



Delft University of Technology

Citation (APA)

Sun, G., Kroesen, M., & Rezaei, J. (2025). Anchoring Bias in Value Function Elicitation Within Multiattribute Value Theory. *Decision Analysis*, 22(4), 284-304. <https://doi.org/10.1287/deca.2024.0308>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.

Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

**Green Open Access added to [TU Delft Institutional Repository](#)
as part of the Taverne amendment.**

More information about this copyright law amendment
can be found at <https://www.openaccess.nl>.




Otherwise as indicated in the copyright section:
the publisher is the copyright holder of this work and the
author uses the Dutch legislation to make this work public.

Anchoring Bias in Value Function Elicitation Within Multiattribute Value Theory

Geqie Sun,^{a,*} Maarten Kroesen,^a Jafar Rezaei^a

^a Faculty of Technology, Policy and Management, Delft University of Technology, 2628 BX Delft, Netherlands

*Corresponding author

Contact: g.sun-3@tudelft.nl,  <https://orcid.org/0000-0001-9847-4762> (GS); m.kroesen@tudelft.nl,  <https://orcid.org/0000-0001-6623-9848> (MK); j.rezaei@tudelft.nl,  <https://orcid.org/0000-0002-7407-9255> (JR)

Received: December 10, 2024

Revised: April 4, 2025; May 15, 2025

Accepted: June 19, 2025

Published Online in Articles in Advance:
July 31, 2025

<https://doi.org/10.1287/deca.2024.0308>

Copyright: © 2025 INFORMS

Abstract. Anchoring bias refers to the human tendency to rely heavily on an initial piece of information when making judgments. This bias has significant implications for decision analysis methods that rely on human judgments. This study examines the influence of anchoring bias in the value function elicitation step of the multiattribute value theory, specifically within the midvalue splitting procedure. We hypothesize that the starting point provided by the analyst during elicitation creates a bias in decision makers' judgments, leading to distorted value functions and ultimately affecting decision outcomes. We also hypothesize that counter-anchoring and avoiding the use of anchors mitigate the effect of anchoring bias. To test the hypotheses, we designed an experiment and collected data from 320 subjects. The findings show that the starting point in the midvalue splitting procedure could change the attribute-specific value functions and, consequently, the overall value of the alternatives. Additionally, two debiasing strategies, counter-anchoring and avoiding the use of anchors, were found to be effective in reducing the effect of anchoring bias. The implications of this study can extend to other structured value function elicitation methods.

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/deca.2024.0308>.

Keywords: anchoring bias • value function • midvalue splitting procedure • multiattribute value theory • debiasing

1. Introduction

Cognitive bias refers to the systematic error in judgment that occurs when individuals process and interpret information (Tversky and Kahneman 1974). It is a well-documented phenomenon in psychology and emerges from the human tendency to use heuristics, or mental shortcuts, to simplify decision making in complex environments. Although such shortcuts can be efficient in reducing cognitive load, they could lead to flawed judgments (Kahneman 2011). Over the past decades, a growing body of research has revealed the pervasive influence of cognitive bias across various fields (Ariely 2008, Thaler and Sunstein 2008, Bazerman and Moore 2012, Bushong and Gagnon-Bartsch 2024). In financial contexts, for instance, overconfidence bias can lead investors to overestimate their knowledge or

abilities to predict outcomes, causing them to make ill-informed investment choices and suffer significant financial losses (Barber and Odean 2001). In managerial decision making, the framing effect, where people respond differently to the same situation depending on how they are framed, can lead to inconsistent results. For instance, managers may opt for riskier choices when the same scenario is framed as avoiding losses rather than achieving gains (Tversky and Kahneman 1981).

Given these widespread implications, cognitive bias has become a central concern in behavioral decision research (BDR) that is divided into two major branches: descriptive BDR and prescriptive BDR (Montibeller and von Winterfeldt 2024). Descriptive BDR aims at developing theories or models to describe the biases in

judgments, whereas prescriptive BDR focuses on developing debiasing strategies to correct those biases, improving the quality of decisions. Notably, different from the descriptive view above, cognitive bias in prescriptive research is defined as “a systematic discrepancy between the ‘correct’ answer in a judgmental task, given by a formal normative rule, and the decision-maker’s or expert’s actual answer to such a task” (Montibeller and von Winterfeldt 2015, p. 1231). This definition highlights the role of the normative rule, and, by extension, the role of methodological design, in prescriptive BDR.

In descriptive BDR, cognitive bias is generally assessed by measuring the deviation of human judgment from either the true value or the coherence of preferences in an unsupported decision-making environment (Montibeller and von Winterfeldt 2024). Conversely, the prescriptive view introduces a medium, the decision analysis method, an aspect considered peripheral or noncentral in the descriptive view. Within the decision analysis field, a variety of methods have been developed to support decision makers (DMs) in structuring and quantifying their preferences, in order to evaluate and choose among alternatives characterized by multiple, often conflicting, attributes (Keeney and Raiffa 1976, Belton and Stewart 2012). The procedures of these methods typically involve a series of structured questions posed by an analyst to elicit the DM’s preferences. These preferences are inferred based on the value judgments of DMs. The underlying assumption of these methods is that DMs hold coherent and stable preferences unaffected by irrelevant information during these judgmental tasks (Fischhoff 1991) and that these methods can reduce human biases in decision making (Keeney 2004). However, this assumption has been extensively challenged by evidence of cognitive bias in decision analysis methods in prescriptive BDR.

The methodological design of the elicitation procedures plays an important role in prescriptive BDR. This is primarily because preferences are not merely revealed but constructed during the elicitation process (Slovic 1995). Empirical research has shown that preferences can be shaped by contextual factors and the elicitation task itself (Payne et al. 1992). This constructive view assigns critical importance to the elicitation method: when thoughtfully and carefully devised, the method can reduce potential biases during this process; however, if the methodological design neglects cognitive

phenomena, it may inadvertently exacerbate existing biases. So, the goal of prescriptive BDR is not only to identify these deviations, but also to develop procedures to mitigate these biases and ensure that decision analysis methods remain effective in real-world contexts.

This study investigates how anchoring bias, one of the most important cognitive biases (Tversky and Kahneman 1974, Kahneman 2011, Montibeller and von Winterfeldt 2015), influences the multiattribute value theory (MAVT), particularly the value function elicitation step. MAVT is a well-known and widely applied decision analysis method developed by Keeney and Raiffa (1976). The value function elicitation step involves determining attribute-specific value functions, which represent the DM’s value judgments over different attribute levels. This step not only allows us to directly observe how cognitive bias may distort the representation of preference, but it also provides a quantifiable structure through which we can assess the impact of bias on the value function.

There are different ways to elicit a value function, including the standard difference procedure, the lock-step procedure, direct rating, successive comparison, and curve fitting, among others (Fishburn 1967, Keeney and Raiffa 1976, Watson and Buede 1987, Beinat 1997). One of the most commonly used, and the focus of this study, is the midvalue splitting procedure. In this procedure, multiple midvalue points will be identified by the DM, which are then used to form the value function. To begin the process, the analyst provides a starting point for the DM to identify whether it is the first midvalue point. However, because of the effect of anchoring bias, the DM might adjust insufficiently away from the starting point, which results in a biased midvalue point. This will, in turn, affect the value function and, ultimately, the final outcome of the decision-making problem. Through this investigation, we aim to understand how anchoring bias impacts the value function shape and the final decision outcomes, as well as to identify effective strategies to mitigate its effects. To achieve this, we develop a few hypotheses and design an experiment to test them. Data were collected from 320 participants, who completed the MAVT procedure under varying anchoring conditions during the value function elicitation step. To assess the effects of anchoring bias on value function elicitation and the

overall value of the alternatives, we conducted both parametric and nonparametric statistical analyses.

This study makes two main contributions to the field of decision analysis. First, it provides a detailed analysis of how anchoring bias affects the value function and, consequently, the decision results of MAVT. It enhances our understanding of anchoring bias in value assessment, offering insights that can extend to other value-based decision methods. Second, it develops effective debiasing strategies to mitigate the anchoring effect. These strategies aim to improve the consistency and reliability of value function elicitation procedure, thereby enhancing the credibility of MAVT and related decision analysis methods in real-world applications. Although our focus is on the mid-value splitting procedure in the value function elicitation step, the implications can be extended to other value function elicitation methods, which will be discussed in Section 6.4.

The remainder of this paper is organized as follows: Section 2 introduces anchoring bias and its implications in decision analysis. Section 3 describes the MAVT, with a particular emphasis on the midvalue splitting procedure. Section 4 presents the research hypotheses. Section 5 outlines the experimental design. Section 6 discusses the results of the study. Finally, Section 7 concludes the paper and suggests directions for future research.

2. Anchoring Bias

Anchoring bias refers to the human tendency to rely heavily on an initial piece of information and to insufficiently adjust away from it when making judgments; see Furnham and Boo (2011) for a review. This initial information, known as the “anchor,” can be provided externally or generated internally (Epley and Gilovich 2005). External anchors are reference points originating outside the DM’s judgment, such as information provided by other people, text, pictures, or other contextual stimuli. Internal anchors, on the other hand, are based on a DM’s prior beliefs, experiences, or memories. For instance, if someone is asked to estimate a product price, an external anchor could be a price tag, whereas an internal anchor might be her memory of the price of a similar product.

Empirical studies of anchoring effects typically use one of two paradigms: estimation tasks, where anchoring

distorts accuracy relative to true value; and valuation tasks, where it affects the internal coherence of preferences in the absence of a known correct answer. The foundational study by Tversky and Kahneman (1974) illustrates anchoring in an estimation task. Two groups of participants are asked to estimate (in five seconds without a calculator) the product of a sequence of numbers presented in reverse order ($8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1$ versus $1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8$). Despite both sequences containing the same numbers, the group starting with a higher initial partial product (i.e., the product of the first few numbers) gave a significantly higher median estimate than the group starting with a lower initial product (2,250 versus 512, respectively). Even though both groups produced estimates much lower than the true value (40,320), the difference in their estimates (due to the initial anchor) demonstrates how anchoring bias distorts individual assessments.

Beyond this seminal work, anchoring effects have been documented across a wide range of applied domains. In estimation tasks, Prava et al. (2016) explore how different anchors, such as an ignorance prior, can influence the probability assessments in surveys. Anchoring bias occurs when participants do not adjust sufficiently from such anchors, leading to partition dependence and carryover biases, distorting the probability assessments. In valuation tasks, anchoring bias is equally pervasive. In negotiations, an initial offer can serve as a powerful anchor, influencing the final agreement terms, even if the offer is arbitrary or exaggerated (Galinsky and Mussweiler 2001). In marketing, the first price that consumers encounter, such as an initial product price or a reference price, can anchor their perceptions of value. This initial price can influence how consumers perceive the value of any subsequent prices, such as a discounted price, ultimately affecting their willingness to pay (Wansink et al. 1998). In healthcare, anchoring bias can lead to admission control bias and path-dependent feedback, where initial diagnostic uncertainty disproportionately influences subsequent decisions and resource allocation, potentially impacting patient outcomes (Kim and Tong 2024). Anchoring also plays a critical role in legal settings, where initial sentencing recommendations—whether reasonable or not—can influence the final sentencing decisions of judges (Englich et al. 2006).

