Taking Probabilistic Choice Seriously

Elizaveta Konovalova[1*] and Thorsten Pachur[2,3]

[1*]Warwick Business School, University of Warwick, Scarman Road, Coventry, CV4 7AL, United Kingdom.
[2]School of Management, Technical University of Munich, Arcisstraße 21, Munich, 80333, Germany.
[3]Center for Adaptive Rationality, Max Planck Institute for Human Development, Lentzeallee 94, Berlin, 14195, Germany.

*Corresponding author(s). E-mail(s): elizaveta.konovalova@wbs.ac.uk;
Contributing authors: pachur@tum.de;

Abstract

Much behaviour and cognition involves choosing between alternatives—whether in the context of consumer decisions, voting, or memory. A fundamental aspect of choice is that it is probabilistic: An organism faced with the same evidence will not always make the same choice. Despite widespread acknowledgement that probabilistic choice is a necessary element in models of cognition, it has long been treated as distracting from the core process driving behaviour. Recently, however, researchers across various fields in the cognitive and behavioural sciences have started to recognise the vital role of probabilistic choice for understanding cognition. This article reviews and synthesizes these developments, distinguishing three main origins of probabilistic choice and proposing future research avenues that integrate probabilistic choice into theories of cognition.

Keywords: Probabilistic choice; Stochasticity; Behavioural variability; Models of choice

Probabilistic choice becomes a focus of research

Much behaviour and cognition involves choosing between alternatives. When people formulate their preferences—whether between consumer goods, economic alternatives or political parties—they evaluate, compare and finally choose one of the available alternatives. When people retrieve information from memory, they select content from several candidate items. When they classify objects, they pick a response from a set of possible categories.

A fundamental feature of choice is that it is probabilistic: When given the same evidence on different occasions, people do not always make the same choice. The *probabilistic nature of choice* has been noted across a wide variety of disciplines in the cognitive and behavioural sciences, including economics, political science, biology, psychology, and neuroscience. Despite the widespread acknowledgement that **probabilistic choice** (see Glossary) needs to be included in models of cognition, it is usually treated as peripheral and as a nuisance factor with little theoretical relevance. In formal treatments it is often relegated to an auxiliary, residual term, and no principled assumptions are made about the internal and external processes that contribute to choice being probabilistic (see Box 1).

In recent years, however, interest in probabilistic choice has been growing and researchers in psychology, neuroscience and economics have made it the focus of their investigations. This article reviews this emerging trend and elaborates on its implications for models of cognition. First, we highlight that probabilistic choice may be a feature rather than a bug of the cognitive system. We review findings suggesting that it is an indispensable tool for learning, confers an advantage in competition and is necessary for the efficient processing of information in the face of cognitive constraints and environmental uncertainty.

Second, we review work illustrating how the measurement of probabilistic choice can be exploited to test and refine behavioural theories. We discuss how cognitive models can be compared and validated based on the observable patterns of probabilistic choice and how a more principled approach toward accommodating **stochasticity** in cognitive models can help explain seemingly irrational behavioural tendencies.

Finally, we synthesize existing ideas regarding the origins of probabilistic choice into a framework by distinguishing three main sources: variability of internal processing, uncertainty in the environment, and stochasticity in behavioural implementation. Based on our proposed framework, we identify new research avenues that focus on both an enhanced understanding and systematic measurement of probabilistic choice. We believe that a more targeted investigation of probabilistic choice can in turn inform a better understanding of behaviour.

Probabilistic choice is a feature, not a bug

Although probabilistic choice can indeed be a nuisance in some cases (e.g., decision-making by judges or medical professionals, where consistency is important), its very ubiquity gives cause for pause. In addition to asking how probabilistic choice can be eliminated, it is also worth asking why probabilistic choice exists in the first place. Existing literature suggests that variability in behaviour can have adaptive benefits. One comes from the organism's deliberate decision to deviate from a previous course of action even when faced with a similar situation. Such behavioural volatility could be optimal when the organism needs to learn about a novel environment; it can also be advantageous against competitors. A second type of benefit comes from stochasticity in information processing (which is usually not under the deliberate control of the organism). This stochasticity allows the organism to adapt to an uncertain environment under cognitive constraints. We next discuss both types of benefits in more detail.

Strategic use of probabilistic choice

Under some circumstances, it can be beneficial for an organism confronted with the same situation to make a different choice. While some organisms may simply have a preference for variety [1] or randomization [2], varying choices across occasions is an important strategy for organisms learning about their surroundings for the first time. The simplest approach is to start by making random choices and see what feedback the environment provides. This strategy has been shown to be the basis for learning songs in some bird species [3], motor control in humans [4], and consumers getting into new markets in online environments [5]. An organism engaging in exploration might employ other strategies: they can choose an alternative deliberately (i.e., "directed exploration" in contrast to "random exploration") or use a mix of the random and directed approaches [for review, see 6].

Probabilistic choice can also be beneficial in a competitive world. In fact, economic game theory is rooted in this idea [7]. To outsmart opponents in zero-sum games with a mixed-strategy equilibrium, the decision maker has to act unpredictably [8]. In the animal world, prey often adopts unpredictable or "Protean" behaviour to avoid being preyed on [9, 10]. If many competitors in the population imitate each other, it can be adaptive for agents to distribute their choices across alternatives when competing for resources [11].

