Resource-Rational Allocation of Metacognitive Resources in Value-Based Multi-Attribute Choice

[1, 2, 3]Trent N. Cash (ORCID: 0000-0001-6624-9616)

[2, 3]Daniel M. Oppenheimer (ORCID: 0000-0002-2363-4220)

[1]School of Psychology, University of Waterloo, Canada

[2]Department of Social and Decision Sciences, Carnegie Mellon University, USA

[3]Department of Psychology, Carnegie Mellon University, USA

Author Note

Correspondence concerning this article should be addressed to Trent N. Cash, Department of Psychology, University of Waterloo, 200 University Ave. West, PAS 4042, Waterloo, ON, N2L3G1, PAS 4042. E-Mail: tcash@uwaterloo.ca, Phone: 330-451-9972. The study reported here was pre-registered: https://aspredicted.org/vzkv-7rfk.pdf. Data, supplementary materials, and analysis code for this manuscript are available at: https://osf.io/z2ryg/?view_only=4571e45137dd404e8fc96e7de0f20360. This work was supported by a Doctoral Dissertation Research Improvement Grant from the National Science Foundation (Award #2333553), and an Academic Research Grant from Sawtooth Software, Inc. that provided us with free access to their Lighthouse Studio CBC software. No grantor had any role other than financial support. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and do not necessarily reflect the views of the grantors.

This is an author original manuscript of an article published by Taylor & Francis in *Thinking & Reasoning*, available upon publication at: http://dx.doi.org/10.1080/13546783.2026.2722961

Abstract

Decision makers overcome computational limitations by allocating limited cognitive resources to the most important information. Here, we investigate whether decision makers also efficiently allocate *metacognitive* resources when making computationally complex decisions. We randomly assigned participants (n = 815) to complete a value-based, multi-attribute choice task that included four or eight attributes. Participants then self-reported the weight they placed on each attribute. We compared the alignment between their stated weights and weights estimated from their choices – a measure of metacognitive knowledge. Participants in the 8-Attribute condition had less accurate metacognitive knowledge. This effect was driven by less accurate metacognitive knowledge only for low-importance attributes. Metacognitive knowledge for high-importance attributes was equally accurate across conditions. This suggests that, as tasks become more resource intensive, decision makers efficiently allocate their limited metacognitive resources to the most important attributes. These findings open the door to future research into the mechanisms underlying this resource-rational behavior.

Keywords: Metacognition; Resource-Rationality; Multi-Attribute Choice; Decision Making

Introduction

A grandmaster playing chess spends minutes contemplating complex mid-game positions but only seconds on obvious opening moves. If she tried to analyze every move with equal depth, she'd run out of time and lose the game. Like a chess player managing time, our minds manage cognitive effort. We can't scrutinize every decision equally, so we ideally invest the most cognitive resources in the decisions that matter most. However, decision making requires both cognitive *and* metacognitive resources. As tasks grow more demanding, our ability to monitor our own thinking often declines (Maniscalco et al., 2017; Maniscalco & Lau, 2015; Oppenheimer & Monin, 2009). So, it is reasonable to ask: do we allocate our limited metacognitive resources in a way that reflects rational trade-offs, or merely random failures? This paper explores how metacognition adapts to increasing task complexity to minimize errors given the constraints of our (meta-)cognitive system.

This conversation begins with acknowledging that, although humans are highly sophisticated creatures that can successfully complete a wide variety of cognitive tasks, our cognitive resources are limited. We can only keep a limited amount of information in our working memory (Miller, 1956), we have limited attention (Smith & Ratcliff, 2009), and we become fatigued when we expend mental effort (Boksem & Tops, 2008). As a result, decision makers often make less optimal choices when faced with decisions that include too much information to consider (Hastings & Weinstein, 2008; Phillips-Wren & Adya, 2020) or too many alternatives to choose from (Iyengar & Lepper, 2000; Timmermans, 1993). These complex decisions simply overwhelm our computational capacities.

The application of our cognitive resources is also influenced by other psychological and environmental factors. For example, the literature suggests that humans are cognitive misers who

attempt to minimize the cognitive effort they expend (Fiedler et al., 2019; Fiske & Taylor, 1991; Stanovich, 2018). As such, decision makers will default to using low-effort cognitive processes, unless they are sufficiently motivated to expend additional mental effort to engage in more cognitively burdensome reasoning processes (Evans & Stanovich, 2013; Stanovich, 2020). Similarly, time pressure may prevent a decision maker from maximizing the application of their cognitive resources (Ackerman, 2014; Ahituv et al., 1998). These additional factors suggest that even when a decision maker may have the mental horsepower to approach a decision rationally, their cognitive resources may be limited by extenuating circumstances.

Resource-Rational Decision Making

Despite these constraints, there is substantial evidence that decision makers engage in adaptive strategies that allow us to navigate complex decisions without sacrificing a great deal of optimality. These strategies often include using heuristics – mental shortcuts that allow decision makers to evaluate simplified versions of the complex decisions (Gilovich et al., 2002; Tversky & Kahneman, 1974). Similarly, decision makers may also engage in non-compensatory strategies in which they only consider a subset of the available information (Swait, 2001; Timmermans, 1993). These approaches are faster and require less mental effort than fully rational reasoning, and thus reduce the cognitive constraints associated with complex decision making (J. W. Payne et al., 1993; Shah & Oppenheimer, 2008).

Decision makers may also overcome complex decisions by using intuitive (i.e., System I), rather than deliberative (i.e., System II), decision processes (De Neys, 2012; De Neys, 2023; Evans, 2010; Evans & Stanovich, 2013; Klein, 2015). In these cases, decision makers may rely on past experiences (Betsch, 2008; Klein, 2015), expertise (Kahneman & Klein, 2009; Salas et al., 2010), or gut instincts (Sadler-Smith & Shefy, 2004) to minimize the effort they put into their

decision processes. Notably, however, this does not necessarily mean that these decision makers engage in less complicated cognitive processes. For example, Glöckner & Betsch (2008) and Brusovansky et al. (2018) both demonstrated that when participants were asked to make multi-attribute choice decisions under extreme time pressure, many still showed evidence of engaging in a complex attribute-weighting process. When participants engage in these kinds of intuitive decision processes, they are likely to be less sensitive to computational complexity and resource constraints (Betsch & Glöckner, 2010).

Though these lower-effort strategies do not always result in optimal decisions and can lead to biases (De Neys, 2012; Gilovich et al., 2002; Haselton et al., 2015; Tversky & Kahneman, 1974), they often result in sufficiently satisfactory outcomes (Schwartz et al., 2002). In line with the concepts of satisficing (Simon, 1956), bounded rationality (Camerer, 1998; Simon, 1955), and naturalistic decision making (Klein, 2015), many scholars have argued that humans often make the best possible decisions, given the constraints they face. When considering the resource-based limitations, some scholars claim that humans are resource-rational; we engage in adaptive strategies that allow us to make the most out of our limited cognitive resources (Bhui et al., 2021; Gigerenzer & Goldstein, 1996; Lieder & Griffiths, 2020).

Metacognition and Resource-Rationality

A growing body of research in the field of metareasoning – metacognition as applied to judgment, decision making, and problem solving (Ackerman & Thompson, 2017) – has demonstrated that metacognitive skills are essential for engaging in resource rational decision making, particularly when faced with complexity. *Metacognitive monitoring* processes allow decision makers to assess how complex decisions feel through the evaluation and aggregation of various metacognitive cues that hint at the difficulty of the task (Alter, 2013; Alter &

Oppenheimer, 2009; Beller, 2017; Koriat, 1997; Scheiter et al., 2020; Schwarz et al., 2021; Undorf & Bröder, 2020). Based on their interpretation of these cues, the decision maker can then adjust their cognitive strategies – a process known as *metacognitive control* (Ackerman & Thompson, 2017; Karpicke, 2009; Lieder & Griffiths, 2017). For example, if the task is difficult but worthwhile, the decision maker can increase the amount of resources they allocate to the task by spending more time on the task (Ackerman et al., 2023; D. G. Lee & Daunizeau, 2021) or by engaging in deeper, more effortful cognitive processes (Alter, 2013; Alter et al., 2013; Evans & Stanovich, 2013). However, if the decision maker does not feel that the decision warrants more effort, they may implement a less effortful decision strategy (Evans & Stanovich, 2013; Shah & Oppenheimer, 2008) or give up on the task entirely (Law et al., 2022; S. J. Payne & Duggan, 2011). From a resource-rational perspective, it is reasonable for a decision maker to metacognitively shift their limited resources to decisions – or aspects of a decision – that maximize their utility (Bhui et al., 2021; Lieder & Griffiths, 2020; Shenhav et al., 2021).

However, metacognitive processes also represent a form of mental effort that may be limited (D. G. Lee & Daunizeau, 2021; Maniscalco et al., 2017; Norman, 2020). As such, it is reasonable to ask: When decision makers are faced with complex decisions, do they also efficiently allocate their *metacognitive* resources to maximize the benefits of their effort? To evaluate this question, we must first address the question of whether metacognition is effortful and therefore limited. One key distinction in this discussion is between monitoring and control. Most theories of metacognition agree that metacognitive control is an effortful process (Ackerman & Thompson, 2017; Nelson & Narens, 1990). Metacognitive monitoring, on the other hand, is often perceived as an automatic process (Ackerman & Thompson, 2017; Koriat, 1997; Nelson & Narens, 1990; Son & Schwartz, 2002). This is because monitoring judgments

can arise from noetic feelings – such as the feeling of knowing (Koriat, 1993; Thompson et al., 2013) or fluency (Alter, 2013; Alter & Oppenheimer, 2009; Schwarz et al., 2021) – that occur as a byproduct of cognitive processes and thus require minimal additional cognitive effort.

