizing the unknowns of their environment. On the one hand, people's cognition is often explained in probabilistic terms: People are said to *sample* outcomes from the environment and from their memory (Fiedler, 2000; Stewart et al., 2006), to learn *probability distributions* (Griffiths & Tenenbaum, 2006; Sanborn & Beierholm, 2016), to perceive *randomness* (Ayton & Fischer, 2004; Bar-Hillel & Wagenaar, 1991; Hahn & Warren, 2009; Nickerson, 2002; Reimers et al., 2018), and to judge the *likelihood* or *probability* of events (Busemeyer et al., 2011; Costello & Watts, 2014; Kahneman & Tversky, 1979; Tversky & Kahneman, 1973). In a similar vein, such explanations are often complemented by references to the statistical structure of the environment (Simon, 1956). Accordingly, they often describe such environments in terms of correlational and causal statistical structure of events (Griffiths & Tenenbaum, 2005; Lagnado et al., 2005; Pearl, 2000; Waldmann & Holyoak, 1992), such as cues and outcomes (Brunswik, 1955; Juslin et al., 2003; Karelaia & Hogarth, 2008; Pleskac & Hertwig, 2014), or rewards and punishments (Dayan & Niv, 2008; Rescorla & Wagner, 1972; Sutton & Barto, 1998).

Yet, while scientific tools can sometimes serve as useful inspirations for theorizing (Gigerenzer, 1991), it is important to understand where such metaphors break. In the current case, the problem is that, since complex systems do not actually behave probabilistically, the application of probability theory nontrivially depends on the arguments and background knowledge as to why we can *pretend* that outcome variability is probabilistic—as we briefly outlined in the case of the die throw. Yet, the application of probabilistic concepts to real-world phenomena is often treated as self-evident allowing the explanation of how people take these necessary steps to be overlooked in psychological theories.

Overview of Our Approach

Our goals in this article are to demonstrate the potential pitfalls and problematic consequences of introducing probabilistic concepts into psychological explanations and to introduce a way toward the resolution of these issues.¹ The general approach we are taking to unpack these arguments is to use probabilistically generated experimental environments as a demonstrative tool: Because people's learning and decision-making in the face of environmental unknowns is predominantly studied in such environments, they provide a natural context to explore common problems and to map out potential solutions.

We start with a theoretical analysis of a standard experimental design through which we illustrate how the typical use of probabilistic concepts in psychological explanations may obscure essential aspects of learning and decision-making. Based on this analysis, we outline an alternative approach for studying how people deal with environmental unknowns that does not necessarily invoke probabilistic concepts. We then provide a case study of this approach in two novel experiments.

To foreshadow, our main argument is that the use of probabilistic concepts allows intermediary steps of dealing with unknowns to be left implicit, and thus these steps may be introduced as assumptions into the explanation—while leaving the psychological mechanisms motivating those assumptions unexplained. To avoid the resulting issues, we suggest that probabilistic concepts do not take on central roles in psychological explanations and that instead the questions of how people attempt to create and apply models of the environment should be prioritized.

Hidden Steps of Probabilistic Explanations: A Theoretical Analysis

Although the mathematical foundation of probability theory is relatively uncontroversial, the way in which the theory should be applied to physical phenomena has historically been and continues to be a thorny issue (for summaries of this debate see, e.g., Hájek, 2019; Schwarz, 2018). The main difficulty stems from the long-recognized problem that the structure of the environment is not manifest through observation (Popper, 1963; more recently in psychology, Brehmer, 1980; Brette, 2019; Felin et al., 2017; Szollosi & Newell, 2020). Consequently, it is not trivial how to represent environmental unknowns probabilistically.

¹ Even though the use of probabilistic concepts differs substantially across psychological theories (e.g., rational, approximate, or “as-if” uses), we argue that the problems that we are highlighting emerge irrespective of interpretation. Therefore, we introduce the main problem without clearly differentiating between these uses but provide a more detailed analysis of each of them in the General Discussion section.

In the first part of this section, we provide a theoretical analysis of how probabilistic concepts are usually used in learning and decision-making research. We do this by showing how two common representations that include probabilistic concepts are applied to explaining performance in a typical decision-making task. In the second part, we highlight how in these applications—and when using probabilistic concepts in psychological explanations more generally—many important steps of the application which connect the representation to the environment can be often left implicit and, therefore, left out of the substantive explanation of how people deal with the problem. Lastly, we sketch and argue in favor of an alternative framework that does not start from the assumption of probabilistic concepts and demonstrate how it avoids the problems associated with them. In this alternative framework, we explain the way people deal with unknowns in their environments in more general terms of the generation and evaluation of hypotheses,² under which the probabilistic approach is only one of many possibilities.

Representing Probabilistically Generated Experimental Environments

Structure of Experimental Environments

Studies of learning and decision-making often model real-world unknowns as probabilistically generated environments in the lab (providing yet another example of the extent to which probabilistic concepts have become an integral part of psychological research). The key similarity across such experimental settings is the use of environments in which events of interest are generated according to a quasistochastic rule and the participants' task is to figure out some parameter of this rule. This connection ties together a large body of learning and decision-making research, from the study of the effects of punishments and reinforcements on animal and human behavior (Rescorla & Wagner, 1972; Sutton & Barto, 1998), to research investigating how people think about probabilities (Estes, 1950; Tversky & Kahneman, 1973), correlational and causal relationships (Brunswik, 1955; Waldmann & Holyoak, 1992), and about random processes in general (Bar-Hillel & Wagenaar, 1991; Nickerson, 2002). Based on this similarity, we can treat the experimental environments used in such studies as isomorphic for our current purposes.

As a demonstrative example, we use a repeated-choice gambling task (Barron & Erev, 2003; Hertwig et al., 2004)—a simple variation of probabilistic experimental designs. An illustration of this task can be seen in Figure 1A. In this hypothetical experiment, participants are asked to choose repeatedly between two options that provide outcomes in an (to the participants) unknown manner. The participants are instructed to lose as few points as possible. Clicking on one of the options reveals the outcomes of both options and the loss associated with the chosen option. By repeating this choice across a number of trials, the participants are expected to learn about the process that generates the outcomes. The outcomes are generated by a pseudorandom sampling process, which provides a loss of 20 points with a probability of .50 (a relatively bad outcome) and 0 points otherwise (a relatively good outcome) on the risky option (left option on Figure 1A), and a loss of 10 points with a probability of 1 on the safe option (right option on Figure 1A).

In this context, the unknowns that participants are expected to learn about are the outcomes of the risky option. To illustrate how

people's learning and decision-making under such conditions is often explained, we will present two probabilistic representations that could be used to learn about these outcomes in the following subsections. We will apply these representations to the first 14 trials of the hypothetical experiment (Figure 1B) and use them to derive expectations about the occurrence of the bad outcome (−20). These are simplified variations of two commonly assumed representations in psychological explanations of such tasks: temporally static and dynamic representations. In the former, the temporal structure of the environment is entirely ignored, and the focus is placed only on the probability or frequency with which different outcomes occur. In the latter, the temporal structure of the environment is considered, for example, by keeping track of the number of trials between occurrences of a particular outcome (e.g., the bad one). Although most psychological theories contain features of both representations, the hidden steps required for the application of probabilistic concepts can be demonstrated more clearly when they are examined separately. We also note that some of the mathematical tools that we introduce in these sections will also be used for the analysis of our experimental results, which will be presented in subsequent sections.

Temporally Static Representations

One way to probabilistically represent such environments is based on the assumption of a stochastic generation process. In this case, unknown aspects of the environment (the probability with which the outcomes of the risky option occur) are assumed to be the result of a random draw from the statistical distribution of 0 and −20 outcomes. Under these assumptions, the temporal structure of the outcomes can be ignored, and expectations about the outcomes (i.e., about their probability) can be based on tracking the frequency with which they occurred before, which can then be used to inform choice.

Mathematically, the tracking of outcome probabilities of the risky option can be expressed by using a cumulative average formula to update the estimate of the probability of each of the outcomes, p_i , on trial t :

$$p_i^t = \frac{p_i^{t-1}(t-1) + R_i^t}{t}, \quad (1)$$

where i indicates the outcome being tracked (e.g., −20 and 0), and R_i is 1 for the outcome that occurred and 0 for other outcomes. The idea is that, in the long run, tracked p_i will converge on the real p value with which the outcomes were generated. This probability value can then be used to calculate the expected value (EV) of an option. This can be calculated, for example, for the risky option as:

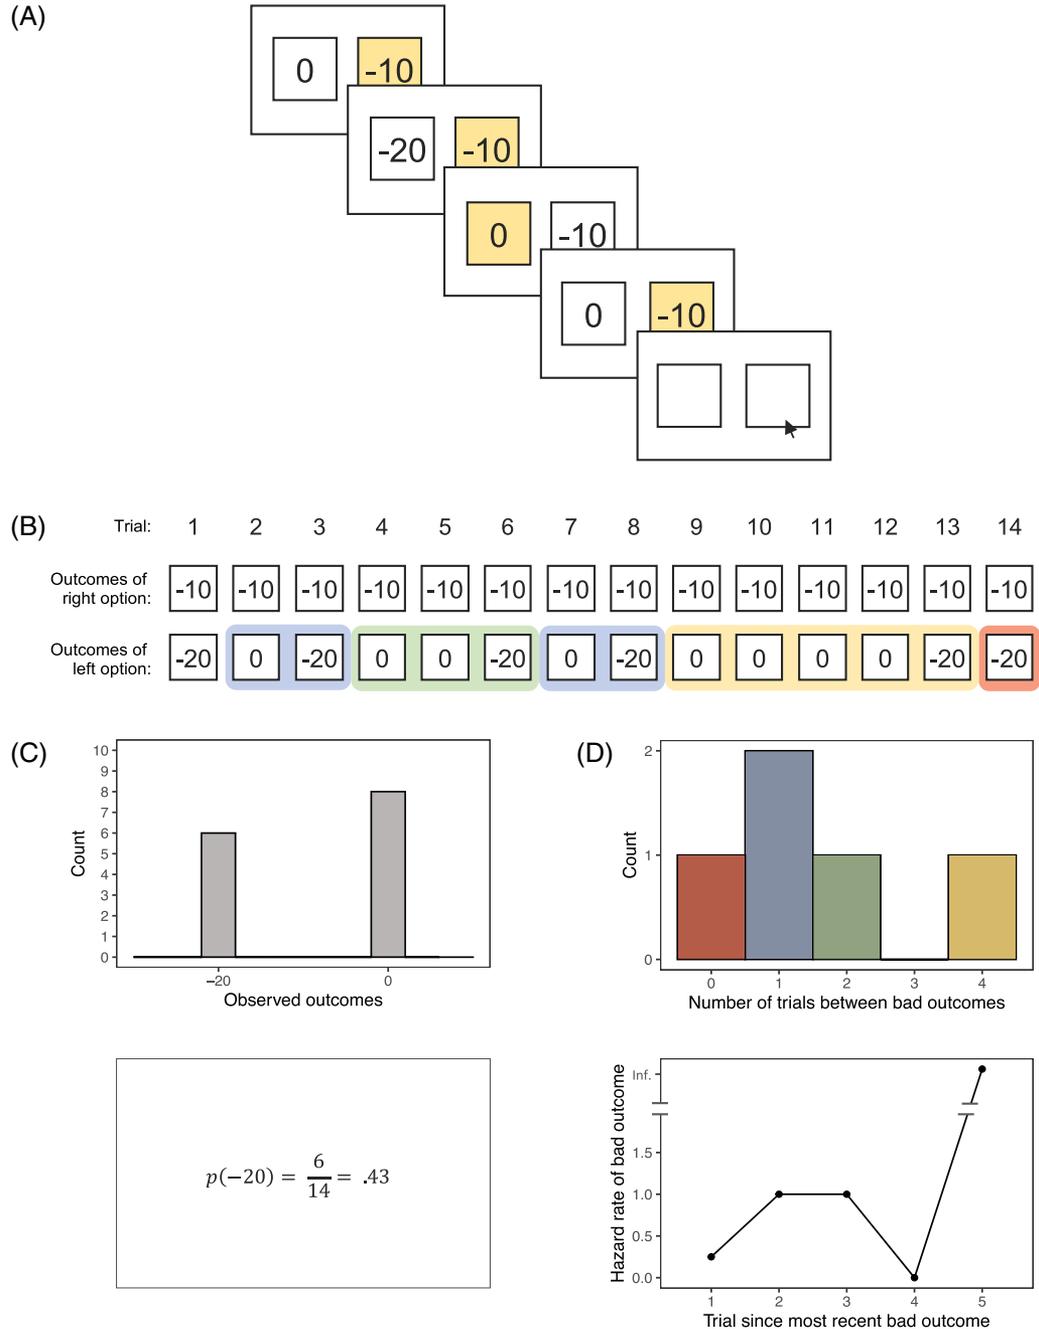
$$EV_{risky} = \alpha \sum_i p_i \cdot V_i, \quad (2)$$

where V_i is the set of outcome values, and α can be used to introduce different subjective utility weights (e.g., Kahneman & Tversky, 1979),

² Although it is common to use the word “hypothesis” as though it inherently refers to probabilities, this need not be the case: People's hypotheses may or may not include references to probabilistic concepts. The difference is important because whereas “probabilistic hypotheses” must be unambiguous and coherent according to the axioms of probability theory, people's actual hypotheses can be ambiguous or incoherent. We explain this difference in more depth in the General Discussion section.

Figure 1

Illustration of a Hypothetical Repeated-Choice Gambling Task and Its Potential Probabilistic Representations



Note. (A) Illustration of the experimental task. Participants choose between unmarked options and receive feedback of chosen (highlighted) and nonchosen outcome. This panel illustrates the history of the task for four trials, and the unmarked options of the current trial. (B) History of outcomes on the first 14 trials of the experiment. (C) Expectations for the occurrence of the bad outcome (−20) after 14 trials, based on a temporally static representation of the task. Expectation of this outcome (in this case probability, p , shown in bottom panel) is calculated based on the number of times it occurred over 14 trials (shown in top panel). (D) Expectations for the occurrence of the bad outcome (−20) after 14 trials, based on a temporally dynamic representation of the task. Expectation of this outcome (hazard rate, shown in bottom panel) for any trial can be derived based jointly on the number of trials since the bad outcome last occurred (x -axis) and on the distribution of previously observed sequence lengths (i.e., the number of good outcome trials between bad outcome occurrences; top panel color matched with panel B). See main text for more details and specific examples. See the online article for the color version of this figure.

and then the option with the highest *EV* on any given trial will be chosen.