Stochasticity in information processing

Recent work in neuroscience has highlighted another benefit of probabilistic choice: It allows an organism to adapt optimally to a changing environment while dealing with cognitive constraints. Consider a foraging animal. Foraging for food is associated with a high level of uncertainty: the number of patches can change, the amount of food in the patches can change, and something completely unforeseen might happen, such as a new predator arriving in the area. Ideally, the animal should make choices based on a rich representation of its environment that considers different types of uncertainty and the probability of unexpected events. However, because the animal's cognitive resources are constrained, precise estimations are unfeasible. What kind of behaviour would be adaptive under such circumstances? A purely exploitative strategy (i.e., picking what is currently the best food patch) is appropriate only when the probabilities or values of alternatives are stable. In volatile environments, switching between alternatives allows the organism to balance the maximization of rewards against the prevailing cognitive constraints [12]. By switching between patches, the animal can adjust optimally to a surprising outcome (i.e., a predator) without dedicating resources to monitoring for such an event [13]. Switching also allows the animal to adapt to the possibility of the availability of food changing without having to explicitly incorporate this possibility into the inference process [14].

Note that in the research discussed above, it is implied that the stochasticity in behaviour is governed by processes that are not under deliberate control. Specifically, the cognitive flexibility that allows the organism to quickly adapt to its environment [15, 16] is thought to result from variability in brain activity (i.e., the variance in different types of neural recordings). As Kloosterman and colleagues put it: the "neural system avoids locking into a stereotypical, rhythmic pattern of activity, while instead continuously exploring its full dynamic range to better prepare for unpredictably occurring events" [16, p. 2].

Behavioural variability as a diagnostic marker in cognitive modelling

In cognitive modelling, probabilistic choice (and the associated stochasticity in cognitive processes) has commonly been treated as something that dilutes and obscures the actual behavioural process under investigation [17]. Yet recent developments suggest that this approach may throw the baby out with the bathwater, and that acknowledging stochasticity as a substantive element of cognition can have genuine value for understanding cognition [18]. In what follows, we showcase how a more principled approach to behavioural variability can be used to compare and evaluate models, and can help explain behavioural biases.

Using probabilistic choice for model comparison and development

Probabilistic choice can be used to compare theoretical predictions about specific choice patterns with the observed empirical behaviour. In the context of violations of **transparent dominance** in repeated choices, for example, it has been pointed out that the predictions of some models of preferential choice do not correspond to the observed rates: Some models overestimate the rates; others underestimate them [19], indicating that the models' assumptions accommodating stochasticity may be inadequate. Similarly, Birnbaum used true and error theory [20] to estimate the rate of violations of **stochastic dominance** in people's choices. He found that the empirical rates of violations were inconsistent with the predictions of several of the existing models of risky choice [21, 22].

Moreover, specific patterns of probabilistic choice can point to descriptive limitations of a model. For example, Loomes and Sugden noted that several prominent models, such as the random preference model, cannot account for asymmetric preference reversals [23]—the empirical observation that people are more likely to switch from a risky alternative to a safe one than the other way around. A similar dependency between consecutive decisions was found by Kubovy and Healy [24] in the context of categorization, where subsequent categorization responses were conditional on whether the current response was correct or not.

Using response stochasticity to distinguish between judgement strategies

In 1955, Egon Brunswik proposed a distinction between intuitive and analytical cognitive processes in judgement [25]. Intuitive processes are characterized by noisy encoding and processing of perceptual information; analytical processes are based on a deterministic algorithm that implements an explicit rule—for example, a mathematical equation. In an experiment where participants were asked to estimate the size of objects, the two types of process resulted in markedly different distributions of errors: Intuitive processing led to few precisely correct answers but to small errors on average, whereas analytical processing led to precisely correct answers but also to large errors. Recently, this approach was revived by Sundh and colleagues, who developed it into a computational model [26]. Not only were they able to validate their model within Brunswik's original setting, but they also used it to distinguish between the two types of process in a different set of tasks. Thus, observed response stochasticity can be highly valuable for identifying cognitive processes and imply that a common approach of modelling using normally distributed errors may misspecify the process underlying the response distribution. Ignoring these important differences can lead to incorrect conclusions about the underlying processes.

Stochasticity in information sampling can explain apparent behavioural biases

Recent research has demonstrated that a more principled approach to stochastic components of cognitive models can provide a simpler and more unifying explanation for human judgement and choice in a variety of psychological tasks. For instance, it has been shown that stochasticity in internal information sampling can explain patterns of intertemporal choice [12]. To illustrate, temporal discounting may be the result of a person engaging in a "noisy" simulation of future rewards (i.e., sampling of possible outcomes). To estimate the value of the future rewards, the person combines the results of the simulations with their prior beliefs about the rewards. Because the future is associated with uncertainty, the resulting valuation will rely more heavily on the prior information, resulting in discounting of the reward value. Adopting this perspective, Gershman and Bhui showed that the magnitude effect in intertemporal choice—the phenomenon that people are more patient when faced with options involving higher rewards—might be due to people investing higher mental effort when faced with higher rewards, which increases precision (i.e., reduces noise) during sampling, which in turn leads to less discounting.