However, these monitoring processes are necessarily error prone and their resultant judgments often are not diagnostic of performance on the cognitive task being monitored (De Bruin et al., 2017; Fleming & Lau, 2014; Lauterman & Ackerman, 2024; Moore & Schatz, 2017; Xue et al., 2021). As such, it is plausible that there exist metacognitive monitoring strategies that are not effortful, but such strategies may sacrifice accuracy, particularly when they are used in a computationally complex context (Shekhar & Rahnev, 2021).This perspective is supported by research in cognitive fluency, which has shown that people will spontaneously discount low-effort fluency-based cues when they believe they are not the most valid cues (Oppenheimer, 2004; Schwarz, 1998; Schwarz et al., 1991; Whittlesea & Williams, 1998). However, this kind of cognitive correction does not occur when people are under cognitive load (Oppenheimer & Monin, 2009), suggesting that metacognitive monitoring – and the process of parsing and interpreting multiple, competing cues – can be effortful.

Another reason we would expect metacognition to be subject to capacity constraints is that metacognitive processes draw upon more basic cognitive abilities, such as working memory (Conte et al., 2023) and attention (Peng & Tullis, 2021) that are known to be limited and subject to resource rational use (Lieder & Griffiths, 2020; Lieder et al., 2026). Indeed, empirical evidence indicates that metacognitive sensitivity declines with fatigue induced by completing cognitive tasks and recovers with rest, suggesting that accurate monitoring is effortful and dependent on limited cognitive resources (Maniscalco et al., 2017; Maniscalco & Lau, 2015). Furthermore, metacognitive sensitivity declines with age, suggesting that accurate monitoring

relies on effortful computational abilities that are known to decline among the elderly (Palmer et al., 2014). Thus, any task that overwhelms the resources associated with our basic cognitive abilities will likely also force our metacognitive systems to take lower-effort shortcuts.

The application of limited metacognitive resources also seems to be constrained by time-based (Ackerman, 2014) and motivational factors (Ackerman et al., 2023; Frederick, 2005; Kahan, 2013). In conjunction, these findings strongly suggest that at least some metacognitive processes are effortful, and that metacognition is a limited, computationally constrained resource. As such, it is reasonable to suspect that when decision makers are faced with complex decisions that exceed their metacognitive constraints, they will (consciously or not) implement resource-rational strategies to maximize the benefits they reap from their limited metacognitive resources. Surprisingly, however, there is a paucity of evidence in the literature supporting this possibility.

Metacognitive Knowledge as a Test Domain

To test the extent to which decision makers efficiently allocate their limited metacognitive resources, we explored participants' *metacognitive knowledge* – the knowledge they hold about how their cognitive processes function, which is acquired through metacognitive monitoring (Basu & Dixit, 2022; Cash & Oppenheimer, 2025; Colombo et al., 2010; Dunlosky & Metcalfe, 2009; Flavell, 1979; Morris et al., 2025). More specifically, we homed in on decision makers' metacognitive knowledge of the weight they place on various attributes when making value-based, multi-attribute choice decisions.

One reason to focus on this construct is that there exists a modern, validated paradigm that can efficiently and precisely estimate the accuracy of participants' metacognitive knowledge of attribute weights – called the *Knowledge of Weights* paradigm (Cash & Oppenheimer, 2025).

In this paradigm, participants complete a choice-based conjoint survey in which they make a series of value-based multi-attribute choice decisions, then self-report the weight they placed on each attribute. These self-reported weights can then be compared to weights estimated from the participants' choices to assess the accuracy of their metacognitive knowledge. The *Knowledge of Weights* paradigm has been tested in a wide variety of domains (Cash & Oppenheimer, 2024, 2025) and, critically, one can easily manipulate the computational complexity of the task being measured simply by changing the number of attributes included in the choice tasks. We will leverage this manipulation to investigate three research questions:

Question #1: Does computational complexity decrease the accuracy of participants' metacognitive knowledge? If so, we would expect the accuracy of decision makers' metacognitive knowledge to be lower when choosing between alternatives that vary on a greater number of attributes than when choosing between alternatives that vary on fewer attributes. This should be evident in both their relative and absolute metacognitive accuracy.

Question #2: When computational complexity is increased, do decision makers efficiently allocate their limited metacognitive resources to the most important information? If so, we would expect the detrimental effect of increasing the number of attributes to be greater for less-important attributes than more-important attributes.

Question #3: If decision makers allocate their limited metacognitive resources to the most important information, does this result in resource-rational outcomes? If so, we would expect behavioral (i.e., choice) accuracy to remain relatively stable, even when the relative and absolute accuracy of metacognitive knowledge for less-important attributes decreases.

Methods

Overview

Participants were randomly assigned to complete variants of the *Knowledge of Weights* paradigm that included either eight or four attributes. We then compared metacognitive knowledge – assessed as the alignment between participants' self-reported attribute weights and the weights revealed from their choices – across conditions to evaluate the impact of increasing the number of attributes. This study was pre-registered: https://aspredicted.org/vzkv-7rfk.pdf.

Participants

To determine the appropriate sample size for each condition, we conducted a simulation in which different numbers of artificial agents completed each choice survey, selecting random alternatives. We then measured the standard error of the utility estimates for these artificial agents and evaluated how the errors were impacted by the number of survey completions. Best practices (Halverson, 2020) suggest that these estimates should be less than .05. The simulation process indicated that this could be achieved with about 150 participants per condition. However, to increase precision we rounded our target sample size up to 200 participants per condition, which reduced the average standard error of utility estimates to about .042 for the simulated agents. This larger sample size also ensured that we would have at least 80% power to detect small-medium effects (d = .30; required n = 176 per condition) when comparing correlations across samples (Faul et al., 2007). To account for any possible participant exclusions across our four conditions (see *Materials and Manipulation)*, we recruited a total of 820 Prolific participants, 815 of whom completed the study (n = 200-207 in each condition). Demographic characteristics of the sample are reported in the Supplemental Materials. All participants passed at least three out of five attention checks (83.1% passed all five), so, following our pre-registration, no participants were excluded from our analyses. Participants were paid $2.00.

Procedures

In line with the *Knowledge of Weights* paradigm (Cash & Oppenheimer, 2025), participants first completed a Choice Based Conjoint (CBC) survey in which they were instructed to imagine that they were buying a new home. At the start of the CBC survey, participants were given a brief description of the attributes on which the homes would vary (n = 4 or 8). These descriptions laid out the five possible levels of each attribute, provided context for typical real-world values for each attribute, and explained the scale on which each attribute was described. Participants were then shown 14 sets of three hypothetical homes and asked to pick which one they would be most interested in purchasing.

After completing the CBC survey, participants were tasked with reporting the weight that they believed they had put on each attribute when making the hypothetical choices offered during the CBC task. These self-report items were *retrospective reflections* about the participants' past decisions – thus making them measures of metacognitive knowledge of their cognitive processes, rather than just a second method of eliciting preferences (Cash & Oppenheimer, 2025). However, the retrospective nature of this judgment makes it difficult to identify the extent to which inaccuracies in metacognitive knowledge result from online metacognitive errors (e.g., inability to learn behavioral patterns) or failures of memory (i.e., forgetting). Participants self-reported their attribute weights in two formats, presented in a random order. In one format, participants were tasked with self-reporting the *percentage* of their decision-making process that was based on each attribute (i.e., their attribute weights). Their responses were required to sum to 100% across attributes. We refer to these percentages as Stated Attribute Weights or **SAW**s.

In the other format, participants rated the importance of each attribute to their decision-making process on a Likert-style scale (*1 = Not At All Important; 9 = Extremely Important*). We will refer to these ratings as Attribute Importance Ratings, or **AIRs**. To establish incentive compatibility, participants were told that the 5% of participants whose self-reported weights most closely aligned with their revealed weights would earn a $0.75 bonus. After completing each self-report measure, participants also self-reported their confidence in the accuracy of their metacognitive ratings (see Supplemental Materials) and completed a demographic survey.

Materials and Manipulation

We randomly assigned participants to complete the *Knowledge of Weights* paradigm with either four or eight attributes. To tease apart the effects of increasing the number of attributes from the effects of the attributes themselves, we created three different four-attribute conditions that included different subsets of the attributes from the eight-attribute condition. Each of these conditions is described below.

8-Attribute condition: We used the six attributes validated by Cash and Oppenheimer (2025): 1) Commute time (to work); 2) Home size; 3) Mortgage (cost as % of income); 4) School district quality; 5) Attractiveness of the home (as rated by buyers); and 6) Lot size). We then added two new attributes that are typically provided in advertisements for homes: 7) Decade built; and 8) Parking. We created five discrete levels for each attribute which were used to create the alternatives in the CBC survey (see Supplemental Materials for details). These levels were designed to reflect what would be seen in typical homes. We then ran a small pilot study (n = 30) in which participants completed the CBC task and the SAW/AIR items with these eight attributes, which provided insight as to the relative weight participants placed on each attribute. We used these insights to design our various four attribute conditions, described below:

4-High Condition: In this condition, participants saw the four attributes that received the *highest* SAW ratings in the pilot study: 1) Commute time; 2) Mortgage; 3) School district quality; and 4) Parking.

4-Low Condition: In this condition, participants saw the four attributes that received the *lowest* SAW ratings in the pilot study: 1) Home size; 2) Attractiveness rating; 3) Lot size; and 4) Decade built.

4-Mixed Condition: In this condition, participants saw the two attributes that received the highest SAW ratings: 1) Mortgage; and 2) School district quality; and the two attributes that received the lowest SAW ratings: 3) Lot size; and 4) Decade built.