Within decision analysis, studies of anchoring bias have primarily focused on valuation tasks. Research has shown that anchoring bias significantly affects weight elicitation—the process by which the relative importance of the attributes (or scaling constants) is elicited. Buchanan and Corner (1997) explored the impact of anchoring bias in two interactive decision analysis methods. In the Zoints and Wallenius method, the DM is usually provided with a fixed starting point as part of the method's initialization, whereas in the free search interactive method, there is no fixed starting point. Instead, it allows the DM to explore the feasible region freely. The experiment results show that the starting solution in the Zoints and Wallenius method had a significant impact on the decision outcome, demonstrating the anchoring effect. In contrast, the anchoring effect was not significant in the free search method. An insight gain from this study is that designing decision analysis methods with flexible starting points and iterative feedback can help reduce the impact of anchoring.

Jacobi and Hobbs (2007) focused on value-tree-induced biases in weight elicitation methods, particularly the splitting bias. Because of splitting bias, decomposing an attribute into multiple subattributes leads to a higher global weight of that attribute compared with directly assessing its relative importance. They suggest that the use of anchor-and-adjustment heuristics is the main cause of such bias. DMs initially allocate equal weights across attributes within each tree partition. The equal allocation serves as an anchor for their judgments. DMs then insufficiently adjust these weights to align with their preferences, resulting in weights being more uniform than they should be.

Collectively, these findings demonstrate the pervasive influence of anchoring bias and highlight the necessity of prescriptive interventions to mitigate its impact. In the context of decision analysis, this points to a critical, yet often overlooked, aspect: although these methods are mathematically sound and intended to guide rational decision making, their structural frameworks can unintentionally induce cognitive bias, such as anchoring bias, by default. Therefore, we should re-evaluate how these methods are designed and implemented and, when necessary, improve them to address cognitive aspects, ensuring their effectiveness.

Generic approaches for reducing cognitive bias, such as consulting with experts, providing regular feedback,

and increasing awareness of biases, have been shown to be somewhat effective in reducing anchoring bias (Gavirneni and Xia 2009, Lilienfeld et al. 2009, Adame 2016, Fasolo et al. 2024). For prescriptive BDR, we need more targeted strategies that focus on the heuristics underlying anchoring bias. This study incorporates two such targeted approaches:

Avoiding anchors: A straightforward, yet effective, strategy is to design decision-making processes that avoid reliance on predefined starting points (Montibeller and von Winterfeldt 2015). Practical implementations include randomizing the presentation order of attributes, structuring decision tasks in a way to prevent the introduction of starting points, or encouraging DMs to independently generate their assessments without external prompts. For example, the free search interactive method, which removes fixed starting points, enables DMs to explore alternatives freely, without being influenced by anchoring bias. However, this strategy has been primarily designed for the external anchors. Internal anchors, which are inherently subjective and often viewed as more credible by DMs (Wilson et al. 1996, Mussweiler and Strack 1999, Chapman and Johnson 2002), require different interventions. For instance, approaches such as monetary rewards for being correct and forewarning individuals about possible biases in judgment have been shown to be effective (Epley and Gilovich 2005).

Consider-the-opposite strategy: This strategy encourages DMs to critically evaluate their initial judgments by considering contradictory information or alternative scenarios (Lord et al. 1984, Mussweiler et al. 2000). In the decision analysis context, counter-anchors are a practical way to operationalize this strategy. For instance, whereas traditional methods such as Simple Multi-Attribute Rating Technique (SMART) and Swing rely on unidirectional anchors, the Best-Worst Method (BWM) (Rezaei 2015) incorporates a bidirectional evaluation process. The BWM requires pairwise comparisons of the best attribute against others and the others against the worst attribute, which helps to balance the evaluation by reducing the influence of a single anchor (Rezaei et al. 2024). However, the use of counter-anchors in decision analysis in practice is not without its challenges. First, if counter-anchor values are too implausible, DMs may recognize the potential for bias and overcorrect their judgments (Brewer et al. 2007). Second, the repeated use of counter-anchors can complicate the

procedure and increase cognitive load. This may lead to mental fatigue, reduced engagement, and potentially less reliable responses from the DM (Kahneman 2011).

Building on the two strategies, avoiding anchors and employing counter-anchors were designed specifically for mitigating anchoring bias in the midvalue splitting procedure. Detailed implementations will be introduced in Section 5.

3. Multiattribute Value Theory

Multiattribute value theory is a well-established decision analysis method developed by Keeney and Raiffa (1976). This theory is grounded in the principles of utility theory, where the objective is to represent DMs' preferences through a mathematical function that aggregates the values or utilities of different attributes into a single score. MAVT not only aids in the systematic evaluation of the alternatives, but also enhances transparency in the decision-making process by clearly articulating the rationale behind the choices made (Anthes et al. 2019, Höfer and Madlener 2020). MAVT has been widely applied and shown great success in various decision fields, such as environmental management (Hostmann et al. 2005), healthcare (Labijak-Kowalska et al. 2023), and policy making (Ferretti 2016). MAVT consists of the following five steps (Keeney and Raiffa 1976; Keeney 2002, 2009):

Identification of the objectives, attributes, and alternatives: First, the problem must be structured, which includes (i) the objectives; (ii) attributes the evaluators used to evaluate how well alternatives meet the objectives; and (iii) alternatives, the possible solutions available for the decision problem. The structuring process often uses tools like value trees to organize these elements hierarchically (Belton and Stewart 2012).

Value function elicitation: Attribute-specific value functions should be developed. An attribute-specific value function translates the performance of an alternative on a specific attribute into a value score, typically on a scale from zero to one, where zero is assigned to the least preferred level of the attribute and one is assigned to the most preferred level.

Weight elicitation: This step involves eliciting scaling constants for each attribute that determine their influence in the overall value function. The Tradeoff procedure is commonly used in this step, where the DM is asked to compare two attributes at a time and

determine how much of one attribute they would sacrifice to gain more of another.

Aggregation: Each alternative's aggregated score is calculated using an additive value function:

$$v(\mathbf{a}_i) = \sum_{j=1}^n w_j v_j(a_{ij}), \quad (1)$$

where $v(\mathbf{a}_i)$ is the overall value of alternative i , scaled from zero to one. $v_j(a_{ij})$ is the attribute-specific value representing the performance of alternative i with respect to attribute j , and w_j is the scaling constant (or weight) of the attribute j . To use an additive value function, the two primary conditions must be satisfied (mutual preference independence and difference independence) (Keeney and Raiffa 1976, Dyer and Sarin 1979).

Definition 1. The attributes X_1, \dots, X_n are mutually preferentially independent if any subset of attributes is preferentially independent of the remaining attributes.

Definition 2. The attribute X_j is difference-independent of the remaining attributes if the preference difference between two levels of X_j is not affected by the fixed levels on the other attributes.

If these conditions are not satisfied, other forms, such as a multiplicative aggregation model, can be used (Keeney 1974).

Ranking and selection: The alternatives are ranked based on their aggregated scores. Specifically, for any two alternatives \mathbf{a}_k and \mathbf{a}_l , the preference relation is expressed as follows:

$$\begin{cases} v(\mathbf{a}_k) > v(\mathbf{a}_l) \Leftrightarrow \mathbf{a}_k > \mathbf{a}_l, \mathbf{a}_k \text{ is strongly preferred to } \mathbf{a}_l \\ v(\mathbf{a}_k) = v(\mathbf{a}_l) \Leftrightarrow \mathbf{a}_k \sim \mathbf{a}_l, \mathbf{a}_k \text{ is indifferent to } \mathbf{a}_l \\ v(\mathbf{a}_k) \geq v(\mathbf{a}_l) \Leftrightarrow \mathbf{a}_k \succeq \mathbf{a}_l, \mathbf{a}_k \text{ is weakly preferred to } \mathbf{a}_l. \end{cases} \quad (2)$$

Among the five steps, this study focuses on the value function elicitation step and examines how anchoring bias can distort value function elicitation and, in turn, distort the evaluation of alternatives. Value functions transform the performance of alternatives into a standardized scale, allowing for consistent comparison across different attributes. In MAVT, both the value function and attribute weights influence the final decision. However, to isolate the impact of value function elicitation, this study fixes the attribute weights,

allowing for a more focused examination of how the elicitation process affects decision results. A detailed explanation of the experiment design used is provided in Section 5.

The midvalue splitting procedure, commonly used in the value function elicitation step, relies on defining a monotonic value function that represents the DM's preferences. When a higher attribute value is preferred, an increasing function is applied. First, a formal definition (Kirkwood and Sarin 1980):

Definition 3. x_1 is said to be the midvalue of the interval $[x_0, x_2]$ if the decision maker will give up the same amount of some other attribute to go from x_0 to x_1 as from x_1 to x_2 .

The procedure involves the following steps: (i) *Assigning initial values:* The analyst assigns a value score of zero to the lowest attribute level x_{lowest} and a value score of one to the highest attribute level x_{highest} . (ii) *Determining the first midvalue point:* The analyst provides an initial point x_1 between x_{lowest} and x_{highest} . The DM is then asked whether they perceive the change from x_{lowest} to x_1 to be equally desirable as the change from x_1 to x_{highest} .¹ If the DM indicates indifference between these changes, x_1 is considered the first midvalue point and is assigned a value score of 0.5. If the DM is not indifferent between the changes, the analyst will propose an alternative point x_1 , and this process will be repeated until the point of indifference is found. (iii) *Determining subsequent midvalue points:* Using the same procedure, the DM defines the indifference point x_2 between x_{lowest} and x_1 . Once x_2 is identified, it is assigned a value score of 0.25. Similarly, the indifference point x_3 is defined between x_1 and x_{highest} , which is then assigned a value score of 0.75. (iv) *Consistency check:* We check if the DM is indifferent between the two changes from x_2 to x_1 and from x_1 to x_3 . If not, the DM will be asked to revise x_1 . (v) *Plotting the value function:* After obtaining the midvalue points (x_1, x_2, x_3) , the value function can be plotted, providing a clear graphical representation of the DM's preferences across the attribute range. In cases when a lower attribute value is preferred, a decreasing monotonic function is used. The value score of one is assigned to the lowest attribute level x_{lowest} , and zero is assigned to the highest attribute level x_{highest} , ensuring that the function accurately represents inverse preferences. The remaining

steps are similar to the case of an increasing function. Notably, more midvalue points can be identified if we want to ensure a more detailed presentation of preference using the value function. It is also important to note that during this procedure, the levels of the other attributes are at a specified fixed level, and mutual preferential independence is given.