Sampling and the stochasticity associated with it have also been proposed as an explanation of biases in probability judgement. An approach called Probability Theory plus Noise (PT+N) assumes that probability judgements largely follow the basic laws of probability theory but are distorted due to noisy information retrieval during sampling [27]. PT+N) provides a unifying account of a variety of biases in probability judgements. For example, conservatism in probability judgements—the phenomenon that people are reluctant to produce probabilities of 0 and 1—may be the result of the erroneous retrieval of the event complementary to the event in focus: When estimating the probability of an event A, the person retrieves instances of ¬A, which leads to the estimated probability of A not being 0 (or 1). Building on this approach, Zhu and colleagues [28] showed that the predictions of PT+N can be improved for conditional probability judgements when, instead of sampling instances of A∪B (i.e., when both A and B are true) and B (i.e., when B is true), instances of A|B are sampled (i.e., whether A is true or not conditional on B being true).[1]

Other work has focused on how stochasticity in behaviour might be linked to the stimulus input and how this link might help explain biases in human judgement of averages [30] and magnitude sensitivity in value-based decisions [31]. Prat-Carrabin and Woodford [30] proposed that people calculating averages weigh numbers differentially and in a non-linear manner due to the noisy encoding of the stimuli. They argued that the amount of stochasticity during the encoding of a stimulus might be linked with the probability of the stimulus occurring during the experiment, and that less likely stimuli will be encoded with more stochasticity. If so, the distribution of participants' estimates should depend on the prior distribution of stimuli. The authors found empirical support for this dependence. In a similar vein, Pirrone and colleagues [31] have suggested a simpler explanation for the magnitude sensitivity observed in value-based choices (i.e., that response times are lower when rewards are higher). Previous research has suggested that relaxing assumptions of linear

[1] Furthermore, Zhu and colleagues showed that the predictions of the PT+N approach are largely equivalent to an implementation of a Bayesian model where the posterior is approximated using the collected samples and the sampled information is adjusted for the (usually) small sample size by using a prior that reflects the inherent uncertainty of sampling and "a conception of probability estimates in a more general [...] sense" [29, p. 2844]. This approach provides a unifying account of a variety of biases in probability judgements.

utility or linear cost of time can account for this magnitude sensitivity. According to Pirrone and colleagues, the simpler explanation is that the stochasticity associated with information accumulation is linked with the magnitude of the rewards, and that higher rewards are associated with higher stochasticity.

Three origins of probabilistic choice

Despite the increasing interest in probabilistic choice (and behavioural variability in general), the respective research is scattered across several fields and tends not to take a comprehensive approach to probabilistic choice, that acknowledges its multiple possible origins. In pursuit of a more encompassing perspective, we propose an organizing framework that distinguishes and synthesizes three possible sources of probabilistic choice that have been discussed in the literature: stochasticity in internal processing, reaction to uncertainty in the environment, and implementation stochasticity.

We sketch the framework around a general description of the cognitive process (see Figure 1). First, the organism perceives information relevant to the current task. Second, they process this information, using memory and making any other necessary computations. Third, they make a choice based on the processed information and finally implement it behaviourally by committing to some action. In what follows, we describe these three sources of probabilistic choice and how they affect each stage of the cognitive process.

Stochasticity in internal processing

Research in psychology and neuroscience usually distinguishes two types of stochasticity that arise from internal processing within the cognitive system (this is similar to what is sometimes referred to as "Thurstonian" uncertainty; [32]): perceptual stochasticity and computational stochasticity [33]. The right

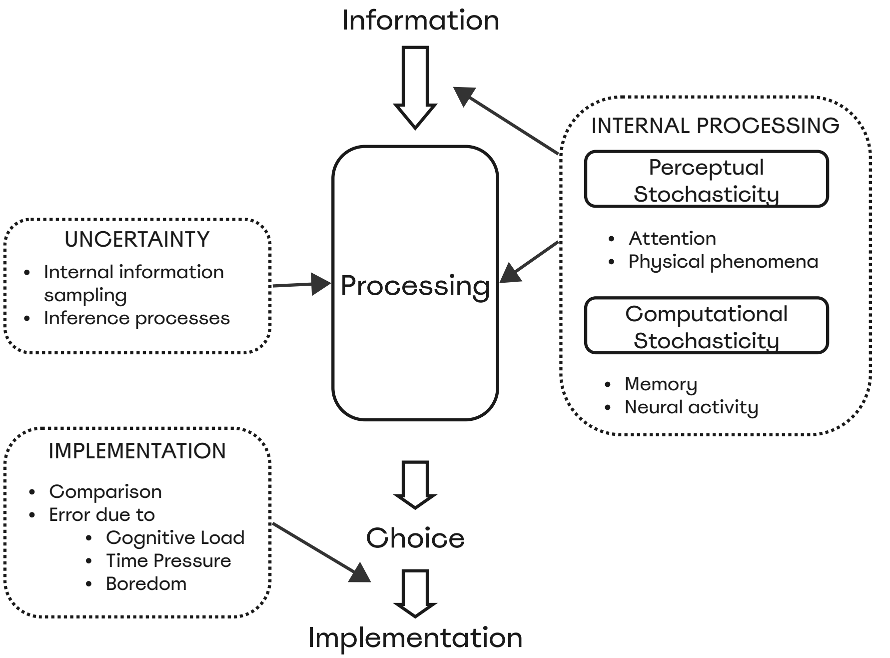

Fig. 1 Three origins of probabilistic choice.

side of Figure 1 shows how these two types relate to different stages of the cognitive process.