An illustrative use of this representation can be shown in the hypothetical experiment in Figure 1. For example, expectation about the risky option can simply be expressed as the function of the number of times an outcome occurred out of the number of all trials that a participant has seen by that point. Early on, this value will vary substantially from trial to trial, but as shown in panel C, by the 14th trial, the *p* value would start to converge on .50 (the value will be .43). This value can then be used in a choice rule to make predictions about people's behavior (e.g., using Equation 2).

Temporally Dynamic Representations

Another possible probabilistic way to represent the experimental environment is through a model that takes the temporal structure of outcome occurrences (i.e., the historical order in which they appeared throughout the trials) into account. Unlike temporally static representations, temporally dynamic representations do not assume that individual trials are independent. Instead, unknowns are assumed to result from a random process that draws *sequences of outcomes* from a distribution of such sequences, and thus expectations about outcomes can be based on the history of the order in which those outcomes occurred on previous trials. For the example displayed in Figure 1, we will consider outcome sequences for the risky option, and take as a sequence a run of good outcomes (0) ended by the occurrence of a bad outcome (−20), including that outcome (example sequences highlighted with color in Figure 1B).

A potential way to take the sequential structure of the outcomes of the risky option into account is to use the so-called *hazard rate* of the outcomes (cf. Grabenhorst et al., 2019). For the current purposes, the hazard rate can be interpreted as an expectation measure for the occurrence of an outcome (in this case the −20 outcome) at any given time point in a sequence, in part, based on the history of outcome sequences. This history can be mathematically represented as the failure density function $f(t)$, which is the mass function of the lengths of the sequences previously observed, in units of trials t . In other words, it is the distribution of the number of trials it takes for an outcome to repeat. In the example sequence, $f(t)$ for the bad outcome is displayed in the top panel of Figure 1D. This histogram shows the mass function of the observed sequences' lengths in the first 14 trials of the experiment (Figure 1B; the sequences are color matched with the corresponding bar in the histogram in the top panel of Figure 1D).

Once $f(t)$ is known, the hazard rate, $h(t)$, can be calculated according to:

$$h(t) = \frac{f(t)}{1 - F(t)}, \quad (3)$$

where $F(t)$ is the cumulative distribution function of $f(t)$. The denominator, $1 - F(t)$, introduces a type of urgency, so that $h(t)$ also takes into account the temporal position of the current trial relative to the longest possible sequence: The smaller the number of potentially longer sequences, the higher the expectation of the outcome's occurrence. The bottom panel of Figure 1D represents $h(t)$ for the bad outcome based on the first 14 trials. To summarize, this can be interpreted as an expectation value of that outcome to occur on the current trial, and it is based on the number of trials since

its last occurrence and the frequency of the possible (i.e., observed) lengths of the sequences.

The potential relevance of $h(t)$ for choices can be easily seen in the above example: When the expectation of a bad outcome is higher, people should be more inclined to choose the safe option; when the expectation of a bad outcome is lower, they should be more inclined to choose the risky option (this value can also be introduced into an *EV* calculation after normalization). For example, if we continued our hypothetical experiment, at Trial 16 (assuming the bad outcome did not occur on Trial 15) the expectation based on $h(t)$ is relatively higher than at Trial 18 (assuming it did not occur on Trials 15–17) because during the first 14 trials, we observed alternations (0 −20) twice, but not sequences with exactly three good outcomes (0 0 0 −20). On Trial 19, however, the $h(t)$ value is infinity (assuming the bad outcome did not occur on Trials 15–18) because sequences longer than five were not observed (and thus the bad outcome needs to occur). Taken together, a representation based on the hazard rate demonstrates a probabilistic way in which temporal expectations could be considered in a repeated-choice experiment.

Applying Probabilistic Models

Through the above examples, we aimed to illustrate how probabilistic concepts are typically used in explanations of people's behavior on repeated-choice gambles—and in studies of learning and decision-making more generally. In many cases, demonstrating that a set of such probabilistic formulae capture people's behavior to an arbitrary accuracy (sometimes accompanied by an explanation of the aspects of the environment the formulae characterize) is accepted as a sufficient psychological explanation. For these examples, one might say that people learn the probability of outcomes or the probability of sequences (or that there is a mixture of people using either of these strategies). However, on closer examination, it becomes apparent that the application of these models depends on numerous hidden steps—arguments, background knowledge, and assumptions not made explicit in the explanation—that nonetheless play an important part in the related psychological claims.

Hidden Steps

To illustrate what these hidden steps are, let us briefly return to our example of how probability theory is applied to model games of chance and compare it with how psychological models are applied. In the case of throwing a six-sided die, an argument could be, events occur in a way that makes them impractical or infeasible to predict perfectly (e.g., lack of sensitive measurement tools, or time or knowledge to build them), and so one may represent the situational unknowns (e.g., force of throw, air pressure, properties of the table's material) with the statistical concept of randomness. However, this can only be done due to the arguments and background knowledge that specify how the physical environment works, how and what can be measured about it, and with what precision, and why the statistical concept of randomness provides an approximation of the aspects that can only be imprecisely measured. Although these parts of the application can be and often are left unsaid, the eventual probabilistic model of the die throw is only a relatively minor part of it.

The same kind of hidden background knowledge and arguments are present in the application of probabilistic representations to the hypothetical experimental environment. Most of these come from knowledge about how the experimental environment is set up, what the researcher's intentions were, or what they thought the participant inferred about their intentions. For example, two simple assumptions in both the static and dynamic representations are that the environment remains stable over time (i.e., that it will behave in a similar way to how it behaved before), and that there are only two possible outcomes that can occur on the risky option and one on the safe option because these have been programmed into the experiment. Another assumption is the a priori determination of the units that should be taken as observations (i.e., individual outcomes for static, sequences of outcomes for dynamic representation) because this is what the experimenter thought the participant will pay attention to. A more complex assumption still is that both representations rely on the concept of random sampling of outcomes or outcome sequences. Such reliance is usually based on the knowledge that the experimental outcomes are determined by a pseudo-random generator—that is, they are random with respect to the participants' presumed ability to predict them.

Of course, not all theories that invoke probabilistic concepts rely on these same exact commitments, even those that are generally concerned with behavior on these types of tasks. Yet such theories must always make use of similar steps, which are often left implicit in probabilistic explanations for convenience (e.g., because other researchers can be expected to possess the relevant background knowledge for the application of those models). This practice can obscure the theoretical commitments to the corresponding psychological mechanisms.

The Problems Posed by Hidden Steps

The hidden steps we just described are not trivial in any kind of meaningful application of a probabilistic model because the mathematical formulae of the probabilistic models make little sense without the accompanying arguments, assumptions, and background knowledge. This means that psychological explanations include these hidden steps in some form whenever they invoke probabilistic concepts. Leaving them unsaid poses a problem because this way the researcher's arguments, assumptions, and background knowledge can be uncritically incorporated into the psychological explanation—while sidestepping the issue of how people themselves take these steps to deal with unknowns. There are two main ways in which such problems can arise.

One way these hidden steps are incorporated in psychological explanations is to argue that people deal with unknowns by (literally) representing them probabilistically. Such an argument must assume that people take the same hidden steps as the researcher in how they apply probabilistic concepts. This is problematic due to the complex nature of the requisite knowledge and the difficulty associated with applying it to the physical environment. Even in the two simple examples we have provided, there is a complex chain of arguments that allows one to pretend that environmental outcomes are probabilistically generated. Although it is undeniably possible that some people with the requisite knowledge (e.g., researchers) do represent these unknowns using probabilistic concepts, it cannot be the general way in which people attempt to deal with unknowns.

Another way to bring these hidden steps into psychological explanations is to explain how similar or equivalent steps are taken by people without assuming they have the requisite knowledge. Such an argument often takes the form of postulating a set of heuristics that provide approximations to the relevant probabilistic models. In our example, for instance, a win-stay-lose-shift strategy would create choice behavior with temporal dynamics, while a strategy to rely on the average of recent observations (e.g., by sampling them) would lead to more static-looking choices. However, in such solutions, the problem of not explaining which of the many possible probabilistic representations of the task should be considered is merely replaced by not explaining which of these probabilistic representations should be approximated. In other words, this approach similarly takes the argumentative steps, background knowledge, and assumptions as self-evident and thereby does not even attempt to explain how individuals take those steps.³

In both the literal and approximate cases, when probabilistic concepts are introduced, it becomes hard to tell what they mean in terms of the psychological explanation. When people are said to learn probability distributions, perceive randomness, judge likelihoods, or sample outcomes it is not clear why and how probabilities, randomness, likelihoods, and sampling (or their approximations) became part of people's repertoire to deal with unknowns in the environment. When they are said to infer the statistical structure, it is not clear how they chose what to include in their models of causes and outcomes, or rewards and punishments. And more generally, it is not clear why the (itself questionable) starting assumption that the unknowns in the environment *should be* represented using probabilistic concepts should be an integral part of a psychological explanation.

Lack of Manifest Structure

The way in which probabilistic concepts should be applied (if at all) to model environmental unknowns is not self-evident and not part of probability theory but needs to be part of the psychological explanations relying on them.⁴ This is because outcome variability in real environments is not produced by genuinely random processes and scientists' choice to develop models that assume it depends on numerous additional considerations. Invoking probabilistic concepts in a psychological explanation without clearly explaining how these additional considerations—the relevant argumentative

³ A variation of this approach explicitly designates the explanation of these steps as a problem that the scientist needs to solve (e.g., as an exercise in developing an optimal probabilistic model) and thereby not making the explanation of how people solve the problem part of the psychological explanation. Such an approach often relies on the argument that the individual does not need to solve this task (or even be conscious of probabilistic or heuristic knowledge) because it already has been solved by evolution. Although not our focus in the present article, our understanding is that heuristics directly resulting from evolutionary processes are rather rigid (in clear contrast with the many potential ways in which people appear to be able to represent environments) and so would be surprised if they played a role in higher-level psychological processes.

⁴ Although one might argue that it is possible to represent these hidden steps probabilistically (e.g., background knowledge as priors, or argumentation as the potential change in likelihoods), such a representation would only be a description of these steps and not an explanation of how and why they are made in the first place. This is because invoking further probabilistic concepts to explain how probabilistic concepts are applied would lead to an infinite regress.

steps, background knowledge, and assumptions—feature in it renders the psychological explanation of how people deal with unknowns incomplete.

Toward Nonprobabilistic Explanations of How People Deal With Unknowns

The problems we outlined with current probabilistic approaches demonstrate the value in at least attempting to explain learning and decision-making without giving probabilistic concepts a major role. An alternative, not (necessarily) probabilistic way to construe what people do when they face environmental unknowns is in terms of *hypothesis generation and evaluation* (e.g., McKenzie, 2004). In general terms, this means that people attempt to deal with the problem of unknowns in their environments by generating hypotheses based on their background knowledge and then evaluating these hypotheses by at least in part using the feedback provided by their environments. Again, here we use the word “hypothesis” in a broad sense, which can include probabilistic hypotheses (when appropriate background knowledge is present), but also nonprobabilistic ones (e.g., deterministic hypotheses, simple heuristics). Thus, in this framework, probabilistic hypotheses are one of an infinite number of possible kinds of hypotheses people can create.

One key difference between this and primarily probabilistic approaches (at least in our interpretation) is that the central research question is moved from attempting to identify (a set of) individual representation(s) to understanding how such representations develop. That is, because people are capable of representing the environment in an infinite number of different ways, it is their capacity to develop various representations and not the particular representation(s) that they develop that remains invariant across settings (Donkin et al., in press; van Rooij & Baggio, 2021). In this section, we discuss the key methodological, empirical, and analytic consequences of this approach and its differences from probabilistic approaches.

Measuring Task Representation and the Role of Feedback

Developing representations of the environment is central to the hypothesis generation and evaluation framework which places a renewed focus on more *sensitive measurement of people’s task representation* in psychological experiments. Although this can be (and sometimes is) done under probabilistic frameworks, it can also be (and often is) avoided because participants’ task representations are typically independently introduced into the probabilistic model as assumptions or “inferred” retrospectively from people’s choices. This can be seen in our example where a preselected set of potential representations (i.e., static and dynamic) is introduced without measuring what representations participants actually entertained. Similarly, the researcher’s choice as to which representation the participant “really” based their decisions on would typically be made based on which representation fits participants’ choice patterns best. Determining people’s representations this way is problematic because they can create an infinite number of hypotheses (not just variation of probabilistic ones), a lot of which are consistent with any one task or observed choice behavior (e.g., Dulany et al., 1984; Dunwoody, 2009; Lovibond & Shanks, 2002; Shanks & St. John, 1994).

Clear and sensitive measurement of task representations can provide a partial solution to this problem by reducing the reliance on unfounded or outright mistaken assumptions to a large degree (Newell & Shanks, 2014). One way toward sensitively measuring participants’ task representations is to consider what kind of background knowledge they might have of the task and to develop measures on that basis. Background knowledge can come in many forms, but perhaps the two more relevant sources for psychology experiments are knowledge that people acquire in their everyday life (e.g., mathematical knowledge acquired through formal education), and knowledge provided via the experimental instructions. Important but often ignored information in the latter is introduced through (potentially unintentional) communicational acts of the experimenter (Hilton, 1995; Sher & McKenzie, 2006). Information communicated in this form can be interpreted by the participants through the knowledge of shared communicational norms (Grice, 1975). For example, the communicational norm of “relevance” can turn the way in which certain stimuli are presented into a communicative act (as we will argue is also the case in the example probabilistically generated experiments). Considering what background knowledge participants can be expected to possess may constrain the measurement of their task representation, of which we will provide an example in our case study experiments below.

Measuring task representations sensitively is also useful to determine the effects that the experimental feedback can be expected to have. Since the environmental structure is not manifest, the information that it provides can only be a function of the participants’ representations of it (Brehmer, 1980; Estes, 1976). For example, as we have discussed in our examples above, when people interpret a probabilistically generated environment through a temporally static representation, they would consider the frequency with which outcomes occur (providing no information about the temporal distribution of the outcomes); when they interpret it through a temporally dynamic representation, they would also consider the sequential distribution of these same outcomes. As this is true not only for probabilistic representations but more generally, understanding the effect of environmental feedback requires a clearer picture of participants’ representations of experimental tasks.

