Received: 21 June 2022 | Revised: 7 October 2022 | Accepted: 7 October 2022
DOI: 10.1002/arcp.1082

REVIEW

CPR CONSUMER PSYCHOLOGY REVIEW
SCP SOCIETY FOR CONSUMER PSYCHOLOGY

Relevance insensitivity: A framework of psychological biases in consumer behavior and beyond

Yang Yang[1] | Xilin Li[2] | Christopher K. Hsee[3]

[1]Warrington College of Business, University of Florida, Gainesville, Florida, USA
[2]China Europe International Business School, Shanghai, China
[3]Booth School of Business, University of Chicago, Chicago, Illinois, USA

Correspondence
Yang Yang, Warrington College of Business, University of Florida, Gainesville, Florida, USA.
Email: yang.yang@warrington.ufl.edu

Abstract

In judgment and choice, consumers show a variety of biases, from the sunk cost fallacy and projection bias to usage frequency neglect and erroneous price–quality inferences. This article explains these seemingly disparate biases and predicts new biases using an overarching framework based on the relevance insensitivity theory proposed by Hsee et al. (2019). According to the theory, many biases arise because people are insufficiently sensitive to the relevance (i.e., weight) of a cue variable to the target variable (the dependent variable). The direction of the bias depends on the normative relevance of the cue—people over-rely on the cue when it is normatively irrelevant and under-rely on the cue when it is normatively highly relevant. We show that ostensibly unique and universal biases are neither unique nor universal: All are manifestations of relevance insensitivity, and each bias attenuates or reverses as the cue variable's relevance changes.

KEYWORDS
biases, diagnosticity, insensitivity, integrative theory, noise, overgeneralization, relevance insensitivity, reverse biases

1 | INTRODUCTION

Many important and fascinating phenomena in consumer behavior are biases in the sense that they are deviations from normative standards or optimal solutions. For decades, research has documented both general biases such as the sunk cost fallacy (Arkes & Ayton, 1999), non-regressive prediction (Kahneman & Tversky, 1973), anchoring effect (Epley & Gilovich, 2006; Tversky & Kahneman, 1974), and base rate neglect (Kahneman & Tversky, 1973) and biases that are specific to consumer behavior, such as the effort heuristic (Chinander & Schweitzer, 2003), projection bias (Loewenstein et al., 2003), and specification-seeking effect (Hsee et al., 2009). These biases often are treated as disconnected from each other, each having its own explanation(s).

According to the relevance insensitivity theory (Hsee et al., 2019), many seemingly unrelated biases share the same underlying cause, namely, that people are insufficiently sensitive to the relevance of a cue variable to a target variable. In their original paper, Hsee et al. use the relevance insensitivity theory to elucidate four classic biases: the sunk cost fallacy, non-regressive prediction, anchoring effect, and base rate neglect. For each bias, they manipulate the relevance of the cue variable, and they find that each original bias emerges near only one end of the relevance spectrum. At the other end of the spectrum, the original bias reverses.

In the present paper, we argue that the relevance insensitivity theory can explain many psychological biases beyond the four classic biases discussed in Hsee et al. (2019). We first review the relevance insensitivity theory and then use it as an overarching framework to explain many documented biases and to predict situations in which reverse biases may occur.

2 | RELEVANCE INSENSITIVITY

All judgments are based on cues. According to the theory of relevance insensitivity, which was first proposed by Hsee et al. (2019), many seemingly unrelated biases have a common cause: People are insensitive to the relevance (i.e., weight) of a cue variable to the target variable (the value to be judged). Notably, "relevance insensitivity" does not mean that people treat the relevance (weight) of the cue as zero; rather, it means that people are insufficiently sensitive to the normative (correct) relevance of the cue. As a result, we predict that the direction of the bias depends on how the normative relevance of the cue in the given situation compares to the cue's relevance in the situations people usually encounter.

Specifically, relevance insensitivity theory posits that people learn the relevance of a cue variable in situations they usually encounter. Of course, a cue's relevance varies situationally, but our theory argues that people are insensitive to these variations. Instead, they overgeneralize the cue's relevance in usually encountered situations to all situations. When the relevance of the cue variable in a particular situation is lower than in usually encountered situations, people will behave as if the relevance were not so low, so they over-rely on the cue variable when judging the target variable. Conversely, when the relevance of the cue variable in a particular situation is higher than in usually encountered situations, people will behave as if the relevance were not so high, so they under-rely on the cue variable when judging the target variable.

To illustrate, consider the relationship between the price (a cue variable) and quality (a target variable) of wine. Suppose that the relationship between the price and quality of wine differs from region to region. You live in a region in which the relevance of price to quality is moderate (so we will call your region the "moderate-relevance region"), and you are adept at predicting quality based on price in your region. Imagine that you are now visiting a region where the relevance of price to quality is low (the "low-relevance region"), such that the price of wine is not at all predictive of wine quality. (The price depends solely on non-quality-related factors such as advertising costs and tariffs.) How accurate would you be at predicting the quality of wine in the low-relevance region? Our theory predicts that you would over-rely on price when judging quality in the low-relevance region, so you would overestimate the difference in the quality of cheap wines and expensive wines.

According to our theory, you over-rely on price when judging quality in the low-relevance region because you are insensitive to the relevance of price to quality. You may know that the price of wine in the low-relevance region does not depend on quality, and yet, consciously or unconsciously, you apply the knowledge you learned in the moderate-relevance region to the low-relevance region, believing that if the price is high, the quality must also be high.