The first type of stochasticity originates in the perceptual system and is the result of physical phenomena (e.g., the activity of photons) or early perceptual systems. It thus affects the information before it is processed. Several cognitive models incorporate this type of stochasticity by assuming that probabilistic choice is a result of attention processes. Specifically, the idea is that people shift their attention stochastically between the alternatives [34] or between the features of the alternatives [35, 36]. Other models identify the noisy perception of information as the source of probabilistic choice. For example, it has been suggested that probabilistic choice in the domain of risky choice can be explained by stochasticity in the encoding of numerical magnitudes [37, 38].

The second type of stochasticity from internal processing originates in computational processes—that is, in how sensory input is mapped onto an internal representation after it has been encoded. Recent research in the domain of risky choice has connected variability in the stochasticity of neural activity (e.g., variability in blood oxygenation level-dependent activity) to the amount of behavioural variability exhibited by participants [39, 40]. Kurtz-David and colleagues have argued that stochasticity in neural processes results in distortions in people's estimations of the value of alternatives. Memory processes, which are known to be variable, may also contribute to this computational stochasticity [41]. Some models suggest that alternatives are evaluated based on samples drawn from memory representations of the alternatives [42]. Because the random nature of memory sampling leads to different samples of retrieved memories across occasions, the estimated value is likewise variable.

Reaction to uncertainty in the environment

A second origin of probabilistic choice is the structure of the environment. The information available about the value of the alternatives is usually uncertain to some extent (this is sometimes referred to as "Brunswikian" uncertainty; [32]). Higher uncertainty is associated with more behavioural variability [43] and with lower maximization of reward [44]. The connection between environmental uncertainty and cognitive processing is shown on the left side of Figure 1.

Models that implement Bayesian computations using sampling illustrate how environmental uncertainty might feed into probabilistic choice. Consider again a foraging animal deciding which food patch to choose based on previous experience. According to Bayes' rule, the choice should be based on posterior beliefs that combine prior beliefs about the patches with the animal's experience. Estimating such posterior beliefs becomes intractable when the number of patches increases. It has therefore been argued that rather than attempting to compute the posterior precisely, an animal approximates it by using random sampling of relevant information from memory or the hypothesis space [28]. Such sampling could lead to probabilistic choice. This suggestion is supported by recent findings showing that uncertainty about visual stimuli is encoded by the width of the probability distribution over the possible outcomes [45], represented by either individual neuron spike behaviour across trials [46] or a combined distribution of neuron pulls [47]. Furthermore, Prat-Carrabin and colleagues [48] compared human inferences in a learning task with that of an optimal Bayesian model. They found that although participants did not make inferences according to the Bayesian model, their responses were qualitatively consistent with it and that the behavioural variance was similar to that predicted by the Bayesian model. Moreover, they showed that human inferences are best explained by a model where the posterior is approximated by some form of sampling (e.g., particle filters).

As another example, it has been proposed that probabilistic choice during learning is the result of imprecise inference processes that help organisms adapt to a changing and thus uncertain environment [13, 14, 49]. Consider again the foraging example. An animal that approaches a patch and finds food there will incorporate this information in their current representation of the patch. The proposal is that the new information is incorporated only imprecisely (e.g., the learning rate varies; [50]). As a consequence, there is some imprecision in the estimated value of the alternative, leading to the possibility of a different (possibly suboptimal) alternative being selected when the choice is repeated. It has been argued that this imprecision helps to balance cognitive resources and accuracy in the face of uncertainty. Specifically, the random element stemming from the imprecision allows the animal to react to changes in the environment without having to invest cognitive resources in the monitoring of such possible changes.

Stochasticity in behavioural implementation

The final origin of probabilistic choice involves variability that occurs when a mental representation is mapped onto a choice (see bottom left of Figure 1). This type of stochasticity is sometimes referred to as response noise [33]. Two mechanisms have been discussed in the literature. First, probabilistic choice can be the result of comparative processes between the alternatives [51], with the degree of stochasticity being a function of how similar the computed values for the alternatives [e.g., **Luce Choice Rule**; 52].

Second, choice can be probabilistic because the organism makes an implementation error. In other words, the choice following from the computed values is not implemented correctly. A common approach implements this by introducing a parameter that expresses the probability that the implementation fails [i.e., **trembling hand error**; 53]. Several factors that might lead to stochasticity in behavioural implementation (e.g., pressing the wrong button) have been discussed in the literature, including cognitive load [54], time pressure [53], and boredom [55].