4. Hypotheses Development

The midvalue splitting procedure, as outlined in Section 3, begins with the analyst asking the DM whether she agrees with the midvalue point provided by the analyst. If this starting point is at the lower end of the attribute range, regardless of a monotonic increasing or decreasing value function, it is called a low anchor, whereas a starting point at the upper end is called a high anchor. When the DM is provided with a low anchor by the analyst, the insufficient adjustment from the anchor by the DM could lead to the first midvalue point being smaller than that of when a high anchor is given ($m_L < m_H$). The analyst's selection of the starting point can be explained by his own judgment of the shape of the value function, a random choice, or other reasons. It is clear that using a low versus a high starting point can affect determining the midvalue point. Similarly, when the analyst uses the midperformance point $((x_{\text{lowest}} + x_{\text{highest}})/2)$ as the starting anchor, anchoring bias may still lead to distortions. For a linear value function, the midvalue point is exactly the midperformance point, but people might adjust away from the anchor (the midperformance point) and thus distort the elicited midvalue point. For nonlinear value functions, the midperformance point is acting as the low or high anchor, depending on the shape of the value function. Therefore, the same arguments for the low and high anchors apply here.

The initial deviation might not only affect the first midvalue point, but also is expected to spread to subsequent midvalue points, ultimately changing the shape of the attribute-specific value function. Although the consistency check in the midvalue splitting procedure is intended to ensure consistent preferences, the systematic deviation across all midvalue points undermines its effectiveness in identifying such errors. We define such consistency as intraconsistency, meaning that the DM maintains internal consistency across all midvalue points within a single value function.

However, the DM may produce different intraconsistent value functions when presented with different anchors. We call such differences a violation of interconsistency. Although there is no objective “true” value function against which to compare outcomes because preferences are constructed rather than preexisting (Slovic 1995), the pattern of violating interconsistency in value function across different anchors offers meaningful insights into the effects of anchoring bias.

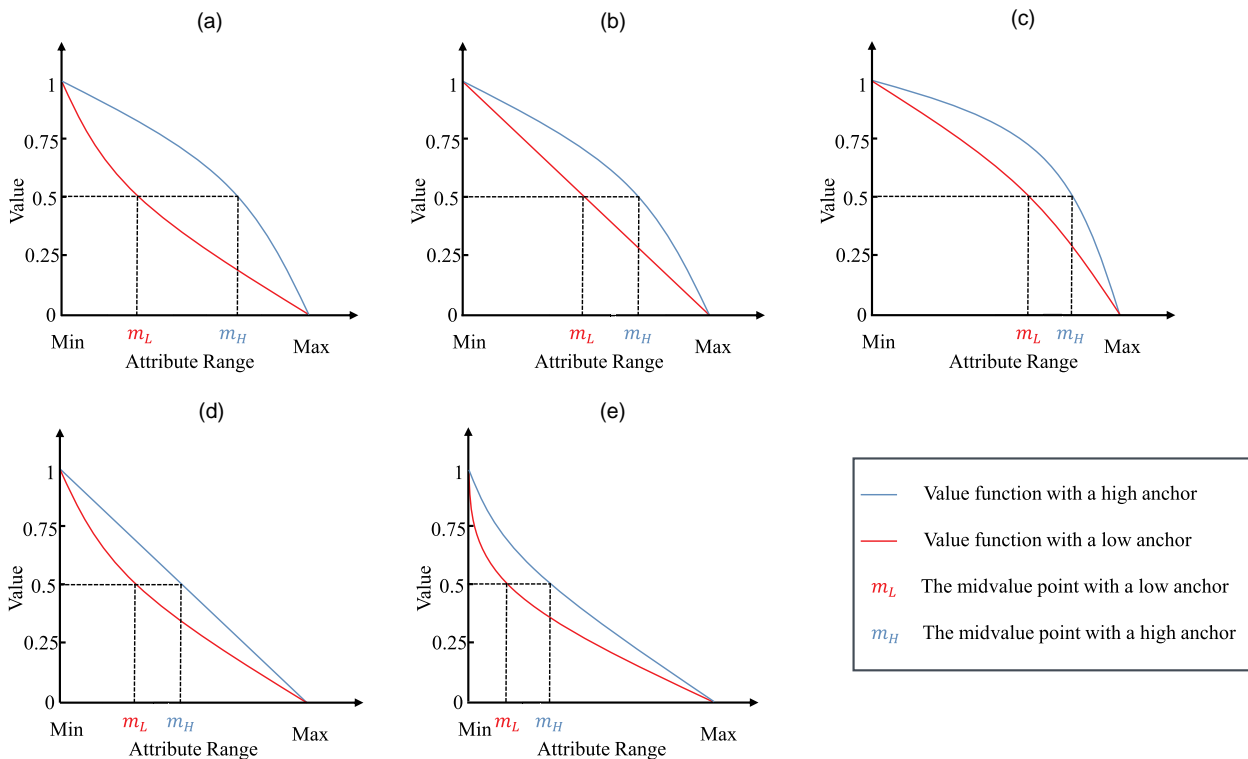
For monotonically decreasing attribute-specific value functions, as illustrated in Figure 1, anchoring bias can alter their shapes in various ways. Specifically, when a high anchor is provided, the resulting value function could be different compared with a low anchor as follows: (a) a concave shape to a convex shape; (b) a concave shape to a linear shape; (c) an extreme concave shape to a concave shape; (d) a linear shape to a convex shape; and (e) a convex shape to an extreme convex shape. In all cases, the area under the curve (AUC) for the value function that

resulted from a low anchor is *smaller* than that of a high anchor. For the monotonically increasing attribute-specific value functions, we could think of five similar situations (convex to concave, convex to linear, extreme convex to convex, linear to concave, and concave to extreme concave), for all of which the AUC for the value function resulted from a low anchor is *larger* than that of a high anchor. These shifts indicate the degree to which judgmental interinconsistencies arise due to anchoring. Therefore, we propose the following hypothesis:

Hypothesis 1. *A low anchor, compared with a high anchor, for the first midvalue point, results in smaller midvalue points and a smaller (larger) area under the curve for a decreasing (increasing) attribute-specific value function.*

To mitigate the effect of anchoring bias, two widely used approaches are (i) avoiding anchors, and (ii) employing counter-anchors. In the context of the midvalue splitting procedure, where the starting point

Figure 1. (Color online) The Possible Effect of Low and High Anchors on Value Function Shape



Notes. (a) A concave shape to a convex shape. (b) A concave shape to a linear shape. (c) An extreme concave shape to a concave shape. (d) A linear shape to a convex shape. (e) A convex shape to an extreme convex shape.

often serves as an anchor, these strategies can be applied in different ways. One approach is to avoid the starting point entirely, allowing the DM to express mid-value points without any prespecified/suggested mid-value by the analyst. Alternatively, counter-anchoring can be implemented by introducing both a low and a high starting point sequentially, in either a low-high or high-low order, with the intention of balancing the influence of each anchor.

When no-anchor or counter-anchoring is employed, the expectation is that the midvalue points will fall between the low and high anchor midvalue points. As a result, the shape of the attribute-specific value function is anticipated to lie between the low and high anchor attribute-specific value functions. Based on this rationale, we hypothesize the following:

Hypothesis 2. *No-anchor and counter-anchoring reduce the impact of anchoring bias, leading to an attribute-specific value function with an AUC in between the AUC of low and high anchor attribute-specific value functions.*

In order to investigate the impact of anchoring bias on the overall value of alternatives, we develop a third hypothesis. The overall value of an alternative is calculated through the aggregation of weights and attribute-specific value functions (see Equation (1)). To isolate the anchoring effect of attribute-specific value function on the overall value, we control the weights through a carefully designed experimental setup (see Section 5).

For a monotonically decreasing value function, the value (v_L) for an attribute level derived from the low-anchored value function will be smaller than the value, v_H , derived from a high-anchored value function. Conversely, for a monotonically increasing value function, the value obtained from the low-anchored value function will be larger than the value obtained from a high-anchored value function. This relationship is illustrated in Figure 2.

Considering a situation with two attributes and fixed weights, if both attributes are monotonically decreasing and have low-anchored value functions, the alternative may receive a lower overall value compared with when both are high-anchored. Conversely, if both attributes are monotonically increasing and low-anchored, the overall value of an alternative may increase compared with the high-anchored scenario. For cases where one attribute is increasing and the other is

decreasing, the opposing effects of anchoring may partially cancel each other out when both attributes are anchored in the same direction, making it challenging to draw a straightforward conclusion.

This reasoning can be extended to scenarios with more than two attributes, though the complexity significantly increases, and experimental analysis becomes more difficult. In the next section, a carefully designed experiment (both attributes are monotonically decreasing, with fixed weights) is used to test the effect of changes in the attribute-specific value functions on the overall value of the alternatives. So, although it is clear how each attribute-specific value function contributes to the overall value, when it comes to experimental analysis, it is more convenient to design situations where the contributions are in the same direction. Specifically, we hypothesize that:

Hypothesis 3. *A low anchor, compared with a high anchor, for the first midvalue point of a decreasing (increasing) attribute-specific value function will contribute to a decrease (increase) in the overall value of an alternative.*

5. Experiment Design

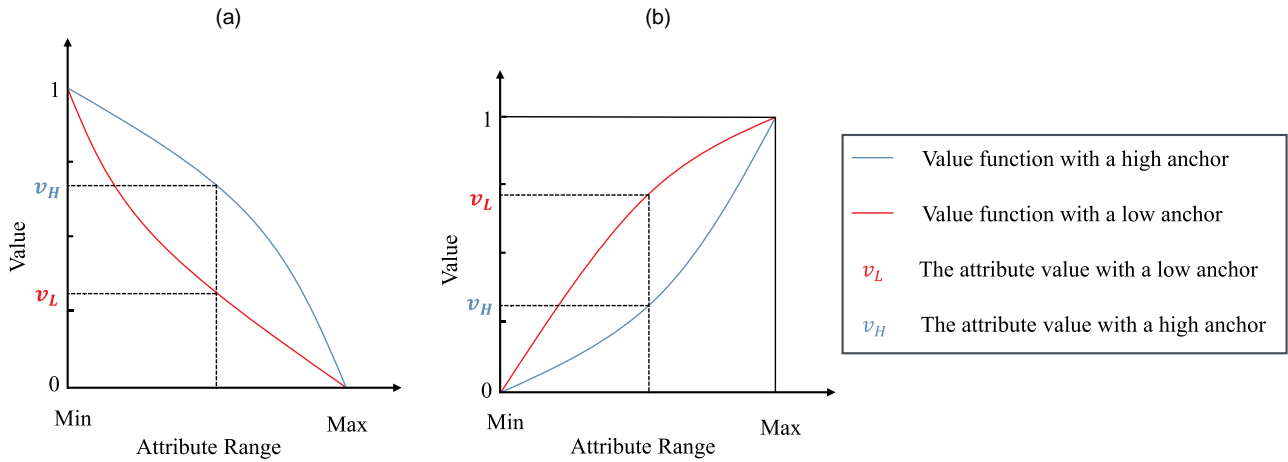
5.1. The Experiment Overview

The experiment was designed following the MAVT steps to test the hypotheses. Before distributing this questionnaire, two pilot studies were conducted to refine the questions and instructions in the questionnaire. Participants were presented with a structured questionnaire including five parts.