For the same reason, you may exhibit a reverse bias when you visit a region where the relevance of price to quality is high (the "high-relevance region"), such that the price of wine is highly predictive of wine quality. Specifically, you may under-rely on price when judging quality in the high-relevance region, leading to an underestimation of the difference in the quality of cheap wines and expensive wines.

It is important to note that the above example assumes that people usually encounter the cue variable when it is moderately relevant (as was true of the relevance of a wine's price to its quality in the moderate-relevance region). If people usually encounter the cue variable when it is highly relevant or irrelevant, then the relevance insensitivity theory predicts that people will exhibit a bias in only one direction, not in the other. For instance, while you live in the moderate-relevance region in the above example, other wine connoisseurs live in the high-relevance region. Thus, they are accustomed to a high relevance of price to quality, so they should exhibit an over-reliance bias when visiting both the moderate-relevance and low-relevance regions—but they will never exhibit an under-reliance bias. By the same logic, wine connoisseurs in the low-relevance region are used to a low relevance of price to quality, so they should display an under-reliance bias when visiting both the moderate-relevance and high-relevance regions, but they will never exhibit an over-reliance bias. In short, a bias may not always have a reverse bias.

Our prior work suggests that many cue variables have moderate relevance in usually encountered situations. All four classic biases investigated in Hsee et al. (2019)—the sunk cost fallacy, non-regressive prediction, anchoring effect, and base rate neglect—have reverse biases, indicating that the usually encountered relevance of the cue variables is neither very high nor very low. Similarly, we suspect that many other biases also have reverse biases on the other end of the relevance spectrum. In this review, we use the relevance insensitivity theory to explain many other biases and predict their potential reverse biases. Future research is needed to investigate whether the reverse biases indeed exist.

The notion of relevance insensitivity builds on several streams of existing literature. First, relevance insensitivity is a manifestation of situation insensitivity (Yang et al., 2021). Human behaviors are adapted but not perfectly adaptive (Arkes & Ayton, 1999; Bar-Anan et al., 2006; Hsee et al., 2013; Klayman & Brown, 1993). People learn to respond appropriately to the situations they encounter often, but people are not highly sensitive to variations in situations. When people encounter a situation that differs from their usually encountered situation, they may not be agile enough to adapt their response fully. Instead, people continue to behave as if they were, more or less, in their usually encountered situation. We call this tendency *situation insensitivity* (Yang et al., 2021). People lack sensitivity to many situational factors. For instance, people are not sensitive to the amount of personal control they have over an outcome, so they overestimate their control in situations in which they have little control, and they underestimate their control in situations in which they have plenty of control (Gino et al., 2011; Langer, 1975). Our relevance insensitivity theory focuses on one general type of situational factor: the relevance of a cue variable to a target variable.

Second, relevance insensitivity is a type of overgeneralization—the tendency to overgeneralize from one situation to another situation. Numerous authors have mentioned overgeneralization when explaining their findings (Arkes & Ayton, 1999; Bolton & Alba, 2012; Chinander & Schweitzer, 2003; Eidelman et al., 2010; Haws et al., 2016; Hsee et al., 2015; Kahneman & Tversky, 1973; Park et al., 2022; Raghunathan et al., 2006; Risen et al., 2007; Yang et al., 2013; Yeung & Soman, 2007; Zhu et al., 2019). For example, Tversky and Kahneman (1974) argue that judgment heuristics such as the availability and representativeness heuristics are overgeneralized

TABLE 1 Summary of known biases and proposed new biases

Cue variable	Target variable	Level of relevance	Bias
Anchor	Target number	Low	Anchoring effect (Epley & Gilovich, 2006; Tversky & Kahneman, 1974)
		High	Reverse bias (Hsee et al., 2019)
Past cost	Future cost	Low	Sunk cost fallacy (Arkes & Ayton, 1999)
		High	Reverse bias (Hsee et al., 2019)
Past outcome	Future outcome	Low	Non-regressive prediction (Kahneman & Tversky, 1973)
		High	Reverse bias (Hsee et al., 2019)
Base rate	Target probability	Low	Reverse bias (Hsee et al., 2019)
		High	Base rate neglect (Kahneman & Tversky, 1973)
Specifications	Consumption utility	Low	Specification-seeking (Hsee et al., 2009)
		High	To be tested
Effort	Quality	Low	Effort heuristic, duration heuristic, input bias, labor illusion, etc. (Buell & Norton, 2011; Chinander & Schweitzer, 2003; Kim & Labroo, 2011; Kruger et al., 2004; Yeung & Soman, 2007)
		High	To be tested
Price	Quality or healthiness	Low	The more expensive = the higher quality intuition (Cronley et al., 2005; Kardes et al., 2004), the healthy = expensive intuition (Haws et al., 2016)
		High	Diversification bias (Read & Loewenstein, 1995)
Current preferences	Future preferences	Low	Projection bias (Loewenstein et al., 2003)
		High	To be tested
Duration (of existence)	Quality	Low	Longer is better heuristic (Eidelman et al., 2010)
		High	To be tested
Category boundary crossing	Distance between two stimuli	Low	Irmak et al. (2011); Donnelly et al. (2022); Thomas and Morwitz (2005)
		High	To be tested
Product name	Risk	Low	Yang et al. (2013)
		High	To be tested
Healthiness	Tastiness	Low	The unhealthy = tasty intuition (Raghunathan et al., 2006)
		High	To be tested
One algorithm	Another algorithm	Low	Zhang and Yang (2022); Longoni et al., (forthcoming)
		High	To be tested
Usage frequency	Consumption utility	Low	To be tested
		High	Usage frequency neglect (Goodman & Irmak, 2013; Mittelman et al., 2020)
Option availability	Utility	Low	Shin and Ariely (2004)
		High	To be tested
Urgency	Utility	Low	Mere urgency effect (Zhu et al., 2018)
		High	To be tested
Reactions	Utility	Low	Mere reaction effect (Hsee et al., 2016)
		High	To be tested
Deadline length	Goal difficulty	Low	Mere deadline effect (Zhu et al., 2019)
		High	To be tested
Observing others' performance	Skill acquisition	Low	Illusion of skill acquisition (Kardas & O'Brien, 2018)
		High	To be tested

