How The Bias-Variance Tradeoff Shapes Human Strategy Selection

Maximilian Maier (max.maier@wbs.ac.uk)[1,2], Falk Lieder[3], & Thomas L. Griffiths[4]
[1]Behavioural Sciences Group, Warwick Business School, University of Warwick
[2]Department of Experimental Psychology, University College London
[3]Department of Psychology, University of California, Los Angeles
[4]Department of Psychology and Department of Computer Science, Princeton University

Abstract

The extent to which biased decision-making is adaptive is one of the longest-debated questions in cognitive science. We investigate whether people adaptively select decision strategies following an optimal bias-variance trade-off: decision strategies that integrate (imperfect) estimates of many uncertain attributes may be unbiased but suffer from high variance, so simpler (biased) decision strategies can be more accurate when uncertainty is high. Using a new multi-attribute choice task, in which we manipulate uncertainty by varying the reliability of the estimates of different attributes across conditions, we provide an empirical demonstration that people adaptively shift their decision strategies in response to environmental uncertainty: participants were more likely to use single-attribute decision-making under high uncertainty and to integrate all attributes under low uncertainty. These results provide evidence that the bias-variance trade-off is a key principle shaping human strategy selection.

Keywords: decision-making; bounded rationality; strategy selection; bias-variance trade-off; heuristics

Introduction

Are humans mostly rational or mostly fallible decision makers? On the one hand, much research has shown that human decision-making is fallible and violates basic requirements of rationality, such as transitivity (Tversky, 1969; Tversky & Kahneman, 1974). Consistent with this view, research on human decision-making has identified a huge number of cognitive biases and distortions in judgment (the current Wikipedia entry on cognitive biases lists more than 150 such biases). On the other hand, proponents of ecological rationality have argued that, while these biases reflect deviations from classical definitions of rationality, they are useful inductive biases in real-life decisions (Gigerenzer & Brighton, 2009; Hertwig et al., 2018; Todd & Gigerenzer, 2012). Laboratory tasks are often not representative of decisions outside the lab, consequently biases that are maladaptive in an experimental task may be adaptive in the real world (Brunswik, 1955).

The perspective of ecological rationality postulates that people use "fast and frugal heuristics", which are adapted to the structure of the environment and perform better than more complicated decision strategies *independent of computational cost* (Gigerenzer, 2004; Gigerenzer & Kurzenhäuser, 2008; Todd & Gigerenzer, 2012). One of the theoretical arguments advanced for why simple heuristics can perform well is the bias-variance trade-off (Brighton & Gigerenzer, 2012; Gigerenzer & Brighton, 2009). The bias-variance trade-off, also often described in terms of the trade-off between underfitting and overfitting, is a concept from statistical learning theory that describes how the ability of a model to generalize to new data depends on its complexity and the amount of available information (Geman et al., 1992). Specifically, the mean squared error (MSE) of a statistical model on new data can be decomposed into bias (systematic deviation of the estimated function from the true function) and variance (sensitivity of the model to arbitrary, noisy changes in the training data):

$$\text{MSE} = \text{Bias}^2 + \text{Variance} \tag{1}$$

Which term dominates depends on the complexity of the model and the amount of available data. Fitting too many parameters to a small dataset can lead to overfitting and high variance as the predictions become very sensitive to noise. In contrast, a model with too few parameters has low variance (i.e., the predictions do not vary much with small changes in the data), but high bias and may be unable to capture a relevant signal. When fitting statistical models, it is important to strike the right balance between bias and variance, extracting signal from the data without fitting a model that captures noise. In introductory statistics classes, this issue is often illustrated with fitting polynomial regressions: A model with too many predictors relative to the data can always account for the data well, but risks overfitting the data and thus generalizing poorly to novel datasets. In contrast, a model with too few predictors may miss relevant signal in the data, also causing it to generalize poorly. Statisticians use techniques such as cross-validation or information criteria (Akaike, 1998) to select models that strike the right balance between overfitting (high variance) and underfitting (high bias).

In human decision-making, the bias-variance trade-off has been invoked to explain why simpler rules can often outperform more sophisticated decision strategies (Brighton & Gigerenzer, 2012; Gigerenzer & Brighton, 2009). Complex decision strategies,[1] such as cost-benefit reasoning, which assign probabilities and utilities to all possible outcomes (or

[1]When discussing the complexity of a decision strategy we only consider the number of attributes that need to be integrated or the number of parameters to be estimated, each of which may lead to more variance. A strategy could also be made more complex without using more data or increasing variance, for instance, by using a meta-strategy that averages multiple simple heuristics (Herzog & von Helversen, 2018). Studying this type of complexity is an interesting future direction for understanding how people could productively employ additional cognitive resources under high uncertainty.