5.1.1. Providing Informed Consent. Participants were first informed about the study's purpose, procedures, potential risks, and benefits. They were then asked to voluntarily provide their consent to participate by signing an informed consent form, acknowledging their understanding of the study and their rights as participants.

5.1.2. Presenting a Hypothetical Decision Problem. We designed an apartment renting decision problem with two attributes: monthly rent and commute distance, making it relatable for most participants. This case also provides measurable attributes. Additionally, choosing these two monotonically decreasing attributes helps to test the third hypothesis (as explained in Section 4). The

Figure 2. (Color online) An Example of Possible Changes in Value Under High vs. Low Anchoring



Notes. (a) Decreasing value function. (b) Increasing value function.

decision problem is presented to the participants as follows:

“Suppose you find yourself at the end of your current rental contract and are actively searching for a new apartment. The rental agency has provided you with a list of options, all of which share identical features: each apartment is a 40 square meter studio with a standard 2-year lease renewal term, the layout, amenities and other features are also the same. However, the monthly rent and commuting distance differ for each option.

Monthly rent (euro): It is the amount of money one has to pay each month to rent the apartment.

Commute distance (kilometer): This is the proximity of the apartment to your workplace.”

The attribute range for “Rent” has been set as [700, 1,500], reflecting current rental market conditions in the countries of the experiment (Eurostat 2023, Statistics Netherlands 2024). This range ensures that the values used in the study are realistic and relevant to participants’ experiences in the housing market. Similarly, the attribute range for “Commute Distance” is defined as [5, 20] kilometers. This range is chosen to minimize the possibility of a nonmonotonic value function, which could occur if DMs have personal preferences for certain commute distances. For example, some individuals may prefer a moderate commute distance rather than living too close to or too far from their workplace (Redmond and Mokhtarian 2001). By setting these ranges,


the study aims to capture realistic decision-making behaviors, while ensuring a structured and logical approach to value function elicitation.

5.1.3. Eliciting and Fixing the Weights. This study examines how the effects observed during the value function elicitation step ultimately influence the decision results. In the additive aggregation model, the overall value of an alternative is determined by both the weights and the attribute-specific value functions (see Equation (1)). To isolate the impact in the value function elicitation step and ensure that the results are unaffected by variations in the weight elicitation step, we first elicit and fix the weights. We achieve this by altering the attribute ranges used to elicit the attribute-specific value functions. Specifically, the procedure works as follows.

Let X and Y be two decreasing attributes² with predefined ranges: $X \in [\bar{x}, \underline{x}]$, where \bar{x} is the minimum (best) performance score and \underline{x} is the maximum (worst) performance score of attribute X ; and $Y \in [\bar{y}, \underline{y}]$, where \bar{y} is the minimum (best) performance score and \underline{y} is the maximum (worst) performance score of attribute Y . Figure 3 presents the hypothetical value functions of attribute X and Y .

Step 1: We provide two alternatives, $\mathbf{a}_1 \equiv (\underline{x}, \bar{y})$ and $\mathbf{a}_2 \equiv (\bar{x}, \underline{y})$, and ask if the DM equally prefers them. Note that $X \in [\bar{x}, \underline{x}]$, $Y \in [\bar{y}, \underline{y}]$.

If the DM is indifferent between the two alternatives, following Equation (2), it implies $w_X v(\underline{x}) + w_Y v(\bar{y}) =$

Figure 3. Value Functions for Attributes X and Y

Notes. (a) Value function for attribute X. (b) Value function for attribute Y.

$w_X v(\bar{x}) + w_Y v(\underline{y})$. Because $v(\underline{x}) = v(\underline{y}) = 0$, $v(\bar{x}) = v(\bar{y}) = 1$, we get $w_Y = w_X$. Given $w_X + w_Y = 1$, we derive $w_X = w_Y = 0.5$ and proceed to Step 2.

If the DM prefers \mathbf{a}_2 to \mathbf{a}_1 , it indicates $w_X > w_Y$ when $X \in [\bar{x}, \underline{x}]$, $Y \in [\bar{y}, \underline{y}]$. We then ask the DM to replace \bar{x} by x' , where $x' \in (\bar{x}, \underline{x})$, such that she equally prefers the two alternatives. Later, we change the range of attribute X to $[x', \underline{x}]$ when eliciting the value function. Because the full range of attribute X is now $[x', \underline{x}]$, it means $v(x') = v(\bar{y}) = 1$. Therefore, the weights of attributes X and Y are still 0.5 when $X \in [x', \underline{x}]$, $Y \in [\bar{y}, \underline{y}]$.

Conversely, if the DM prefers \mathbf{a}_1 to \mathbf{a}_2 , it implies $w_X < w_Y$ when $X \in [\bar{x}, \underline{x}]$, $Y \in [\bar{y}, \underline{y}]$. We ask the DM to replace \bar{y} by y' , where $y' \in (\bar{y}, \underline{y})$, so that she equally prefers the two alternatives. We then adjust the range of attribute Y to $[y', \underline{y}]$, when eliciting the value function, maintaining $w_X = w_Y = 0.5$ when $X \in [\bar{x}, \underline{x}]$, $Y \in [y', \underline{y}]$.

For now, we elicited and fixed the weights of the two attributes to 0.5 by changing the attribute ranges (see Table 1). Consequently, in doing so, the overall values of a_1 and a_2 are also fixed to 0.5. To test Hypothesis 3 regarding the effect on the overall value of an alternative, we need another alternative with attribute levels not at the extreme levels of the attribute range.³

Step 2: For participants in the first and last scenarios (see Table 1), we provide a_3 with a random value within the attribute value range of X, like \dot{x} , and ask the DM to choose a value for Y, like \dot{y} , such that she equally prefers the three alternatives. For participants in the second scenario, we provide a_3 with a random value within the attribute value range of Y, like \dot{y} , and ask the

DM to choose a value for X, like \dot{x} , such that she equally prefers the three alternatives.

For the first scenario $w_X = w_Y$, when $X \in [\bar{x}, \underline{x}]$, $Y \in [\bar{y}, \underline{y}]$:

$$\begin{aligned} v(\mathbf{a}_1) &= w_X v(\underline{x}) + w_Y v(\bar{y}) = w_Y \\ v(\mathbf{a}_2) &= w_X v(\bar{x}) + w_Y v(\underline{y}) = w_X \\ v(\mathbf{a}_3) &= w_X v(\dot{x}) + w_Y v(\dot{y}). \end{aligned}$$

For the second scenario $w_X > w_Y$, when $X \in [x', \underline{x}]$, $Y \in [\bar{y}, \underline{y}]$:

$$\begin{aligned} v(\mathbf{a}_1) &= w_X v(\underline{x}) + w_Y v(\bar{y}) = w_Y \\ v(\mathbf{a}_2) &= w_X v(x') + w_Y v(\underline{y}) = w_X v(x') \\ v(\mathbf{a}_3) &= w_X v(\dot{x}) + w_Y v(\dot{y}). \end{aligned}$$

For the third scenario $w_X < w_Y$, when $X \in [\bar{x}, \underline{x}]$, $Y \in [y', \underline{y}]$:

$$\begin{aligned} v(\mathbf{a}_1) &= w_X v(\underline{x}) + w_Y v(y') = w_Y v(y') \\ v(\mathbf{a}_2) &= w_X v(\bar{x}) + w_Y v(\underline{y}) = w_X \\ v(\mathbf{a}_3) &= w_X v(\dot{x}) + w_Y v(\dot{y}). \end{aligned}$$

In all scenarios, \mathbf{a}_1 and \mathbf{a}_2 are constructed using extreme attribute levels from the corresponding attribute

Table 1. Scenarios of Different Attribute Ranges for Fixing Weights

| Scenario | Attribute X | Attribute Y |
|-------------|----------------------------|----------------------------|
| $w_X = w_Y$ | $[\bar{x}, \underline{x}]$ | $[\bar{y}, \underline{y}]$ |
| $w_X > w_Y$ | $[x', \underline{x}]$ | $[\bar{y}, \underline{y}]$ |
| $w_X < w_Y$ | $[\bar{x}, \underline{x}]$ | $[y', \underline{y}]$ |

ranges. The values of these extreme attribute levels are always zero (for the worst) or one (for the best), regardless of any changes in the shape of the attribute-specific value functions. Therefore, the overall values of \mathbf{a}_1 and \mathbf{a}_2 remain fixed at 0.5 based on the calculation above. We design \mathbf{a}_3 in a way that its attribute levels \dot{x} and \dot{y} are within the attribute ranges; thus, any changes in the shape of the attribute-specific value functions due to anchoring bias can affect the values of \dot{x} and \dot{y} . After eliciting the value functions for attributes X and Y in different experiment groups, we can calculate a new overall value of \mathbf{a}_3 , and Hypothesis 3 can be tested by comparing $v(\mathbf{a}_3)$ in different anchor groups. The equal preference among three alternatives serves as a benchmark, enabling us to attribute any deviations in $v(\mathbf{a}_3)$ to anchoring bias. The experimental design establishes that $v(\mathbf{a}_3)$ should theoretically be lower than 0.5 for the low anchor group, higher than 0.5 for the high anchor group, and equal to 0.5 for the three mitigation groups. This setup is intentionally structured to test the anchoring effect on decision-making outcomes, as well as the effectiveness of the mitigation strategies in neutralizing this bias.

5.1.4. Determining the Value Functions Using the Midvalue Splitting Procedure. In this part, participants are presented with a toy example to help them understand the midvalue splitting procedure. Then, they are randomly assigned to one of five groups: low anchor group, high anchor group, low-high counter-anchor group, high-low counter-anchor group, and no-anchor group. The first two groups are to test the anchoring bias, and the last three groups are to test the debiasing strategies. To ensure consistency, all anchored groups receive two consecutive anchors spaced at 20% of the total attribute range. See Table 2 for the detailed information on anchors given in each group. A between-

subject design is essential for this study, which ensures that the impact of each anchor group can be clearly isolated and measured without interference from prior conditions (Shadish et al. 2002). This design eliminates the risk that participants may adjust their responses based on prior exposure to a different anchor (e.g., shifting from a low to high anchor or vice versa), which could distort the observed effects of the anchors.

Take the rent attribute value function elicitation in the high anchor group as an example; to identify the first midvalue point, the participants are presented with two possible rent drops: “Suppose you can get a lower rent for the apartment by increasing the commute distance. Suppose the drop in monthly rent would be either from 1500 euros to 1340 euros or from 1340 euros to 700 euros. For which drop in price would you accept a larger increase in commute distance?” Here, 1,340 euros serves as the high anchor. Regardless of the participants’ answer, in the second question, we will replace 1,340 euros with another high value, 1,180 euros, to the DMs and ask if they are indifferent between the two latest rent drops. After that, participants will directly give a value for x_1 such that they are indifferent between the rent drops from 1,500 euros to x_1 euros and from x_1 euros to 700 euros. This value of x_1 becomes the first midvalue point and is assigned a value score of 0.5. We understand that in a real-world scenario, an analyst would adjust the values in their following questions based on the DM’s response and would continue asking questions until a midvalue point is reached. But for this experiment, we focus on the effect of anchors, and we stop after two questions for the sake of convenience and consistency (in terms of the number of adjustments) across all participants.⁴

This procedure is consistent across all groups, with only the starting points differing. Specifically, in the no-anchor group, no starting point is given to the participants; they directly give the first midvalue point. For the subsequent midvalue points, for all five groups, each subject is asked directly to give the points. After obtaining all three midvalue points, a consistency check is done by asking if the participants are indifferent between the change from $x_{0.25}$ to $x_{0.5}$ and from $x_{0.5}$ to $x_{0.75}$. If they are not indifferent between the two changes, then they will give a new score to $x_{0.5}$, and this new score becomes the first midvalue point.