(Continues)

TABLE 1 (Continued)

Cue variable	Target variable	Level of relevance	Bias
Initial experience	Subsequent experience	Low	Billeter et al. (2010)
		High	To be tested
One's own feelings	Others' feelings	Low	Desensitization bias (Campbell et al., 2014)
		High	To be tested

rules of thumb. Likewise, Arkes and Ayton (1999) attribute the sunk cost fallacy to an overgeneralized waste-not rule. However, overgeneralization is a rather loose and vague concept. People may overgeneralize myriad things including observations, responses, heuristics, and the relevance of a cue variable. Relevance insensitivity is a specific form of overgeneralization in which the "thing" that people overgeneralize is the relevance (i.e., the weight) of a cue variable to a target variable.

Third, relevance insensitivity is a type of regression in that people behave as if the relevance of a cue variable is less extreme than it actually is (Fiedler & Unkelbach, 2014). The typical regression phenomenon occurs when the response scale is bounded, and there is noise in the response, such that the observed responses regress toward the middle of the response scale. For example, when people are asked to estimate an extreme probability, their responses tend to regress toward 50%. The regression effect occurs because the probability scale is bounded by 0% and 100%. If the actual probability (normative benchmark) is low, perhaps 1%, then there is more room for people to err by overestimating the actual probability than by underestimating it; likewise, if the actual probability is high (e.g., 99%), then there is more room for people to underestimate it than to overestimate it. Various cognitive biases have been explained by endpoint-induced regression (Costello & Watts, 2014; Erev et al., 1994; Hilbert, 2012). The relevance insensitivity theory is both more general and more specific than endpoint-induced regression. Our theory is more general because, as we explain in the next section, relevance insensitivity can explain biases that occur when the response scale has no specific endpoints. Our theory is more specific because we argue that people's responses regress not toward the midpoint of a bounded scale but rather toward responses in usually encountered situations.

3 | REVIEW, INTEGRATION, AND PREDICTION

People base their judgments and decisions on a variety of cues: anchors, past costs, numerical product specifications, price, the amount of effort needed to make the products, and so on. These cues often are somewhat relevant to the target variables, but their relevance fluctuates greatly from situation to situation—a cue might be highly relevant in some situations but completely irrelevant in other situations. Our theory predicts that people are insensitive to cue relevance: They over-rely on the cue when it is irrelevant and under-rely on it when it is highly relevant.

In what follows, we review some well-documented biases in the judgment and decision-making literature, explain how they fit our theory, and predict when they may attenuate or even reverse (Table 1). Despite their apparent disparateness, all biases reviewed below are consistent with our theory that many judgments and decisions are either influenced by irrelevant information (i.e., information that is moderately relevant in most situations but is irrelevant in the present situation) or are not sufficiently influenced by highly relevant information (i.e., information that is moderately relevant in most situations but is highly relevant in the present situation). In other words, all the documented biases we review can be characterized as either an over-reliance on information that is presently irrelevant or an under-reliance on information that presently is highly relevant—so the biases can be attributed to relevance insensitivity.

We wish to make two clarifications regarding the purpose of this review. First, it is intended to be illustrative, not exhaustive. Second, we do not aim to rule out or discredit other explanations for the biases we review. We recognize that a bias can be explained at different levels and by different accounts. Our review offers a level of analysis that is general enough to encompass many seemingly disparate phenomena and specific enough to produce testable predictions.

3.1 | Anchors

One of the best-known psychological biases is the anchoring effect, in which a person's judgment of a target variable (typically a numerical one) is influenced by an externally given or self-generated number, the "anchor." First discovered by Tversky and Kahneman (1974), the anchoring effect has been studied extensively in many contexts (Ariely et al., 2003; Jung et al., 2016; Simonson & Drolet, 2004; Stewart, 2009). A typical finding is that people estimate or predict higher numbers if they encounter a high (vs. low) anchor in the decision process, even when the anchor is clearly irrelevant to the actual judgment, such as a number generated by a wheel of fortune (Tversky & Kahneman, 1974) or the last two digits of one's social security number (Ariely et al., 2003).

Explanations for the anchoring effect include anchoring and adjustment (Tversky & Kahneman, 1974), selective accessibility (Strack & Mussweiler, 1997), scale distortion (Frederick & Mochon, 2012), and noise (Lee & Morewedge, 2022), but relevance insensitivity also can explain the anchoring effect. We argue that people are insensitive to the relevance of the anchor (i.e., the cue variable) to the numerical judgment (i.e., the target variable). Specifically, we predict that if an anchor has little to no relevance to the numerical judgment, people will over-rely on the anchor, exhibiting the typical anchoring effect. If an anchor is highly relevant to the target judgment, however, then we predict that people may under-rely on the anchor, exhibiting a reverse bias.