To summarize, sensitively measuring people’s environmental representations—as opposed to assuming and/or “inferring” them from choices—is critical to evaluate the potential impact of experimental feedback and is also important to evaluate how representations vary and change over time (discussed in the next subsection). Although such measurements can be implemented under probabilistic approaches, due to the customary way of allowing researchers’ knowledge to freely enter psychological explanations of people’s representations, these measurements are not as critical as under the framework we are advocating.

Measuring Representational Variability

A perhaps even more important practical consequence of the hypothesis generation and evaluation framework comes from its implications about the relevant invariances in learning and decision-making. Specifically, if we accept that people can create any representation, the relevant invariances to measure will not be which particular representation(s) they entertain, but rather how these representations develop. The representation(s) participants build are of no special significance (although not irrelevant, see below)

because these representations are mostly a function of their background knowledge, which is not invariant over time (i.e., it changes due to changes in cultural knowledge, formal education, etc.). Thus, what should be the key interest of psychological studies of learning and decision-making are the processes that underlie hypothesis generation and evaluation.

To enable the development of hypotheses, these processes rely on concurrently increasing and decreasing variability of entertained hypotheses (by generating and evaluating them). Therefore, empirically, such processes can be expected to influence the *variability of the representations* people develop. For example, experimental instructions constrain how people adopt their background knowledge to create a representation of the experimental environment, and it is potentially further constrained by experimental feedback. It is easy to understand why this would be the case: If experimental instructions do not clearly constrain the set of representations participants can develop, people's representations can be expected to be more variable than when instructions provide clearer constraints. Similarly, if feedback is uninformative with respect to the entertained hypotheses, we can expect a greater variety of representations than when feedback is relevant and informative.

Although the emphasis here is moved from the task representation to the processes responsible for its development, this does not mean that the former is irrelevant. On the contrary, determining people's task representations is useful to document changes over time and to measure representational variability, and it is also important to derive expectations about the potential effects of feedback—only in this scheme determining the actual task representation becomes a means to understanding hypothesis development instead of being an end in itself. This shift in emphasis also has implications for data analysis and presentation, which we will demonstrate in the subsequent case study experiments (roughly, it implies the depiction and comparison of the variability across and within people's responses, instead of comparing group means or other model-fit estimates).

The experimental approach we are proposing stands in stark contrast with the prevailing approach under probabilistically informed theories (and more generally in psychology) in which the focus is on the evaluation of one or a set of specific representations. Often this is done by comparing people's performance to a predetermined (set of) representation(s), or by estimating what proportion of people are best described by the various predetermined representations. Variability in this framework is assumed to be produced by random processes (guiding both people's choice on the task and their choice between predetermined representations) and so is only estimated based on predetermined probability distributions. Instead, we argue that response variability (across and within individuals) is the critical dependent measure because it emerges from the invariant processes responsible for developing representations.

Beyond Probabilistic Psychological Explanations

To summarize, our proposed theoretical approach differs from primarily probabilistically approaches in two main practical ways. We argued that *sensitively measuring people's task representations* is essential under our approach because it helps to document changes over time and to measure representational variability, and it is also important to derive expectations about the potential effects of feedback. However, since representations can change over time, we also argued that their measurement cannot be the main aim

of experimental studies, and that instead *representational variability* should be focused on. This is because only the processes that are responsible for hypothesis development (i.e., increase and decrease their variability) remain invariant over time.

In the following main section, we provide two experimental case studies of how this approach could work in practice. Continuing with the experimental task from our theoretical analysis, we revisited the common finding that people attempt to identify and exploit the sequential structure of repeated-choice gambling tasks in “probabilistically” generated environments (e.g., Plonsky et al., 2015; Szollosi et al., 2019).

Experiment 1

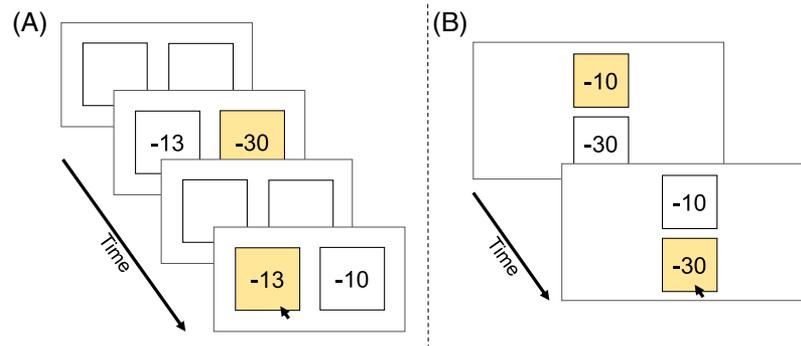
Experimental Approach

Experiment 1 is an illustration of how to study learning and decision-making without relying on probabilistic concepts, using a repeated-choice gambling task in which we primarily focus on measuring people's task representation. To inform how we measure this, as a first step, we consider the relevant background knowledge that people can be expected to hold and that they can adapt to form their task representation. For instance, it is unlikely that they assume that events in the experiment are stochastically independent because even if they had heard about concepts such as randomness or stochastic independence, they are unlikely to have a deep enough understanding of these concepts to apply them in novel circumstances (e.g., Gal & Baron, 1996). On the other hand, there are many reasons for the participants to think that the temporal aspect of the task should not be ignored. For example, the fact that the task has a trial-by-trial structure could be a sufficient reason in itself: A participant might reasonably think that if the experimenter wanted them to ignore the temporal structure, they would have presented the stimuli all at once. Taken together, without the precise knowledge of how the environment was generated, people have good reasons to think that the experimenter wants them to learn about the sequential structure of the task. In this case, they would presumably try to come up with hypotheses with respect to this aspect of the task by trying to identify informative aspects of the temporal structure and to keep track of them using simple mathematical heuristics.

Based on these considerations, we set out to measure the temporal aspect of people's task representation. We did this by asking participants to generate the aspects of interest once they had completed the “actual” experiment, which is a generally promising method to measure people's task representation. We adapted a method used by Tran et al. (2017). In their study, they measured how well people incidentally (i.e., without being instructed) learned the spatial structure of some stimuli by asking them to generate the spatial locations similar to what they observed in the experiment. Here, we adapted this method to measure people's knowledge of the temporal structure of the repeated-choice task and how this knowledge informed their choices.

We also focused our data analysis on the variability in the hypotheses people came up with, even though in this experiment we relied on the standard design of repeated-choice gambles and so did not attempt to manipulate this variability. Relatedly, while our explanation of what people might do did not refer to probabilistic concepts, we still used some such concepts to describe and to visualize both the experimental environment and participants' behavior.

Figure 2
General Structure of the Experiments



Note. (A) Repeated choice task. Participants make repeated choices between a risky and a safe option on 200 trials (two hypothetical trials of the $p = .15$ condition pictured). They are shown the outcome of both the chosen and the nonchosen options (chosen option highlighted). (B) Sequence generation task. Participants are shown the possible outcomes of the risky option. They generate a sequence of good and bad outcomes by clicking on the potential outcome (chosen option highlighted) in the desired order on 200 trials (two hypothetical trials of the $p = .15$ condition pictured). See the online article for the color version of this figure.

Method

Participants, Open Materials, and Data Availability

Ninety students, at the University of New South Wales, participated in Experiment 1. Participants received course credit and performance-contingent payment in exchange for their participation (see the Procedure section, for more details). They also consented to participation in the experiment. The UNSW School of Psychology Ethics Committee (Ref. #2909) approved all experimental work. Materials and data from all experiments, as well as the analysis code, are available in the Supplemental Materials.

Procedure

The experiment consisted of two parts: a repeated-choice gambling task and a sequence-generation task (Figure 2). Participants were only told about the generation task after they completed the repeated-choice task. In the repeated-choice gamble (Figure 2A), participants were asked to choose between a risky and a safe option for 200 trials. The safe option provided the same outcome every time, while the risky option provided two different outcomes: a good and a bad outcome relative to the safe outcome. When participants chose an option, they received full feedback of the payoffs (i.e., feedback on the outcome of both the chosen and nonchosen options), but they were not told what the possible outcomes are in advance. Participants started the experiment with 5,000 points (equivalent of AUD 5.00) and by choosing between the options they lost some number of points. They were told that they will receive a monetary payoff based on the number of points they have left by the end of the experiment.

Participants were randomly assigned to three between-subject conditions, in which we manipulated the probability (frequency) of the bad outcome across three levels (.15, .25, and .50), with the good outcome changing to keep the expected values equivalent to the safe

option. All participants within the conditions observed the same gamble, the payoff structure of which is presented in Table 1.

We aimed to provide a random-like order of the risky outcomes, and so we generated the outcome sequences such that their sequence-length distributions retained a monotonically decreasing shape (an important feature of random sequences). We determined a fixed sequence of good and bad outcomes for each frequency condition that had this property (summarized in Figure 3A, gray bars), and, within each condition, every participant observed this same sequence (see Supplemental Materials, for the actual sequences). The rationale for determining a fixed sequence was that spurious sequential patterns (e.g., in the form of nonmonotonically decreasing sequence-length distributions) can easily emerge in pseudorandomly generated sequences consisting of a relatively small number of trials.⁵

In the second part of the experiment (Figure 2B), participants were asked to generate the sequence of outcomes the risky option would provide if the experiment continued (based on Tran et al., 2017). We asked participants to do the following:

Now we are interested in determining how much of an intuitive sense you have gained for the order in which the [outcome 1] and [outcome 2] numbers occurred on the [risky-side] button. Please show us this by generating a new sequence of [outcome 1] and [outcome 2] numbers. If you think there were any other patterns in the original sequence, please try to generate a sequence that reflects those patterns, too. Do not worry about mimicking the actual order of the original sequence. Just try to produce a sequence of numbers which is as much like the original sequence as you can make it. You may receive additional payment based on how much your sequence is like the original in the first part of the experiment.

⁵ To be more precise, sequential patterns emerge in any finite sequence. Whether the generation of a sequence of outcomes is well approximated by the mathematical random model (where no sequential patterns emerge, under the assumption of an infinitely long sequence) depends on the knowledge of the modeler (e.g., about the computer program that provides the randomization).

Table 1
Pay-Off Structure of Gambles in Experiment 1

$p(\text{bad})$	Good	Bad	Safe
.15	-10	-30	-13
.25	-8	-28	-13
.50	-6	-20	-13

Note. $p(\text{bad})$ indicates frequency of bad outcomes.

Participants were paid an additional AUD 1.00, if the proportion of good and bad outcomes that they generated matched that of the observed outcomes within a range of $\pm 5\%$. Participants generated a sequence of 200 outcomes.

Results

We present the results in the following order: We start with the analysis of the sequences that participants generated, then proceed to analyze people's choices, and lastly we look at the connection between the two. The reason for leading with the analysis of the sequence-generation task is so that we can better explain the types of temporal patterns we should expect in people's choices based on their task representation. Additionally, since visual inspection of the data revealed bimodality in people's responses in the sequence-generation task, our analyses focus strongly on the modal aspects of these responses.

Sequence-Generation Task

The histograms in Figure 3A depict the sequence-length distribution that participants produced (yellow bars) overlaid on what they observed (gray bars) in the sequence-generation task. More specifically, the histograms show the lengths of the sequences (i.e., the number of good outcomes) between repeats of bad outcomes over the course of the whole experiment, similar to our explanation in the Temporally Dynamic Representation section. To reiterate through an example, a hypothetical outcome sequence of B G G G B B G B (G—good outcome, B—bad outcome) would load on the positions 0 (B B), 1 (B G B), and 3 (B G G G B) of the histogram. In the analysis of the experimental results, the sequence lengths were counted starting from the first observed/generated bad outcome.

Visual inspection of these histograms reveals some differences between the generated (Figure 3A, yellow bars) and the observed sequence-length distributions (gray bars) at the group level. Specifically, in the $p = .15$ and $.25$ conditions the generated sequence-length distributions were bimodal, compared to the unimodal distributions that the participants observed. In contrast, in the $p = .50$ condition participants generated relatively accurate distributions (i.e., unimodal, monotonically decreasing). Additionally, in all conditions, participants failed to generate the rare longer sequences they observed, apparent from the relative flatness of the tail of the generated distributions.

The differences between observed and generated distributions at the aggregate level could have resulted from subgroups of participants behaving differently. For example, the bimodality in the leftmost panel of Figure 3A could have resulted from a subgroup

of participants producing a lot of immediate repeats (a sequence of B B; position 0 in the histogram), and another group producing a lot of longer sequences (e.g., a sequence of B G G G B; position 3 in the histogram). Thus, to get a better understanding of individual-level responding, we analyzed the two most-frequent sequences that each individual participant generated (Figure 3B, 3C).