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

all outcomes that the decision maker can consider), require estimating many parameters (e.g., the probabilities and utilities of the different outcomes), which leads to high variance when information is scarce (i.e., uncertainty is high, Bennis et al., 2010). Thus, using complex decision strategies that take many noisily measured attributes into account can lead to worse performance than simpler heuristics with fewer inputs.

The bias-variance dilemma has been postulated as an explanation for "less-is-more effects", where strategies that integrate all available information are outperformed by frugal heuristics such as *tallying* (which assigns equal weights to all attributes) or *take-the-best* (which ignores all attributes except for the most reliable one, Gigerenzer and Gaissmaier, 2011). These frugal heuristics perform worse than using linear regression weights on the same data used to estimate those weights. However, they can outperform linear regression when predicting new data. This is because simple heuristics are less likely to overfit the training data (Brighton & Gigerenzer, 2012).

The bias-variance trade-off has also been used to understand the costs and benefits of different learning strategies (Doroudi & Rastegar, 2023). For instance, Glaze et al. (2018) shows that participants fall along a continuum with respect to how much they rely on earlier experiences. Less history-dependent (high variance) strategies can adapt quickly to new environments but might also adapt to arbitrary noise. In contrast, less flexible, more history-dependent (lower variance) strategies take longer to adjust to new environments but are also less sensitive to random noise (Glaze et al., 2018). Further, Briscoe and Feldman (2011) showed how the bias-variance trade-off can help understand the costs and benefits of exemplar-based versus prototype-based category learning.

So far, the bias-variance trade-off has mostly been postulated as an ad-hoc explanation for why simple heuristics perform well in real-world decisions, but no research has tested whether people indeed systematically adjust which decision strategies they rely on based on the perceived uncertainty in the environment. Previous research on adaptive strategy selection has shown that people adjust their reliance on decision strategies in line with whether the application of a given strategy led to good outcomes in similar past situations (Lieder & Griffiths, 2017; Maier et al., 2025; Rieskamp & Otto, 2006). An open question is whether people can use information about environmental uncertainty to select a strategy with the right level of complexity, even if they have not been reinforced for applying this type of strategy to this type of problem.

To tackle this question, we designed a task where people learn about the uncertainty associated with different attributes so that either a simple strategy (single-reason decision-making) or a more complex strategy (integrating multiple predictors) would be optimal. Importantly, people did not receive any feedback about whether the application of either decision strategy was more successful; instead, they only learned about the measurement error of each strategy's inputs. We adjusted the uncertainty about each choice option's attributes (inputs) so single-reason decision-making is optimal in one environment with high uncertainty and integration of all predictors is optimal in another environment with lower uncertainty.

Manipulating the Bias-Variance Tradeoff in a Simple Multi-Attribute Decision Task

To test whether people adaptively adjust their reliance on decision strategies based on the uncertainty in line with the bias-variance trade-off, we designed a multi-attribute choice task: a hypothetical job selection task where participants decided between two candidates based on different predictors of job performance (i.e., a candidate's attributes; for more details see Method). Importantly, those attributes were noisily measured, and this noise differed across the two between-subjects conditions. In one condition, it was optimal to rely only on a single attribute, whereas in the other condition, it was optimal to integrate information from all three attributes.

Specifically, participants selected candidates based on three attributes that were equally important for job performance. True attribute values are simulated from a multivariate normal distribution with $\mu = 0$ and $\sigma = 1$ and were correlated with $\rho = .7$. The observed attribute values (interview scores) were calculated by adding normally distributed measurement errors to the true attribute values. Each measurement error was sampled from a normal distribution with mean 0 but their variances differed. Crucially, attribute 1 was always measured precisely ($SD = 0$), while attributes 2 and 3 were measured with different levels of precision depending on the experimental condition ($SD = 0.5$ in the low-noise condition vs. $SD = 2.5$ in the high-noise condition).