Hsee et al. (2019) empirically tested these hypotheses. In one study, participants were told that 100 men and 100 women were attending a dance party. The hostess of the party paired one man with one woman either randomly or according to their height (i.e., the tallest man with the tallest woman, the second tallest man with the second tallest woman, and so on). Participants were asked to predict the height of the man (i.e., the target variable) in a certain pairing based on the height of his dance partner, who was either 4 ft 11 in. or 5 ft 11 in. (i.e., the anchor). The results align with the relevance insensitivity account: When the height of the woman was unrelated to the height of the man (i.e., random pairing), participants over-relied on the height of the woman when predicting the height of the man, replicating the typical anchoring effect. Conversely, when the height of the woman was perfectly correlated with the height of the man (i.e., pairing by height), participants under-relied on the height of the woman, exhibiting the reverse bias.

3.2 | Past cost

A sunk cost is a cost that was incurred in the past and cannot be recovered (e.g., a nonrefundable ticket and money spent on a project). Rationally speaking, people should ignore sunk costs when making decisions, but many people fail to do so, falling prey to the sunk cost fallacy (Arkes & Ayton, 1999; Thaler, 1999).

The sunk cost fallacy has been explained by loss aversion (Kahneman & Tversky, 1984), self-justification (Staw, 1976), and—of the most relevance to our theory—the misapplication of the norm against waste (Arkes & Ayton, 1999). According to our theory, the sunk cost fallacy occurs because many people learn that past costs usually are moderately relevant to future costs, but people are insensitive to situational variation in the relevance. On the one hand, when the past cost is irrelevant to future costs (i.e., irrecoverable), people fail to fully ignore the past cost when making decisions, so they exhibit the typical sunk cost fallacy. On the other hand, when the past cost is highly predictive of future costs (i.e., fully recoverable), people fail to fully incorporate the past cost into their decisions, so they exhibit a reverse bias.

Hsee et al. (2019) attempted to elicit the reverse bias with a scenario in which participants had just signed a lease on an apartment and then found a different apartment that they liked better. All participants imagined that they had paid a $300 fee for the original apartment, and the fee was either a nonrefundable one-time processing fee (i.e., irrelevant to future costs) or a deposit that would be counted toward their rent if they moved in as planned (i.e., highly relevant to future costs). Participants needed to decide how much they would be willing to pay for the new apartment. When the fee was nonrefundable (irrelevant), participants failed to ignore the fee, replicating the sunk cost fallacy. When the fee was a recoverable deposit (highly relevant), however, participants did not fully incorporate it into their willingness to pay for the new apartment. In other words, when the relevance of the cue variable was high, participants exhibited the reverse sunk cost fallacy as predicted by the relevance insensitivity account.

3.3 | Past outcome

One common statistical phenomenon is the regression toward the mean effect, in which extreme outcomes are often followed by moderate outcomes (Kahneman & Tversky, 1973). Despite the prevalence of the phenomenon, people often have trouble recognizing it: they tend to underpredict the extent to which extreme outcomes will regress toward the mean (Kahneman, 2011; Kahneman & Tversky, 1973). The bias has been explained by the representativeness heuristic and people's tendency to come up with unjustified speculations.

The underprediction of regression is also compatible with our relevance insensitivity theory. We argue that people underpredict regression because they are not sensitive to the relevance of past outcomes (i.e., the cue variable) to future outcomes (i.e., the target variable). Indeed, Hsee et al. (2019) both replicated the underprediction bias and demonstrated the reverse: When past outcomes were not diagnostic of future outcomes (low relevance), participants over-relied on past outcomes, underpredicting the regression. However, when past outcomes were highly diagnostic of future outcomes, participants under-relied on past outcomes, overpredicting the regression.

3.4 | Base rate

When judging the probability of uncertain events, people tend to focus on individuating information and neglect the base rate. In a classic demonstration of base rate neglect, participants estimated the probability that a person was an engineer, and they relied heavily on whether the person's personality sketch was consistent with engineer stereotypes but not on whether the base rate of engineers was 30% or 70% (Kahneman & Tversky, 1973).

Base rate neglect has been attributed to the use of the representativeness heuristic, but we argue that relevance insensitivity may also play a role: People are insensitive to the relevance of the base rate to the target judgment. Thus, they under-rely on the base rate when it is highly relevant to the target judgment (base rate neglect) but over-rely on the base rate when it is not very relevant (a reverse bias). We empirically tested both hypotheses and found supporting evidence in Hsee et al. (2019).

3.5 | Specifications

Virtually, all consumer products carry quantitative specifications (e.g., the resolution of a camera and the number of airbags in a massage chair). Hsee et al. (2009) found that consumers rely on specifications even when the specifications are spurious and are no more informative than the consumer's direct experience. In one study, for example, participants were asked to generate specifications based on their actual experiences, so the specifications could not possibly provide additional information beyond their experiences. Specifically, participants drew two circles to represent the softness of two towels and then generated softness specifications by either estimating the diameters of the two circles (e.g., 3 and 2 cm) or the areas of the two circles (e.g., 7.07 and 3.14 cm^2). While neither specification offered information beyond participants' experiences, participants who estimated areas (such that the products had a larger difference in the specifications) were more likely to choose the softer towel than participants who estimated diameters (a smaller difference). The result is evidence of "specification seeking": Participants relied on specifications even though they were irrelevant to the task at hand.