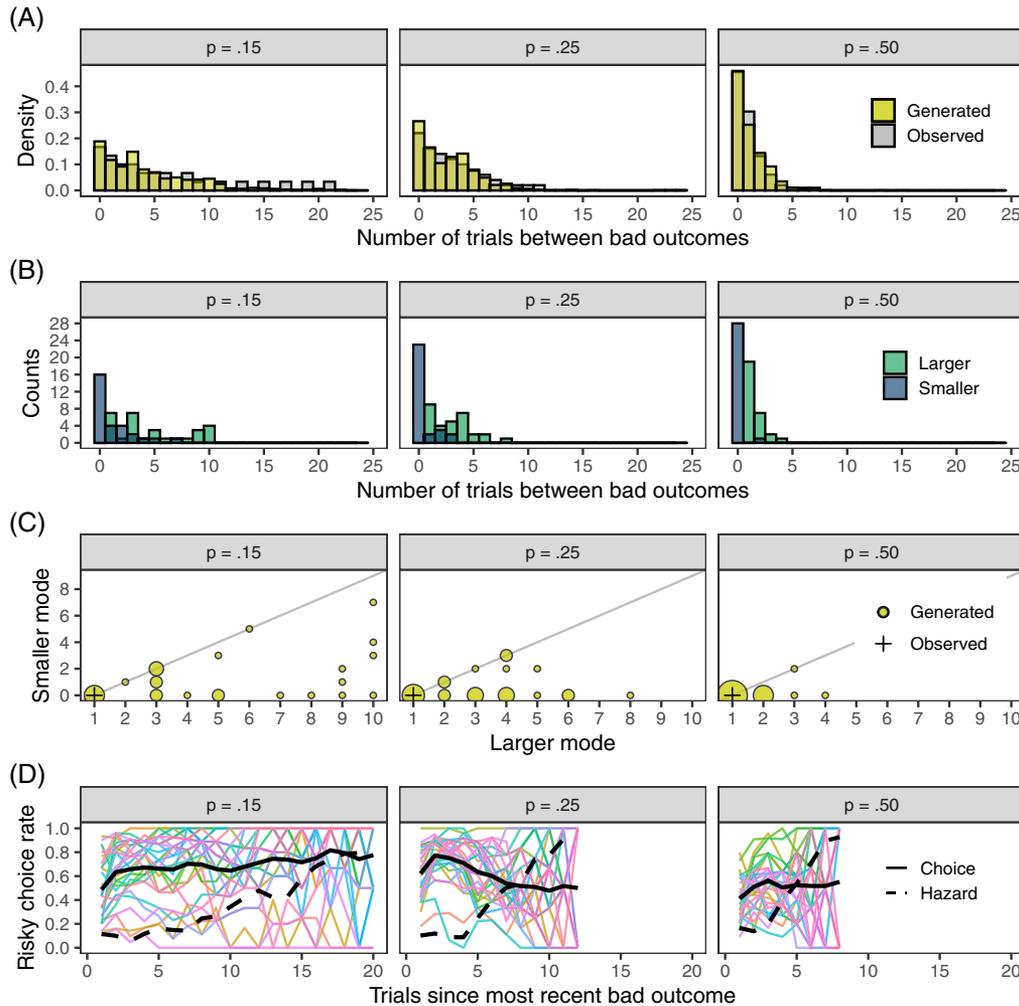
Figure 3B shows the two most-frequent sequences that participants generated, the blue bars representing the numerically smaller and the green bars the numerically larger modes (in case of multiple modes, we selected the numerically smaller). This analysis reveals whether only a subgroup of participants was responsible for the apparent bimodality at the aggregate level, or whether this bimodal pattern holds for individuals. In the case of the former possibility, we should see the two modes distributed relatively evenly (i.e., bimodality in both distributions) because for some participants shorter sequences would be more likely to be modal, while for others longer sequences would be modal (and so *both* of each participant's modes would be plotted at those respective positions). In the latter case, however, we would expect separation for these two modes (i.e., blue and green unimodal distributions) because each participant would have modal shorter and modal longer sequences (and so there would be separation between each participant's modes).

Inspecting Figure 3B, we see that the general bimodal pattern that we observed at the group level holds for at least some participants at the individual level. This can be seen by the clustering of the smaller modes (blue bars) at position 0 (indicating a large number of immediate repeats), while the larger modes (green bars) were more spread out (indicating various longer sequences). This pattern generally held in all conditions, albeit less pronounced in the $p = .50$ condition. However, it is also apparent from the figure that some participants' modal sequence lengths were 0 and 1 (i.e., immediate repeats B B, and alternations B G B), which were the actual most-frequent sequence lengths in the experiment. Thus, it seems that the group-level bimodal pattern masked some participants who generated unimodal distributions (or at least the modes that they produced coincided with the actual modes that they observed); but there were also participants that generated two distinct modes.

To reveal the proportion of participants in each of these subgroups, we looked at the degree of separation between the two modes for each individual. This separation can be seen in Figure 3C, where we plot the covariance of the two most-frequent modes for each individual (yellow circles). The plot also indicates the covariance of the two actual most-frequent modes in the experiment (i.e., immediate repeats B B, and alternations B G B; indicated by crosses in the figure). In all conditions, a substantial number of participants generated sequences with such adjacent modes (24% in the $p = .15$, 30% in the $p = .25$, and 66% in the $p = .50$ condition). Fewer participants generated adjacent modes at other positions, which can be seen by the relatively small number of participants plotted on the diagonal line which represents adjacent frequent modes (17% in the $p = .15$, 17% in the $p = .25$, and 3% in the $p = .50$ condition).

The remaining participants' most-frequent sequences had varying levels of separation between their modes (59% in the $p = .15$, 53% in the $p = .25$, and 31% in the $p = .50$ condition; the distance between

Figure 3
Descriptive Results of Experiment 1



Note. (A) Density histograms of observed (gray bars) and generated (yellow bars) sequence-length distributions (i.e., number of good outcomes between subsequent bad outcomes) in the respective conditions for all participants. (B) Counts of the two most-frequent (numerically smaller—blue bars; numerically larger—green bars) sequence lengths in individuals' sequence-length distributions. (C) Scatterplot of observed and generated smaller and larger modes of sequence distributions for each participant. Size of yellow circles represent number of individuals for that mode pairing. Gray diagonal line represents potential adjacent modes (greater distance from the line represents greater separation of modes). (D) Average risky-choice proportions for individuals and groups (colored and black solid lines, respectively) based on the temporal distance from the most-recent bad outcome. Dashed line represents average hazard rate based on generated sequences. See the online article for the color version of this figure.

these modes is represented by the distance from the gray diagonal). For example, looking at the $p = .25$ condition (the second panel of Figure 3C), the abundance of yellow circles in the bottom row of the figure indicates that the two most-frequent sequence lengths for many participants were immediate repeats (B B) and three- and four-trial long sequences (B G G G B, and B G G G G B, respectively).

Choice Task

We focus on the sequential nature of participants' choices, in light of the results of the sequence-generation task and earlier

findings demonstrating that participants take (supposed) sequential features of the task into account in their decisions (e.g., Plonsky et al., 2015; Szollosi et al., 2019). However, we do note that, at the group level, we replicated the general finding that participants "underweight" rare outcomes in such tasks (Hertwig et al., 2004), revealed by their average preference for the risky option in conditions with rare outcomes (i.e., group-level risky-choice rates of 65% and 67% in the $p = .15$ and $p = .25$ conditions, respectively; there was no rare event in the $p = .50$ condition—the risky-choice rate was 47%).

Figure 3D depicts the temporal structure of participants' risky-choice behavior aggregated at the group level (solid black lines), and

at the individual level (colored lines). On this figure (and in the subsequent analyses), we averaged participants' choices on each trial based on that trial's temporal distance from the most-recent bad outcome (similarly to the analysis of the generated sequences). At the group level, we can see that participants chose as if they thought there was a predictable temporal structure in the task. This can be seen in the recency pattern of participants' choices in the first two trials following the observation of a bad outcome: In each condition, risky-choice rates are always lowest on the first trial, and this is followed by an increase in this rate on the second trial. Additionally, a reversing of this pattern (a decrease in risky-choice rates) can be seen in later trials: Most markedly in the $p = .25$ condition, where risky choices can be seen to taper off gradually after the initial increase. A less pronounced version of this same pattern is apparent in the $p = .50$ condition. In the $p = .15$ condition, the general decreasing pattern is not evident at the group level, which might be the result of the individual-level differences. For some participants, the decrease following the initial increase in risky choices began at different points in the sequence, after which some returned to choosing the risky option, whereas others increased their risky choices consistently after the bad outcome's occurrence.

Relationship Between Choices and Generated Sequences

Lastly, we analyzed the relationship between the risky outcome sequences that participants generated and their choices. To test whether participants' choices were related to their expectations about the outcome sequences, we first calculated a hazard rate for the bad outcomes (according to Equation 3) for each participant based on the outcome sequences that they generated. We did this to get a (admittedly crude) measure of participants' temporal expectation about the outcomes of the risky option. We calculated the hazard rate based on the sequences that participants generated (i.e., using the empirical mass function of the sequence lengths for the bad outcome; yellow bars in Figure 3A).

There were some problems associated with using the raw values of the hazard rates, thus, to ensure some level of comparability between participants, we modified the hazard rate calculation in three ways. First, because the sequence-length distributions had an upper bound for each participant, the hazard rate for this upper bound (the longest sequence that participants generated) was always mathematical infinity (i.e., the bad outcome was predicted with complete certainty). To be able to plot this number, we replaced it with either two times (the value of two was arbitrarily chosen) the otherwise maximum value of $h(t)$, or, in case the largest value was 0 (i.e., if a participant generated only one type of sequence over and over), we replaced it with 1. Second, to convert hazard rate to a probability, we normalized the values for each participant according to $\frac{h(t) - \min(h(t))}{\max(h(t)) - \min(h(t))}$ so that the rates would be bound between 0 and 1. Third, when connecting the normalized hazard rates to the choices, it could sometimes be the case that participants observed sequences in the choice task that were longer than the longest sequence they produced (i.e., when there were no values in the hazard rates for such long sequences). In these cases, we set the hazard rate value to 1 (i.e., the largest possible number), to predict the imminent occurrence of the risky option.

We plotted the group-level average of the (modified) hazard rate in Figure 3D (dashed line). To reiterate, the values of this measure

can be interpreted as one aspect of the expectation about the occurrence of the bad outcome at any given point in the experiment based on participants' retrospectively generated sequences. Thus, we could predict that when expectations of the bad outcome were higher, risky-choice rates would decline, and when its expectation was lower, risky-choice rates would increase. We found that this measure tracked risky-choice rates to some degree. Yet, as we have seen in the case of the generation task, individual-level effects might be masked at the group level.

Therefore, to demonstrate the connection between these measures at the individual level, we fitted generalized linear models (with binomial error models) in which we predicted each participant's choices from their (modified) hazard rates.⁶ The estimated β coefficients reflect the association between participants' hazard-based expectation of the bad outcome and their choice of the risky option, with negative values reflecting consistency (i.e., a tendency to choose the risky option when the bad outcome is not expected). We found that the proportion of participants with negative coefficients was 34%, 80%, and 76% in the $p = .15$, $.25$, and $.50$ conditions, respectively.

Discussion

Through Experiment 1, we aimed to illustrate a way to sensitively measure people's task representation in a repeated-choice gamble and a way to analyze the variability of these representations which arises from hypothesis generation and evaluation.

The analysis of people's task representation revealed well-known temporal features of typical task representation of repeated-choice gambles. Specifically, we found that many participants generated mainly two different types of sequences (indicated by distinct longer and shorter modes in their sequence-length distribution), and their choices appeared to preempt similar patterns of sequential dependencies. In contrast with typical interpretations of such findings, the explanation that we gave implies that these sequential features are not invariant, but rather arise from the background knowledge participants bring to the task and received from the experimental instructions.

In addition, there seemed to be a large degree of individual variability in the extent and type of hypotheses that participants entertained about the sequential structure. Although both participants' background knowledge and the experimental instructions would have reasonably led participants to suspect that they need to attend to the temporal structure, neither were unambiguous with

⁶ There are several reasons to consider the hazard rate as a measure of people's expectation of the occurrence of the bad outcome too crude. For example, the hazard rate was calculated based on retrospective estimates of the sequence lengths, the perception of which was almost certainly different, and possibly gradually changing during the task. The hazard rate also cannot express expectations about trials that are part of sequences longer than a participant generated (and thus had to fall back on a heuristic solution). Finally, the temporal structure may not be the only information that participants considered in their choices. The reason why we decided to analyze the data using this method was only to demonstrate that there was some level of coherence between the temporal effects in people's choices and in the sequences that they generated. High levels of correlation between these measures can show when the participants clearly made their choices based on the perceived temporal dependencies; however, lower levels of correlation between these measures are nondiagnostic because they cannot rule out alternative sources of variability that still take perceived temporal dependencies into account (such as more complicated decision strategies). See the Discussion section of both experiments for a more detailed explanation.

respect to which aspect of the temporal structure that requires attention. Thus, participants might have attempted to generate hypotheses about the temporal structure but came up with different variations of these hypotheses. Additionally, the environmental feedback did not provide sufficiently clear constraints to rule out many of these potential hypotheses, and thus could contribute little to the reduction of this variability. The same arguments can be made for the hypotheses that people entertained about how to make decisions based on their representation of the task structure.

The apparent bimodality in some participants' sequence-length distribution presents an interesting exception to this representational variability because it shows some extent of convergence in the hypotheses participants generated. This could have been the result of these participants forming a hypothesis implying the existence of longer and shorter sequence lengths categories. Generally, convergence in task representations suggests that similar hypothesis selection processes took place, which we will further explore in the next experiment.

Overall, this experiment and its results show an alternative perspective on repeated-choice gambles that does not start with the premise that participants represent the task the same way it was generated—probabilistically—or indeed in any fixed way. Instead, the detailed measurement of participants' task representations revealed wide representational variability but also important similarities, which jointly illuminate the processes responsible for the development of these representations.

Experiment 2

Experimental Approach

In Experiment 2, we further aimed to illustrate how to study hypothesis generation and evaluation. Specifically, in addition to using the same measures of task representation, we aimed to selectively influence the variability of people's hypotheses. We attempted to do this by making experimental feedback more relevant for the hypotheses that participants entertained about the environment. Therefore, we manipulated the sequential structure of the task in the following way. First, inspired by participants' apparent attention to modal aspects of the sequences observed in Experiment 1, in Experiment 2, we showed them both unimodal and bimodal sequence distributions. Additionally, inspired by Tran et al.'s (2017) study, we manipulated the discernibility of these modes by selectively adding "noise sequences" that differed in length from the modal sequences. Lastly, we manipulated the frequency of the bad outcome to be able to show participants shorter and longer sequences, since that also seemed to affect their representations in Experiment 1. This resulted in three main manipulations and thus three between-subjects conditions: (a) modality of the sequence-length distribution (unimodal or bimodal), (b) noise of the sequence-length distribution (noisy or discrete), and (c) frequency of the bad outcome ($p = .10$ or $p = .25$).

The expectation was that the manipulations will selectively decrease the ambiguity of the feedback that participants get in the task by targeting the possible hypotheses that participants entertain. More specifically, we expected the manipulations to increase the constraints on hypotheses that participants might entertain (specifically bimodal more ambiguous than unimodal, noisy more ambiguous than discrete, and longer sequences more ambiguous than shorter sequences).

Method

Participants

We recruited 241 participants at the University of New South Wales. Participants received course credit and performance-contingent payment for their participation.