Operationalization of the Bias-Variance Trade-Off in the Experimental Paradigm. Here, the bias of a decision strategy s in assessing a candidate can be defined as the difference between the mean prediction of that strategy and the true score for a given applicant across K assessments, differing in noise:

$$\text{Bias}_s = \frac{1}{K}\sum_{k=1}^{K}\hat{y}_{k,s} - y. \tag{2}$$

The variance of a strategy is simply the variance of its predictions across assessments for a given candidate. Finally, the mean squared error (MSE) of the strategy is:

$$\text{MSE}_s = \frac{1}{K}\sum_{k=1}^{K}\left(\hat{y}_k^{(s)} - y_k\right)^2 = \text{Bias}_s^2 + \text{Var}_s, \tag{3}$$

For the assessment of the skill difference between two candidates we extend this decomposition by replacing y and $\hat{y}$ in the equations above with $\Delta\hat{y}$ and Δy, which denote the difference in predicted performance between the two candidates when each of them is evaluated using a decision strategy s and the difference in true performance.

To evaluate how the bias-variance trade-off shapes the performance of different decision strategies in our task, we simulated 1000 applicant pairs based on the underlying distribution of true skills (correlated skills with $\mu = 0$, $\sigma = 1$, & $\rho = .7$).

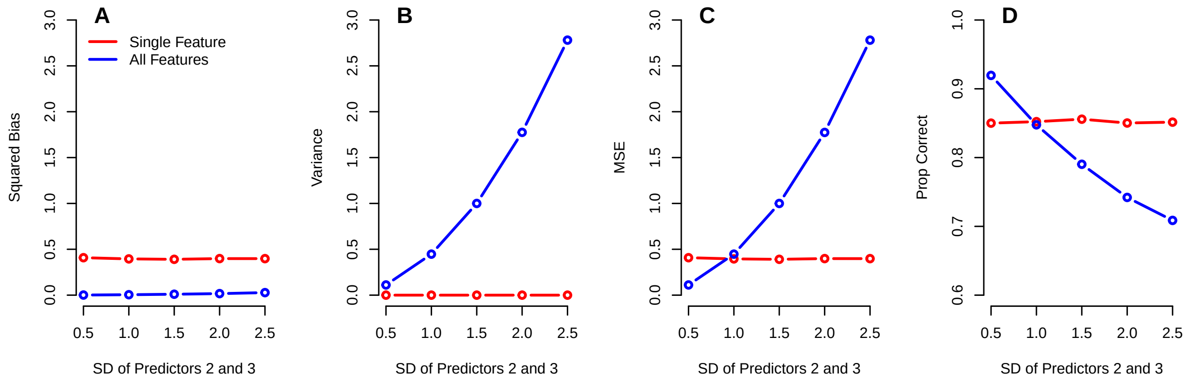

Figure 1: Bias-variance decomposition for the difference between two candidates and implications for proportion of correct choices when following each decision strategy. Simulated using 1000 iterations each with different error terms (horizontal axis) for 1000 different candidates, which were simulated from a multivariate normal with $\mu = 0$, $\sigma = 1$, and correlation between dimensions of $\rho = .7$. Panels A-C show the squared bias, variance, and MSE of the estimated skill difference between two candidates using each strategy, respectively. Panel D shows the proportion of times each strategy chooses the better candidate in a binary decision.

For each of these 1000 simulated candidates, we then simulated 1000 assessments with five different levels of measurement error for skills 2 and 3 (SD $\in \{0.5, 1, 1.5, 2, 2.5\}$) and evaluated bias, variance, and MSE for each decision strategy (single-reason decision-making versus using all attributes).

Figure 1 shows the average values across all simulation conditions. The results indicate that the single-attribute strategy has constant bias (Figure 1A) and no variance (Figure 1B). The reason is that attribute 1 is always measured precisely; therefore, the choices of this strategy are not stochastic as they do not depend on the SD of attributes 2 and 3. In contrast, the strategy of using all attributes has no bias, because, on average across all simulations, it recovers the true value. However, with increasing SD of attributes 2 and 3, it has increasing variance, which under high uncertainty leads to a larger MSE than single-attribute decision-making (Figure 1C). Finally, Figure 1D shows the proportion of correct decisions according to each decision strategy. Following the MSE, we see a crossover whereby using all attributes performs better under low uncertainty, but when the measurement error with which the true scores are assessed increases, singe-reason decision-making starts to perform better. Overall, the experimental task captures the key idea of adaptive strategy selection according to the bias-variance trade-off. We next test whether people indeed shift their decision strategy based on the difference in uncertainty with which attributes 2 and 3 are assessed.