Table 2. Anchor Values in Different Groups

| Groups | Starting points rent (euro) | Starting points commute distance (km) |
|-------------------------|-----------------------------|---------------------------------------|
| Low anchor | 860 and 1,020 | 8 and 11 |
| No anchor | None | None |
| Low-high counter-anchor | 860 and 1,340 | 8 and 17 |
| High-low counter-anchor | 1,340 and 860 | 17 and 8 |
| High anchor | 1,340 and 1,180 | 17 and 14 |

5.1.5. Collecting Demographic Information. The final section of the questionnaire gathers demographic data such as age, gender, and education level. Additionally, information about the participant's current living space size, monthly rent or mortgage, and commuting distance were also collected as control variables. This information provides a more comprehensive profile of the participants (Creusen 2010, Hammer 2011).

To implement the experiment in an online survey, we used the advanced and user-friendly Qualtrics platform, which allows for dynamic question flows, flexible layouts, and a variety of question types (Couper 2000, Dillman et al. 2014). This platform enhances participant interaction by providing a seamless and intuitive survey experience.

5.2. Participants

Participants were recruited from six European countries: the Netherlands, Germany, France, Belgium, Denmark, and Luxembourg. These countries were chosen because they share similar cultures and apartment renting situations (Eurostat 2023). Data collection was conducted using the online platform Prolific, which offers a large, diverse, and reliable participant pool (Palan and Schitter 2018). Prolific's features for prescreening and response verification enhance the quality of the data. The prescreening function allowed us to limit participants based on their nationality and level of English proficiency. Additionally, the response verification process enabled us to reject incomplete answers or those that failed the midvalue splitting procedure.

Participants received a small monetary incentive for their participation, which has been shown to improve response rates, enhance the quality of responses, and accelerate data collection in research studies (Singer and Ye 2013). This study was not incentivized in the traditional sense of providing performance-based rewards. Although economists often use monetary incentives, arguing that they can elicit more realistic and effortful responses from participants, this is not universally accepted as the best practice, especially in psychology and behavioral decision-making fields (Hertwig and Ortmann 2001). Hascher et al. (2021) suggested that the use of incentives in low-stakes valuation tasks may not be necessary and could potentially be counterproductive. The primary aim of our study is to investigate anchoring effects and value function

elicitation within a controlled experimental setup. Introducing performance-based incentives could have influenced participants' responses, potentially leading to strategic behavior rather than authentic expressions of their preferences (Camerer and Hogarth 1999).

In our study, participants spent an average of 16 minutes and 38 seconds, with a standard deviation of 8 minutes and 59 seconds, completing the experiment. This time includes participation in two experiments in one questionnaire, one of them being the present study. A total of 440 participants were recruited. After data cleaning, 36 responses were excluded because they did not complete the questionnaire. Additionally, in order to maintain logical consistency, 84 participants were excluded from the analysis due to failing the midvalue splitting procedure, as they provided values outside the specified ranges or identical values for all three midvalue points. This indicated either inattention to the questionnaire or a lack of understanding of the questions. This highlights a potential drawback of online questionnaires, where the lack of direct interaction may lead to misunderstandings or errors in following the procedures, resulting in inaccurate responses. Additionally, there were 19 participants who failed to give qualified answers when defining the initial indifference ranking, but all other responses adhered to the task logic and value ranges. This suggests that they understood the instructions but may have had momentary lapses on this specific question. Therefore, they were given a second chance to answer the initial ranking question, and they provided qualified answers.

The final sample included 320 participants, whose data were used for the statistical analysis. Table 3 presents demographic information for the participants. Ethical approval for the study was obtained from the Institutional Review Board of Delft University of Technology.

6. Results and Discussion

This section presents the findings from the experimental analysis designed to test the influence of anchoring bias on value function elicitation within the context of MAVT. The results are structured around the hypotheses; both parametric and nonparametric statistical methods were used.

The nonparametric tests identify consistent directional shifts caused by anchoring bias, irrespective of

Table 3. Demographic Characteristics of Participants ($n = 320$)

| Characteristics | Levels | Percent (%) |
|-----------------|-------------------|-------------|
| Gender | Male | 64 |
| | Female | 35 |
| | Other | 1 |
| Age | [18,24] | 25 |
| | [25,34] | 50 |
| | [35,44] | 14 |
| | >44 | 11 |
| Education | High school | 14 |
| | Bachelor’s degree | 32 |
| | Master’s degree | 36 |
| | Other | 18 |

the scale or distribution of the data. Parametric tests quantify the magnitude of the difference between groups, offering insight into the size and practical significance of the anchoring effect. Together, the two distinct, but complementary, tests provide converging evidence of anchoring bias, demonstrating both its systematic influence on DM judgments and the extent of the effect. Notably, we conducted analysis of covariance tests to examine whether subjects’ current living space, housing costs, commuting distance, and demographics influenced the main results of anchoring bias in this study. The results indicated that none of these variables had a statistically significant effect on the primary dependent variables (all $p > 0.05$).

6.1. Hypothesis 1: Impact of Low vs. High Anchors on Midvalue Points and Value Function Shape

Hypothesis 1 posited that low starting points would result in smaller midvalue points and a smaller (larger) area under the curve compared with high starting points for decreasing (increasing) attribute-specific value functions. To test this, two key indicators were analyzed: the first midvalue point and the AUC.

The first midvalue point x_1 was selected because it represents the initial response most directly influenced by the starting point. A significant difference in this point would suggest that the DM’s initial adjustments were affected by anchoring bias. Notably, the attribute ranges during the value function elicitation step may vary across participants due to the experimental setup. To enable meaningful comparisons and analysis, we normalized the attribute ranges, which also adjusted the scores of the midvalue points accordingly.

The AUC, a measure of the total area under the curve relative to a reference axis, was used to quantify the overall shape of the value function. AUC was chosen over alternative indicators, such as the sign of the second derivative for a value function, because AUC captures the nuanced changes in value function shape between groups. Specifically, whereas for a decreasing value function, a high anchor (compared with a low anchor) may shift a value function from extreme convex to convex, such changes cannot be fully captured by merely categorizing functions as convex or concave. AUC thus provides a more meaningful representation of these shifts and was recognized as the most appropriate measure for this study.

6.1.1. Rent Attribute. For the rent attribute, descriptive statistics showed that participants in the low anchor group produced lower midvalue points and smaller AUC values compared with those in the high anchor group, as detailed in Table 4.

An independent-samples t -test was conducted to compare the first midvalue point for rent between the low and high anchor groups. Levene’s test for equality of variances was not significant, $F(1,139) = 1.721$, $p = 0.192$, indicating that the assumption of equal variances was met. The results showed a statistically significant difference in the first midvalue point between the two groups, $t(139) = -3.496$, $p < 0.001$, confirming that participants in the low anchor group consistently provided lower midvalue points compared with the high anchor group. The effect size (Cohen’s d) was 0.589, representing a medium effect magnitude. To further explore these results, a nonparametric Mann-Whitney U test was performed. This test, which compares the ranks rather than means of the midvalue points, also revealed a significant difference between the two groups, $Z = -3.586$, $p < 0.001$. The Mann-Whitney U test further confirmed that participants in the low anchor group consistently identify their

Table 4. Midvalue Points and AUC for Low and High Anchor Groups (Rent)

| Anchor group | Normalized first midvalue point mean (SD) | AUC value mean (SD) |
|--------------|---|---------------------|
| Low | 0.42 (0.11) | 0.45 (0.07) |
| High | 0.49 (0.12) | 0.49 (0.08) |

midvalue points lower than those in the high anchor group, providing additional evidence of a systematic directional shift caused by anchoring bias.

The anchoring effect on the shape of the value function was examined using the AUC values. The independent-samples t -test for AUC values revealed that the low anchor group produced significantly lower AUC values, $t(139) = -3.276$, $p < 0.001$, with a corresponding medium effect size (Cohen's $d = 0.552$). The Mann-Whitney U test also indicated a significant difference in the ranks of AUC values between the two groups, $Z = -3.147$, $p = 0.002$, with the low anchor group having generally lower ranks compared with the high anchor group. These results imply that the influence of anchoring bias extends beyond the first midvalue point, affecting the overall shape of the value function for the Rent attribute.

6.1.2. Commute Distance Attribute. The descriptive statistics for the attribute commute distance showed a similar pattern: the low anchor group reveals lower first midvalue points and smaller AUC values than the high anchor group (see Table 5).

An independent-samples t -test was conducted to compare the first midvalue point for commute distance between the low and high anchor groups. Levene's test for equality of variances was not significant, $F(1, 139) = 1.574$, $p = 0.212$, indicating that the assumption of equal variances was satisfied. The t -test results showed a statistically significant difference in midvalue points between the low and high anchor groups, $t(139) = -2.504$, $p = 0.007$. The effect size (Cohen's d) was 0.422, representing a medium effect magnitude. Similarly, a Mann-Whitney U test further validated the significant difference between the distributions of midvalue points across the two groups, $Z = -2.605$, $p = 0.009$.

The t -test for AUC values showed significant difference ($t(139) = -2.271$, $p = 0.012$), with a medium effect size (Cohen's $d = 0.383$). The Mann-Whitney U test

Table 5. Midvalue Points and AUC for Low and High Starting Point Groups (Commute Distance)

| Anchor group | Normalized first midvalue point mean (SD) | AUC value mean (SD) |
|--------------|---|---------------------|
| Low | 0.40 (0.13) | 0.43 (0.08) |
| High | 0.45 (0.11) | 0.45 (0.07) |

further corroborated the results, with $Z = -2.191$, $p = 0.028$. These results indicate that anchoring bias also affects the value function elicitation process for commute distance, though its effect is smaller than that of rent. This suggests that commute distance, as a non-monetary attribute, might introduce greater individual variability than rent, thus weakening the presence of bias. This aligns with evidence that nonmonetary attributes reduce preference reversals compared with monetary attributes (Slovic et al. 1990).

6.2. Hypothesis 2: Effectiveness of Debiasing Strategies

Hypothesis 2 examines whether the use of no-anchor or counter-anchoring could mitigate the effects of anchoring bias. This hypothesis was tested by comparing the first midvalue points and AUC values generated by the low-high, high-low, and no-anchor groups with those from the low and high anchor groups to see if the debiasing groups produced less extreme values compared with the low and high anchor groups.

6.2.1. Rent Attribute. According to Table 6, the first midvalue points and AUC values for the mitigation strategies (no-anchor, low-high, and high-low groups) fell between the extremes of the low and high anchor groups, suggesting that these strategies helped to mitigate the anchoring effect by reducing the reliance on one anchor.