Concluding remarks and future perspectives

Choice is one of the basic processes in cognition, and organisms provided with the same information about the available alternatives will not always make the same choice. Although the phenomenon of probabilistic choice is generally recognized as an important element of any model of cognition, it has commonly been treated as a peripheral, unsystematic factor that interferes with and distracts from the process proper. Yet it appears that this view is slowly changing. In this article, we have reviewed recent developments across multiple fields in the cognitive and behavioural sciences that acknowledge probabilistic choice and the associated stochasticity of cognitive processes to be a multi-faceted and theoretically interesting element that can play a substantive role in models of cognition. We have discussed both the possible functionality of probabilistic choice for decision-making and how it can be harnessed to test and refine models of cognition. Moreover, we have sketched a conceptual framework distinguishing three main origins of probabilistic choice: stochasticity within internal processing, reaction to uncertainty in the environment, and implementation stochasticity. Our review offers several new avenues for researchers interested in probabilistic choice as well as general suggestions for cognitive modellers.

The framework presented in the previous section illuminates the multiple sources that can contribute to probabilistic choice. More effort should be directed at understanding the multi-faceted nature of behavioural variability—and especially the nature of the processes that lead to probabilistic choice as well as their potential adaptive value. For example, although much work has elaborated how environmental uncertainty might be reflected in cognitive and neural processes, thus giving rise to probabilistic choice, it remains an open question to what extent different types of uncertainty might result in different patterns of probabilistic choice, trigger different cognitive processes, or require different types of adaptation (see Box 2). See the Outstanding Questions for more suggestions.

On a more general level, we argue that measuring and analyzing patterns of probabilistic choice (and associated stochasticity) should become a core part of behavioural research. First, on an empirical level, researchers should take into account the complexity of probabilistic choice and its various possible origins when designing experiments, and make sure that the aspects important for the model(s) in question can be controlled. Second, on a theoretical level, researchers should analyze the response distributions, as they can help to distinguish between different cognitive mechanisms. Finally, on a methodological level, researchers should take into account observed patterns of probabilistic choice and contrast them with the model's predictions. It's high time to take probabilistic choice seriously.

Acknowledgements

We are grateful for the discussions with and comments by the members of the Center for Adaptive Rationality, Graham Loomes, and Nick Chater. We also thank Susannah Goss for editing this manuscript.

References

[1] Wallin, A., Swait, J. & Marley, A. A. Not just noise: A goal pursuit interpretation of stochastic choice. *Decision* **5**, 253–271 (2018). https://doi.org/10.1037/DEC0000077 .

[2] Agranov, M. & Ortoleva, P. Revealed preferences for randomization: An overview. *AEA Papers and Proceedings* **112**, 426–30 (2022). https://doi.org/10.1257/PANDP.20221093 .

[3] Olveczky, B. P., Otchy, T. M., Goldberg, J. H., Aronov, D. & Fee, M. S. Changes in the neural control of a complex motor sequence during learning. *Journal of Neurophysiology* **106**, 386–397 (2011). https://doi.org/10.1152/jn.00018.2011 .

[4] Wu, H. G., Miyamoto, Y. R., Castro, L. N. G., Ölveczky, B. P. & Smith, M. A. Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nature Neuroscience* **17**, 312–321 (2014). https://doi.org/10.1038/nn.3616 .

[5] Schulz, E. & Gershman, S. J. The algorithmic architecture of exploration in the human brain. *Current Opinion in Neurobiology* **55**, 7–14 (2019). https://doi.org/10.1016/J.CONB.2018.11.003 .

[6] Wilson, R. C., Bonawitz, E., Costa, V. D. & Ebitz, R. B. Balancing exploration and exploitation with information and randomization. *Current Opinion in Behavioral Sciences* **38**, 49–56 (2021). https://doi.org/10.1016/J.COBEHA.2020.10.001 .

[7] Icard, T. Why be random? *Mind* **130**, 111–139 (2021). https://doi.org/10.1093/MIND/FZZ065 .

[8] Leyton-Brown, K. & Shoham, Y. *Essentials of Game Theory* (Springer International Publishing, 2008).

[9] Humphries, D. A. & Driver, P. M. Protean defence by prey animals. *Oecologia* **5**, 285–302 (1970). https://doi.org/10.1007/BF00815496 .

[10] Winter, G., Wirsching, L. & Schielzeth, H. Condition dependence of (un)predictability in escape behavior of a grasshopper species. *Behavioral Ecology* 741–750 (2023). https://doi.org/10.1093/BEHECO/ARAD047 .

[11] Schulze, C., van Ravenzwaaij, D. & Newell, B. R. Of matchers and maximizers: How competition shapes choice under risk and uncertainty. *Cognitive Psychology* **78**, 78–98 (2015). https://doi.org/10.1016/j.cogpsych.2015.03.002 .

[12] Gershman, S. J. & Bhui, R. Rationally inattentive intertemporal choice. *Nature Communications* **11**, Article 3365 (2020). https://doi.org/10.1038/s41467-020-16852-y .

[13] Findling, C., Skvortsova, V., Dromnelle, R., Palminteri, S. & Wyart, V. Computational noise in reward-guided learning drives behavioral variability in volatile environments. *Nature Neuroscience* **22**, 2066–2077 (2019). https://doi.org/10.1038/s41593-019-0518-9 .

[14] Findling, C., Chopin, N. & Koechlin, E. Imprecise neural computations as a source of adaptive behaviour in volatile environments. *Nature Human Behaviour* **5**, 99–112 (2020). https://doi.org/10.1038/s41562-020-00971-z .