Related Paradigms. Our paradigm is similar to previous work on strategy selection learning (e.g., Binz et al., 2022; Rieskamp & Otto, 2006) and heuristic decision-making, specifically in the context of the single-attribute heuristic (Hogarth & Karelaia, 2005, 2007; Karelaia & Hogarth, 2008) or the take-the-best heuristic (Bröder & Eichler, 2006; Bröder & Schiffer, 2003; Newell et al., 2003). However, in previous experiments that involved learning, participants learned about the outcomes of their decisions directly. In our task, participants do not receive feedback about the outcomes of their decisions. Thus, in our paradigm, strategy selection cannot be based on outcome learning but must be based solely on learning about the environment's structure. Another related literature is multiple-cue probability learning (Cook & Schipper, 1985; Jones et al., 1978; York et al., 1987), where participants integrate multiple noisy cues to choose between two options. However, tasks in this tradition usually do not provide information about the cues' reliability in the form we do here (i.e., uncertainty about a predictor) but in terms of conditional probabilities of the event given the cue (which confounds validity and reliability, Cook and Schipper, 1985; Jones et al., 1978). A notable exception is York et al. (1987), who study multiple-cue probability learning and deconfounded validity and reliability; however, they provide outcome feedback after each trial rather than feedback about uncertainty as in our task.

Method

Participants. We recruited 150 US participants using Prolific. Twelve participants failed an attention check asking them how candidate skills were scored, and 24 participants failed an attention check asking what the goal of the task was (i.e., hiring candidates that would perform well in the performance review). After these exclusions, we were left with a final sample of 117 participants: 66 male, 49 female, and two for whom gender information was unavailable. The mean age was 43.50 (SD = 13.59). The study received ethics approval from the University of Warwick Social Science Research Ethics Committee (HSSREC 32/25-26).

Design, Materials, and Procedure. After providing informed consent, participants received the following information about the hypothetical hiring task: "Imagine you are the head of an innovative aerospace company, 'Future Habitats Inc.', that specializes in designing sustainable liv-

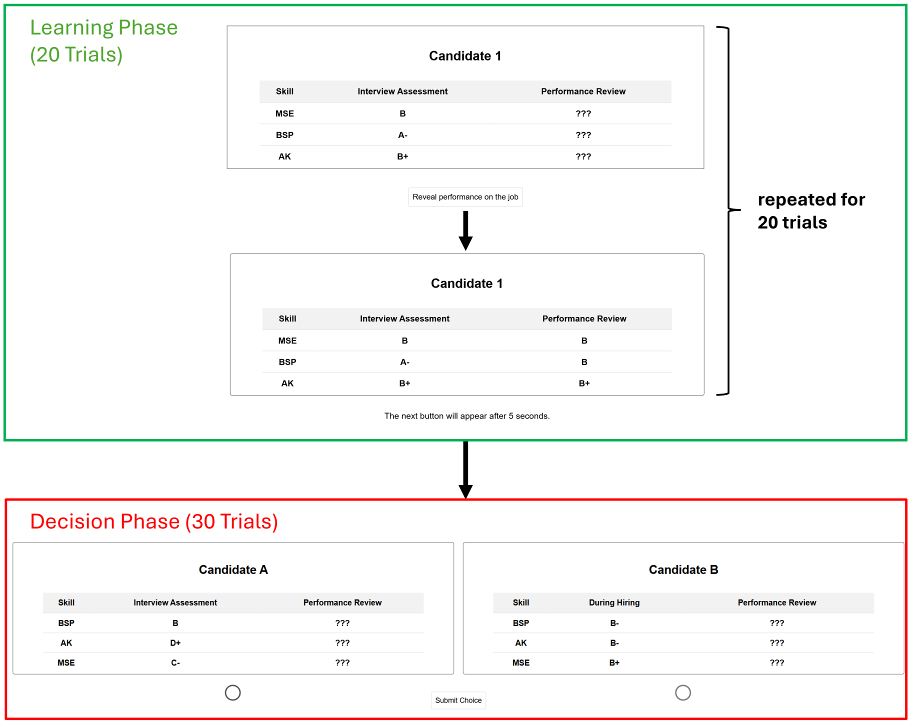

Figure 2: Experimental task. Participants choose the candidate who would have the best performance review based on a noisy assessment of skills. In a separate learning phase, participants learn how well each of the skills predicts performance.

ing environments for extraterrestrial settings as a first step in building human civilizations on other planets. You are currently in the process of hiring a new Exoplanetary Habitat Designer. This role is crucial for the company's upcoming project, which involves creating a self-sustaining habitat for humans on Mars." Participants were then instructed to hire the best candidates based on three (invented) skills important for habitat design (Astroengineering Knowledge, Biospheric Science Proficiency, & Materials Science Expertise) together with a short description of each skill.