In order to examine the anchoring effect across different conditions and to compare specific group differences, an analysis of variance (ANOVA) test and post hoc analyses were used. The ANOVA test for the first midvalue point showed significant differences between the five groups ($F(4, 315) = 3.705$, $p = 0.006$). Post hoc comparisons identified two key findings regarding the effectiveness of debiasing strategies: the no-anchor group and low-high counter-anchor group produced significantly smaller midvalue points than the high anchor group ($p = 0.005$ and $p = 0.038$, respectively), and the high-low counter-anchor group resulted in significantly larger midvalue points than the low anchor group ($p = 0.047$). Importantly, there were no significant differences between the three debiasing strategy groups, indicating that these strategies produce similar midvalue points.

Table 6. Midvalue Points and AUC for Five Anchor Groups (Rent)

| Measure | Low anchor Mean (SD) | No-anchor Mean (SD) | Low-high anchor Mean (SD) | High-low anchor Mean (SD) | High anchor Mean (SD) |
|---------------------------------|-------------------------|------------------------|------------------------------|------------------------------|--------------------------|
| Normalized first midvalue point | 0.42 (0.11) | 0.43 (0.12) | 0.45 (0.10) | 0.46 (0.11) | 0.49 (0.12) |
| AUC value | 0.45 (0.07) | 0.45 (0.08) | 0.46 (0.06) | 0.47 (0.07) | 0.49 (0.08) |

For AUC values, the ANOVA test indicated significant differences between the five groups with ($F(4, 216) = 3.601$) at $p = 0.007$. Post hoc analysis showed that the no-anchor group and low-high counter-anchor group had significantly smaller AUC values compared with the high anchor group ($p = 0.008$ and $p = 0.026$, respectively); the high-low counter-anchor group had significantly larger AUC values than the low anchor group ($p = 0.027$). Similar to the results for the first midvalue point, the AUC values for the three mitigation strategies were nonsignificant, suggesting that they produce similar AUC values. This indicates that the three mitigation strategies can reduce anchoring bias to a similar extent.

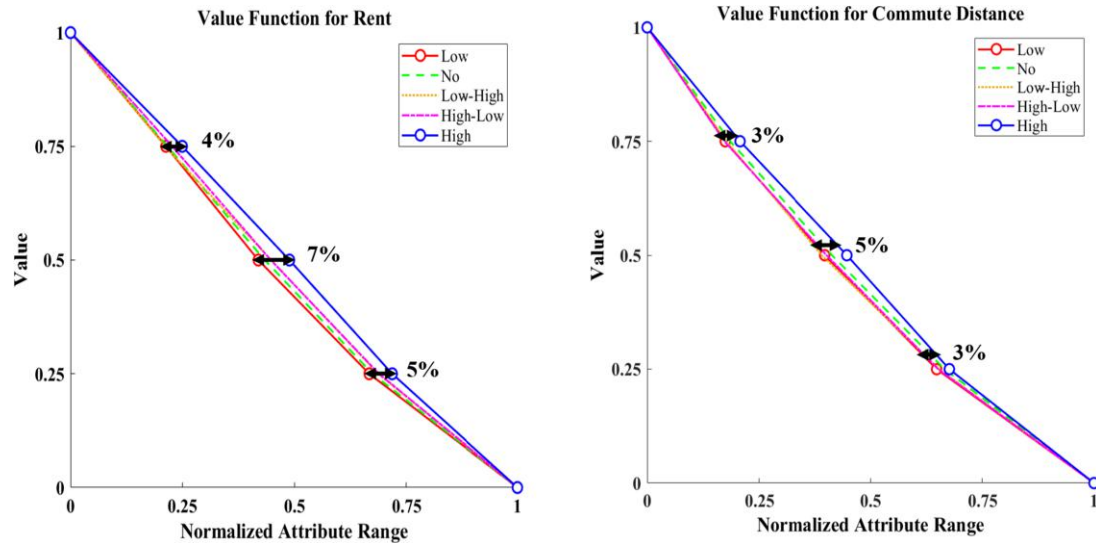
6.2.2. Commute Distance Attribute. Table 7 presents the mean and standard deviation of the first midvalue points and the AUC values for the low and high anchor groups, alongside the three mitigation groups. Although the ANOVA for the first midvalue points ($F(4, 315) = 2.365$, $p = 0.053$) and for the AUC values ($F(4, 315) = 1.830$, $p = 0.123$) did not reach conventional levels of significance, these results do not imply that anchoring bias is absent in the low or high anchor groups. Rather, they suggest that the commute distance attribute is less susceptible to anchoring bias overall. The post hoc analysis revealed important findings for debiasing: the low-high counter-anchor group and high-low counter-anchor group produced significantly lower values than the high anchor group for both the first midvalue point and AUC values ($p = 0.008$, $p = 0.031$ for first midvalue point, respectively; $p = 0.036$ and $p = 0.037$ for AUC values, respectively). Besides,

the three mitigation groups were not significantly different from each other, suggesting that they are equally effective in reducing anchoring bias.

To better visualize these findings, we generated representative value functions (see Figure 4) for each experimental group by normalizing the attribute range and using the average midvalue points. For rent attribute, the representative value functions associated with the debiasing strategies generally fell between those observed in the low and high anchor conditions. The pattern for commute distance attribute was a bit different. That is, whereas the no-anchor condition produced value functions that lay intermediate to the low and high anchor groups, the counter-anchoring strategies did not consistently yield intermediate values. These results suggest that although anchoring bias is evident in the value function elicitation procedure for both attributes, and although the debiasing strategies are all effective for both attributes, the degree of effectiveness varies for the two attributes. Notably, although the normalization of the attribute ranges in Figure 4 visually compresses these differences, their practical impact remains nontrivial, especially considering the statistically significant differences and the medium effect sizes. The anchoring-induced shifts in midvalue points (7% for rent and 5% for commute distance between the low and high anchor groups) could significantly influence decision outcomes, especially in decisions involving critical factors such as financial commitments, safety, security, and health. In the next section, we will show how the decision results could change due to the changes in the value functions as a result of different anchors.

Table 7. Midvalue Points and AUC for Five Anchor Groups (Commute Distance)

| Measure | Low anchor Mean (SD) | No anchor Mean (SD) | Low-high anchor Mean (SD) | High-low anchor Mean (SD) | High anchor Mean (SD) |
|---------------------------------|-------------------------|------------------------|------------------------------|------------------------------|--------------------------|
| Normalized first midvalue point | 0.40 (0.13) | 0.41 (0.13) | 0.39 (0.11) | 0.40 (0.11) | 0.45 (0.11) |
| AUC value | 0.43 (0.08) | 0.44 (0.08) | 0.43 (0.07) | 0.43 (0.07) | 0.45 (0.07) |

Figure 4. (Color online) The Representative Value Functions

6.3. Hypothesis 3: Impact of Anchoring on the Overall Value of Alternatives

Hypothesis 3 suggests that the anchoring bias introduced in value function elicitation would also affect the overall value of alternatives and that the proposed mitigation strategies would be effective in reducing this bias. Specifically, low anchors were expected to lead to lower overall values compared with high anchors for the same alternative in decreasing attribute-specific value functions. The overall value of alternative 3 was used to test this hypothesis.

As shown in Table 8, the group exposed to the low anchor produced the smallest overall value, whereas the high anchor group produced the largest overall value. The three mitigation strategies (no-anchor, low-high, high-low) resulted in overall values that fall between the extremes of the low and high anchor groups, aligning with the expectations of the hypothesis.

ANOVA test results indicated significant differences among the five groups, $F(4, 315) = 4.805, p < 0.001$. Post hoc comparisons revealed that the low anchor group provided significantly lower overall values compared with the high anchor group, with $p < 0.001$. Additionally, all three debiasing groups produced significantly lower overall values than the high anchor group ($p = 0.003, p = 0.005, \text{ and } p = 0.039$, respectively).

The high-low counter-anchor group produced a significantly higher value than the low anchor group ($p = 0.048$). Noticeably, no significant differences were observed among the three mitigation groups, suggesting that they have similar effects on the overall value.

Supporting these findings, the Kruskal-Wallis H test observed significant differences among the five groups ($\chi^2(4) = 22.217, p < 0.001$). The Mann-Whitney U test for the overall value of a_3 revealed similar effects as with the ANOVA post hoc analysis results. All three debiasing groups produced significantly lower overall values than the high anchor group ($p < 0.001, p = 0.006, \text{ and } p = 0.025$, respectively). The high-low counter-anchor group produced a significantly higher value than the low anchor group ($p = 0.024$). No significant differences were observed among the three mitigation strategies, which shows their mitigation effect is similar.

To examine how anchoring bias deviates participants from their initial equal preferences within each group, we tested if the overall value of a_3 is different from 0.5 across the five groups. One-sample t -test results show that all groups are significantly above 0.5 ($p < 0.05$ for all), a result that initially seems counterintuitive. Ideally, one would expect the low anchor group to yield values below 0.5, the high anchor group to produce

Table 8. Overall Value for a_3 in Different Anchor Groups

| Measure | Low anchor Mean (SD) | No-anchor Mean (SD) | Low-high anchor Mean (SD) | High-low anchor Mean (SD) | High anchor Mean (SD) |
|----------|-------------------------|------------------------|------------------------------|------------------------------|--------------------------|
| $v(a_3)$ | 0.52 (0.09) | 0.54 (0.09) | 0.54 (0.11) | 0.56 (0.09) | 0.59 (0.10) |

values above 0.5, and the debiasing groups to be closer to 0.5.

A closer inspection of individual-level data (see Table 9) provides important context beyond the group means. Although group means are all above 0.5, the distribution of responses shows systematic differences. Notably, for $v(a_3) < 0.5$, the low anchor group has the highest percentage of participants (33%), whereas the high anchor group has the lowest percentage (13%). The debiasing groups fall between the extremes. For $v(a_3) > 0.5$, the high anchor group shows the highest percentage (82%). In contrast, the low anchor group shows the lowest percentage (51%). The remaining groups fall in between. A Chi-square test indicated significant differences across the five groups, with $\chi^2(8,320) = 19.392$ and $p = 0.013$. Further, Chi-square analysis among the group pairs showed that the low anchor group, no-anchor group, and the low-high counter-anchor group are significantly different from the high anchor group ($p < 0.001$, $p = 0.015$, and $p = 0.038$, respectively).

In summary, although the overall means across groups are elevated (see Table 8), the individual-level analysis reveals clear, systematic differences that align with our theoretical expectations: low anchor leads to more values below 0.5, high anchor leads to more values above 0.5, and the debiasing groups produce values in between the extremes.