Procedure

The general structure of Experiment 2 was similar to that of Experiment 1 (Figure 2): Participants first made choices in a repeated gamble task between safe and risky options and then recreated the sequence of outcomes they observed on the risky option. The experiments differed in the outcome structure of the risky option. We manipulated this structure in a $2 \times 2 \times 2$ design across the following factors: (a) the modality of the sequence-length distribution of the good and bad outcomes (unimodal or bimodal); (b) the noise around the modes of the sequence-length distributions (noisy or discrete); and (c) the frequency of bad outcomes ($p = .10$ and $p = .25$; for the respective payoff structure, see Table 2). The sequence-length distributions for the respective conditions are displayed in Figure 4A (gray bars). Every participant within each condition observed the same sequence of good and bad outcomes that we sampled from the respective distribution (i.e., one sequence for each of the eight conditions; see Supplemental Materials, for the full sequences). We inspected these distributions to ensure that the main structural patterns are evenly distributed throughout the experiment (we wanted to avoid issues such as, in the bimodal conditions, the shorter sequences being presented disproportionately early in the experiment). Additionally, to ensure consistency, each condition started off with a bad outcome (this outcome was added as an "extra" bad outcome, thereby slightly increasing the frequency of that outcome in each condition).

Compared to Experiment 1, the sequence distributions in Experiment 2 were simplified: Instead of monotonically decreasing distributions, we introduced a single or two distinguishable modes. The positions of these modes were dependent on the frequency of outcomes: For more-frequent bad outcomes, the lengths of the modal sequences were shorter, while for less-frequent bad outcomes, the length increased. For an example, consider the sequential structure of the $p = .25$, unimodal, discrete condition. The mode of the sequence-length distribution was 3, which means that the bad outcome occurred on every fourth trial (i.e., there were 3 good outcomes between each bad outcome, thus each sequence was B G G G B). In contrast, in the $p = .10$, unimodal, discrete condition, the mode of the sequence-length distribution was 9, meaning that the bad outcome occurred on every tenth trial (i.e., a sequence of B G G G G G G G G B).

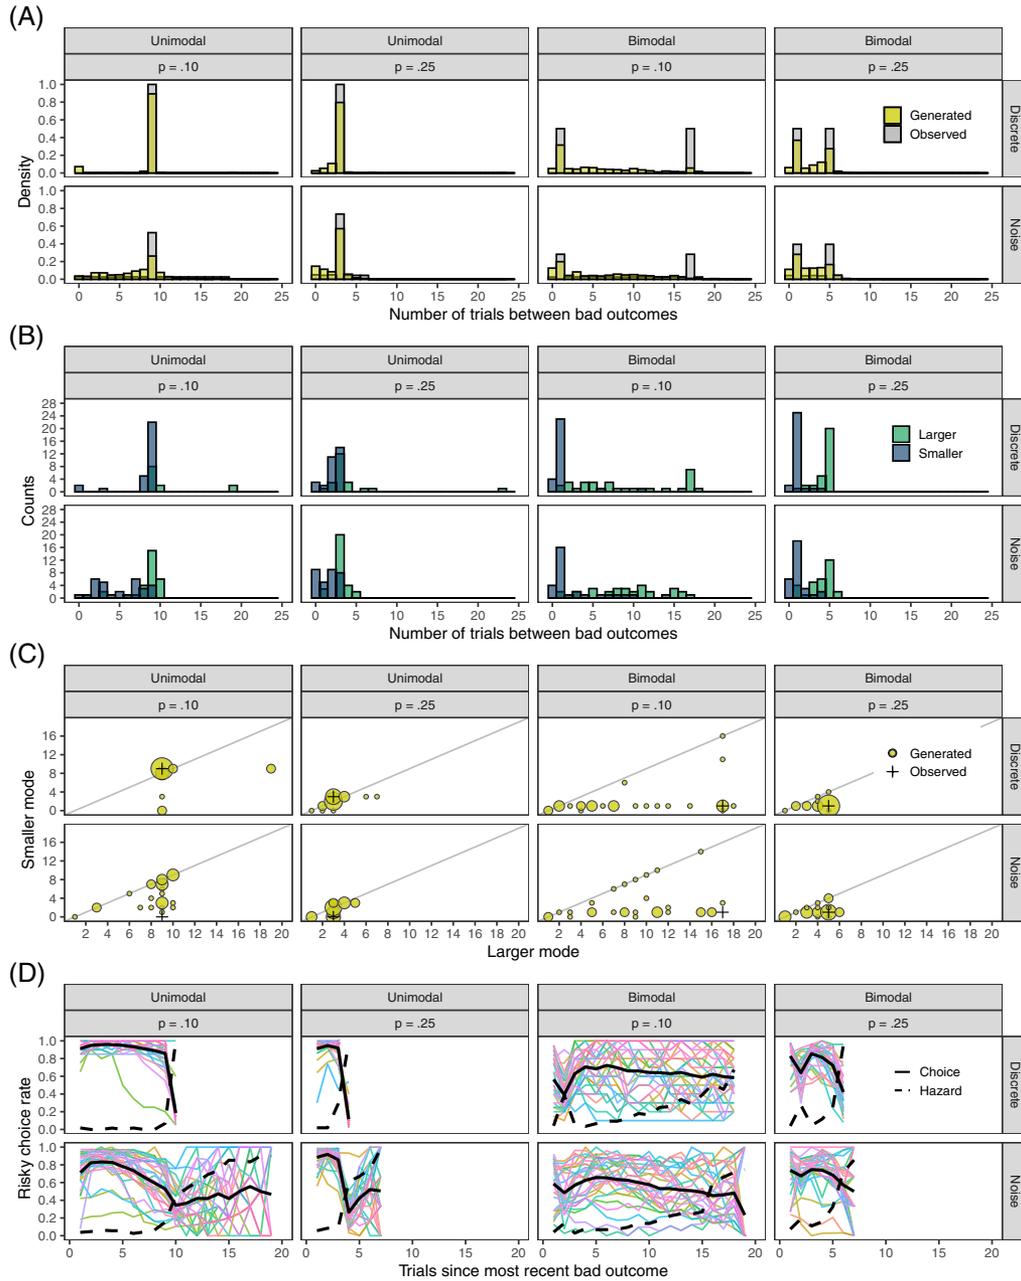
To introduce variance, we added a number of "noise" sequences with various lengths to some conditions. Whenever present, these

Table 2
Pay-Off Structure of Gambles in Experiment 2

$p(\text{bad})$	Good	Bad	Safe
.10	-10	-40	-13
.25	-8	-28	-13

Note. $p(\text{bad})$ indicates frequency of the bad outcome.

Figure 4
Descriptive Results of Experiment 2



Note. (A) Density histograms of observed (gray bars) and generated (yellow bars) sequence-length distributions (i.e., number of good outcomes between subsequent bad outcomes) in the respective conditions for all participants. (B) Counts of the two most-frequent (numerically smaller—blue bars; numerically larger—green bars) sequence lengths in individuals' sequence-length distributions. (C) Scatterplot of observed and generated smaller and larger modes of sequence distributions for each participant. Size of yellow circles represent number of individuals for that mode pairing. Gray diagonal line represents potential adjacent modes (greater distance from the line represents greater separation of modes). (D) Average risky-choice proportions for individuals and groups (colored and black solid lines, respectively) based on the temporal distance from the most-recent bad outcome. Dashed line represents average hazard rate based on generated sequences. See the online article for the color version of this figure.

noisy sequences followed a uniform distribution, such that the number of longer sequences was equivalent to the number of sequences between the mean of the distribution and zero. For the noisy $p = .10$ conditions, each noise sequence was shown once; for the noisy $p = .25$ conditions, each noise sequence was shown three times. One consequence of introducing the noise manipulation was that the number of trials increased in certain conditions. We decided that it was preferable to have an unequal number of trials across conditions, so that the respective modal sequences could be presented the same number of times in all conditions. The number of trials were as follows: 200 trials in all discrete conditions; 390 trials in the $p = .10$, bimodal, noise condition; 284 trials in the $p = .25$, bimodal, noise condition; 380 trials in the $p = .10$, unimodal, noise condition; 272 trials in the $p = .25$, unimodal, noise condition.

Results

In the presentation of the results, we followed the same general logic as in Experiment 1. We start with the analysis of the sequence-generation task. We then derive expectations from participants' responses on this task to evaluate the sequential effects in their choices.

Sequence-Generation Task

Figure 4A displays the group-level sequences that participants produced in the generation task (yellow bars) overlaid on the sequences that they observed in the choice task (gray bars). Visual inspection of the figures suggested that learning accuracy decreased with the increased ambiguity in the sequence distributions. This decrease in accuracy between the conditions can be seen in the heights of the yellow bars relative to the gray bars at the modes: Relatively smaller yellow bars indicate less accurate learning. However, participants were reasonably accurate when recreating the general modal structure of the distributions: The unimodal and bimodal structure of the sequence lengths distributions are readily apparent at the group level in all but the most ambiguous (bimodal, $p = .10$, noisy and discrete) conditions.

Similar to our analysis in Experiment 1, we aimed to clarify these findings by looking at the two most-frequent sequences that each individual generated (the first two modes of each individual's sequence-length distribution; plotted in Figure 4B, 4C). We observed the same general pattern in individuals as we did at the group level, with two notable differences. First, at this level of analysis, it seems that participants learned the general bimodal structure of the task even in the most ambiguous (bimodal, $p = .10$, noisy and discrete) conditions. This result is revealed in Figure 4B by the relatively greater spread of the green bars, and in Figure 4C by most participants' distance from the diagonal (the distance indicates the degree of separation between the modal sequences) in those conditions (third panel from the right in the top and bottom rows).

Another pattern that is important to point out is that learning was more accurate for shorter sequences. This can be seen in that in all bimodal conditions, the smaller mode (alternating sequences) was better learned. Although this was not the case for the unimodal discrete conditions (learning accuracy was at ceiling in both conditions), it was true in the unimodal noisy conditions (the mode in the $p = .25$ condition was better learned than in the $p = .10$ condition). A related finding was that errors (sequences that were

not the modal length) were almost exclusively made in the direction of creating shorter sequences. This can be seen in Figure 4B in that the sequence lengths that were different from the modal length(s) are almost always located at the left of the largest (or only) mode.

Relationship Between Choices and Generated Sequences

Figure 4D shows participants' group (black solid line) and individual-level (colored lines) average risky-choice rates based on the temporal distance from the most-recent bad outcome (for more details on the visualization, see the description for Experiment 1). Participants' choices were connected to their representations of the sequential structure of the environment: Participants' expectation of the bad outcome (based on the hazard rate calculated from participants' retrospectively generated sequences; see the calculations for Experiment 1) was coherent with their choices: When the expectation of the bad outcome's occurrence was lower, their choices were riskier, and vice versa. This correlation can be seen in Figure 4D by the negative association between the group-level hazard rate (dashed black line) and the risky-choice rate (solid black line). The clear connection between expectations and choices was even more evident from the individual-level correlation of hazard rates and risky-choice rates. The proportion of people that had a negative correlation was relatively large in all conditions (see Table 3).

Importantly, individual variability in both the generated sequences and choice behavior increased as a function of the ambiguity in environmental feedback (i.e., bimodal more ambiguous than unimodal, noisy more ambiguous than discrete, and longer sequences more ambiguous than shorter sequences). For the generated sequences, the increased individual variability can be seen in that the spread of the circles in Figure 4C was greater in those respective conditions (indicating larger differences in the modal aspects of the generated distributions). For participants' choices, it can be seen in that the spread of the colored lines in Figure 4D was greater in the respective conditions (indicating larger differences in risky-choice behavior).

Discussion

In Experiment 2, we aimed to provide a clearer example of how to study hypothesis evaluation experimentally. Since in Experiment 1 participants seemed to spontaneously generate hypotheses about the temporal structure of the task, in Experiment 2 we manipulated the feedback about this aspect of the task. The manipulations aimed to selectively reduce the variability of participants' responses by proving their hypotheses wrong to varying extents.

Table 3

Proportion of Participants With a Negative Association Between Hazard Rate and Risky Choice Rate Within the Respective Conditions

Condition	Unimodal		Bimodal	
	$p = .10$	$p = .25$	$p = .10$	$p = .25$
Discrete	97%	100%	57%	93%
Noise	90%	94%	83%	77%

It was apparent from these results that when feedback is relevant and diagnostic for participants' hypotheses, it had a clear effect both on their representation of the temporal structure of the task and on the way in which they made choices. One way in which this can be seen is that participants were better at recreating the temporal structure of this experiment than that of Experiment 1. Participants' learning was generally good for the relevant modal aspects: In most conditions they recreated the modal sequences at the correct positions, and even when they did not, they recreated the modality well. These results can be explained in terms of how ambiguous the environmental feedback was with respect to the hypotheses that participants may have entertained. Specifically, since the modal aspects of the temporal structure of the outcomes were clearer (i.e., simpler and frequently repeating), many of the potential hypotheses that participants possibly entertained could have been ruled out more easily.

Another way the effect of feedback can be seen is through the finding that individual variability decreased with the decrease of ambiguity (similar to Experiment 1). We found that the less ambiguous the environmental feedback was (with regards to the possible hypotheses that participants had), the smaller individual variability seemed to be in both participants' choices and sequence-length distributions. This was presumably the case because different individuals were able to rule out similar hypotheses in the conditions with clear feedback, but not in the conditions with ambiguous feedback. Thus, in general, participants' learning of the aspects of the environment they thought to be important improved as a result of providing more relevant and diagnostic feedback about those aspects.

The reduction of variability can also be seen in the relationship between the generated sequences and participants' choices. In the least ambiguous conditions (unimodal), the hazard function was strongly associated with participants' choices, and this connection was more variable in the more ambiguous conditions (in the bimodal conditions, and also in Experiment 1). A potential explanation for this is that people converge more on what is an advantageous choice strategy in the less ambiguous environment, while the possibility to entertain other choice strategies is greater in more ambiguous environments (e.g., because worse strategies are harder to rule out using the experimental feedback).

Although most of participants' inaccuracies in the sequences they generated (deviations from the actual structure shown to them) can be explained by how much constraint the environmental feedback provided for their hypotheses, not all of them can be. Instead, some inaccuracies seemed to have attentional sources. Specifically, the finding that participants made more errors the longer the modal sequences they needed to track were (e.g., apparent in the variability of the second mode for the bimodal, $p = .10$ condition, and the first mode for the unimodal with noise, $p = .10$ condition), could be explained as a result of errors that people can make when counting the lengths of those longer sequences during the repeated-choice task (i.e., the longer the sequence is that they need to count, the more opportunities they have to become distracted).