Participants were informed that all skills were assessed using a letter grade between A+ and D- based on the candidate's interview assessment by the HR team. We transformed the z-scores sampled from the multivariate Gaussian model described in the previous section to letter grades as follows: z-scores above 1.645 were recorded as A+, z-scores below -1.645 were recoded as D-, and the space in between was divided into 10 equal-sized intervals corresponding to the 10 remaining letter grades (A, A-, B+, B, B-, ..., D). Participants were further informed that the HR team's assessments were only based on short interviews and prone to error. Consequently, there may be a discrepancy between a candidate's apparent competence according to the interview assessment and their actual performance on the job later, as assessed in a performance review. Participants in both conditions were informed that this probabilistic mismatch between interview impressions and job performance was more common and more pronounced for some skills than for others. The measurement error across skills varied based on experimental condition: attribute 1 was always measured precisely ($SD = 0$), while the standard deviation of the measurement error of attributes 2 and 3 was five times as high in the high-noise condition ($SD = 2.5$) as in the low-noise condition ($SD = 0.5$).

Figure 2 summarizes the experimental procedure. To learn how well different skills can be assessed, participants first completed 20 practice trials in which they viewed a participant's interview assessment, and, after clicking "reveal performance on the job", the candidate's actual performance for comparison. Participants needed to stay on the page showing the interview assessment and performance review for at least 5 seconds before they could continue to the next trial.

In the main phase of the experiment, participants made a hiring decision between two candidates based on interview assessment alone. To incentivise accuracy without allowing for learning about the success of different strategies directly from the outcomes, participants were told that they would receive a bonus of 0.05 for each correct decision at the end of the task, but there was no trial-by-trial feedback. Decision trials were pre-filtered such that on each trial, single-reason decision-making and using all predictors of job performance would lead to opposing results. This allowed us to infer strategies directly from participants choices based on whether they chose the option that single reason decision making would chose or the option that averaging of all features would choose.

However, because participants saw letter grades, rather than the underlying scores, the two decision strategies nevertheless made the same predictions in 14.98% of trials after the grades had been assigned due to rounding. We removed those trials, leaving us with an average of 25.50 (SD = 3.14) trials per participant (14 trials for the participants with the least trials).[2]

After completion of the task, participants answered two multiple-choice attention check questions ("What were you expected to do in this experiment?" [correct answer: "Hiring candidates that you think would have the best evaluation in the performance review."], and "How were candidate skills scored?" [correct answer: "On a letter scale from A+ to D-"]), some exploratory measures, and open questions regarding how they answered in the task.

Results

Strategy Selection is Driven by Uncertainty. Figure 3 shows that, as predicted, participants in the high-noise (HN) condition were more likely to rely on single-reason decision-making than participants in the low-noise (LN) condition (HN condition: 44.4%, 95% CI [32.8%, 56.6%]; LN condition: 19.4%, 95% CI [12%, 29.8%]). We tested this using a logistic mixed-effects model with random intercepts for participants and a fixed effect for condition and found a statistically significant difference ($\chi^2(1) = 9.42, p = .002$).

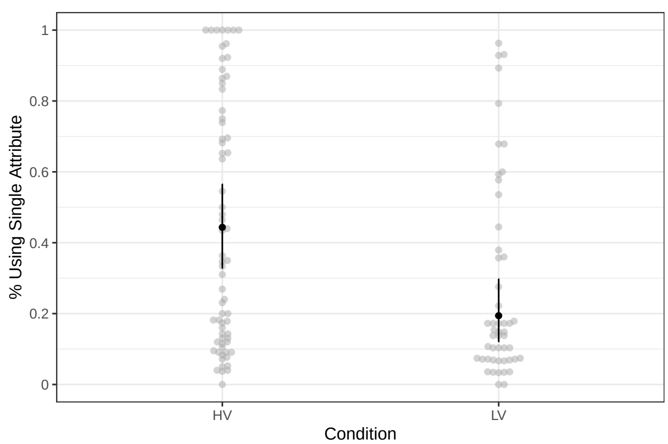

Figure 3: Participants in the high-noise condition are more likely to use single-attribute decision-making than in the low-noise condition. HN = high-noise condition, LN = low-noise condition. Error bars represent 95% CIs.