Table 9. Distribution of Individual-Level Overall Values for a_3 Across Anchor Groups

| Anchor group | $v(a_3)$ | | | |
|-----------------|------------------|---------------------|------------------|--------------|
| | Below 0.5 (%) | Equal to 0.5 (%) | Above 0.5 (%) | Total (%) |
| Low anchor | 24 (33) | 12 (16) | 37 (51) | 73 (100) |
| No-anchor | 12 (20) | 11 (18) | 37 (63) | 60 (100) |
| Low-high anchor | 16 (28) | 6 (10) | 36 (62) | 58 (100) |
| High-low anchor | 13 (21) | 6 (10) | 42 (69) | 61 (100) |
| High anchor | 9 (13) | 3 (4) | 56 (82) | 68 (100) |

The experimental setting, where three alternatives were initially equally preferred, was designed to fix the weights and thereby isolate the effect of anchoring bias on the value function. In normal decision-making scenarios, where weights are not fixed, the effect of anchoring bias can be even more complex and pronounced due to its combined influence on both the weight elicitation (Rezaei et al. 2024) and value function elicitation processes. Furthermore, in real-world decision problems, even when the alternatives are not similar with respect to the attributes, the overall values of the alternatives are often close to each other, making the ranking highly sensitive to the inputs (i.e., value functions and weights). When the value functions or weights shift even slightly due to anchoring bias, the aggregated scores can cause significant and complex changes to the rankings.

6.4. Implications for Similar Procedures

Although this study focuses on the midvalue splitting procedure, its implications can be extended to other value function elicitation methods as well. Below, we discuss three such methods and their susceptibility to anchoring bias. Future research could investigate the vulnerability of these methods to anchoring bias and examine the effectiveness of the proposed debiasing strategies.

6.4.1. Standard Difference Procedure (Beinat 1997). This procedure divides the attribute range into equal value-spaced subintervals and plots the value function by solving a system of linear equations. When eliciting the value function for an increasing attribute X_j , the analyst will first define a value of $x_1 \in (\underline{x}, \bar{x})$ and then ask the DM to identify a value of $x_2 \in (x_1, \bar{x}]$ such that she is indifferent between the subintervals (\underline{x}, x_1) and (x_1, x_2) . The analyst continues to ask the DM to identify a value of x_N until $x_N = \bar{x}$. Consequently, the value function is defined using these indifference points—for example, $v_j(x_1) = v_j(x_N)/N = 1/N$. Below, we discuss how

anchoring bias could affect the value function elicited using this procedure.

During this process, the first indifference point x_1 set by the analyst can act as an anchor, distorting subsequent judgments in identifying the indifference points in two key ways: (i) Beginning this procedure at the lower versus upper end of the range can result in different value functions. When DMs identify subsequent indifference points x_2, x_3, \dots, x_N , they might adjust insufficiently from x_1 due to anchoring bias. If the procedure starts at the lower end, the increments between indifference points may be smaller than they should be. Conversely, starting at the upper end can result in smaller decrements. Like the midvalue splitting procedure, this effect on the indifference points can lead to an overall effect on the shape of the value function due to the iterative structure. In case of an increasing value function, beginning this procedure at the lower end could then lead to a value function with a larger AUC compared with beginning at the upper end. (ii) Beginning the procedure with a smaller or a larger value of x_1 can result in different value functions. The identification of x_2, x_3, \dots, x_N can be affected by x_1 due to anchoring bias. Beginning this procedure with a smaller value of x_1 could lead to smaller increments between subsequent indifferent points compared with beginning this procedure at a larger value of x_1 . In case of an increasing value function, this could also lead to a value function with a larger AUC when beginning this procedure with a smaller value of x_1 . Building on the counter-anchoring strategy, debiasing could involve conducting the procedure first at the lower end and then repeating it at the upper end, or vice versa. By eliciting indifference points from both directions, the influence of any single starting point is reduced, and the final value function can be derived by averaging or synthesizing the results. Similarly, the procedure could be conducted using different starting values within the range to further mitigate anchoring effects.

6.4.2. Lock-Step Procedure (Keeney and Raiffa 1976).

This procedure constructs the attribute-specific value functions by iteratively defining indifference points between two attributes. Specifically, the worst outcomes for both attributes are defined at first, $v(\underline{x}, \underline{y}) = v_X(\underline{x}) = v_Y(\underline{y}) = 0$. Second, the analyst selects $x_1 > \underline{x}$ and sets $v_X(x_1) = 1$ to fix the “unit” of value for

attribute X. Third, the analyst will ask the DM to identify y_1 such that she is indifferent between (x_1, \underline{y}) and (\underline{x}, y_1) and define $v_Y(y_1) = 1$. This process continues until sufficient indifferent points (e.g., $x_2, y_2, x_3, y_3, \dots, x_k, y_k$) are defined to satisfy chains of indifference (e.g., $(x_2, \underline{y}) \sim (x_1, y_1) \sim (\underline{x}, y_2)$), with values incremented sequentially ($v_X(x_k) = v_Y(y_k) = k$). In this process, the analyst’s selection of x_1 can also serve as an anchor, distorting subsequent judgments in identifying the indifference points in a similar way as explained in the standard difference procedure. Both the starting point in the attribute range (lower versus upper end) and the size of the initial value (x_1) affect the resulting value function. Counter-anchoring debiasing strategies—such as bidirectional elicitation and multiple starting values for x_1 —can help mitigate these biases.

6.4.3. Direct Rating (Fishburn 1967). This procedure directly assigns 0 to the worst performance level and 100 to the best performance level and then asks the DM to assign scores to other performance levels. However, the order (rating from high-to-low performance level versus from low-to-high performance level) can also serve as anchors and affect the elicited values. For example, starting with high-performance levels may anchor DMs to overvalue subsequent midrange levels. Building on the no-anchor strategy, debiasing could involve randomizing the presentation order. Building on the counter-anchoring strategy, debiasing could involve presenting performance levels in an alternating sequence, starting with the highest, then the lowest, followed by the second highest, then the second lowest, and so on.

In summary, these elicitation procedures share similar characteristics: the analyst’s suggestions or presentation orders (anchors) that might affect the identification of indifference points or elicited values, and iterative procedures that can carry this effect to the value function. These common features allow the implications of this study to be extended to other value function elicitation methods.

7. Conclusion

In this study, we theorized and empirically tested the effect of anchoring bias on the midvalue splitting procedure, a common approach for value function elicitation, within the context of multiattribute value theory.

We found that the starting point in the midvalue splitting procedure can serve as an anchor, influencing attribute-specific value functions. Consequently, this anchoring effect alters the overall value of the alternatives. The debiasing strategies tested in this study (no-anchor and counter-anchoring) effectively mitigate anchoring bias by producing less extreme results compared with the low and high anchor conditions. Moreover, the proposed mitigation strategies do not introduce additional bias when an attribute is less prone to bias, demonstrating their robustness.

The results of this study align with and extend the existing literature on anchoring bias and its impact on decision-making processes (Tversky and Kahneman 1974, Englich et al. 2006). Our findings contribute to this field by showing its significant effects on the shape of elicited value functions and results in MAVT. In the context of MADM, whereas research has been primarily focused on the effect of anchoring bias in weight elicitation (Buchanan and Corner 1997, Jacobi and Hobbs 2007, Rezaei 2021, Rezaei et al. 2024), this study provides empirical evidence in another essential part, the value function elicitation. The findings raise concerns about the reliability of commonly used elicitation procedures, as the characteristics of the value function directly affect the evaluation of the alternatives. Anchoring bias strongly influenced the overall value of alternatives, even when only two attributes were considered. In real-world decision-making contexts, where the number of attributes and alternatives is typically much larger, the impact of anchoring bias can result in even stronger distortions in the evaluation process, potentially leading to systematically biased decisions.

The findings of this study carry important implications for analysts employing MAVT and similar methods to support decision makers. First, analysts must recognize the susceptibility of these methods to anchoring bias in value function elicitation. To reduce such bias, analysts should remain neutral throughout all procedures. This includes avoiding the introduction of starting points that may unintentionally serve as anchors. When initial values or reference points are necessary, providing counter-anchors can help reduce their impact. Additionally, these insights extend beyond individual analysts to the design and implementation of decision-support systems that rely on value function elicitation. Such systems should be developed with

mechanisms to counteract anchoring bias, ensuring that their structure and interfaces do not inadvertently introduce or reinforce biased reference points.

This study enhances our understanding of anchoring bias in value function elicitation and provides important implications for both theoretical and practical improvements of MAVT and other decision analysis methods. However, several limitations should be acknowledged. Despite refinements through two pilot experiments, some participants still failed the midvalue splitting procedure. This might suggest that the procedure imposes considerable cognitive demands on participants, potentially leading to fatigue and errors. The sample's cultural and geographical specificity may limit the generalizability of the findings (Norenzayan et al. 2007, Ma-Kellams 2020). The simplified two-attribute decision problem may not fully capture real-world complexity. The degree of effectiveness of the debiasing strategies varies for the two attributes examined in this study, indicating different susceptibility to anchoring bias. These limitations highlight the need for replication studies to confirm the robustness of the observed effects.

Future research could build on these findings by examining the influence of anchoring bias in complex decision problems involving more than two attributes and different types of attribute value functions (e.g., a mix of increasing and decreasing value functions). Additionally, it would be valuable to investigate how anchoring bias affects other value function elicitation procedures and whether the proposed mitigation strategies remain effective in those contexts. We introduced some general ideas in Section 6.4, which future research can further develop and extend. Furthermore, conducting the experiment in a more controlled environment, such as a supervised laboratory setting, could help improve data quality by minimizing inattentive responses and enabling better identification of potentially invalid data, such as unusually fast responses on specific tasks. Finally, expanding the experiments to different cultural or geographical contexts could help generalize the study's findings.

Endnotes

¹ Given that the additive assumption is validated, it is easy and natural to ignore the changes in values of the other attributes when answering these questions (Smith and Dyer 2021).

² In this paper, we use \underline{x} and \bar{x} to show the worst and best performance scores of attribute X . For a decreasing attribute, then, the range is shown as $X \in [\bar{x}, \underline{x}]$, whereas for an increasing attribute the range is shown as $X \in [\underline{x}, \bar{x}]$.

³ The additive assumptions for the aggregation model were verified between attributes through relevant questions in the questionnaire.

⁴ See, for an example, Online Appendix A, the part of the questionnaire related to midvalue splitting procedure for participants in the high anchor condition of the first scenario.