For each condition, we generated 300 versions of the CBC survey. Each version contained a unique set of 14 choices, each of which included three hypothetical homes that were random permutations of the possible levels of each attribute. Each participant experienced a unique set of choices. Survey versions were generated using Lighthouse Studio (Sawtooth Software, Inc., 2024).

Analysis & Results

Estimating Revealed Attribute Weights (RAWs)

To assess participants' metacognitive knowledge of their attribute weights, we first estimated the weights revealed by the choices that participants made during the CBC task. We refer to these as Revealed Attribute Weights, or **RAWs**. The first step in calculating these RAWs was to estimate the utility that each participant placed on each level of each attribute. To do so, we implemented Hierarchical Bayes (HB) estimation, using the HB function in Lighthouse

Studio (Sawtooth Software, Inc., 2024). HB estimation is generally considered the gold standard for analyzing CBC survey data (Eggers et al., 2022; Hein et al., 2020; Orme, 2002).

The HB estimation process began with the assumption of a null preference model in which the utility of each level of each attribute was set to zero. The model then randomly generated a new set of utilities for each level of each attribute for each participant. These utilities were then entered into a probability model that estimated the likelihood that the participant would have made the choices that they did if they had used the new set of utilities. If the participant's choices were more likely given the new set of utilities than the old set of utilities, then the new set of utilities was probabilistically retained as the current estimate and used to inform the generation of the next set of utilities. This continuous updating process – known as a Markov Chain Monte Carlo – allows the model to converge on accurate utility estimates. For this study, we allowed the model to generate 20,000 iterations, the last 10,000 of which were averaged to generate point estimates of utilities.

When the HB estimation process was completed, we were left with point estimates for the utility of each attribute for each participant (4 attributes x 5 levels = 20 utilities in the 4-attribute conditions; 8 attributes x 5 levels = 40 utilities in the 8-attribute conditions). These utility estimates were then converted into attribute-level RAWs using the following formula (Cash & Oppenheimer, 2025; Eggers et al., 2022; Orme, 2002), where U_j is a vector containing the utility values for the five levels of attribute *j* (the attribute for which the RAW is being calculated), U_i is a vector containing the utility values for the five levels of attribute *i,* and the set from which attribute *i* is pulled includes all attributes from the CBC, including *j:*

$$RAW_j = 100 * \frac{\max(U_j) - \min(U_j)}{\sum_{i=1}^{N}(\max(U_i) - \min(U_i))}$$

This process calculates RAWs separately for each participant, thus allowing for different preference orders across the levels of each attribute (e.g., one participant may prefer newer houses, while another may prefer older houses). This method of estimating weights was validated in a psychological context by Cash & Oppenheimer (2025). A more technical overview of the HB function can be found in Sawtooth Software, Inc. (2021).

Validating RAWs

We first sought to validate that our estimation procedure was generating RAWs that accurately described participants' choice behavior, To evaluate this, we individually analyzed each choice that each participant made during the CBC task (n = 815 participants x 14 choices = 11,410 choices) and estimated the utility they would have assigned to each of the three homes by multiplying their RAW for each attribute by the level of that attribute (scored as 1-5) and summing across the four or eight attributes. The attribute level scores were transformed to reflect each participant's rank-order preferences for the five levels of each attribute, with their highest-utility level being scored as a 5 and their lowest-utility level being scored as a 1. We then calculated how often RAWs predicted that the participant would give the highest utility to the alternative that they actually chose.

RAWs correctly predicted the participant's choice in 94.83% of choice tasks (95% CI = 93.95 – 95.60%) in the 8-Attribute condition; 89.68% of choice tasks (95% CI = 88.50 – 90.75%) in the 4-High condition; 88.50% of choice tasks (95% CI = CI: 87.25 – 89.65%) in the 4-Low condition; and 89.64% of choice tasks (95% CI = 88.44% - 90.73%) in the 4-Mixed condition. An omnibus four-sample test of equality of proportions indicated that these proportions were significantly different from one another, $X^2(3, N_{choices} = 11{,}410) = 81.68$, $p < .001$. Pairwise comparisons indicated that accuracy was significantly higher in the 8-attribute

condition than all other conditions (ps < .001), and this effect was robust to a Bonferroni correction for six pairwise comparisons (.05/6 = .008 = α). None of the 4-attribute conditions were significantly different from one another (ps ≥ .16). These results suggest that our RAWs were good estimates of participants' decision weights (~90% accurate). The greater accuracy in the 8-Attribute condition was unsurprising, given that having more attributes increases the number of data points in the model.

Descriptives

As a manipulation check, we evaluated the mean RAWs, SAWs, and AIRs for each attribute in the 8-Attribute condition (see Table 1). We found that the average weight for the four high-importance attributes was greater than the average weight for the four low-importance attributes, regardless of metric. Mean RAWs, SAWs, and AIRs for the 4-attribute conditions are reported in the Supplemental Materials.

Data Transformation.

RAWs and SAWs are both compositional in nature – that is, for each participant, the values are required to sum to the constant value of 100. As such, any shift in the RAW or SAW for a single attribute necessitates a compensatory shift in the RAWs or SAWs for one or more remaining attributes (i.e., a closure constraint). This results in non-independence between attributes, typically reflected in artificially negative correlations (Aitchison, 1982; Smithson & Broomell, 2024). Formally, this non-independence arises because compositional data live within a multi-dimensional simplex, rather than the unrestricted Euclidean space that is required for most statistical analyses, including regressions and correlations (Smithson & Broomell, 2024). Compositional data can be moved to Euclidean space by conducting one of a variety of log-ratio

transformations. Here, we use a Centered Log-Ratio Transformation in which we calculate the logarithm of the ratio between the RAW or SAW for each attribute and the geometric mean of all RAWs or SAWs for that participant (Cash & Oppenheimer, 2025; Smithson & Broomell, 2024). These transformed RAWs and SAWs (henceforth tRAWs and tSAWs) exist in Euclidean space, mitigating spurious dependence induced by closure constraints. Thus, tRAWs and tSAWs can be analyzed without violating the assumptions of standard statistical methods. This transformation was not pre-registered but follows best practices for analyzing compositional data (for a dedicated review of compositional data analysis methods, see Smithson & Broomell, 2024) and aligns with prior work using the KoW paradigm (Cash & Oppenheimer, 2025).

Table 1: Weights for the 8-Attribute Condition.

Attribute	Mean RAW (sd)	Mean SAW (sd)	Mean AIR (sd)
High Importance Attributes			
Mortgage	18.46 (9.46)	22.84 (19.38)	6.84 (2.04)
Commute	16.42 (8.59)	16.55 (15.53)	5.85 (2.53)
School Quality	16.26 (10.22)	11.50 (16.48)	4.47 (2.97)
Parking	11.75 (7.26)	12.84 (14.40)	5.47 (2.65)
Mean (4 High Only)	**15.72**	**15.93**	**5.59**
Low-Importance Attributes			
Attractiveness Rating	12.05 (7.49)	9.84 (11.39)	4.83 (2.31)
Home Size	9.65 (5.73)	11.93 (10.95)	5.34 (2.26)
Lot Size	8.02 (4.29)	8.79 (9.67)	4.89 (2.36)
Decade Built	7.38 (3.67)	5.71 (7.72)	3.97 (2.41)
Mean (4 Low Only)	**9.28**	**9.07**	**4.76**

Participant-Level Correlations between tRAWs and tSAWs

This section and the following two sections will seek to answer our first two research questions. In each section, we will test: 1) Whether metacognitive knowledge is less accurate in the 8-attribute condition than the 4-atttribute conditions; and 2) Whether reductions in the accuracy of metacognitive knowledge in the 8-attribute condition are larger for low-importance attributes than high-importance attributes. We first investigated participants' metacognitive knowledge by calculating the Pearson correlations between each participant's four or eight pairs of tRAWs and tSAWs. Participants who gave identical SAWs for each attribute (n = 9) were excluded from this analysis because their correlations were undefined. We then averaged these individual-level correlations for each condition. Power analyses using G-Power (Faul et al., 2007) indicated that t-tests comparing these average correlations across conditions would have 80% power (α = .05, ns = 200 per condition) to detect small to medium effects (d = .28).