We hypothesize that the specification-seeking effect occurs not because consumers generally overweight specifications but because consumers are insensitive to the relevance of specifications (i.e., the cue variable) to consumption utility (i.e., the target variable). In Hsee et al. (2009), specifications were not predictive of consumption utility (i.e., relevance was low), and consumers over-relied on specifications. We predict that when specifications are highly predictive of consumption utility (i.e., relevance is high), consumers will under-rely on specifications, such that they underestimate the difference in consumption utility between products with low and high numerical specifications.

3.6 | Effort

Consumers often evaluate products and services based on the amount of effort or time it takes to produce the product or render the service (Buell & Norton, 2011; Chinander & Schweitzer, 2003; Kim & Labroo, 2011; Kruger et al., 2004; Yeung & Soman, 2007). Prior research has yielded converging evidence that consumers over-rely on effort, such that consumers are biased in favor of products and services associated with more effort even when the effort is irrelevant (Kim & Labroo, 2011) or counterproductive (Buell & Norton, 2011; Yeung & Soman, 2007). In one study by Yeung and Soman (2007), participants were asked to judge an unlocking service. Rationally speaking, a faster unlocking service is more efficient and probably better than a slower service. However, participants judged the price of the slower unlocking service as more reasonable than the price of the faster service because participants relied on effort (time) as a cue.

Consistent with the relevance insensitivity theory, some researchers have reasoned that the "more-effort-is-better" effect occurs because people over-rely on inputs as a proxy for outcomes (Chinander & Schweitzer, 2003; Kim & Labroo, 2011; Kruger et al., 2004; Yeung & Soman, 2007). We agree—it seems that consumers are insufficiently sensitive to the relevance of effort (the cue variable) to quality (the target variable), so consumers over-rely on effort to infer quality when quality does not depend on effort (i.e., low relevance), demonstrating the "more-effort-is-better" effect. We also predict the reverse: When quality heavily depends on effort (i.e., high relevance), people may under-rely on effort to predict quality. For example, the taste of a Chinese old-mother-chicken soup depends heavily on how long it is cooked, such that the relevance of effort to quality is higher than in usually encountered situations. Thus, we would expect people to under-rely on effort (time) when predicting the taste of the soup.

3.7 | Price

The relevance of a product's price to its attributes, such as quality and healthiness, is context specific. Previous work has found that even when price is not related to quality, consumers still perceive more expensive products as higher-quality than cheaper products (Cronley et al., 2005; Kardes et al., 2004). Similarly, people assume that more expensive products are healthier even when price is not associated with healthiness (Haws et al., 2016). In both cases, consumers demonstrate an over-reliance on the product's price when judging its attributes.

The price–quality and price–health heuristics documented in previous literature are consistent with our theory: People are insensitive to the relevance of price (i.e., the cue variable) to product attributes such as quality and healthiness (i.e., the target variable). When price is irrelevant to quality and healthiness, people over-rely on price information, as documented in the literature. We also predict that when price is highly relevant, people may under-rely on price information when judging quality and healthiness. For example, the price of construction materials (e.g., carbon steel, stainless steel, and alloy) are highly predictive of quality (e.g., toughness and resistance to heat and corrosion), but people may underestimate the quality difference between materials with different prices.

3.8 | Current preferences

When predicting future preferences, people tend to anchor on their current preferences, believing (often incorrectly) that their future preferences will be similar (Gilbert et al., 2002; Loewenstein et al., 2003; Sayette et al., 2008; VanEpps et al., 2016). In one study, people were asked to choose between two snacks—one healthy (fruit) and the other unhealthy (junk food), and they were informed that they would receive the chosen snack in a week. Normatively speaking, the participant's current hunger should not influence their choice of a snack to be consumed in a week. However, results showed a "projection bias": Participants who made their choice in a hungry state were more likely to choose the unhealthy snack than those who made their choice in a satisfied state (Read & van Leeuwen, 1998).

The projection bias has been attributed to insufficient adjustment (Gilbert et al., 2002) and the hot–cold empathy gap (Loewenstein, 2005), but it also is consistent with the relevance insensitivity theory. We propose that the projection bias occurs because people are insensitive to the relevance of their current preferences (i.e., the cue variable) to their future preferences (i.e., the target variable). People over-rely on current preferences when they are inconsistent with future preferences, as documented in Loewenstein et al. (2003), and we predict that people will under-rely on current preferences (a reverse projection bias) when current preferences are highly consistent with future preferences. Indeed, research has found that when choosing what to consume in the future, consumers may choose too little of their favorite option due to an erroneous belief that they will become satiated with it (Read & Loewenstein, 1995).

3.9 | Duration (of existence)

People tend to evaluate an object or a practice more positively when the object/practice has existed for a longer duration, even when duration is unrelated to quality (Eidelman et al., 2010; Kupor et al., 2021). For example, in one study by Eidelman et al. (2010), participants viewed a relatively unknown painting and judged its aesthetics. Participants who were informed that the painting was from 1905 rated it as more aesthetically pleasing than participants who were informed that the painting was from 2005.