General Discussion

In this article, we set out to evaluate the role of probabilistic concepts in psychological explanations of learning and decision-making and argued that the use of these concepts can obscure the steps that are needed to solve the problem of how the structure of an

unknown environment is identified. For an illustration, we revisited experiments on people's learning and decision-making in probabilistically generated environments. Our aim with this reassessment was twofold. First, we used these experimental environments to demonstrate how probabilistically inspired models can conceal the steps that are necessary to form and apply any kind of representation in unknown environments. Second, we outlined and provided an example of an alternative approach to study how people deal with environmental unknowns in their learning and decision-making that does not rely on probabilistic concepts.

In both experiments, we aimed to demonstrate an alternative approach that does not start from the assumption that people represent their environment using tools of probability theory. We argued that under this alternative framework it is essential to sensitively measure people's task representations and illustrated a way to do this in both experiments. These measures revealed a large extent of variability in how people construed the task, but also important similarities. We explained these results as a joint product of people's hypotheses based on their background knowledge (including task instructions) and the experimental feedback—higher levels of variability as the result of hypotheses being under constrained, while lower levels of variability as the result of increased constraints. We also showed how to analyze and visualize data in these terms.

To reiterate, the general finding that people try to identify and exploit sequential patterns in randomly generated environments is not the novel aspect of this research. There is a long history of documenting this pattern of behavior in studies of learning and decision-making (Bar-Hillel & Wagenaar, 1991; Estes, 1976; Gallistel & Gibbon, 2000; Navarro et al., 2016; Plonsky et al., 2015). Instead, what was novel about our studies was the way we explained and treated these results: We provided an explanation of participants actual task representation as only a transient feature of their psychology, which is based on the current knowledge they have, and as such, their identification is not the main result of our study but a tool that illuminates how people generate and evaluate hypotheses. This is radically different from the customary approach which aims for and stops at the supposed identification of one or a set of representations.

Taken together, these studies serve as an example of some of the beneficial influences of this perspective on empirical work, apparent in the experimental design (e.g., sensitive probing of people's task representation, providing specific feedback on aspects that are relevant to them), the analytical technique (e.g., taking response variability as a crucial measure), and the interpretation of the results (e.g., not assuming that representations are stable over time, not assuming extensive background knowledge, deeper consideration of invariant features of learning and decision-making). In the following section, we outline a general theoretical framework for hypothesis generation and evaluation and use it to clarify the status of probabilistic models and to provide a more detailed explanation of the current experimental results.

Intuitive Scientists

Our position is connected to the research tradition that considers laypeople's learning and decision-making similar to what scientists do, generally suggesting that people should be considered a type of intuitive scientists whose behavior is guided by their explanations (or theories) of their environments (e.g., Carey & Spelke, 1996;

Gopnik & Meltzoff, 1997; Kuhn, 1989; Murphy & Medin, 1985). Our view differs from some of these accounts in that we identify the similarities between scientists and people not in the use of any of science's past or current methodologies (such as experimental testing or statistical modeling), but rather in the general way in which they attempt to explain their environments (Szollosi & Newell, 2020).

Explanations, for the current purposes, can be considered statements that provide answers to why and how questions (cf. Keil, 2006; Lombrozo, 2006): They state what they attempt to explain, how they work, and why they work that way, at an arbitrary level of accuracy. We note here again that the way in which we used the words "theories" or "hypotheses" (i.e., to refer to people's explanations) deviates from the common usage in psychology, where these concepts are often used interchangeably with probabilistic models. There are two main advantages of not equating the two: (a) whereas probabilistic hypotheses need to be unambiguous and coherent with the axioms of probability, explanations can be ambiguous or incoherent and (b) probabilistic hypotheses rely on very specific forms of background knowledge (as we have outlined in detail), whereas this need not be the case for explanations (they can rely on less complex and even incomplete background knowledge).

The Role of Theories

How do people attempt to provide and improve explanations of their environment? To reiterate, the general answer that many probabilistic models (and some variations of the intuitive scientist theories) imply is that theories are derived from or are suggested by the environment (or that at least in part they are). This is often suggested to be done on the basis of correlations between cues and outcomes, or presumably obvious effects in the environment, or some sampling process. These ideas presuppose that there is some self-evident structure in the environment for what counts as a cue and an outcome, what effects there are to look for, or what to sample—yet, as we have argued, no such structure is ever self-evident in physical environments. Instead, we argued that theories always need to be generated before they can be applied to the environment.

This solution is based on Popper's (1963; see also Deutsch, 2011) idea of knowledge generation. In this conception, theories are generated via creative recombination or reorganization of the parts of existing theories and are subsequently evaluated through different forms of criticism. These criticisms can come in the form of other theories (e.g., ideas about what makes that theory good) or from observations (interpreted through theories). These processes can be thought of as similar to evolutionary variance increasing (e.g., via genetic mutations) and decreasing (e.g., via natural selection) processes, one of the main differences being that these processes take place in people's minds (Popper, 1972). The core idea here is to avoid delegating the role of structure identification to the environment because, as the structure is not manifest, this would essentially be equivalent to leaving it unexplained.

The specific aspects of learning and decision-making that we highlighted in this article—background knowledge and environmental feedback—can be easily integrated into this framework. Background knowledge provides the basis for the creative construction of the environmental representation: It can be modified to make it fit for the situation that people are confronted with and

when some criticism shows that it does not fit for that purpose, it can be modified again. Environmental feedback informs these theories, by enabling a specific form of criticism: empirical testing. These tests can show people whether their theories about the environment are inaccurate, if their theories about it allow for such criticism (e.g., they are specific enough, or include counterfactuals of what cannot happen).

Dealing With Unknowns

Let us now turn back to our original question of how people deal with unknowns in their environment and answer it through the intuitive scientist framework. To illustrate, we can now give a more detailed explanation of how laypeople behave in a probabilistically generated experimental environment. The background knowledge for most participants (in the present cultural context) presumably constitutes of knowledge of communicational norms (e.g., people emphasize relevant information in their communication) and some basic mathematical principles (e.g., counting, keeping track of numbers, etc.). They could start by enacting only minor modifications to this background knowledge to fit the present circumstances because they are likely not motivated enough to attempt to think of better ways of representing the environmental unknowns in such contexts. This low level of motivation combined with the relevant communicational norms applied to the current situation (e.g., "the experimenter set up the experiment in such a repeated structure, because that structure is relevant") and the relevant (coarse) mathematical background knowledge applied to the current situation (e.g., counting and tracking apparently relevant features of the environment) could lead people to attempt to detect and exploit sequential patterns in the experiments.

Note that there was no reference to probabilistic concepts in this explanation. However, there could be for some people. One way in which probabilistic models can feature in this explanation is actually the historical account of the development of probability theory and its applications (Hacking, 2006). The way in which this mathematical framework and its applications were developed took years of research and guesswork. At the end of this process, the people who were involved in developing this framework could use probabilistic representations to represent unknowns. A less complicated way to represent the environment using probabilistic concepts is by learning about these concepts (e.g., in school) and applying them to a context because the requisite arguments for the application are also present (e.g., the experimenter explicitly told the participant to use them, and the participants believed the experimenter as to why they should use them). To summarize, the use of probabilistic concepts to represent the environment is a possible way to represent (some) unknowns in the environment for knowledgeable people, but it is not a fundamental feature of learning and decision-making.

Potential Objections, Alternative Explanations

In this section, we consider existing counterarguments and potential alternative explanations relating to some of the points in our critique of probabilistic models. We structure these counterarguments around the three different uses of probabilistic concepts—as optimal, approximate, and "as-if" models—common in psychological theories.

Optimal Solutions to Computational-Level Problems

The first counterargument to our position might be that probabilistic models are not used as theories of cognition, but rather as a way to provide insights into different analytical levels of cognition—often building on Marr’s (1982) levels of analyses. On the one side of this view, probabilistic models supposedly provide insights into the *computational level* of the problems that people face (e.g., Griffiths et al., 2012). This level is often defined as the abstract characterization of the problem that people face in the environment of interest. The idea is that using this abstract model, one can provide an objectively optimal solution for the task and can then check how people’s actual behavior measures up to this standard. Using this framework, researchers often develop a probabilistic formulation of a given environment (in some cases taking environmental or cognitive constraints into consideration to an arbitrary degree, see below, and Lieder & Griffiths, 2020), and then they designate the solution to this formal problem as the optimal solution in that environment. Much of the literature on whether or not learning and decision-making are rational is grounded in comparisons between people’s behavior and such computational-level solutions to certain problems.⁷

The problem with these lines of arguments is that such solutions only reflect a snapshot of the researcher’s current representation of the environment, rather than an objective optimal solution. This is because it is always possible that there is another yet undiscovered aspect of the environment that can be considered and exploited but is not contained in the currently designated optimal solution. In other words, the researcher may have found an optimal solution for the abstract computational problem that is an arbitrary approximation of the experimental environment—but whether this solution is optimal depends on how well the abstract problem matches the real one (similar issues emerge in the application of probabilistic models for scientific problems; Kellen et al., 2021; Navarro, 2018; Szollosi & Donkin, 2019).

Is it possible to provide an optimal solution to problems in physical environments? We argue that it is not because there always exist elements in the environment that people may consider outside of the narrow abstract characterization of the experimental environment. These elements can be found even in probabilistic experimental environments. One such element that we already mentioned is that the information encoded in the communicational aspect of the experimental instruction is often ignored in optimal solutions (Grice, 1975; Hilton, 1995; McKenzie, 2003; Sher & McKenzie, 2006). But people can go even further than that if they are motivated and knowledgeable about other aspects of the environment. For example, they could consider a broader contextual element in which the experiment is taking place: That it runs on a computer. Background knowledge in this area may allow them to devise novel ways to make advantageous choices in the experiments (e.g., by opening the source code of the experiment and finding out about which outcomes will occur before making a choice) or to make the experimenter believe that they did (e.g., by modifying the program to display a high total). The point is that, in physical environments, there is no limit on new elements that people may choose to include in their representations, and it is impossible to account for all of these possibilities in optimal models.

More generally, probabilistic models do not have a special status as models of the environment. As we have already explained, the use

of such models is only sensible when there exist arguments that imply that they are the current best solutions (or approximation) for the specific problem at hand. Such arguments can be made, for example, in the case of playing games of chance at a casino. Using probabilistic models may be considered best for that specific problem situation because developing or using more sensitive measurement devices that can predict the outcome of the games with better accuracy is not allowed in such places. However, in alternative gambling scenarios in which the rules do not prohibit the use of more sensitive measurement tools people might decide to use those tools instead of (or in conjunction with) relying on probabilistic models (relatively simple and common examples include, e.g., counting cards in a Blackjack game, adjusting the cost of car insurance based on a person’s age, etc.).

To summarize, even though we agree that problem specification is an important part of cognition, researchers cannot provide a complete solution to how this is executed because it is always possible for both the participants and the researcher to consider novel aspects of the physical environment in this specification. Therefore, optimal probabilistic models are better off treated either as merely “as-if” or descriptive models (e.g., Tauber et al., 2017; see below) or as explicanda (i.e., treating the researcher’s capacity to develop such representations as behavior to be explained by a psychological theory).

Bounded Rational Approximations

The other side of the previous criticism might be that our point is moot because no one claims that people represent unknowns probabilistically at the *algorithmic level* (Marr, 1982). In contrast with the computational level’s focus on optimal solutions, algorithmic-level investigations are concerned with identifying the cognitive processes that are involved in actually solving the problem (Griffiths et al., 2012). One important consideration for this level is the idea of *bounded rationality* (Simon, 1956)—that the rationality of people’s behavior should be understood in terms of ecological and cognitive constraints. The common framing of this idea is that people can only approximate more complex optimal representations to deal with the problems of environmental unknowns, thus theories of cognitive processes need to be psychologically and ecologically plausible (i.e., less computationally demanding, while taking environmental constraints into account). Examples for such cognitive processes include heuristics (Gigerenzer & Todd, 1999), simple linear models (Juslin et al., 2008), or intuitive physical theories (Battaglia et al., 2013).

An important recognition of this bounded rational approach is that there are several environmental (Fiedler, 2000; Pleskac & Hertwig, 2014; Todd & Gigerenzer, 2007) and cognitive constraints (Juslin et al., 2009; van Rooij, 2008) that theories of cognitive processes need to consider—an idea that has been integrated to some extent even in computational-level analyses (e.g., Lieder & Griffiths, 2020).

⁷ There is little that matches or mismatches between people’s behavior and optimal models can tell us (Szollosi & Newell, 2020): Matches could mean that both people’s and researcher’s optimal models correspond to the same aspect of the environment, but it is possible that they do not have the same content at all (i.e., participants’ representations are in some respect confounded with the assumptions of the researcher); mismatches could mean that either the researcher’s or the participant’s model is (more) incorrect (McKenzie, 2003).

Another important recognition is that probabilistic models are not essential in psychological explanations because they may not provide good approximations to certain kinds of environmental unknowns, such as future cultural and technological advances (e.g., Knight, 1921, Luce & Raiffa, 1957; see Kozyreva & Hertwig, 2021, for a summary).

Yet, such theories still often feature probabilistic concepts. People are assumed to use more or less complex sampling methods (Brown & Steyvers, 2009; Hau et al., 2008; Sanborn & Chater, 2016), and the environment is assumed to provide cues based on correlational structures (Brunswik, 1955; Pleskac & Hertwig, 2014). Falling back on these concepts is indicative of the more general problem that we already highlighted: That probabilistic concepts are only substitutes for the process by which the environmental representation develops.

To illustrate the problem on probabilistically generated experiments again, a potential bounded rational account would presumably give an explanation for why people consider temporal structure in their decisions similar to this: “People’s natural environments often contain events that are temporally autocorrelated. A sequence-tracking heuristic can be used to exploit these statistical structures in a computationally efficient way. Although such a heuristic is useful in natural environments, it can be misapplied in an artificial experimental environment (e.g., Ayton & Fischer, 2004).” It is easy to see how the environmental specification is entered as an assumption into such an explanation: How does such a heuristic “know” what constitutes a sequence? Are all possible temporal sequences already prespecified by the set of heuristics? How is it possible that some people use a much more complex probabilistic representation instead of such a heuristic? From our point of view, a deeper explanation for how and why people would use such a heuristic would need to account not only for environmental and cognitive constraints, but for the interplay between those and people’s background knowledge, while also accounting for possible more complex representations of the environment (similar to our explanation of how some participants *could have* used probabilistic representations).