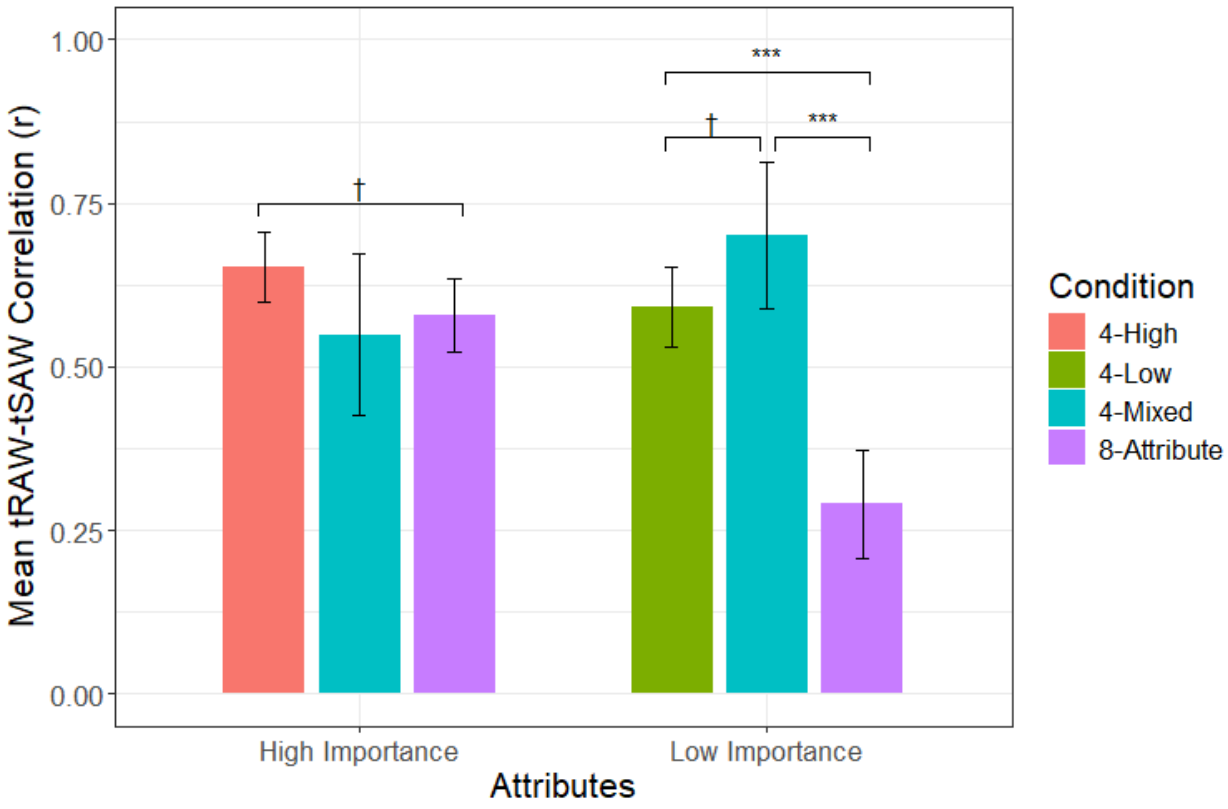
The mean participant-level correlations were 0.53 (sd = 0.30) in the 8-Attribute condition, 0.65 (sd = 0.39) in the 4-High condition, 0.59 (sd = 0.43) in the 4-Low condition, and 0.65 (sd = 0.39) in the 4-Mixed condition. A one-way ANOVA indicated that these mean correlations were significantly different from one another, $F(3, 802) = 4.86$, $p = .002$. Pairwise post-hoc comparisons corrected using Tukey's HSDs indicated that the mean tRAW-tSAW correlation in the 8-attribute condition was significantly lower than the mean correlation in the 4-Mixed condition ($p = .008$, $BF_{01} = 0.02$) and the 4-High condition ($p = .005$, $BF_{01} = 0.03$), but not the 4-Low condition ($p = .32$, $BF_{01} = 2.12$). None of the 4-Attribute conditions were significantly different from one another ($ps > .37$, $BFs_{01} \geq 3.13$). These results suggest that, on average, increasing the number of attributes decreased the accuracy of participants'

metacognitive knowledge of their attribute weights, though the effect did not reach significance when comparing the 8-Attribute condition to the 4-Low condition.

Next, we separately calculated correlations for the high- and low-importance attributes in the 8-Attribute condition. The mean correlation when only the four high-importance attributes were considered was 0.58 (sd = 0.41), which was marginally lower than the mean correlation in the 4-High condition, $t(406.82) = 1.85$, $p = .07$, $BF_{01} = 1.76$. The Bayes Factor suggested that this marginal difference was more supportive of the null hypothesis than the alternative. The mean correlation when only the four low-importance attributes were considered was 0.29 (sd = 0.59), which was significantly lower than the correlation in the 4-Low condition, $t(349.85) = 5.75$, $p < .001$, $BF_{01} < .001$. This effect was robust to a Bonferroni correction for six comparisons.

We then calculated separate correlations for the high- and low-importance attributes in the 4-Mixed condition. The mean correlation for high-importance attributes in the 4-Mixed condition was 0.55 (sd = 0.84). This was not significantly different than the mean correlation for the four high-importance attributes in the 8-Attribute condition, $t(255.69) = 0.44$, $p = .66$, $BF_{01} = 8.02$, or the mean correlation for the 4-High condition, $t(248.47) = 1.52$, $p = .13$, $BF_{01} = 2.66$. The mean correlation for the low-importance attributes in the 4-Mixed condition was 0.70 (sd = 0.72). This was significantly higher (robust to Bonferroni correction) than the mean for the four low-importance attributes in the 8-Attribute condition, $t(307.72) = 5.79$, $p < .001$, $BF_{01} < 0.001$ and marginally higher than the mean correlation in the 4-Low condition, $t(249.3) = 1.68$, $p = .09$, $BF_{01} = 1.93$. These results (summarized in Figure 1) suggest that the less accurate metacognitive knowledge of attribute weights in the 8-Attribute condition was primarily driven by the low-importance attributes.

Figure 1: Mean tRAW-tSAW Correlations for High- and Low-Importance Attributes



†$p < .10$; ***$p < .001$; *Note: Error bars depict 95% confidence intervals.*

Participant-Level Correlations between tRAWs and AIRs

We next evaluated whether the pattern of results was similar when evaluating participant-level correlations between tRAWs and AIRs (n = 14 excluded for giving identical AIRs to all attributes). These analyses had the same power as the comparable analyses of tRAW-tSAW correlations. The mean correlations were 0.53 (sd = 0.29) in the 8-Attribute condition, 0.64 (sd = 0.39) in the 4-High condition, 0.56 (sd = 0.48) in the 4-Low condition, and 0.67 (sd = 0.37) in the 4-Mixed condition. A one-way ANOVA indicated that these mean correlations were significantly different from one another, $F(3, 797) = 6.01$, $p < .001$. Pairwise post-hoc

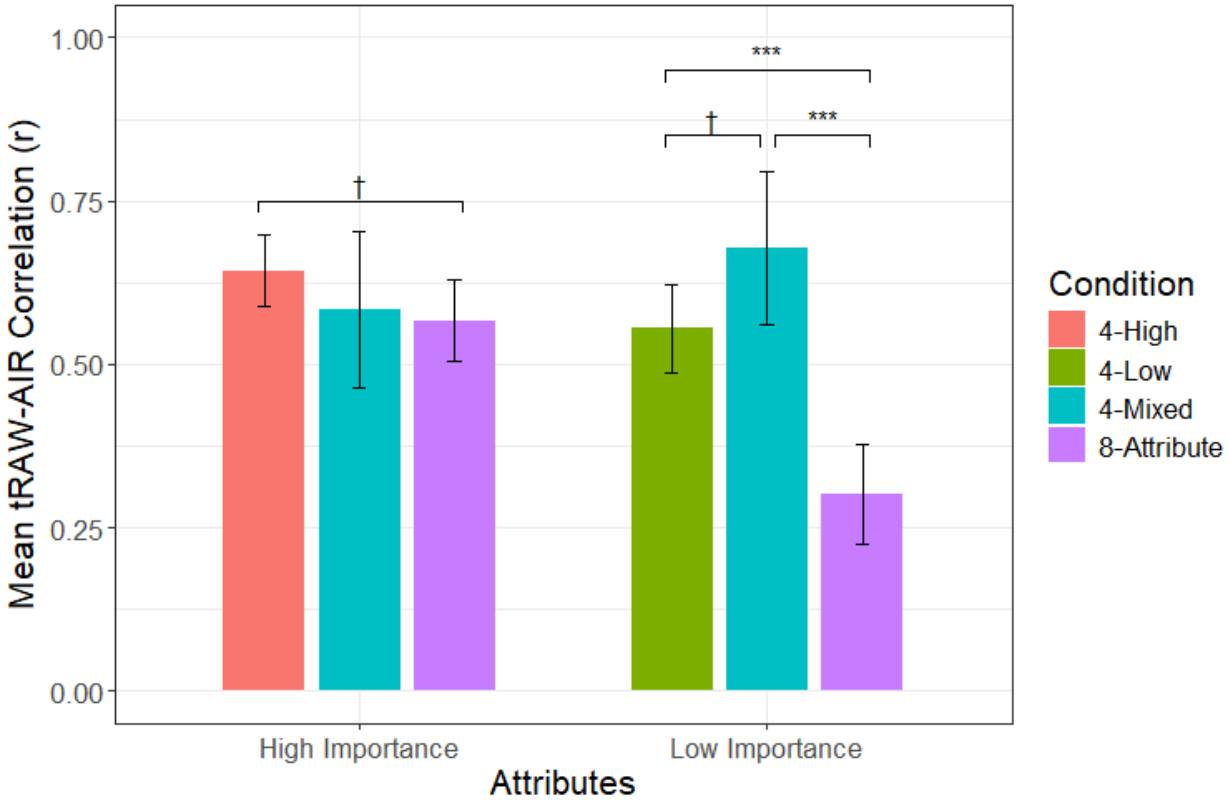
comparisons using Tukey's HSDs indicated that the mean correlation in the 8-attribute condition was significantly lower than the mean correlation in the 4-Mixed condition ($p = .002$, $BF_{01} = 0.002$) and the 4-High condition ($p = .02$, $BF_{01} = 0.04$), but not the 4-Low condition ($p = .90$, $BF_{01} = 7.24$). The mean correlation in the 4-Low condition was significantly lower than the mean correlation in the 4-Mixed condition ($p = .02$, $BF_{01} = 0.35$). None of the other differences between 4-attribute conditions were significant ($ps > .10$, $BFs_{01} \geq 1.25$).