Consistent with our theory, Eidelman et al. (2010) attributed their effect to the overapplication of the "longer is better" heuristic. More generally, we propose that the bias occurs because people are insensitive to the relevance of the duration of existence (i.e., the cue variable) to quality (i.e., the target variable). Eidelman et al. (2010) found an over-reliance bias when duration was unrelated to quality, and we predict that when duration is highly predictive of quality, people may under-rely on duration when judging quality. For example, the longevity of hedge funds is highly predictive of performance because hedge funds with bad performance usually do not survive for long in such a competitive market. We predict that people may under-rely on longevity when predicting the performance of hedge funds.

3.10 | Category boundary crossing

Often, stimuli in different categories are more different than stimuli within the same category, so people tend to infer the difference between two stimuli based on their category membership—even when the category membership is irrelevant. For example, a pair of cities in different states are perceived as farther apart (e.g., Sacramento, CA, and Eugene, OR) than a pair of cities in the same state (e.g., Sacramento, CA, and San Diego, CA) even though the actual distance is the same (Irmak et al., 2011). Time periods feel longer when they cross more time categories (e.g., from 1:45 p.m. to 3:15 p.m. vs. from 1:15 p.m. to 2:45 p.m.; Donnelly et al., 2022). Number pairs seem more different if they have different leftmost digits (e.g., 299 and 300) than if they share the leftmost digit (e.g., 298 and 299; Sokolova et al., 2020; Thomas & Morwitz, 2005).

We argue that these biases occur because people are insensitive to the relevance of category boundary crossing (i.e., the cue variable) to the distance between two stimuli (i.e., the target variable). Category boundary crossing was irrelevant in the aforementioned studies, so the studies demonstrated an over-reliance bias. We predict that an under-reliance bias will occur when category boundary crossing is highly diagnostic. For example, suppose that people are asked to estimate the distance between Sacramento, CA, and Miami, FL, and the distance between Sacramento, CA, and San Francisco, CA. Then, the fact that Sacramento and Miami are in different states while Sacramento and San Francisco are in the same state is a highly valid cue that the former distance is much greater than the latter. In this case, we predict that people will under-utilize the cue of category boundary crossing and will underestimate the difference between the two distances.

3.11 | Product name

The name of a product often, but not always, provides hints about its characteristics and value. Research has shown that consumers consider a product's name even when it is not very informative. For example, consumers have a significantly lower willingness to pay for a probabilistic good if it is labeled a "lottery ticket," "raffle," "coin flip," or "gamble" (vs. "gift certificate" or "voucher") because the names are associated with high (vs. low) risk (Yang et al., 2013).

From the perspective of the relevance insensitivity theory, the cue variable is the product name, and the target variable is the actual risk associated with the probabilistic good. In line with our theory, the authors attributed the labeling effect to an over-reliance on the name of a probabilistic good as a proxy for the risk associated with the good. We believe that an under-reliance bias also is possible if the product name is highly diagnostic of the product risk. For instance, the names of casino games are highly diagnostic of the odds of winning: Blackjack has the best odds of winning (with a house edge of 1% in most casinos), whereas slot machines have the worst odds of winning (with a house edge of 10% or more; Matarese, 2018). However, the increasing popularity of slot machines suggests that consumers may not use the names of casino games as proxies for the odds of winning.

3.12 | Healthiness

Consumers learn through personal experience and observations that unhealthy food often tastes better than healthy food. Overgeneralizing this experience, consumers judge the same food item as tastier when it is portrayed as less healthy (Raghunathan et al., 2006). In one study, consumers were presented with two new brands of cheddar-flavored snack crackers and rated the tastiness of each cracker. Both crackers contained 13 g of fat per serving, but one cracker contained more unsaturated fat (generally considered healthy), while the other cracker contained more saturated fat (unhealthy). Although the two crackers should have tasted the same given that they had the same total fat content, consumers reported that the cracker with more saturated fat was tastier than the one with more unsaturated fat.

The explanation that people tend to overgeneralize the "unhealthy = tasty" intuition (Raghunathan et al., 2006) is compatible with the relevance insensitivity theory: People are insensitive to the relevance of healthiness (i.e., the cue variable) to tastiness (i.e., the target variable). While Raghunathan et al. (2006) documented an over-reliance bias when healthiness is not predictive of tastiness, we predict that an under-reliance bias will occur when healthiness is highly predictive of tastiness, such that consumers may underestimate the difference in tastiness between healthy and unhealthy food items.

3.13 | Algorithmic bias

Zhang and Yang (2022) investigated the consequences of raising public awareness of racial bias in algorithms. They found that after learning from the news that some algorithms exhibit racial bias, people overgeneralize the information and behave as if other algorithms are also racially biased. Consequently, raising awareness of algorithmic racial bias can deter Black consumers, but not White consumers, from using "good" (i.e., fair and beneficial) algorithms. Relatedly, Longoni et al., (forthcoming) found that people are more likely to generalize algorithmic failures than human failures.

The authors theorized that most consumers have a poor understanding of how algorithms work, so they tend to overgeneralize any knowledge they gain about one algorithm. We agree—it seems that people are insensitive to the relevance of one algorithm (i.e., the cue variable) to another algorithm (i.e., the target variable). In Zhang and Yang (2022), consumers learned about an algorithm (racially biased) that was not at all diagnostic of the target algorithm (fair and beneficial) and then over-relied on their knowledge when judging the target algorithm. We also predict that people may under-rely on their knowledge about a biased algorithm when the biased algorithm is highly diagnostic of the target algorithm, such that the biased and target algorithms use the same methodology, model specifications, training data, and so on.