To summarize, we broadly agree with the claims that environmental and cognitive constraints need to be taken seriously, and that people often rely on representations that are not that computationally resource demanding if certain constraints (environmental, computational, motivational) prevent them from thinking more. But such representations are only part of how people deal with unknowns: It is true that sometimes people use such simple rules of thumb, but sometimes they create much more elaborate representations of their environments. So instead of determining which of the many potential approximations people use, the way in which interesting aspects of the environment are specified at varying levels of complexity need to take a major role in its explanation. Bounded rational accounts often point to evolutionary, developmental, and social origins of this specification, but such accounts are in our opinion not sufficiently explained (Szollosi & Newell, 2020).

“As-If” Models

Lastly, one could argue that probabilistic concepts are not used in a literal sense in psychological explanations, but only in a more noncommittal sense, as “as-if” models (e.g., Tauber et al., 2017; sometimes also referred to as descriptive models). In this view,

people are not assumed to use (neither literally nor approximately) probabilistic models for their learning and decision-making. Instead, the models merely provide a mathematical redescription of some aspect of what people do. Although sometimes it seems difficult to separate the descriptive mathematical model from the psychological explanation—as has been the case with iterations of Bayesian models (Bowers & Davis, 2012; Jones & Love, 2011)—we agree that probabilistic concepts can occasionally be useful in this sense.

What we are concerned about, however, is that providing a probabilistic representation of some aspect of people’s responses is often accepted as a sufficient psychological explanation. For example, describing people’s learning as noisy sampling of evidence or the variability in their choices as stochastic deviation from some ideal model is relatively irrelevant without an explanation of how they chose what constitutes as evidence and what led to that variability. Similarly, descriptions of the environment in terms of noisy correlational cue–outcome structures or variably occurring environmental events are described as a result of stochastic generative processes is not that meaningful without some explanation of how people might decide to consider something a cue or an outcome or an event. Stopping before a psychological explanation is given is problematic because the probabilistic model is only an approximate mathematical redescription of what people have been observed to do in the experiment, rather than an explanation of how people deal with unknowns.

What further limits the role of probabilistic concepts is that they cannot meaningfully describe the argumentative steps that we argued for throughout the article (although they can redescribe them) because those steps are not part of probability theory. Indeed, whenever psychological explanations actually attempt to answer these questions, the parts of the explanation that answer these questions come not from probability theory, but from arguments and knowledge external to it (currently verbal parts of psychological theories). In other words, as we have argued throughout the article, probabilistic models only provide a good description of any phenomenon in conjunction with the accompanying arguments and background knowledge that explain when they can be used. While probabilistic concepts offer a convenient way to mathematically express some aspects of people’s behavior and their environment (e.g., our use of density histograms and hazard rates only redescrbed aspects of the environment and what people did in the experiment in a simple manner), they cannot in themselves fully describe how people deal with unknowns.

The Scope of Our Argument

To clarify the scope of our argument, let us review how we aimed to preempt some potential challenges to our position. We argued that irrespective of interpretation specifying how and why probabilistic concepts are applied is important because—as demonstrated by our theoretical and experimental examples—there is no trivial way to create representations of environments. We suggested that there needs to be either a clear explanation for where probabilistic knowledge comes from (in literal and bounded rational interpretations), or a clear explanation for why probabilistic models provide appropriate approximations for phenomena known to be nonprobabilistic (in “as-if” uses). Although we focused our discussion on higher-level psychological processes such as learning and decision-making, these arguments hold for any psychological explanation

that invokes concepts from probability theory, including lower-level psychological processes (e.g., those associated with perception).

To provide a brief sketch of how our argument would apply to lower-level psychological processes, consider research on how people use similarity in categorization (e.g., Ashby & Alfonso-Reese, 1995). Probabilistic concepts are often used in models of such psychological processes—for instance, one might argue that the similarity of a novel stimulus to older stimuli can be modeled as a probability distribution. Such use of probabilistic concepts can be interpreted as an “as-if” model, meaning that the probabilistic model is considered useful to summarize how aspects of similarity affects people’s categorization judgments, rather than considering such processes truly probabilistic. However, as we argued, this use of probability theory can still easily obscure much of what is interesting in such decisions. For example, such a probabilistic model does not explain how and why particular aspects of the environment are used to judge the similarity between stimuli when a range of potential candidates may have been possible.⁸

How to use probabilistic concepts in literal or bounded rational senses to explain lower-level psychological processes is less clear-cut. In higher-level processes, we argued that, for these uses, the source of probabilistic knowledge can either be explained in terms of the acquisition of abstract knowledge about probability theory and its application, or in terms of the acquisition of heuristic strategies that exploit some aspects of the mechanics of probabilistic models. We expect that explanations of lower-level processes would be more similar to the latter—for instance, they can be stated in terms of how evolutionary processes produced heuristics that in some way work similarly to probability theory. In these cases, it would be important to clarify why such explanations should not be better considered “as-if” interpretations.

In this section, we aimed to clarify how our argument applies to common uses of probability theory. While we think probabilistic concepts can be usefully applied in certain cases and suggested necessary conditions for how this can be done, we expect that reliance on such concepts will become increasingly less prevalent as more of the unexamined underlying assumptions are made explicit.

Future Directions

By highlighting often-neglected aspects of learning and decision-making, our aim was to start a discussion about, rather than provide definitive answers to, how to understand people’s treatment of unknowns without probabilistic concepts. We anticipate that this framework can not only extend the range of questions that we can ask about learning and decision-making, but it could change the way we ask those questions and the questions we find interesting. In this section, we offer a selection from the exciting questions that this framework could lead to.

As we have already said, this view opens up new ways to think about the role of environmental feedback in improving people’s theories of unknowns. An important implication is that feedback is only meaningful in terms of people’s specific knowledge of that environment. One interesting line of investigations could be to understand how this applies to environments that “provide” different types of feedback about their unknown aspects (Hogarth et al., 2015; Kahneman & Klein, 2009). As we have argued, in real-life people can always improve their representation. Such studies, however, often make it difficult for people to do this: They can either stay

within the narrow confines of the experiment and develop a simple satisfactory strategy (e.g., detecting sequential patterns); going further than that (e.g., accessing the source code of the software), however, requires a substantial amount of additional motivation and background knowledge. Researchers may consider supplementing the experiments with mechanisms that allow people more ways to find out about the experimental environment. One way to do this is to allow them to figure out which aspect of the environment is important via experimentation (Gureckis & Markant, 2012). Another way is to make the experiments more enticing by making them more closely resemble actual physical environments: Most physical environments that people interact with contain many unknowns that are not “probabilistic” (i.e., are not well described by probabilistic models), which should be considered when designing experimental environments (e.g., by making the deterministic aspects of experiments more meaningful).

A connected implication of our present analysis is the need to study how people generate hypotheses. There have been promising recent attempts in this direction in the form of compositional generative models (Bramley et al., 2017; Lake et al., 2017; Piantadosi, 2021). Such models rely to a large extent on probabilistic concepts. However, their core commitments are that people generate hypotheses by using a finite set of primitives and rules that specify compositional combinations. Thus, the core claims are broadly consistent with the approach advocated here. It would be interesting to see whether the probabilistic aspects of the present compositional generative models can be replaced by more psychologically meaningful mechanisms. That is, since variability in people’s hypotheses is not generated by real or pseudorandom generators, the substantive explanation for how such variability arises should be given in nonprobabilistic terms (even if it turns out that these processes are well approximated by probability models). Generally, however, we believe that some form of the compositional structure such models invoke is necessary in any explanation that aims to account for the flexibility in people’s capacity to create new hypotheses.

A closely related question concerns the role of motivation. We have already alluded to the importance of this in our explanation and when explaining what other potential theories people may construct. For example, what makes people pursue having deeper, more complex explanations, instead of being satisfied with simple heuristic or satisficing solutions? Such a question implies that the way in which people’s theories develop cannot be sufficiently explained without concepts that make people inclined to pursue potential improvements, such as intrinsic motivation or curiosity (Berlyne, 1966; Deci, 1972; Loewenstein, 1994).

Another interesting area to which our arguments relate are the measurement of learning and decision-making traits, such as risk preferences (Charness et al., 2013; Frey et al., 2017; Loewenstein et al., 2001).

⁸ Stimuli can be considered similar across many different features (e.g., Tversky, 1977), which can be easily demonstrated by the process of constructing similarity measures. To construct such measures, one needs to make several choices, for example, decide which features to include and which ones to disregard, as there are an infinite number of physical features that may be considered (commonly considered features, such as brightness, color, or shape, are only a very limited set of the possible ones—other features like hue, emitted nonvisible light, whether they are a primary color, are rarely or never considered). Such measures can either be constructed by a person (for higher-level processes) or by evolutionary processes (for lower-level processes).

One important question in this area of research is whether risk preferences are constructed when people make decisions or there exist a stable trait that is revealed by people's behavior (Lichtenstein & Slovic, 2006; Pedroni et al., 2017; Slovic, 1995). The results of the current experiment pose additional difficulties to this debate—for example, under which particular representation of the environment can a participant's choice be considered to be risky? If a participant thinks that there is an exploitable temporal structure, their choices of the risky option do not reveal a preference trait, but rather some (perhaps mistaken) knowledge of the environment. This also makes it difficult to compare the risk preference of people, since their representations of the same environments may differ substantially. Studies into risk attitudes—especially those in the lab—need to consider how people construct their environmental representation, for example, by attempting to introduce more sensitive measures of what hypotheses people entertain of their environments.

Conclusion

Due to the success of scientific applications of probability theory, probabilistic concepts took on a dominant role in psychological explanations of how people deal with unknowns. In this article, we critically assessed the role that these concepts may play in such explanations. Through theoretical and experimental analyses of a typical experiment of related research areas, we argued that many intermediary argumentative steps (that are not part of probability theory but are essential for representing unknowns in physical environments) can be left implicit when applying probabilistic concepts. This leads to the problem that such steps are left unexplained in the corresponding psychological explanation. Our hope is that the research approach that we outlined and demonstrated here will lead to psychological theories of learning and decision-making in which probabilistic concepts only play a relatively minor role. Specifically, these concepts should either be treated as explicanda for psychological explanations (i.e., as a mode of representation that people are capable of creating), or as tools for scientists to conveniently represent (a small number of) aspects of people's performance or environment.

References

Ashby, F. G., & Alfonso-Reese, L. A. (1995). Categorization as probability density estimation. *Journal of Mathematical Psychology*, 39(2), 216–233. <https://doi.org/10.1006/jmps.1995.1021>

Ayton, P., & Fischer, I. (2004). The hot hand fallacy and the gambler's fallacy: Two faces of subjective randomness? *Memory & Cognition*, 32(8), 1369–1378. <https://doi.org/10.3758/BF03206327>

Bar-Hillel, M., & Wagenaar, W. A. (1991). The perception of randomness. *Advances in Applied Mathematics*, 12(4), 428–454. [https://doi.org/10.1016/0196-8858\(91\)90029-1](https://doi.org/10.1016/0196-8858(91)90029-1)

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, 16(3), 215–233. <https://doi.org/10.1002/bdm.443>

Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences of the United States of America*, 110(45), 18327–18332. <https://doi.org/10.1073/pnas.1306572110>

Berlyne, D. E. (1966). Curiosity and exploration. *Science*, 153(3731), 25–33. <https://doi.org/10.1126/science.153.3731.25>

Bowers, J. S., & Davis, C. J. (2012). Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin*, 138(3), 389–414. <https://doi.org/10.1037/a0026450>

Bramley, N. R., Dayan, P., Griffiths, T. L., & Lagnado, D. A. (2017). Formalizing Neurath's ship: Approximate algorithms for online causal learning. *Psychological Review*, 124(3), 301–338. <https://doi.org/10.1037/rev0000061>

Brehmer, B. (1980). In one word: Not from experience. *Acta Psychologica*, 45(1–3), 223–241. [https://doi.org/10.1016/0001-6918\(80\)90034-7](https://doi.org/10.1016/0001-6918(80)90034-7)

Brette, R. (2019). Is coding a relevant metaphor for the brain? *Behavioral and Brain Sciences*, 42, Article e215. <https://doi.org/10.1017/S0140525X19000049>

Brown, S. D., & Steyvers, M. (2009). Detecting and predicting changes. *Cognitive Psychology*, 58(1), 49–67. <https://doi.org/10.1016/j.cogpsych.2008.09.002>

Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review*, 62(3), 193–217. <https://doi.org/10.1037/h0047470>

Busemeyer, J. R., Gluth, S., Rieskamp, J., & Turner, B. M. (2019). Cognitive and neural bases of multi-attribute, multi-alternative, value-based decisions. *Trends in Cognitive Sciences*, 23(3), 251–263. <https://doi.org/10.1016/j.tics.2018.12.003>

Busemeyer, J. R., Pothos, E. M., Franco, R., & Trueblood, J. S. (2011). A quantum theoretical explanation for probability judgment errors. *Psychological Review*, 118(2), 193–218. <https://doi.org/10.1037/a0022542>

Carey, S., & Spelke, E. (1996). Science and core knowledge. *Philosophy of Science*, 63(4), 515–533. <https://doi.org/10.1086/289971>

Charness, G., Gneezy, U., & Imas, A. (2013). Experimental methods: Eliciting risk preferences. *Journal of Economic Behavior & Organization*, 87, 43–51. <https://doi.org/10.1016/j.jebo.2012.12.023>

Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences*, 10(7), 287–291. <https://doi.org/10.1016/j.tics.2006.05.007>

Costello, F., & Watts, P. (2014). Surprisingly rational: Probability theory plus noise explains biases in judgment. *Psychological Review*, 121(3), 463–480. <https://doi.org/10.1037/a0037010>

Dayan, P., & Niv, Y. (2008). Reinforcement learning: The good, the bad and the ugly. *Current Opinion in Neurobiology*, 18(2), 185–196. <https://doi.org/10.1016/j.conb.2008.08.003>

Deci, E. L. (1972). Intrinsic motivation, extrinsic reinforcement, and inequity. *Journal of Personality and Social Psychology*, 22(1), 113–120. <https://doi.org/10.1037/h0032355>

Deutsch, D. (2011). *The beginning of infinity: Explanations that transform the world*. Allen Lane.

Donkin, C., Szollosi, A., & Bramley, N. R. (in press). Observing effects in various contexts won't give us general psychological theories. *Behavioral and Brain Sciences*.