We then calculated the participant-level tRAW-AIR correlations in the 8-Attribute condition for the high- and low-importance attributes. When only the high-importance attributes were considered, the mean correlation for the 8-Attribute condition was 0.57 ($sd = .45$). This was marginally lower than the mean correlation in the 4-High condition, $t(395.64) = 1.82$, $p = .07$, $BF_{01} = 1.84$, but again the Bayes Factor suggested that the null hypothesis was more likely than the alternative. When only the low-importance attributes were considered, the mean correlation for the 8-Attribute condition was 0.30 ($sd = 0.59$), which was significantly lower than the mean correlation in the 4-Low condition, $t(384.08) = 4.94$, $p < .001$, $BF_{01} < 0.001$. This significant difference was robust to a Bonferroni correction for six comparisons.

We then calculated participant-level correlations for the high- and low-importance attributes in the 4-Mixed condition. The mean correlation for the two high-importance attributes in the 4-Mixed condition was 0.58 ($sd = 0.81$). This was not significantly different than the mean correlation for the four high-importance attributes in the 8-Attribute condition, $t(266.96) = 0.25$, $p = .80$, $BF_{01} = 8.53$, or the 4-High condition, $t(245.63) = 0.87$, $p = .38$, $BF_{01} = 5.88$. However, the mean correlation for the low-importance attributes was 0.68 ($sd = 0.74$), which was significantly higher (robust to Bonferroni correction) than the mean for the four low-importance attributes in the 8-Attribute condition, $t(274) = 5.32$, $p < .001$, $BF_{01} < 0.001$ and marginally

higher than the mean correlation in the 4-Low condition, $t(250.23) = 1.79$, $p = .07$, $BF_{01} = 1.55$. These findings, summarized in Figure 2, once again suggest that, in the 8-Attribute condition, metacognitive knowledge was primarily hindered for low-importance attributes.

Figure 2: Mean tRAW-AIR Correlations for High- and Low-Importance Attributes



†$p < .10$; ***$p < .001$; *Note: Error bars depict 95% confidence intervals.*

Aitchison Distance between tRAWs and tSAWs

The correlational metrics reported in the previous section functionally capture decision makers' ability to rank-order the attributes by importance – a measure of their *relative accuracy* or *resolution* (Fleming & Lau, 2014). However, it is also of interest to test how closely the exact values of their SAWs match up with their RAWs – a measure of *absolute accuracy,* or

calibration (Fleming & Lau, 2014). Following Cash and Oppenheimer (2025), we assessed absolute accuracy by calculating the Aitchison Distance[1] (i.e., absolute distance) between tRAWs and tSAWs for each participant. For participants in the 8-attribute condition, we calculated Aitchison Distances separately for each 4-attribute subset (4-High, 4-Low, 4-Mixed; see Supplemental Materials for details). We calculated Aitchison Distances using the following formula, where *D* denotes an individual's Aitchison Distance, *i* refers to the different attributes in the set, and *K* refers to the total number of attributes:

$$D = \frac{1}{K}\sqrt{\sum_{i=1}^{K}(tRAW_i - tSAW_i)^2}$$

Since Aitchison Distance is an absolute distance metric, higher values indicate lower absolute accuracy. When evaluating the four high-importance attributes, the mean Aitchison Distance was 0.35 (95% CI = 0.32 - 0.37) in the 4-High condition and 0.44 (95% CI = 0.41 – 0.47) in the 8-Attribute condition. When evaluating the four low-importance attributes, the mean Aitchison Distance was 0.34 (95% CI = 0.32 – 0.36) in the 4-Low condition and 0.52 (95% CI = 0.48 – 0.56) in the 8-Attribute condition. When evaluating the two high-importance and two low-importance attributes used in the 4-Mixed condition, the mean Aitchison Distance was 0.38 (95% CI = 0.35 – 0.41) in the 4-Mixed condition and 0.51 (95% CI = 0.48 – 0.55) in the 8-Attribute condition. A one-way ANOVA comparing all six of these Aitchison Distances was significant, $F(5, 1215) = 25.04$, $p < .001$. Pairwise comparisons using Tukey's HSDs indicated that each Aitchison Distance from the 8-attribute condition was significantly higher than its counterpart from the 4-attribute conditions ($ps \leq .001$; $BFs_{01} < 0.001$). However, the magnitude of the effect

[1] Aitchison Distances are equivalent to Euclidean Distances, but are used for compositional data that has undergone a centered log-ratio transformation (Aitchison et al., 2000)

was smallest when evaluating only the high-importance attributes (d = 0.43), largest when evaluating low-importance attributes (d = 0.75), and in the middle when evaluating the mixed set of attributes (d = 0.57). This supports the conclusion that increasing the number of attributes decreases the accuracy of metacognitive knowledge writ large, but that the effect is stronger for less-important attributes, again suggesting resource-rational behavior.

SAWs Predicting Same Choices as RAWs

To answer our third research question, we turned to a meaningful behavioral measure of metacognitive accuracy – the extent to which participants' SAWs could predict the same choices as their RAWs. If this metric was similar across conditions, it would suggest that allocating additional metacognitive resources to high-importance attributes is a resource-rational behavior that lessens the impact of the additional computational complexity in the 8-attribute condition.

For this analysis, we used the same procedure that we used to validate the RAWs, but instead predicted which alternative would have had the highest utility if participants used their SAWs as their decision weights. We then evaluated the percentage of choice tasks for which SAWs and RAWs gave the highest utility to the same alternative (for validation of this method, see Cash & Oppenheimer, 2025). In the 8-attribute condition, RAWs and SAWs gave the highest utility to the same alternative for 83.64% of choice tasks (k = .75, 95% CI = 82.22% – 84.98%). In the 4-High condition, RAWs and SAWs gave the highest utility to the same alternative for 84.48% of choice tasks (k = .77, 95% CI = 83.08% - 85.79%). In the 4-Low condition, RAWs and SAWs gave the highest utility to the same alternative for 82.51% of choice tasks (k = .74, 95% CI = 81.02% - 83.91%). Finally, in the 4-Mixed condition, RAWs and SAWs gave the highest utility to the same alternative for 84.44% of choice tasks (k = .77, 95% CI = 83.02% - 85.77%). An omnibus four-sample test of equality of proportions indicated that these proportions

were not different from one another, with evidence strongly in favor of the null hypothesis, $X^2(3, N_{choices} = 11{,}132) = 5.16$, $p = .16$. $BF_{01} = 7{,}211.13$. These results suggest that even though participants in the 8-Attribute condition had less accurate metacognitive knowledge than participants in the 4-attribute conditions, their SAWs were just as likely to predict the same choice as their RAWs. This suggests that decision makers allocate their metacognitive resources in a resource-rational way. We find the same pattern of results when evaluating how often SAWs accurately predict decision makers' actual choices (see Supplemental Materials).

Deviations from Pre-Registration

In our initial pre-registration, we indicated that we would compare RAW-SAW correlations at the attribute level. In response to reviewer feedback about best practices in metacognition research (Fleming & Lau, 2014), we instead decided to report participant-level correlations across attributes. We report the attribute-level correlations in the supplemental materials. Furthermore, we indicated in our pre-registration that our analyses would generate metacognitive knowledge metrics for "averages across attributes," but did not specify that we would generate these metrics for subsets of attributes in the 8-Attribute condition. While this was always an intended analytical approach, it should be interpreted with appropriate caution. The centered log-ratio transformation was not pre-registered because we were not familiar with the transformation at the time of pre-registration. We chose to align with best practices rather than our pre-registration, but it is worth noting that the results were qualitatively similar without the transformation. Lastly, the analysis of Bayes Factors was not pre-registered, but was included in response to reviewer feedback. BFs were calculated using the BayesFactor package in R (Morey & Rouder, 2026).

Summary of Results

In sum, the results presented here indicate that when decision makers are faced with computationally complex decision environments, the accuracy of their metacognitive knowledge decreases relative to less-complex decision environments. However, decision makers mitigate the negative impacts of complexity by focusing their limited metacognitive resources on the most important information in the environment. This paid off at the behavioral level, as stated weights provided by participants in the 8-attribute condition were no less predictive of their actual choices than stated weights provided by participants in the 4-attribute conditions. These findings suggest that decision makers allocate metacognitive resources in a resource-rational way that has not been previously documented. This has significant implications for our understanding of how people choose when and how to engage in metacognitive monitoring and control.

Discussion

The decisions we make have implications for our health, happiness, and well-being. But decision making is hard. We are often faced with computationally complex decisions that force us to choose between many alternatives (Iyengar & Lepper, 2000; Timmermans, 1993) and sort through great deals of information (Hastings & Weinstein, 2008; Phillips-Wren & Adya, 2020). Despite this, humans have limited cognitive resources (Boksem & Tops, 2008; Miller, 1956; Smith & Ratcliff, 2009). These constraints prevent humans from being fully rational (Thaler, 2000), but evidence suggests that decision makers can implement various strategies to make the best decision possible with the cognitive resources that they have. In other words, human decision makers are resource-rational (Bhui et al., 2021; Lieder & Griffiths, 2020).

Critically, however, decision making is not just a cognitive process; it is also a metacognitive process (Ackerman & Thompson, 2017; Fiedler et al., 2019). In order to navigate computationally complex decisions, we must monitor the efficacy and efficiency of our decision-

making processes (Ackerman & Thompson, 2017; Alter & Oppenheimer, 2009) and strategically control our allocation of mental resources (Ackerman et al., 2023; Karpicke, 2009; D. G. Lee & Daunizeau, 2021). These metacognitive processes can be effortful and the resources necessary to engage in them are limited (Maniscalco et al., 2017; Maniscalco & Lau, 2015; Shekhar & Rahnev, 2021). As such, it is reasonable to hypothesize that decision makers would allocate their metacognitive resources in a resource-rational way when a decision is computationally complex.