3.14 | Usage frequency

Previous literature shows that when consumers decide whether to purchase a product, they overlook how frequently they will use the product. Thus, consumers undervalue a durable product with high usage frequency (Mittelman et al., 2020) and overvalue a multi-feature product, of which some features have low usage frequency (Goodman & Irmak, 2013). The authors attribute "usage frequency neglect" to the non-accessibility of the frequency information at the choice stage (Mittelman et al., 2020) and a failure to distinguish between *having* product features and *using* product features (Goodman & Irmak, 2013).

According to our theory, however, usage frequency neglect occurs because people are insensitive to the relevance of a product's usage frequency (i.e., the cue variable) to its consumption utility (i.e., the target variable). In Goodman and Irmak (2013) and Mittelman et al. (2020), usage frequency was highly diagnostic of consumption utility, so consumers under-relied on usage frequency. We also predict an over-reliance bias when usage frequency is not very diagnostic of consumption utility. For example, many homeowners do not have fire extinguishers, perhaps because they do not think they will ever need to use one. In other words, when evaluating the consumption utility of a fire extinguisher, homeowners over-rely on the fact that the usage frequency is extremely low.

3.15 | Option availability

People often exert effort to hold onto multiple attractive options (e.g., college majors and romantic partners). Shin and Ariely (2004) conducted a series of experiments using "door games" and found that people invest too much to hold onto options even when the options are not appealing.

Our theory argues that the underlying cause of the effect is that people are insensitive to the relevance of option availability (i.e., the cue variable) to utility (i.e., the target variable). When keeping options available does little to improve the payoff (i.e., relevance is low), consumers waste effort keeping options available, as documented in Shin and Ariely (2004). We also predict that when keeping options available substantially enhances the payoff (i.e., relevance is high), people may not try as hard as they should to keep options available. You probably have heard someone lament their mistake of not holding a stock for long enough.

3.16 | Urgency

The urgency of a task often is moderately diagnostic of its consequences, so consumers may learn to use task urgency as an input when prioritizing tasks. Zhu et al. (2018) found that people seek urgent tasks even when the only consequence of doing so is a suboptimal financial payoff. Specifically, the authors found that online workers were more likely to perform lower-paying tasks over comparable higher-paying tasks when the lower-paying tasks were more urgent (e.g., would expire in 10 min vs. in 24 h). The "mere urgency effect" occurred even when the urgency was illusory and irrelevant (e.g., when everyone could easily complete the task within 5 min).

While the authors argued that the mere urgency effect happened because urgency drew people's attention to the task's time dimension instead of its outcome dimension, we argue that people are insensitive to the relevance of a task's urgency (i.e., the cue variable) to its utility (i.e., the target variable). The urgent task was not associated with a better payoff (i.e., relevance was low) in Zhu et al. (2018), so the likelihood of choosing the urgent task was irrationally high. We expect that people may not put enough effort into the urgent task when task urgency has a strong positive association with the payoff (i.e., relevance is high). For instance, people may decide not to apply to a dream school or dream job for which the deadline is approaching, thereby missing out on the opportunity to realize their full potential.

3.17 | Reactions

Decades of research show that a reaction can reinforce an action if the reaction is a reward (i.e., a priori positive) or carries useful information (Baumeister, 1998; Skinner, 1953). Recently, however, Hsee et al. (2015) found that a reaction can also reinforce the action even if the reaction is a priori non-positive and non-informative. For example, in a field study, the authors ran a donation campaign for homeless cats. People could donate by inserting coins into a special donation box, and people donated more when the donation box emitted an unpleasant sound (vs. no sound) when a coin was inserted.

The authors attributed the "mere reaction effect" to overgeneralization: People internalize the tendency to seek reactions because reactions often (but not always) carry useful information. Similarly, our theory posits that people are insensitive to the relevance of a reaction (i.e., the cue variable) to utility (i.e., the target variable). When reactions offer no benefit (i.e., relevance is low), people seek too many reactions, as shown in Hsee et al. (2015). We also predict that when reactions offer considerable benefit (i.e., relevance is high), people may fail to seek enough reactions, exhibiting the reverse bias. Consistent with this prediction, research in product design has shown that although consumer reactions (e.g., user input and insights) are highly beneficial for product design, designers unfortunately tend not to solicit enough consumer feedback (Norman, 2013).

3.18 | Deadline length

In many situations, more difficult goals have longer deadlines, so people may learn to use deadline length as a predictor of goal difficulty. Zhu et al. (2019) demonstrated that people commit more resources to goals with longer deadlines even when deadline length is incidental (i.e., not relevant to the goal's difficulty or normative resource requirement). For example, in one study, participants imagined that they were organizing an afternoon tea party for their high school best friend, who would visit in either 3 months (the longer deadline) or 1 month (the shorter deadline). Participants with the longer (vs. shorter) deadline expected hosting the party to be more challenging—so they decided to spend more money on it.

Zhu et al. (2019) attributed the "mere deadline effect" to the overapplication of the rule that "a longer deadline = a more difficult task." It seems that people are insensitive to the relevance of the deadline length (i.e., the cue variable) to goal difficulty (i.e., the target variable). In Zhu et al. (2019), the deadline length was not diagnostic of goal difficulty, so people over-relied on the deadline length to judge goal difficulty. We expect that an under-reliance bias will occur when the deadline length is highly diagnostic of goal difficulty. When people need to complete either one project in 1 month or three projects in 3 months, they may not allocate three times as many resources in the latter case as in the former.