Dulany, D. E., Carlson, R. A., & Dewey, G. I. (1984). A case of syntactical learning and judgment: How conscious and how abstract? *Journal of Experimental Psychology: General*, 113(4), 541–555. <https://doi.org/10.1037/0096-3445.113.4.541>

Dunwoody, P. T. (2009). Theories of truth as assessment criteria in judgment and decision making. *Judgment and Decision Making*, 4(2), 116–125.

Estes, W. K. (1950). Toward a statistical theory of learning. *Psychological Review*, 57(2), 94–107. <https://doi.org/10.1037/h0058559>

Estes, W. K. (1976). The cognitive side of probability learning. *Psychological Review*, 83(1), 37–64. <https://doi.org/10.1037/0033-295X.83.1.37>

Felin, T., Koenderink, J., & Krueger, J. I. (2017). Rationality, perception, and the all-seeing eye. *Psychonomic Bulletin & Review*, 24(4), 1040–1059. <https://doi.org/10.3758/s13423-016-1198-z>

Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review*, 107(4), 659–676. <https://doi.org/10.1037/0033-295X.107.4.659>

- Fiedler, K., & Juslin, P. (2006). *Information sampling and adaptive cognition*. Cambridge University Press.
- Frey, R., Pedroni, A., Mata, R., Rieskamp, J., & Hertwig, R. (2017). Risk preference shares the psychometric structure of major psychological traits. *Science Advances*, 3(10), Article e1701381. <https://doi.org/10.1126/scia.dv.1701381>
- Gal, I., & Baron, J. (1996). Understanding repeated simple choices. *Thinking & Reasoning*, 2(1), 81–98. <https://doi.org/10.1080/135467896394573>
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*, 107(2), 289–344. <https://doi.org/10.1037/0033-295X.107.2.289>
- Gigerenzer, G. (1991). From tools to theories: A heuristic of discovery in cognitive psychology. *Psychological Review*, 98(2), 254–267. <https://doi.org/10.1037/0033-295X.98.2.254>
- Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart*. Oxford University Press.
- Gopnik, A., & Meltzoff, A. N. (1997). *Words, thoughts, and theories*. MIT Press.
- Grabenhorst, M., Michalareas, G., Maloney, L. T., & Poeppel, D. (2019). The anticipation of events in time. *Nature Communications*, 10(1), Article 5802. <https://doi.org/10.1038/s41467-019-13849-0>
- Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics: Vol. 3. Speech arts* (pp. 41–58). Academic Press.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, 51(4), 334–384. <https://doi.org/10.1016/j.cogpsych.2005.05.004>
- Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science*, 17(9), 767–773. <https://doi.org/10.1111/j.1467-9280.2006.01780.x>
- Griffiths, T. L., Vul, E., & Sanborn, A. N. (2012). Bridging levels of analysis for probabilistic models of cognition. *Current Directions in Psychological Science*, 21(4), 263–268. <https://doi.org/10.1177/0963721412447619>
- Gureckis, T. M., & Markant, D. B. (2012). Self-directed learning. *Perspectives on Psychological Science*, 7(5), 464–481. <https://doi.org/10.1177/1745691612454304>
- Hacking, I. (2006). *The emergence of probability: A philosophical study of early ideas about probability, induction and statistical inference*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511817557>
- Hahn, U., & Warren, P. A. (2009). Perceptions of randomness: Why three heads are better than four. *Psychological Review*, 116(2), 454–461. <https://doi.org/10.1037/a0015241>
- Hájek, A. (2019). Interpretations of probability. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University.
- Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description-experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, 21(5), 493–518. <https://doi.org/10.1002/bdm.598>
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8), 534–539. <https://doi.org/10.1111/j.0956-7976.2004.00715.x>
- Hilton, D. J. (1995). The social context of reasoning: Conversational inference and rational judgment. *Psychological Bulletin*, 118(2), 248–271. <https://doi.org/10.1037/0033-2909.118.2.248>
- Hogarth, R. M., Lejarraga, T., & Soyer, E. (2015). The two settings of kind and wicked learning environments. *Current Directions in Psychological Science*, 24(5), 379–385. <https://doi.org/10.1177/0963721415591878>
- Jones, M., & Love, B. C. (2011). Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behavioral and Brain Sciences*, 34(4), 169–188. <https://doi.org/10.1017/S0140525X10003134>
- Juslin, P., Karlsson, L., & Olsson, H. (2008). Information integration in multiple cue judgment: A division of labor hypothesis. *Cognition*, 106(1), 259–298. <https://doi.org/10.1016/j.cognition.2007.02.003>
- Juslin, P., Nilsson, H., & Winman, A. (2009). Probability theory, not the very guide of life. *Psychological Review*, 116(4), 856–874. <https://doi.org/10.1037/a0016979>
- Juslin, P., Olsson, H., & Olsson, A.-C. (2003). Exemplar effects in categorization and multiple-cue judgment. *Journal of Experimental Psychology: General*, 132(1), 133–156. <https://doi.org/10.1037/0096-3445.132.1.133>
- Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: A failure to disagree. *American Psychologist*, 64(6), 515–526. <https://doi.org/10.1037/a0016755>
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–292. <https://doi.org/10.2307/1914185>
- Karelaia, N., & Hogarth, R. M. (2008). Determinants of linear judgment: A meta-analysis of lens model studies. *Psychological Bulletin*, 134(3), 404–426. <https://doi.org/10.1037/0033-2909.134.3.404>
- Keil, F. C. (2006). Explanation and understanding. *Annual Review of Psychology*, 57(1), 227–254. <https://doi.org/10.1146/annurev.psych.57.102904.190100>
- Kellen, D., Davis-Stober, C. P., Dunn, J. C., & Kalish, M. L. (2021). The problem of coordination and the pursuit of structural constraints in psychology. *Perspectives on Psychological Science*, 16(4), 767–778. <https://doi.org/10.1177/1745691620974771>
- Knight, F. H. (1921). *Risk, uncertainty, and profit*. Houghton and Mifflin.
- Kolmogorov, A. N. (1950). *Foundations of the theory of probability*. Chelsea Publishing.
- Kozyreva, A., & Hertwig, R. (2021). The interpretation of uncertainty in ecological rationality. *Synthese*, 198(2), 1517–1547. <https://doi.org/10.1007/s11229-019-02140-w>
- Kuhn, D. (1989). Children and adults as intuitive scientists. *Psychological Review*, 96(4), 674–689. <https://doi.org/10.1037/0033-295X.96.4.674>
- Lagnado, D. A., Waldmann, M. R., Hagmayer, Y., & Sloman, S. A. (2005). Beyond covariation: Cues to causal structure. In A. Gopnik & L. Schultz (Eds.), *Causal learning: Psychology, philosophy and computation* (pp. 154–172). Oxford University Press.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, Article e253. <https://doi.org/10.1017/S0140525X16001837>
- Lichtenstein, S., & Slovic, P. (2006). *The construction of preference*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511618031>
- 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, Article e1. <https://doi.org/10.1017/S0140525X1900061X>
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, 116(1), 75–98. <https://doi.org/10.1037/0033-2909.116.1.75>
- Loewenstein, G. F., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as feelings. *Psychological Bulletin*, 127(2), 267–286. <https://doi.org/10.1037/0033-2909.127.2.267>
- Lombrozo, T. (2006). The structure and function of explanations. *Trends in Cognitive Sciences*, 10(10), 464–470. <https://doi.org/10.1016/j.tics.2006.08.004>
- Lovibond, P. F., & Shanks, D. R. (2002). The role of awareness in Pavlovian conditioning: Empirical evidence and theoretical implications. *Journal of Experimental Psychology: Animal Behavior Processes*, 28(1), 3–26. <https://doi.org/10.1037/0097-7403.28.1.3>
- Luce, R. D., & Raiffa, H. (1957). *Games and decisions: Introduction and critical survey*. Dover Publication.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. MIT Press.

- McKenzie, C. R. M. (2003). Rational models as theories—Not standards—Of behavior. *Trends in Cognitive Sciences*, 7(9), 403–406. [https://doi.org/10.1016/S1364-6613\(03\)00196-7](https://doi.org/10.1016/S1364-6613(03)00196-7)
- McKenzie, C. R. M. (2004). Hypothesis testing and evaluation. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 200–219). Blackwell Publishing. <https://doi.org/10.1002/9780470752937.ch10>
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92(3), 289–316. <https://doi.org/10.1037/0033-295X.92.3.289>
- Navarro, D. J. (2018). Between the devil and the deep blue sea: Tensions between scientific judgement and statistical model selection. *Computational Brain & Behavior*, 2(1), 28–34. <https://doi.org/10.1007/s42113-018-0019-z>
- Navarro, D. J., Newell, B. R., & Schulze, C. (2016). Learning and choosing in an uncertain world: An investigation of the explore-exploit dilemma in static and dynamic environments. *Cognitive Psychology*, 85, 43–77. <https://doi.org/10.1016/j.cogpsych.2016.01.001>
- Newell, B. R., & Shanks, D. R. (2014). Unconscious influences on decision making: A critical review. *Behavioral and Brain Sciences*, 37(1), 1–19. <https://doi.org/10.1017/S0140525X12003214>
- Nickerson, R. S. (2002). The production and perception of randomness. *Psychological Review*, 109(2), 330–357. <https://doi.org/10.1037/0033-295X.109.2.330>
- Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198524496.001.0001>
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge University Press.
- Pedroni, A., Frey, R., Bruhin, A., Dutilh, G., Hertwig, R., & Rieskamp, J. (2017). The risk elicitation puzzle. *Nature Human Behaviour*, 1(11), 803–809. <https://doi.org/10.1038/s41562-017-0219-x>
- Piantadosi, S. T. (2021). The computational origin of representation. *Minds and Machines*, 31(1), 1–58. <https://doi.org/10.1007/s11023-020-09540-9>
- Pleskac, T. J., & Hertwig, R. (2014). Ecologically rational choice and the structure of the environment. *Journal of Experimental Psychology: General*, 143(5), 2000–2019. <https://doi.org/10.1037/xge0000013>
- Plonsky, O., Teodorescu, K., & Erev, I. (2015). Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review*, 122(4), 621–647. <https://doi.org/10.1037/a0039413>
- Popper, K. (1963). *Conjectures and refutations: The growth of scientific knowledge*. Routledge.
- Popper, K. (1972). *Objective knowledge: An evolutionary approach*. Clarendon Press.
- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, 20(4), 260–281. <https://doi.org/10.1016/j.tics.2016.01.007>
- Reimers, S., Donkin, C., & Le Pelley, M. E. (2018). Perceptions of randomness in binary sequences: Normative, heuristic, or both? *Cognition*, 172, 11–25. <https://doi.org/10.1016/j.cognition.2017.11.002>
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). Appleton-Century-Crofts.
- Sanborn, A. N., & Beierholm, U. R. (2016). Fast and accurate learning when making discrete numerical estimates. *PLoS Computational Biology*, 12(4), Article e1004859. <https://doi.org/10.1371/journal.pcbi.1004859>
- Sanborn, A. N., & Chater, N. (2016). Bayesian brains without probabilities. *Trends in Cognitive Sciences*, 20(12), 883–893. <https://doi.org/10.1016/j.tics.2016.10.003>
- Schwarz, W. (2018). No interpretation of probability. *Erkenntnis*, 83(6), 1195–1212. <https://doi.org/10.1007/s10670-017-9936-9>
- Shanks, D. R., & St. John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, 17(3), 367–395. <https://doi.org/10.1017/S0140525X00035032>
- Sher, S., & McKenzie, C. R. M. (2006). Information leakage from logically equivalent frames. *Cognition*, 101(3), 467–494. <https://doi.org/10.1016/j.cognition.2005.11.001>
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129–138. <https://doi.org/10.1037/h0042769>
- Slovic, P. (1995). The construction of preference. *American Psychologist*, 50(5), 364–371. <https://doi.org/10.1037/0003-066X.50.5.364>
- Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, 53(1), 1–26. <https://doi.org/10.1016/j.cogpsych.2005.10.003>
- Sutton, R., & Barto, A. (1998). *Reinforcement learning, an introduction*. MIT Press. <https://doi.org/10.1109/TNN.1998.712192>
- Szollosi, A., & Donkin, C. (2019). Neglected sources of flexibility in psychological theories: From replicability to good explanations. *Computational Brain & Behavior*, 2(3–4), 190–192. <https://doi.org/10.1007/s42113-019-00045-y>
- Szollosi, A., Liang, G., Konstantinidis, E., Donkin, C., & Newell, B. R. (2019). Simultaneous underweighting and overestimation of rare events: Unpacking a paradox. *Journal of Experimental Psychology: General*, 148(12), 2207–2217. <https://doi.org/10.1037/xge0000603>
- Szollosi, A., & Newell, B. R. (2020). People as intuitive scientists: Reconsidering statistical explanations of decision making. *Trends in Cognitive Sciences*, 24(12), 1008–1018. <https://doi.org/10.1016/j.tics.2020.09.005>
- Tauber, S., Navarro, D. J., Perfors, A., & Steyvers, M. (2017). Bayesian models of cognition revisited: Setting optimality aside and letting data drive psychological theory. *Psychological Review*, 124(4), 410–441. <https://doi.org/10.1037/rev0000052>
- Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, 10(7), 309–318. <https://doi.org/10.1016/j.tics.2006.05.009>
- Todd, P. M., & Gigerenzer, G. (2007). Environments that make us smart. *Current Directions in Psychological Science*, 16(3), 167–171. <https://doi.org/10.1111/j.1467-8721.2007.00497.x>
- Tran, R., Vul, E., & Pashler, H. (2017). How effective is incidental learning of the shape of probability distributions? *Royal Society Open Science*, 4(8), Article 170270. <https://doi.org/10.1098/rsos.170270>
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84(4), 327–352. <https://doi.org/10.1037/0033-295X.84.4.327>
- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2), 207–232. [https://doi.org/10.1016/0010-0285\(73\)90033-9](https://doi.org/10.1016/0010-0285(73)90033-9)
- van Rooij, I. (2008). The tractable cognition thesis. *Cognitive Science: A Multidisciplinary Journal*, 32(6), 939–984. <https://doi.org/10.1080/03640210801897856>
- van Rooij, I., & Baggio, G. (2021). Theory before the test: How to build high-verisimilitude explanatory theories in psychological science. *Perspectives on Psychological Science*, 16(4), 682–697. <https://doi.org/10.1177/1745691620970604>
- Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General*, 121(2), 222–236. <https://doi.org/10.1037/0096-3445.121.2.222>

Received October 15, 2020

Revision received December 15, 2021

Accepted December 20, 2021 ■