We tested this hypothesis by evaluating how the accuracy of participants' metacognitive knowledge would be influenced by increasing the computational complexity of a decision task. We found that as the number of attributes in a value-based, multi-attribute choice task increased, the accuracy of participants' metacognitive knowledge declined. This suggests that metacognition is limited by computational constraints that prevent a decision maker from fully processing complex decisions. However, the more interesting finding is that this impairment of metacognitive knowledge was neither random nor uniform across attributes. Instead, as complexity increased, decision makers' metacognitive knowledge for low-importance attributes decreased substantially, while their metacognitive knowledge for high-importance attributes saw a smaller decline, or no decline at all. This suggests that decision makers could not maintain accurate metacognitive knowledge for all eight attributes and thus focused their metacognitive resources on the most important attributes.

Interestingly, however, participants in the 8-Attribute condition did not show a deficit in metacognitive accuracy at the behavioral (i.e., choice) level. Stated weights for participants in the 8-Attribute condition were as predictive of their choice behavior as the stated weights generated by participants in the 4-attribute conditions. These findings suggest that allocating limited metacognitive resources to the most important attributes may be an adaptive strategy

when faced with many attributes, as doing so allowed participants in the 8-Attribute condition to maintain behavioral accuracy, despite reductions in the relative and absolute accuracy of their metacognitive knowledge. In conjunction, this set of findings suggests that metacognitive processes can be resource rational (Bhui et al., 2021; Lieder & Griffiths, 2020; J. W. Payne et al., 1993). That is, decision makers efficiently allocate their metacognitive resources to the pieces of information that are most critical to the success of their decision-making process. Evaluating metacognition as a resource-rational process opens new avenues for advancing theory in both metacognition and decision making.

Limitations & Future Directions

In the present study, participants made hypothetical decisions. It is plausible that decision makers in ecologically valid contexts would be willing to expend far more metacognitive resources – particularly when the stakes are high. However, ecologically valid contexts also tend to be much more computationally complex than the choice tasks completed here. Real-life multi-attribute choice environments often require decision makers to evaluate dozens of alternatives, each of which vary on dozens of attributes (Hastings & Weinstein, 2008; Iyengar & Lepper, 2000). As such, it is likely that decision makers in real world domains still must make decisions about how to allocate their limited metacognitive resources, just on a larger scale. Future research should assess the accuracy of metacognitive knowledge in real-world decisions.

Another limitation of the present study is that our results do not allow us to identify the specific mechanism by which the accuracy of metacognitive knowledge for low-importance attributes was reduced in the 8-Attribute condition. On one hand, it is plausible that decision makers in the 8-Attribute condition may not have attended to the low-importance attributes in the first place and therefore may not have encoded metacognitive information about the weight they

placed on those attributes. This hypothesis aligns with findings from the memory literature suggesting that cognitive load during encoding hinders individuals' ability to attend to (Lavie, 2005) and selectively remember (Dubravac & Meier, 2023) task-relevant information. Decision makers may also have found the cognitively demanding nature of the 8-Attribute condition to be demotivating (Iyengar & Lepper, 2000) and applied heuristics to reduce the number of attributes they had to consider (Shah & Oppenheimer, 2008; Tversky & Kahneman, 1974).

On the other hand, it is also plausible that participants did attend to and encode the weights that they applied to the low-importance attributes, but that they lacked sufficient resources to retrieve the weights they placed on all eight attributes (Engle, 2001; Miller, 1956; Oberauer et al., 2016). Given that important and task-relevant information is more likely to be maintained in working memory (Caccamise & Kintsch, 1978; Murphy & Castel, 2021; Walter & Meier, 2014), this could also result in participants in the 8-Attribute condition preferentially recalling the weight they placed on high-importance attributes, while forgetting the weight they placed on low-importance attributes. This memory-based account is particularly difficult to isolate because metacognitive knowledge is, by definition, stored in memory (Flavell, 1979).

Furthermore, our manipulation of computational complexity – changing the number of attributes – may have simultaneously affected working memory load and attentional demands. This is largely unavoidable, as metacognition relies on these more basic cognitive abilities (Conte et al., 2023; Maniscalco et al., 2017; Maniscalco & Lau, 2015; Peng & Tullis, 2021). As such, we remain agnostic about the exact mechanism by which computational complexity increased the metacognitive demands of the task. Despite this, our results clearly suggest that increasing the number of attributes increased metacognitive demands, and that decision makers responded to these increased demands by focusing their metacognitive resources on the most

important attributes. Similarly, we cannot speak to the extent to which the observed resource rational effects were due to conscious strategic deliberation, or automatic responses to complex environments. We encourage further investigations into these mechanism-related questions.

Future researchers should also explore how decision makers regulate their metacognitive resource usage for other metacognitive skills, including metacognitive monitoring and metacognitive control (Ackerman & Thompson, 2017), for decisions in other domains, and for tasks outside of the decision-making context. It would also be of empirical interest to investigate how individual differences influence the implementation of resource-rational metacognitive strategies. For example, more intelligent individuals may be able to process more metacognitive information (Sheppard & Vernon, 2008) and individuals who are high in need for cognition may be willing to spend more time developing their metacognitive knowledge, rather than accepting a resource-rational approach (Cacioppo et al., 1996). Individuals who are more experienced in a domain may also be better able to generate accurate metacognitive knowledge using their limited metacognitive resources – though some evidence suggests that domain familiarity may not improve the accuracy of metacognitive knowledge (Cash & Oppenheimer, 2024).

Conclusion

Our results provide initial evidence that when decision makers are faced with computationally complex tasks, they allocate their limited metacognitive resources in a resource-rational way that maximizes the accuracy of their metacognitive judgments, given their cognitive and metacognitive constraints. This suggests that our metacognitive systems recognize when a decision is computationally complex and adjust accordingly. These findings open the door to new lines of research exploring how resource-rational metacognitive strategies affect our decision making and other cognitive processes.

Acknowledgements

This work was supported by a Doctoral Dissertation Research Improvement Grant from the National Science Foundation (Award #2333553), and an Academic Research Grant from Sawtooth Software, Inc. that provided us with free access to their Lighthouse Studio CBC software. No grantor had any role other than financial support. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and do not necessarily reflect the views of the grantors. We would also like to thank Elizabeth Cash, Nicole Oppenheimer, Sharon Carver, Julie Downs, and Bonnie Nozari for their support throughout the research process.

Funding: This work was supported by a Doctoral Dissertation Research Improvement Grant from the National Science Foundation (Award #2333553), and an Academic Research Grant from Sawtooth Software, Inc. that provided us with free access to their Lighthouse Studio CBC software. No grantor had any role other than financial support. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and do not necessarily reflect the views of the grantors.

Conflicts of Interest: The authors report that there are no competing interests to declare.

Data Availability Statement: Data, materials, and code for all studies presented in this manuscript are available at https://osf.io/z2ryg/?view_only=4571e45137dd404e8fc96e7de0f20360

References

Ackerman, R. (2014). The diminishing criterion model for metacognitive regulation of time investment. *Journal of Experimental Psychology: General*, *143*(3), 1349–1368. https://doi.org/10.1037/a0035098

Ackerman, R., & Beller, Y. (2017). Shared and distinct cue utilization for metacognitive judgements during reasoning and memorisation. *Thinking & Reasoning*, *23*(4), 376–408. https://doi.org/10.1080/13546783.2017.1328373

Ackerman, R., Binah-Pollak, A., & Lauterman, T. (2023). Metacognitive effort regulation across cultures. *Journal of Intelligence*, *11*(9), 171. https://doi.org/10.3390/jintelligence11090171

Ackerman, R., & Thompson, V. A. (2017). Meta-reasoning: Monitoring and control of thinking and reasoning. *Trends in Cognitive Sciences*, *21*(8), 607–617. https://doi.org/10.1016/j.tics.2017.05.004

Ahituv, N., Igbaria, M., & Sella, A. V. (1998). The effects of time pressure and completeness of information on decision making. *Journal of Management Information Systems*, *15*(2), 153–172. https://doi.org/10.1080/07421222.1998.11518212

Aitchison, J. (1982). The statistical analysis of compositional data. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, *44*(2), 139–160. https://doi.org/10.1111/j.2517-6161.1982.tb01195.x

Aitchison, J., Barceló-Vidal, C., Martín-Fernández, J. A., & Pawlowsky-Glahn, V. (2000). Logratio analysis and compositional distance. *Mathematical Geology, 32*, 271-275. https://doi.org/10.1023/A:1007529726302

Alter, A. L. (2013). The benefits of cognitive disfluency. *Current Directions in Psychological Science*, *22*(6), 437–442. https://doi.org/10.1177/0963721413498894

Alter, A. L., & Oppenheimer, D. M. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review*, *13*(3), 219–235. https://doi.org/10.1177/1088868309341564

Alter, A. L., Oppenheimer, D. M., & Epley, N. (2013). Disfluency prompts analytic thinking—But not always greater accuracy. *Cognition*, *128*(2), 252–255. https://doi.org/10.1016/j.cognition.2013.01.006

Basu, S., & Dixit, S. (2022). Role of metacognition in explaining decision-making styles: A study of knowledge about cognition and regulation of cognition. *Personality and Individual Differences*, *185*, 111318. https://doi.org/10.1016/j.paid.2021.111318

Betsch, T. (2008). The nature of intuition and its neglect in research on judgment and decision making. In H. Plessner, C. Betsch, & T. Betsch (Eds.), *Intuition in judgment and decision making* (pp. 3–22). Lawrence Erlbaum Associates Publishers.