Accuracy of Different Strategies and Participants. Figure 4 shows participants' accuracy relative to single-attribute decision-making and averaging all attributes across the two experimental conditions. In the HN condition, single-attribute decision-making performed considerably better than averaging all attributes (75.3%, 95% CI [73.0%, 77.5%] vs. 24.7%, 95%CI [22.5%, 27%]). In contrast, in the LN condition, averaging all attributes performed considerably better than single-attribute decision-making (72.6%, 95% CI [70.0%, 75.0%]vs. 27.4%, 95% CI [25.0%, 30.0%]). Consistent with the finding that more participants relied on averaging in the LN condition than on single-attribute decision-making in the HN condition, participants performed slightly better in the LN condition than in the HN condition (proportion correct LN: 65.9%, 95%CI [60.9%, 70.6%]; proportion correct HN: 50.8%, 95%CI [46.0%, 55.5%], $\chi^2(1) = 17.15, p < .001$). Consistent with adaptive strategy selection, people chose the better candidate more often (58.55%, 95% CI [55.03%, 61.99%]) than they would have if they had always used single-reason decision-making (53.11%, 95% CI [51.24%, 54.99%]) or had always averaged (46.88% 95% CI [45.01%, 48.76%]).

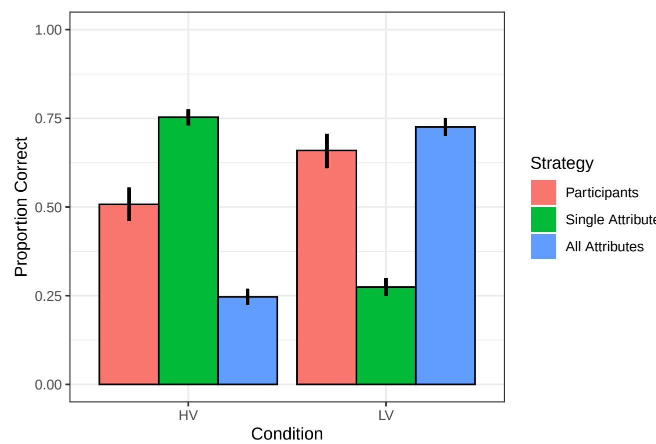

Figure 4: Proportion of correct decisions of participants, single-attribute decision-making, and averaging all attributes. HN = high-noise condition, LN = low-noise condition. Error bars represent 95% CIs.

Response Times. Rational use of cognitive resources would imply that participants 1) respond faster when using single-attribute decision-making than averaging (the latter is computationally more demanding as it requires integration of three attributes) and 2) respond faster in the HN condition, where single-attribute decision making performs better than in the LN condition. Consistent with this prediction, we found that participants responded faster in the HN than the LN condition ($F(1, 114.99) = 5.14, p = .025$). However, surprisingly, participants responded more slowly when using single-attribute decision-making than when using averaging, ($F(1, 93.29) = 8.04, p = .006$). We also observed a significant interaction between condition and single-attribute decision-making, ($F(1, 93.29) = 6.23, p = .014$), with simple contrasts revealing that averaging is associated with faster decision-making in the LN condition ($t(2979) = 5.06, p < .001$; averaging: 6.31s, 95% CI [5.45, 7.30]; single-attribute:

[2] See Figure 4 for the correct proportions of following each strategy by condition taking into account the pre-filtering. Relative to Figure 1 Panel D, the average proportions of correct choices are lower due to the pre-filtering, and the differences between strategies based on uncertainty are somewhat increased.

7.65s, 95% CI [6.55, 8.93]), but not in the HN condition ($t(2979) = 0.10, p = .919$; averaging: 5.49s, 95% CI [4.82, 6.24]; single-attribute: 5.47s, 95% CI [4.80, 6.23]).

Discussion

Our results provide evidence that participants adaptively shift whether they rely on decision strategies with more or fewer parameters based on learning about the uncertainty in the environment. Unlike in previous research on strategy selection learning (e.g., Lieder and Griffiths, 2017; Maier et al., 2025; Rieskamp and Otto, 2006), participants only received reward and feedback for accurate choices at the end of the experiment rather than on a trial-by-trial basis. In other words, they were able to adjust their decision strategies solely based on acquired knowledge about the structure of the environment.

How do these findings relate back to the debate about human rationality and whether people use adaptive "fast and frugal heuristics" versus second-best simplifications due to limited cognitive resources? It is our view that the bias-variance trade-off and the opportunity cost of using limited cognitive resources are two principles that jointly shape adaptive strategy selection. Because our task required the integration of relatively few attributes and did not induce strong resource constraints (participants could take as much time as they wanted), it focused mostly on the bias-variance trade-off as an understudied principle that shapes strategy selection. However, most likely, if we increased the number of attributes, the opportunity cost of integrating all of them would eventually outweigh the benefit in predictive performance, even if low uncertainty would make it adaptive to do so from a predictive perspective (see also Lieder and Griffiths, 2020; Maier et al., under review). By highlighting the optimal bias-variance trade-off as a mechanism of strategy selection in addition to resource constraints, this paper offers various future research directions and predictions regarding the interaction between these two factors. For instance, one would expect that more cognitive resources (e.g., providing people more time to deliberate) would increase strategy complexity only when the environment is characterized by relatively low uncertainty.