[15] Grady, C. L. & Garrett, D. D. Brain signal variability is modulated as a function of internal and external demand in younger and older adults. *NeuroImage* **169**, 510–523 (2018). https://doi.org/10.1016/J.NEUROIMAGE.2017.12.031 .

[16] Kloosterman, N. A., Kosciessa, J. Q., Lindenberger, U., Fahrenfort, J. J. & Garrett, D. D. Boosts in brain signal variability track liberal shifts in decision bias. *eLife* **9**, Article e54201 (2020). https://doi.org/10.7554/ELIFE.54201 .

[17] McAllister, J. W. What do patterns in empirical data tell us about the structure of the world? *Synthese* **182**, 73–87 (2011). https://doi.org/10.1007/S11229-009-9613-X .

[18] Bogen, J. Noise in the world. *Philosophy of Science* **77**, 778–791 (2010). https://doi.org/10.1086/656006 .

[19] Loomes, G. Modelling the stochastic component of behaviour in experiments: Some issues for the interpretation of data. *Experimental Economics* **8**, 301–323 (2005). https://doi.org/10.1007/s10683-005-5372-9 .

[20] Birnbaum, M. H. Testing mixture models of transitive preference: Comment on Regenwetter, Dana, and Davis-Stober (2011). *Psychological Review* **118**, 675–683 (2011). https://doi.org/10.1037/a0023852 .

[21] Birnbaum, M. H. & Bahra, J. P. Separating response variability from structural inconsistency to test models of risky decision making. *Judgment and Decision Making* **7**, 402–426 (2012). https://doi.org/10.1017/S1930297500002758 .

[22] Birnbaum, M. H. Testing transitivity of preference in individuals. *Decision* **10**, 153–180 (2022). https://doi.org/10.1037/DEC0000185 .

[23] Loomes, G. & Sugden, R. Testing different stochastic specifications of risky choice. *Economica* **65**, 581–598 (1998). https://doi.org/10.1111/1468-0335.00147 .

[24] Kubovy, M. & Healy, A. F. The decision rule in probabilistic categorization: What it is and how it is learned. *Journal of Experimental Psychology: General* **106**, 427–446 (1977). https://doi.org/10.1037/0096-3445.106.4.427 .

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

[26] Sundh, J., Collsiöö, A., Millroth, P. & Juslin, P. Precise/not precise (pnp): A Brunswikian model that uses judgment error distributions to identify

cognitive processes. *Psychonomic Bulletin & Review* **28**, 351–373 (2021). https://doi.org/10.3758/s13423-020-01805-9 .

[27] Howe, R. & Costello, F. Random variation and systematic biases in probability estimation. *Cognitive Psychology* **123**, 101306 (2020). https://doi.org/10.1016/J.COGPSYCH.2020.101306 .

[28] Zhu, J. Q., Sanborn, A. N. & Chater, N. The bayesian sampler: Generic Bayesian inference causes incoherence in human probability judgments. *Psychological Review* **127**, 719–748 (2020). https://doi.org/10.1037/REV0000190 .

[29] Sundh, J., Zhu, J.-Q., Chater, N. & Sanborn, A. A unified explanation of variability and bias in human probability judgments: How computational noise explains the mean–variance signature. *Journal of Experimental Psychology: General* **152**, 2842–2860 (2023). https://doi.org/10.1037/XGE0001414 .

[30] Prat-Carrabin, A. & Woodford, M. Efficient coding of numbers explains decision bias and noise. *Nature Human Behaviour* **6**, 1142–1152 (2022). https://doi.org/10.1038/s41562-022-01352-4 .

[31] Pirrone, A., Reina, A. & Gobet, F. Input-dependent noise can explain magnitude-sensitivity in optimal value-based decision-making. *Judgment and Decision Making* **16**, 1221–1233 (2021). https://doi.org/10.1017/S1930297500008408 .

[32] Juslin, P. & Olsson, H. Thurstonian and brunswikian origins of uncertainty in judgment: A sampling model of confidence in sensory discrimination. *Psychological Review* **104**, 344–366 (1997). https://doi.org/10.

1037/0033-295X.104.2.344 .

[33] Sanborn, A. *et al.* Noise in cognition: Bug or feature? *PsyArXiv* Published online on December 05, 2022 (2022). https://doi.org/10.31234/osf.io/438nd .

[34] Yang, X. & Krajbich, I. A dynamic computational model of gaze and choice in multi-attribute decisions. *Psychological Review* **130**, 52–70 (2022). https://doi.org/10.1037/REV0000350 .

[35] Tversky, A. Elimination by aspects: A theory of choice. *Psychological Review* **79**, 281–299 (1972). https://doi.org/10.1037/h0032955 .

[36] Usher, M. & McClelland, J. L. Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review* **111**, 757–769 (2004). https://doi.org/10.1037/0033-295X.111.3.757 .

[37] Khaw, M. W., Li, Z. & Woodford, M. Cognitive imprecision and small-stakes risk aversion. *The Review of Economic Studies* **88**, 1979–2013 (2021). https://doi.org/10.1093/RESTUD/RDAA044 .