Betsch, T., & Glöckner, A. (2010). Intuition in judgment and decision making: Extensive thinking without effort. *Psychological Inquiry*, *21*(4), 279–294. https://doi.org/10.1080/1047840X.2010.517737

Bhui, R., Lai, L., & Gershman, S. J. (2021). Resource-rational decision making. *Current Opinion in Behavioral Sciences*, *41*, 15–21. https://doi.org/10.1016/j.cobeha.2021.02.015

Boksem, M. A. S., & Tops, M. (2008). Mental fatigue: Costs and benefits. *Brain Research Reviews*, *59*(1), 125–139. https://doi.org/10.1016/j.brainresrev.2008.07.001

Brusovansky, M., Glickman, M., & Usher, M. (2018). Fast and effective: Intuitive processes in complex decisions. *Psychonomic Bulletin & Review*, *25*(4), 1542–1548. https://doi.org/10.3758/s13423-018-1474-1

Caccamise, D. J., & Kintsch, W. (1978). Recognition of important and unimportant statements from stories. *The American Journal of Psychology*, *91*(4), 651. https://doi.org/10.2307/1421513

Cacioppo, J. T., Petty, R. E., Feinstein, J. A., & Jarvis, W. B. G. (1996). Dispositional differences in cognitive motivation: The life and times of individuals varying in need for cognition. *Psychological Bulletin*, *119*(2), 197–253. https://doi.org/10.1037/0033-2909.119.2.197

Camerer, C. (1998). Bounded rationality in individual decision making. *Experimental Economics*, *1*(2), 163–183. https://doi.org/10.1023/A:1009944326196

Cash, T. N., & Oppenheimer, D. M. (2024). Parental rights or parental wrongs: Parents' metacognitive knowledge of the factors that influence their school choice decisions. *PLOS ONE*, *19*(4), e0301768. https://doi.org/10.1371/journal.pone.0301768

Cash, T. N., & Oppenheimer, D. M. (2025). Assessing metacognitive knowledge in subjective decisions: The *knowledge of weights* paradigm. *Thinking & Reasoning*, *31*(3), 331–373. https://doi.org/10.1080/13546783.2024.2426543

Colombo, B., Iannello, P., & Antonietti, A. (2010). Metacognitive knowledge of decision-making: An explorative study. In A. Efklides & P. Misailidi (Eds.), *Trends and Prospects in Metacognition Research* (pp. 445–472). Springer US. https://doi.org/10.1007/978-1-4419-6546-2_20

Conte, N., Fairfield, B., Padulo, C., & Pelegrina, S. (2023). Metacognition in working memory: Confidence judgments during an n-back task. *Consciousness and Cognition, 111*, Article 103522. https://doi.org/10.1016/j.concog.2023.103522

De Bruin, A. B. H., Dunlosky, J., & Cavalcanti, R. B. (2017). Monitoring and regulation of learning in medical education: The need for predictive cues. *Medical Education*, *51*(6), 575–584. https://doi.org/10.1111/medu.13267

De Neys, W. (2012). Bias and conflict: A case for logical intuitions. *Perspectives on Psychological Science*, *7*(1), 28–38. https://doi.org/10.1177/1745691611429354

De Neys, W. (2023). Advancing theorizing about fast-and-slow thinking. *Behavioral and Brain Sciences*, *46*, e111. https://doi.org/10.1017/S0140525X2200142X

Dubravac, M., & Meier, B. (2023). Cognitive load at encoding hurts memory selectivity. *Quarterly Journal of Experimental Psychology*, *76*(7), 1515–1538. https://doi.org/10.1177/17470218221132846

Dunlosky, J., & Metcalfe, J. (2009). *Metacognition*. SAGE.

Eggers, F., Sattler, H., Teichert, T., & Völckner, F. (2022). Choice-based conjoint analysis. In C. Homburg, M. Klarmann, & A. Vomberg (Eds.), *Handbook of Market Research* (pp. 781–819). Springer International Publishing. https://doi.org/10.1007/978-3-319-57413-4_23

Engle, R. W. (2001). What is working memory capacity? In H. L. Roediger, J. S. Nairne, I. Neath, & A. M. Surprenant (Eds.), *The nature of remembering: Essays in honor of Robert G. Crowder* (pp. 297–314). American Psychological Association. https://doi.org/10.1037/10394-016

Evans, J. S. B. T. (2010). Intuition and reasoning: A dual-process perspective. *Psychological Inquiry*, *21*(4), 313–326. https://doi.org/10.1080/1047840X.2010.521057

Evans, J. S. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, *8*(3), 223–241. https://doi.org/10.1177/1745691612460685

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*(2), 175–191. https://doi.org/10.3758/BF03193146

Fiedler, K., Ackerman, R., & Scarampi, C. (2019). Metacognition: Monitoring and controlling one's own knowledge, reasoning and decisions. In R. J. Sternberg & J. Funke (Eds.), *The Psychology of Human Thought* (pp. 89–111). Heidelberg University Publishing. https://doi.org/10.17885/heiup.470

Fiske, S. T., & Taylor, S. E. (1991). *Social cognition* (2nd ed.). McGraw-Hill.

Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive–developmental inquiry. *American Psychologist*, *34*(10), 906–911. https://doi.org/10.1037/0003-066X.34.10.906

Fleming, S. M., & Lau, H. C. (2014). How to measure metacognition. *Frontiers in Human Neuroscience*, *8*. https://doi.org/10.3389/fnhum.2014.00443

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, *19*(4), 25–42. https://doi.org/10.1257/089533005775196732

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*(4), 650–669. https://doi.org/10.1037/0033-295X.103.4.650

Gilovich, T., Griffin, D., & Kahneman, D. (Eds.). (2002). *Heuristics and biases: The psychology of intuitive judgment*. Cambridge Univ. Press.

Glöckner, A., & Betsch, T. (2008). Multiple-reason decision making based on automatic processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(5), 1055–1075. https://doi.org/10.1037/0278-7393.34.5.1055

Halverson, C. (2020). *Sample size rule of thumb for choice-based conjoint (CBC)*. Sawtooth. https://sawtoothsoftware.com/resources/blog/posts/sample-size-rules-of-thumb

Haselton, M. G., Nettle, D., & Andrews, P. W. (2015). The evolution of cognitive bias. In D. M. Buss (Ed.), *The Handbook of Evolutionary Psychology* (1st ed., pp. 724–746). Wiley. https://doi.org/10.1002/9780470939376.ch25

Hastings, J. S., & Weinstein, J. M. (2008). Information, school choice, and academic achievement: Evidence from two experiments. *Quarterly Journal of Economics*, *123*(4), 1373–1414. https://doi.org/10.1162/qjec.2008.123.4.1373

Hein, M., Kurz, P., & Steiner, W. J. (2020). Analyzing the capabilities of the HB logit model for choice-based conjoint analysis: A simulation study. *Journal of Business Economics*, *90*(1), 1–36. https://doi.org/10.1007/s11573-019-00927-4

Iyengar, S. S., & Lepper, M. R. (2000). When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology*, *79*(6), 995–1006. https://doi.org/10.1037/0022-3514.79.6.995

Kahan, D. M. (2013). Ideology, motivated reasoning, and cognitive reflection. *Judgment and Decision Making*, *8*(4), 407–424. https://doi.org/10.1017/S1930297500005271

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

Karpicke, J. D. (2009). Metacognitive control and strategy selection: Deciding to practice retrieval during learning. *Journal of Experimental Psychology: General*, *138*(4), 469–486. https://doi.org/10.1037/a0017341

Klein, G. (2015). A naturalistic decision making perspective on studying intuitive decision making. *Journal of Applied Research in Memory and Cognition*, *4*(3), 164–168. https://doi.org/10.1016/j.jarmac.2015.07.001

Koriat, A. (1993). How do we know that we know? The accessibility model of the feeling of knowing. *Psychological Review*, *100*(4), 609-639. https://doi.org/10.1037/0033-295X.100.4.609

Koriat, A. (1997). Monitoring one's own knowledge during study: A cue-utilization approach to judgments of learning. *Journal of Experimental Psychology: General*, *126*(4), 349–370. https://doi.org/10.1037/0096-3445.126.4.349

Lauterman, T., & Ackerman, R. (2024). Initial judgment of solvability: Integrating prior expectations with experience-based heuristic cues. *Thinking & Reasoning*, *30*(1), 135–168. https://doi.org/10.1080/13546783.2023.2214378

Lavie, N. (2005). Distracted and confused?: Selective attention under load. *Trends in Cognitive Sciences*, *9*(2), 75–82. https://doi.org/10.1016/j.tics.2004.12.004

Law, M. K. H., Stankov, L., & Kleitman, S. (2022). I choose to opt-out of answering: Individual differences in giving up behaviour on cognitive tests. *Journal of Intelligence*, *10*(4), 86. https://doi.org/10.3390/jintelligence10040086

Lee, D. G., & Daunizeau, J. (2021). Trading mental effort for confidence in the metacognitive control of value-based decision-making. *eLife*, *10*, e63282. https://doi.org/10.7554/eLife.63282

Lieder, F., Callaway, F., & Griffiths, T. L. (2026). *The rational use of cognitive resources*. Princeton University Press. https://doi.org/10.1515/9780691259932

Lieder, F., & Griffiths, T. L. (2017). Strategy selection as rational metareasoning. *Psychological Review*, *124*(6), 762–794. https://doi.org/10.1037/rev0000075.supp

Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*, e1. https://doi.org/10.1017/S0140525X1900061X

Maniscalco, B., & Lau, H. (2015). Manipulation of working memory contents selectively impairs metacognitive sensitivity in a concurrent visual discrimination task. *Neuroscience of Consciousness*, *2015*(1). https://doi.org/10.1093/nc/niv002