One surprising finding was that even though participants responded more slowly in the low-variance condition (in line with the adaptive use of cognitive resources), they responded more slowly when using single-attribute decision-making. This effect seems paradoxical, as integrating multiple attributes should take more time. The most likely explanation for this finding is that, unlike uncertainty, strategy usage was not directly manipulated. Thus, the relationship between strategy use and response time is confounded by participants' fluency in integrating and comparing attributes across alternatives. It is likely that participants who are faster at these types of operations are also more likely to employ decision strategies that integrate more information, thereby introducing a negative correlation between strategy complexity and response time. This negative relationship would likely disappear when controlling for participants' numerical abilities and processing speed.

Future Directions. In this paper, we contrasted only two decision strategies: single-reason decision-making and equal integration of all cues, which always lead to opposing choices in our task. However, it is likely that people would employ strategies falling between these two extremes, such as using shrinkage to regularize estimates based on information about the distribution of true values from practice trials or using all attributes with variable weighting. We are currently implementing computational models to describe these strategies and comparing them with the two approaches contrasted in this paper. Comparison with these additional models will also help further understand the role of limited cognitive resources: for instance, shrinkage-based methods may perform better than single-reason decision-making but are much more cognitively demanding. Further, the current study inferred choices from participants' decisions. Future work should employ process tracing methods in both the learning and the decision phase of the task to understand the underlying processes.

In addition to advancing descriptive models of participants' learning and decision making, future work could also improve the derivation of the optimal strategy. Previous work has derived the conditions under which single-variable decision-making outperforms other strategies (Hogarth & Karelaia, 2005) and the direct implementation of the bias-variance decomposition for binary choices (Kohavi, Wolpert, et al., 1996). An important next step is to integrate these models with Bayesian models of optimal learning during the practice trials to analytically derive optimal strategies for the task.

This study only manipulated one type of uncertainty: uncertainty about attributes that are used as inputs for a decision strategy. There are many other ways in which uncertainty can operate, such as uncertainty about which attributes are relevant for a decision, or uncertainty about what objective the decision maker should maximize (Kozyreva and Hertwig, 2021). These types of uncertainty might have very different implications for strategy selection. One of the authors is currently working on creating a taxonomy of uncertainty types and their implications for optimal and human decision-making.

Concluding Comments. This paper offers the first evidence that people adaptively shift their reliance on decision strategies based on the structure of the task environment by using simpler strategies when uncertainty is high and more complex strategies when uncertainty is low. This finding highlights the difficulty of generalizing from experimental tasks, which are often characterized by relatively constrained environments with known outcomes and probabilities, to the 'open world' of real-world decisions, which usually involves considerably more uncertainty. To better understand how people make decisions outside the laboratory, an important step for cognitive science is to analyze the uncertainty people face in everyday decisions and design experimental tasks that are representative of this uncertainty.

Data and Materials Availability

Data, materials, and analysis code are available at https://osf.io/j9dmk.

Acknowledgements

We thank Jonas Wolfram for helpful feedback on experimental materials, Vanessa Cheung for helpful comments on the manuscript, and Sean Enderby for setting up the server to host the experiment.

References

Akaike, H. (1998). Information theory and an extension of the maximum likelihood principle. In *Selected papers of hirotugu akaike* (pp. 199–213). Springer.

Bennis, W. M., Medin, D. L., & Bartels, D. M. (2010). The costs and benefits of calculation and moral rules. *Perspectives on Psychological Science*, *5*(2), 187–202. https://doi.org/10.1177/1745691610362354

Binz, M., Gershman, S. J., Schulz, E., & Endres, D. (2022). Heuristics from bounded meta-learned inference. *Psychological review*, *129*(5), 1042–1077.

Brighton, H., & Gigerenzer, G. (2012). Homo heuristicus and the bias–variance dilemma. *Action, Perception and the Brain*, 68–91.