[38] Garcia, M. B., Grueschow, M., Polania, R., Woodford, M. & Ruff, C. Individual risk attitudes arise from noise in neurocognitive magnitude representations. *Nature Human Behaviour* **7**, 1551–1567 (2023). https://doi.org/10.1038/s41562-023-01643-4 .

[39] Chew, B. *et al.* Endogenous fluctuations in the dopaminergic midbrain drive behavioral choice variability. *Proceedings of the National Academy of Sciences of the United States of America* **116**, 18732–18737 (2019). https://doi.org/10.1073/pnas.1900872116 .

[40] Kurtz-David, V., Persitz, D., Webb, R. & Levy, D. J. The neural computation of inconsistent choice behavior. *Nature Communications 2019 10:1* **10**, 1–14 (2019). https://doi.org/10.1038/s41467-019-09343-2 .

[41] Kahana, M. J., Aggarwal, E. V. & Phan, T. D. The variability puzzle in human memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **44**, 1857–1863 (2018). https://doi.org/10.1037/xlm0000553 .

[42] Dai, J., Pleskac, T. J. & Pachur, T. Dynamic cognitive models of intertemporal choice. *Cognitive Psychology* **104**, 29–56 (2018). https://doi.org/10.1016/j.cogpsych.2018.03.001 .

[43] Lee, D. & Coricelli, G. An empirical test of the role of value certainty in decision making. *Frontiers in Psychology* **11**, 574473 (2020). https://doi.org/10.3389/FPSYG.2020.574473 .

[44] Gershman, S. J. Deconstructing the human algorithms for exploration. *Cognition* **173**, 34–42 (2018). https://doi.org/10.1016/J.COGNITION.2017.12.014 .

[45] Bergen, R. S. & Jehee, J. F. Probabilistic representation in human visual cortex reflects uncertainty in serial decisions. *Journal of Neuroscience* **39**, 8164–8176 (2019). https://doi.org/10.1523/JNEUROSCI.3212-18.2019 .

[46] Hénaff, O. J., Boundy-Singer, Z. M., Meding, K., Ziemba, C. M. & Goris, R. L. T. Representation of visual uncertainty through neural gain variability. *Nature Communications* **11**, 1–12 (2020). https://doi.org/10.1038/s41467-020-15533-0 .

[47] Walker, E. Y., Cotton, R. J., Ma, W. J. & Tolias, A. S. A neural basis of probabilistic computation in visual cortex. *Nature Neuroscience* **23**, 122–129 (2019). https://doi.org/10.1038/s41593-019-0554-5 .

[48] Prat-Carrabin, A., Wilson, R. C., Cohen, J. D. & da Silveira, R. A. Human inference in changing environments with temporal structure. *Psychological Review* **128**, 879–912 (2021). https://doi.org/10.1037/REV0000276 .

[49] Findling, C. & Wyart, V. Computation noise in human learning and decision-making: origin, impact, function. *Current Opinion in Behavioral Sciences* **38**, 124–132 (2021). https://doi.org/10.1016/J.COBEHA.2021.02.018 .

[50] Wyart, V. & Koechlin, E. Choice variability and suboptimality in uncertain environments. *Current Opinion in Behavioral Sciences* **11**, 109–115 (2016). https://doi.org/10.1016/J.COBEHA.2016.07.003 .

[51] Woodford, M. Modeling imprecision in perception, valuation, and choice. *Annual Review of Economics* **12**, 579–601 (2020). https://doi.org/10.1146/annurev-economics-102819-040518 .

[52] Luce, R. D. On the possible psychophysical laws. *Psychological Review* **66**, 81–95 (1959). https://doi.org/10.1037/h0043178 .

[53] Olschewski, S. & Rieskamp, J. Distinguishing three effects of time pressure on risk taking: Choice consistency, risk preference, and strategy selection. *Journal of Behavioral Decision Making* **34**, 541–554 (2021). https://doi.org/10.1002/BDM.2228 .

[54] Olschewski, S., Rieskamp, J. & Scheibehenne, B. Taxing cognitive capacities reduces choice consistency rather than preference: A model-based

test. *Journal of Experimental Psychology: General* **147**, 462–484 (2018). https://doi.org/10.1037/xge0000403 .

[55] Yakobi, O. & Danckert, J. Boredom proneness is associated with noisy decision-making, not risk-taking. *Experimental Brain Research* **239**, 1807–1825 (2021). https://doi.org/10.1007/S00221-021-06098-5 .

[56] Waschke, L., Kloosterman, N. A., Obleser, J. & Garrett, D. D. Behavior needs neural variability. *Neuron* **109**, 751–766 (2021). https://doi.org/10.1016/J.NEURON.2021.01.023 .

Glossary

Probabilistic choice: The empirical phenomenon that people given the same evidence do not always make the same choice. Also known as an 'inconsistent' or 'stochastic' choice.

Stochasticity: A property of an element (or collection of elements) in a model contributing unsystematic variability to the model's predictions. It is usually implemented by assuming a random draw from a specified probability distribution.

Transparent dominance: When the worst outcome of alternative A is better than the best outcome of alternative B, then alternative A transparently dominates B. In other words, A is clearly the better alternative.

Stochastic dominance: When every possible outcome of alternative A is at least as likely under alternative B and one outcome is more likely under alternative A, then alternative A stochastically dominates alternative B.