Maniscalco, B., McCurdy, L. Y., Odegaard, B., & Lau, H. (2017). Limited cognitive resources explain a trade-off between perceptual and metacognitive vigilance. *The Journal of Neuroscience*, *37*(5), 1213–1224. https://doi.org/10.1523/JNEUROSCI.2271-13.2016

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63*(2), 81–97. https://doi.org/10.1037/h0043158

Moore, D. A., & Schatz, D. (2017). The three faces of overconfidence. *Social and Personality Psychology Compass*, *11*(8), e12331. https://doi.org/10.1111/spc3.12331

Morey, R. D., & Rouder, J. N. (2026). *BayesFactor: Computation of Bayes Factors for Common Designs* (Version 0.9.12-2) [Computer software]. https://github.com/richarddmorey/bayesfactor

Morris, A., Carlson, R. W., Kober, H., & Crockett, M. (2025). Introspective access to value-based multi-attribute choice processes. *Nature Communications, 16*, 3733. https://doi.org/10.1038/s41467-025-59080-y

Murphy, D. H., & Castel, A. D. (2021). Responsible remembering and forgetting as contributors to memory for important information. *Memory & Cognition*, *49*(5), 895–911. https://doi.org/10.3758/s13421-021-01139-4

Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. *Psychology of Learning and Motivation*, *26*, 125–173. https://doi.org/10.1016/S0079-7421(08)60053-5

Norman, E. (2020). Why metacognition is not always helpful. *Frontiers in Psychology*, *11*, 1537. https://doi.org/10.3389/fpsyg.2020.01537

Oberauer, K., Farrell, S., Jarrold, C., & Lewandowsky, S. (2016). What limits working memory capacity? *Psychological Bulletin*, *142*(7), 758–799. https://doi.org/10.1037/bul0000046

Oppenheimer, D. M. (2004). Spontaneous discounting of availability in frequency judgment tasks. *Psychological Science*, *15*(2), 100–105. https://doi.org/10.1111/j.0963-7214.2004.01502005.x

Oppenheimer, D. M., & Monin, B. (2009). Investigations in spontaneous discounting. *Memory & Cognition*, *37*(5), 608–614. https://doi.org/10.3758/MC.37.5.608

Orme, B. (2002). Interpreting conjoint analysis data. *Sawtooth Software, Inc.* https://leeds-faculty.colorado.edu/ysun/MKTG4825_files/Interpreting_CA_Data.pdf

Palmer, E. C., David, A. S., & Fleming, S. M. (2014). Effects of age on metacognitive efficiency. *Consciousness and Cognition*, *28*, 151–160. https://doi.org/10.1016/j.concog.2014.06.007

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The Adaptive Decision Maker* (1st ed.). Cambridge University Press. https://doi.org/10.1017/CBO9781139173933

Payne, S. J., & Duggan, G. B. (2011). Giving up problem solving. *Memory and Cognition*, *39*(5), 902–913. https://doi.org/10.3758/s13421-010-0068-6

Peng, Y., & Tullis, J. G. (2021). Dividing attention impairs metacognitive control more than monitoring. *Psychonomic Bulletin & Review, 28*(6), 2064–2074. https://doi.org/10.3758/s13423-021-01950-9

Phillips-Wren, G., & Adya, M. (2020). Decision making under stress: The role of information overload, time pressure, complexity, and uncertainty. *Journal of Decision Systems*, *29*(sup1), 213–225. https://doi.org/10.1080/12460125.2020.1768680

Sadler-Smith, E., & Shefy, E. (2004). The intuitive executive: Understanding and applying 'gut feel' in decision-making. *Academy of Management Perspectives*, *18*(4), 76–91. https://doi.org/10.5465/ame.2004.15268692

Salas, E., Rosen, M. A., & DiazGranados, D. (2010). Expertise-based intuition and decision making in organizations. *Journal of Management*, *36*(4), 941–973. https://doi.org/10.1177/0149206309350084

Sawtooth Software, Inc. (2021). *CBC/HB system technical paper V5.6: The CBC/HB system for Hierarchical Bayes Estimation*. https://sawtoothsoftware.com/resources/technical-papers/cbc-hb-technical-paper

Sawtooth Software, Inc. (2024). *Lighthouse Studio* (Version 9.14.2) [Computer software]. https://sawtoothsoftware.com/lighthouse-studio

Scheiter, K., Ackerman, R., & Hoogerheide, V. (2020). Looking at mental effort appraisals through a metacognitive lens: Are they biased? *Educational Psychology Review*, *32*(4), 1003–1027. https://doi.org/10.1007/s10648-020-09555-9

Schwartz, B., Ward, A., Monterosso, J., Lyubomirsky, S., White, K., & Lehman, D. R. (2002). Maximizing versus satisficing: Happiness is a matter of choice. *Journal of Personality and Social Psychology*, *83*(5), 1178–1197. https://doi.org/10.1037/0022-3514.83.5.1178

Schwarz, N. (1998). Accessible content and accessibility experiences: The interplay of declarative and experiential information in judgment. *Personality and Social Psychology Review*, *2*(2), 87–99. https://doi.org/10.1207/s15327957pspr0202_2

Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauer-Schatka, H., & Simons, A. (1991). Ease of retrieval as information: Another look at the availability heuristic. *Journal of Personality and Social Psychology*, *61*(2), 195–202. https://doi.org/10.1037/0022-3514.61.2.195

Schwarz, N., Jalbert, M., Noah, T., & Zhang, L. (2021). Metacognitive experiences as information: Processing fluency in consumer judgment and decision making. *Consumer Psychology Review*, *4*(1), 4–25. https://doi.org/10.1002/arcp.1067

Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: An effort-reduction framework. *Psychological Bulletin*, *134*(2), 207–222. https://doi.org/10.1037/0033-2909.134.2.207

Shekhar, M., & Rahnev, D. (2021). Sources of metacognitive inefficiency. *Trends in Cognitive Sciences*, *25*(1), 12–23. https://doi.org/10.1016/j.tics.2020.10.007

Shenhav, A., Prater Fahey, M., & Grahek, I. (2021). Decomposing the motivation to exert mental effort. *Current Directions in Psychological Science*, *30*(4), 307–314. https://doi.org/10.1177/09637214211009510

Sheppard, L. D., & Vernon, P. A. (2008). Intelligence and speed of information-processing: A review of 50 years of research. *Personality and Individual Differences*, *44*(3), 535–551. https://doi.org/10.1016/j.paid.2007.09.015

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, *69*(1), 99. https://doi.org/10.2307/1884852

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*(2), 129–138. https://doi.org/10.1037/h0042769

Smith, P. L., & Ratcliff, R. (2009). An integrated theory of attention and decision making in visual signal detection. *Psychological Review*, *116*(2), 283–317. https://doi.org/10.1037/a0015156

Smithson, M., & Broomell, S. B. (2024). Compositional data analysis tutorial. *Psychological Methods*, *29*(2), 362–378. https://doi.org/10.1037/met0000464

Son, L. K., & Schwartz, B. L. (2002). The relation between metacognitive monitoring and control. In T. J. Perfect & B. L. Schwartz (Eds.), *Applied Metacognition* (1st ed., pp. 15–38). Cambridge University Press. https://doi.org/10.1017/CBO9780511489976.003

Stanovich, K. E. (2018). Miserliness in human cognition: The interaction of detection, override and mindware. *Thinking & Reasoning*, *24*(4), 423–444. https://doi.org/10.1080/13546783.2018.1459314

Stanovich, K. E. (2020). Why humans are cognitive misers and what it means for the Great Rationality Debate. In R. Viale (Ed.), *Routledge Handbook of Bounded Rationality* (1st ed., pp. 196–206). Routledge. https://doi.org/10.4324/9781315658353-12

Swait, J. (2001). A non-compensatory choice model incorporating attribute cutoffs. *Transportation Research Part B: Methodological*, *35*(10), 903–928. https://doi.org/10.1016/S0191-2615(00)00030-8

Thaler, R. H. (2000). From Homo Economicus to Homo Sapiens. *Journal of Economic Perspectives*, *14*(1), 133–141. https://doi.org/10.1257/jep.14.1.133

Thompson, V. A., Turner, J. A. P., Pennycook, G., Ball, L. J., Brack, H., Ophir, Y., & Ackerman, R. (2013). The role of answer fluency and perceptual fluency as metacognitive cues for initiating analytic thinking. *Cognition*, *128*(2), 237–251. https://doi.org/10.1016/j.cognition.2012.09.012

Timmermans, D. (1993). The impact of task complexity on information use in multi-attribute decision making. *Journal of Behavioral Decision Making*, *6*(2), 95–111. https://doi.org/10.1002/bdm.3960060203

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131. https://doi.org/10.1126/science.185.4157.1124

Undorf, M., & Bröder, A. (2020). Cue integration in metamemory judgements is strategic. *Quarterly Journal of Experimental Psychology*, *73*(4), 629–642. https://doi.org/10.1177/1747021819882308

Walter, S., & Meier, B. (2014). How important is importance for prospective memory? A review. *Frontiers in Psychology*, *5*. https://doi.org/10.3389/fpsyg.2014.00657

Whittlesea, B. W. A., & Williams, L. D. (1998). Why do strangers feel familiar, but friends don't? A discrepancy-attribution account of feelings of familiarity. *Acta Psychologica*, *98*(2–3), 141–165. https://doi.org/10.1016/S0001-6918(97)00040-1

Xue, K., Shekhar, M., & Rahnev, D. (2021). Examining the robustness of the relationship between metacognitive efficiency and metacognitive bias. *Consciousness and Cognition*, *95*, 103196. https://doi.org/10.1016/j.concog.2021.103196