Briscoe, E., & Feldman, J. (2011). Conceptual complexity and the bias/variance tradeoff. *Cognition*, *118*(1), 2–16. https://doi.org/10.1016/j.cognition.2010.10.004

Bröder, A., & Eichler, A. (2006). The use of recognition information and additional cues in inferences from memory. *Acta psychologica*, *121*(3), 275–284. https://doi.org/10.1016/j.actpsy.2005.07.001

Bröder, A., & Schiffer, S. (2003). Take the best versus simultaneous feature matching: Probabilistic inferences from memory and effects of reprensentation format. *Journal of Experimental Psychology: General*, *132*(2), 277–293.

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

Cook, P. J., & Schipper, L. M. (1985). Source reliability in a multiple cue probability processing task. *Bulletin of the Psychonomic Society*, *23*(1), 31–34. https://doi.org/10.3758/BF03329771

Doroudi, S., & Rastegar, S. A. (2023). The bias–variance tradeoff in cognitive science. *Cognitive Science*, *47*(1), e13241. https://doi.org/10.1111/cogs.13241

Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural computation*, *4*(1), 1–58.

Gigerenzer, G. (2004). Fast and frugal heuristics: The tools of bounded rationality. *Blackwell handbook of judgment and decision making*, *62*, 88.

Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, *1*(1), 107–143.

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual review of psychology*, *62*(2011), 451–482. https://doi.org/10.1146/annurev-psych-120709-145346

Gigerenzer, G., & Kurzenhäuser, S. (2008). Fast and frugal heuristics.

Glaze, C. M., Filipowicz, A. L., Kable, J. W., Balasubramanian, V., & Gold, J. I. (2018). A bias–variance trade-off governs individual differences in on-line learning in an unpredictable environment. *Nature Human Behaviour*, *2*(3), 213–224. https://doi.org/10.1038/s41562-018-0297-4

Hertwig, R., Hogarth, R. M., & Lejarraga, T. (2018). Experience and description: Exploring two paths to knowledge [Publisher: Sage Publications Sage CA: Los Angeles, CA]. *Current Directions in Psychological Science*, *27*(2), 123–128. https://doi.org/10.1177/0963721417740645

Herzog, S. M., & von Helversen, B. (2018). Strategy selection versus strategy blending: A predictive perspective on single-and multi-strategy accounts in multiple-cue estimation. *Journal of Behavioral Decision Making*, *31*(2), 233–249. https://doi.org/10.1002/bdm.1958

Hogarth, R. M., & Karelaia, N. (2005). Ignoring information in binary choice with continuous variables: When is less "more"? *Journal of Mathematical Psychology*, *49*(2), 115–124.

Hogarth, R. M., & Karelaia, N. (2007). Heuristic and linear models of judgment: Matching rules and environments. *Psychological Review*, *114*(3), 733–758.

Jones, D. P., Schipper, L. M., & Holzworth, R. J. (1978). Effects of amount of information on decision strategies. *The Journal of General Psychology*, *98*(2), 281–294. https://doi.org/10.1080/00221309.1978.9920882

Karelaia, N., & Hogarth, R. M. (2008). Determinants of linear judgment: A meta-analysis of lens model studies. *Psychological Bulletin*, *134*(3), 404–426.

Kohavi, R., Wolpert, D. H., et al. (1996). Bias plus variance decomposition for zero-one loss functions. *ICML*, *96*, 275–283.

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

Lieder, F., & Griffiths, T. L. (2017). Strategy selection as rational metareasoning. *Psychological Review*, *124*(6), 762.

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*.

Maier, M., Cheung, V., & Lieder, F. (2025). Learning from outcomes shapes reliance on moral rules versus cost-benefit reasoning. *Nature Human Behaviour*. https://doi.org/10.1038/s41562-025-02271-w

Maier, M., Cheung, V., & Lieder, F. (under review). Moral decision-making with bounded cognitive resources and limited information.

Newell, B. R., Weston, N. J., & Shanks, D. R. (2003). Empirical tests of a fast-and-frugal heuristic: Not everyone "takesthe-best". *Organizational Behavior and Human Decision Processes*, *91*(1), 82–96.

Rieskamp, J., & Otto, P. E. (2006). Ssl: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*(2), 207.

Todd, P. M., & Gigerenzer, G. (2012). *Ecological rationality: Intelligence in the world*. OUP USA.

Tversky, A. (1969). Intransitivity of preferences. *Psychological review*, *76*(1), 31. https://doi.org/10.1037/h0026750

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *science*, *185*(4157), 1124–1131.

York, K. M., Doherty, M. E., & Kamouri, J. (1987). The influence of cue unreliability on judgment in a multiple cue probability learning task. *Organizational Behavior and Human Decision Processes*, *39*(3), 303–317. https://doi.org/10.1016/0749-5978(87)90026-4