References

- Adame BJ (2016) Training in the mitigation of anchoring bias: A test of the consider-the-opposite strategy. *Learn. Motivation* 53(2016): 36–48.
- Anthes R, Maier M, Ackerman S, Atlas R, Callahan L, Dittberner G, Yoe J, et al. (2019) Developing priority observational requirements from space using multi-attribute utility theory. *Bull. Amer. Meteorological Soc.* 100(9):1753–1774.
- Ariely D (2008) *Predictably Irrational: The Hidden Forces That Shape Our Decisions* (Harper Collins, NY).
- Barber BM, Odean T (2001) Boys will be boys: Gender, overconfidence, and common stock investment. *Quart. J. Econom.* 116(1):261–292.
- Bazerman MH, Moore DA (2012) *Judgment in Managerial Decision Making* (John Wiley & Sons, Hoboken, NJ).
- Beinat E (1997) *Value Functions for Environmental Management* (Springer, Dordrecht, Netherlands).
- Belton V, Stewart T (2012) *Multiple Criteria Decision Analysis: An Integrated Approach* (Springer, New York).
- Brewer NT, Chapman GB, Schwartz JA, Bergus GR (2007) The influence of irrelevant anchors on the judgments and choices of doctors and patients. *Medical Decision Making* 27(2):203–211.
- Buchanan JT, Corner J (1997) The effects of anchoring in interactive MCDM solution methods. *Comput. Oper. Res.* 24(10):907–918.
- Bushong B, Gagnon-Bartsch T (2024) Failures in forecasting: An experiment on interpersonal projection bias. *Management Sci.* 70(12):8735–8752.
- Camerer CF, Hogarth RM (1999) The effects of financial incentives in experiments: A review and capital-labor-production framework. *J. Risk Uncertainty* 19(1):7–42.
- Chapman GB, Johnson EJ (2002) Incorporating the irrelevant: Anchors in judgments of belief and value. *Heuristics and Biases: The Psychology of Intuitive Judgment* (Cambridge University Press, New York), 120–138.
- Couper MP (2000) Web surveys: A review of issues and approaches. *Public Opinion Quart.* 64(4):464–494.
- Creusen ME (2010) The importance of product aspects in choice: The influence of demographic characteristics. *J. Consumer Marketing* 27(1):26–34.
- Dillman DA, Smyth JD, Christian LM (2014) *Internet, Phone, Mail, and Mixed-Mode Surveys: The Tailored Design Method* (John Wiley & Sons, Hoboken, NJ).
- Dyer JS, Sarin RK (1979) Measurable multiattribute value functions. *Oper. Res.* 27(4):810–822.
- Englich B, Mussweiler T, Strack F (2006) Playing dice with criminal sentences: The influence of irrelevant anchors on experts' judicial decision making. *Personality Soc. Psych. Bull.* 32(2):188–200.
- Epley N, Gilovich T (2005) When effortful thinking influences judgmental anchoring: Differential effects of forewarning and incentives on self-generated and externally provided anchors. *J. Behav. Decision Making* 18(3):199–212.
- Eurostat (2023) Housing in Europe. Accessed October 12, 2023, <https://ec.europa.eu/eurostat/web/interactive-publications/housing-2023>.
- Fasolo B, Heard C, Scopelliti I (2024) Mitigating cognitive bias to improve organizational decisions: An integrative review, framework, and research agenda. *J. Management* 51(6): 2182–2211.
- Ferretti V (2016) From stakeholders analysis to cognitive mapping and multi-attribute value theory: An integrated approach for policy support. *Eur. J. Oper. Res.* 253(2):524–541.
- Fischhoff B (1991) Value elicitation: Is there anything in there? *Amer. Psychologist* 46(8):835.
- Fishburn PC (1967) Methods of estimating additive utilities. *Management Sci.* 13(7):435–453.
- Furnham A, Boo HC (2011) A literature review of the anchoring effect. *J. Socio-Econom.* 40(1):35–42.
- Galinsky AD, Mussweiler T (2001) First offers as anchors: The role of perspective-taking and negotiator focus. *J. Personality Soc. Psych.* 81(4):657–669.
- Gavirneni S, Xia Y (2009) Anchor selection and group dynamics in newsvendor decisions—A note. *Decision Anal.* 6(2):87–97.
- Hammer CS (2011) The importance of participant demographics. *Amer. J. Speech-Language Pathology* 20(4):261.
- Hascher J, Desai N, Krajčich I (2021) Incentivized and non-incentivized liking ratings outperform willingness-to-pay in predicting choice. *Judgment Decision Making* 16(6):1464–1484.
- Hertwig R, Ortmann A (2001) Experimental practices in economics: A methodological challenge for psychologists? *Behav. Brain Sci.* 24(3):383–403.
- Höfer T, Madlener R (2020) A participatory stakeholder process for evaluating sustainable energy transition scenarios. *Energy Policy* 139:111277.
- Hostmann M, Bernauer T, Mosler HJ, Reichert P, Truffer B (2005) Multi-attribute value theory as a framework for conflict resolution in river rehabilitation. *J. Multi-Criteria Decision Anal.* 13(2–3): 91–102.
- Jacobi SK, Hobbs BF (2007) Quantifying and mitigating the splitting bias and other value tree-induced weighting biases. *Decision Anal.* 4(4):194–210.
- Kahneman D (2011) *Thinking, Fast and Slow* (Farrar, Straus and Giroux, New York).
- Keeney RL (1974) Multiplicative utility functions. *Oper. Res.* 22(1):22–34.
- Keeney RL (2002) Common mistakes in making value trade-offs. *Oper. Res.* 50(6):935–945.
- Keeney RL (2004) Making better decision makers. *Decision Anal.* 1(4):193–204.
- Keeney RL (2009) *Value-Focused Thinking: A Path to Creative Decision-making* (Harvard University Press, Cambridge, MA).
- Keeney RL, Raiffa H (1976) *Decisions with Multiple Objectives: Preferences and Value Trade-Offs* (Cambridge University Press, Cambridge, UK).
- Kim S-H, Tong J (2024) Admission control bias and path-dependent feedback under diagnosis uncertainty. *Manufacturing Service Oper. Management* 26(1):117–136.
- Kirkwood CW, Sarin RK (1980) Preference conditions for multiattribute value functions. *Oper. Res.* 28(1):225–232.
- Labijak-Kowalska A, Kadziński M, Spychała I, Dias LC, Fiallos J, Patrick J, Michalowski W, Farion K (2023) Performance evaluation of emergency department physicians using robust value-based additive efficiency model. *Internat. Trans. Oper. Res.* 30(1):503–544.

- Lilienfeld SO, Ammirati R, Landfield K (2009) Giving debiasing away: Can psychological research on correcting cognitive errors promote human welfare? *Perspect. Psych. Sci.* 4(4):390–398.
- Lord CG, Lepper MR, Preston E (1984) Considering the opposite: A corrective strategy for social judgment. *J. Personality Soc. Psych.* 47(6):1231–1243.
- Ma-Kellams C (2020) Cultural variation and similarities in cognitive thinking styles versus judgment biases: A review of environmental factors and evolutionary forces. *Rev. General Psych.* 24(3):238–253.
- Montibeller G, von Winterfeldt D (2015) Cognitive and motivational biases in decision and risk analysis. *Risk Anal.* 35(7):1230–1251.
- Montibeller G, von Winterfeldt D (2024) Behavioral decision research: Descriptive and prescriptive perspectives. Federspiel FM, Montibeller G, Seifert M, eds. *Behavioral Decision Analysis* (Springer, Cham, Switzerland), 15–40.
- Mussweiler T, Strack F (1999) Hypothesis-consistent testing and semantic priming in the anchoring paradigm: A selective accessibility model. *J. Experiment. Soc. Psych.* 35(2):136–164.
- Mussweiler T, Strack F, Pfeiffer T (2000) Overcoming the inevitable anchoring effect: Considering the opposite compensates for selective accessibility. *Personality Soc. Psych. Bull.* 26(9):1142–1150.
- Norenzayan A, Choi I, Peng K (2007) Perception and cognition. *Handbook of Cultural Psychology* (The Guilford Press, New York), 569–594.
- Palan S, Schitter C (2018) Prolific.ac—A subject pool for online experiments. *J. Behav. Experiment. Finance* 17:22–27.
- Payne JW, Bettman JR, Johnson EJ (1992) Behavioral decision research: A constructive processing perspective. *Annu. Rev. Psych.* 43(1):87–131.
- Prava VR, Clemen RT, Hobbs BF, Kenney MA (2016) Partition dependence and carryover biases in subjective probability assessment surveys for continuous variables: Model-based estimation and correction. *Decision Anal.* 13(1):51–67.
- Redmond LS, Mokhtarian PL (2001) The positive utility of the commute: Modeling ideal commute time and relative desired commute amount. *Transportation* 28(2):179–205.
- Rezaei J (2015) Best-worst multi-criteria decision-making method. *Omega* 53:49–57.
- Rezaei J (2021) Anchoring bias in eliciting attribute weights and values in multi-attribute decision-making. *J. Decision Systems* 30(1):72–96.
- Rezaei J, Arab A, Mehregan M (2024) Analyzing anchoring bias in attribute weight elicitation of smart, swing, and best-worst method. *Internat. Trans. Oper. Res.* 31(2):918–948.
- Shadish W, Cook T, Campbell DT (2002) *Experimental and Quasi-Experimental Designs for Generalized Causal Inference* (Houghton Mifflin, Boston).
- Singer E, Ye C (2013) The use and effects of incentives in surveys. *Ann. Amer. Acad. Polit. Soc. Sci.* 645(1):112–141.
- Slovic P (1995) The construction of preference. *Amer. Psychologist* 50(5):364–371.
- Slovic P, Griffin D, Tversky A (1990) Compatibility effects in judgment and choice. *Insights in decision making: A tribute to Hillel J. Einhorn*, 5–27 (The University of Chicago Press, Chicago).
- Smith JE, Dyer JS (2021) On (measurable) multiattribute value functions: An expository argument. *Decision Anal.* 18(4):247–256.
- Statistics Netherlands (2024) Centraal Bureau Voor de Statistiek. Accessed July 8, 2024, <https://www.cbs.nl/>.
- Thaler RH, Sunstein CR (2008) *Nudge: Improving Decisions About Health, Wealth, and Happiness* (Yale University Press, New Haven, CT).
- Tversky A, Kahneman D (1974) Judgment under uncertainty: Heuristics and biases. *Science* 185(4157):1124–1131.
- Tversky A, Kahneman D (1981) The framing of decisions and the psychology of choice. *Science* 211(4481):453–458.
- Wansink B, Kent RJ, Hoch SJ (1998) An anchoring and adjustment model of purchase quantity decisions. *J. Marketing Res.* 35(1):71–81.
- Watson SR, Buede DM (1987) *Decision Synthesis: The Principles and Practice of Decision Analysis* (Cambridge University Press, Cambridge, UK).
- Wilson TD, Houston CE, Etling KM, Brekke N (1996) A new look at anchoring effects: Basic anchoring and its antecedents. *J. Experiment. Psych. General* 125(4):387–402.

Geqie Sun is a PhD candidate at Delft University of Technology, the Netherlands. Her research focuses on behavioral decision making, with a particular interest in how cognitive biases influence preference elicitation in multicriteria decision analysis and how these biases can be mitigated.

Maarten Kroesen is an associate professor of travel behavior research at Delft University of Technology, the Netherlands. He holds a PhD and a Master of Science in Complex Systems Engineering, both from Delft University of Technology. His research focuses on the use of quantitative empirical methods to analyze, explain, and predict travel behavior. Through his work, he aims to promote healthy and sustainable forms of mobility.

Jafar Rezaei is a professor of decision science at Delft University of Technology, the Netherlands. His research focuses on multicriteria decision analysis and behavioral decision making. He is the founder of the Best-Worst Method. He currently serves as the Editor-in-Chief of the Journal of Multi-Criteria Decision Analysis. His work has been published in leading journals in decision science, and he regularly collaborates with a range of organizations to support complex decision making.