Trembling hand error: Trembling hand error assumes that—irrespective of the difference in subjective valuation between the alternatives in a choice set—the alternative with the higher subjective valuation is selected with constant probability $1 - \epsilon$, where ϵ is the probability of making an error or that "the hand trembled."

Luce Choice Rule: This rule derives from the work of Duncan Luce and assumes that the probability of choosing alternative A is a function of the difference between the subjective valuation of alternative A and the sum of the subjective valuations of all alternatives in the choice set (including A).

Box 1: Probabilistic choice in cognitive models

Many formal cognitive models that accommodate the probabilistic nature of choice share a similar general structure, consisting of two components: a *core component* that specifies how the subjective value for each alternative is determined, and a *choice rule* that derives a probability that each alternative is chosen.

While both the core component and the choice rule have been used to incorporate stochasticity, most models use just one of the components—only a few assume stochasticity in both components. Proposals for how stochasticity can be incorporated in the core component vary considerably, from a straightforward Gaussian error term to a random process representing information accumulation. Most models incorporating stochasticity in the choice rule use one of four rules: Luce Choice Rule, softmax, probit, or the trembling hand error. In the first three, the predicted choice probability is a function of the relative evidence for each alternative. In the fourth, it is constant across different sets of compared alternatives.

Both the different stochastic components and the many possible implementations of stochasticity can, in principle, be combined in any way. This has resulted in a great heterogeneity of approaches to probabilistic choice in cognitive models. Often, there is no principled approach as to which implementation of probabilistic choice is employed in a given model. At the same time, this heterogeneity is not uniformly distributed across disciplines and fields: Assumptions and approaches seem to be clustered, further suggesting that they are guided by the specific conventions in a subfield rather than by general principles of cognition. This fragmentation suggests that probabilistic choice is often treated as an auxiliary aspect of the actual substantive theory of cognition. In many cases, models are equipped with a stochastic component simply to "accommodate human choice stochasticity" [44, p. 41], without providing a functional and process-specific rationale, or based on a model comparison of various probabilistic components.

Box 2: Types of uncertainty

Three types of uncertainty can be distinguished: *aleatory uncertainty*, *epistemic uncertainty* and *ambiguity*. Whereas aleatory uncertainty arises from objective physical features of the environment (e.g., the design of a die makes it inherently uncertain), both epistemic uncertainty and ambiguity arise from the agent's limited information about the environment.

Suppose a foraging animal knows that the current patch produces 10 berries and that there is a probability of 80% that they will be eaten by another animal. In other words, the animal knows the possible outcomes and their probabilities. Here, *aleatory uncertainty* would correspond to the animal being uncertain about whether there will be any food at the patch because there is an 80% chance that another animal has already eaten it.

Now, suppose the animal knows that 10 berries can be produced but does not know the probability that they will be eaten by another animal. In other words, the animal knows the possible outcomes but not their probabilities.

Here, *epistemic uncertainty* would correspond to the animal being uncertain about whether the food will be there because there is some possibility it has already been eaten.

Finally, suppose the animal has limited information about the amount of food and the probability it will be eaten by another animal, such that the animal knows neither all possible outcomes nor how likely they are. Here, *ambiguity* corresponds to the animal being uncertain about both the amount of food that might be there and the probability that any food will be there. A related concept to ambiguity is uncertainty about *known/unknown unknowns*. The animal might have a set of hypotheses about what could interfere with the food supply, with each hypothesis being more or less likely. For example, the animal might think it is likely that another animal will eat the food, but that it is less likely that insects will interfere with the food source, and very unlikely that the weather will affect the supply. This is uncertainty about *known unknowns*. It is also possible that hypotheses are extremely unlikely and/or have not even been considered. This is uncertainty about *unknown unknowns*.

Highlights

Probabilistic choice has long been treated as peripheral to the core processes of decision-making. Recent developments challenge that view and instead highlight the possible importance of probabilistic choice for understanding cognition.

First, probabilistic choice has been shown to have adaptive benefits for the organism.

Second, probabilistic choice (and the associated stochasticity) has been shown to serve as a diagnostic tool for testing, comparing, and refining cognitive models.

Toward a more structured perspective on probabilistic choice, we propose a framework that distinguishes three origins of probabilistic choice: stochasticity in the internal processing, in reaction to uncertainty in the environment, and during implementation of internal computations into a behaviour.

Outstanding questions

- Do patterns of probabilistic choice differ depending on the origin of probabilistic choice? In addition to intuitive and analytic strategies of judgement, what other cognitive strategies can be distinguished based on their predicted response distributions?
- Is the stochasticity that is due to internal processing the same that emerges in response to uncertainty in the environment? And how does the amount of uncertainty in the environment affect the amount of stochasticity in internal processing? For example, how does the variability in memory retrieval change under uncertainty? Does memory retrieval become more stochastic, resulting in more probabilistic choice?
- In light of the evidence that higher variability in neuronal firing is sometimes associated with higher task performance [e.g., 56], what exactly is the mechanistic relationship between neural variability and probabilistic choice?
- To what extent do different types of uncertainty result in different patterns of probabilistic choice, trigger different cognitive processes, or require different types of adaptivity?
- What are the empirical rates of implementation stochasticity across different tasks?