3.19 | Observing others' performances

Kardas and O'Brien (2018) found that people who repeatedly watch others perform a skill (e.g., throwing darts, playing online games, and dancing) over-predict their own performance. In one experiment, participants watched a training video of moonwalking either one time (low exposure) or 20 times (high exposure) and then attempted to moonwalk without any practice. Actual performance (as rated by external observers) did not differ between conditions, but participants with high (vs. low) exposure were more confident in their performance.

The authors reasoned that the "illusion of skill acquisition" happens because people do not sufficiently account for the gap between seeing (knowing what the performance looks like) and doing (knowing what the performance feels like). Our theory offers a slightly different account: Observation is moderately helpful for skill acquisition in usually-encountered situations, and people are insensitive to situational variation in the relevance of observation (i.e., the cue variable) to skill acquisition (i.e., the target variable). In Kardas and O'Brien (2018), observing a skill provided little help with skill acquisition, so participants over-relied on the number of observations to predict their own performance. When observing a skill is very helpful for acquiring the skill (e.g., design a website using a website builder), we expect people to under-rely on the number of observations to predict their own performance.

3.20 | Initial experience

Many technology products require skill to use, and Billeter et al. (2010) discovered that people who briefly try a new product underestimate how quickly they could master the skills required for fluent product use. In one study, participants used a new keyboard to complete two rounds of a typing task. After finishing the first round, participants predicted their performance in the second round before completing the round—and their predictions were worse than their actual performance. The authors argued that underestimation happened because unskilled users have difficulty imagining being skilled.

We propose that the bias occurs because people are insensitive to the relevance of their initial experience (i.e., the cue variable) to their subsequent experience (i.e., the target variable). When an initial experience is not diagnostic of subsequent experiences, people over-rely on the initial experience. In Billeter et al. (2010), the initial experience was difficult, so participants underpredicted their speed of learning. (Over-reliance also could manifest as an overprediction of the speed of learning when the initial experience is easy.) We also predict that an under-reliance bias is possible: when an initial experience is highly diagnostic of subsequent experiences, people may under-rely on the initial experience, so they may overpredict (underpredict) their speed of learning when their initial experience is difficult (easy).

3.21 | One's own feelings

People often anticipate the reactions of others in social interactions, and often (but not always), one's own reaction to an experience is a reasonable barometer for others' reactions to the same experience. Campbell et al. (2014) found that repeated exposure to an evocative event (e.g., an image and a joke) leads to an under-prediction of the emotional intensity of first-time viewers. For example, participants were asked to copy a funny joke either once or five times. Afterward, participants rated how funny they found the joke and predicted how funny someone else would find the joke. Participants who had copied the joke five times (vs. only once) not only thought the joke was less funny but also predicted that the joke would be less funny for others.

The authors attributed the "desensitization bias" to an insensitivity to one's own desensitization and a lack of perspective-taking. The relevance insensitivity theory offers a more general explanation: people are insensitive to the relevance of their own feelings (i.e., the cue variable) to others' feelings (i.e., the target variable). In Campbell et al. (2014), after copying the joke five times, participants' apathetic feelings about the joke were not very diagnostic of the feelings of the hypothetical others who would read the joke for the first time. Thus, participants over-relied on their own feelings, leading to an underestimation of the intensity of others' feelings. (If participants' feelings were intense rather than apathetic, we would expect an overestimation of the intensity of others' feelings.) We also predict that an under-reliance bias will occur when a person's feelings are highly diagnostic of others' feelings, such that people will overestimate (underestimate) the intensity of others' feelings when their own feelings are apathetic (intense).

4 | CONCLUSION

This review uses the relevance insensitivity theory proposed by Hsee et al. (2019) to analyze and integrate ostensibly unrelated biases in consumer behavior and beyond. The theory is general enough to encompass and integrate a wide range of biases yet specific enough to identify the cue variable and target variable in each bias, explain the bias, and predict the opposite bias.

Notably, the biases reviewed above are just examples. Other biases that fit the relevance insensitivity framework include the numerosity heuristic (Pelham et al., 1994), neglect of marginal utility (Li & Hsee, 2021), overcommunicating upward counterfactual information (Li et al., forthcoming), the confidence heuristic (Price & Stone, 2004), ease-of-retrieval effect (Schwarz et al., 1991), affect-as-information effect (Schwarz & Clore, 1983), and biased assimilation effect (Lord et al., 1979). For example, in the affect-as-information effect, people incorporate their momentary affective state into judgments. The relevance insensitivity framework posits that people are insensitive to the relevance of their current mood (i.e., the cue variable) to the judgment at hand (i.e., the target variable). We predict that people over-rely on mood when it is irrelevant and under-rely on mood when it is highly relevant.

While reviewing the biases, we noticed an interesting pattern in the literature: The vast majority of documented biases involve an over-reliance on irrelevant cues rather than an under-reliance on highly relevant cues. We speculate that this pattern occurs because it is easier for researchers to construct situations in which a cue is normatively irrelevant than to construct situations in which a cue is highly relevant (more relevant than usual). For example, it is easy to find situations in which anchor numbers are normatively irrelevant (e.g., spinning a wheel), but it is more difficult to find situations in which anchor numbers are of unusually high relevance.

It is also important to note that a vast body of literature in social and consumer psychology has shown that people's judgments, perceptions, and behaviors are susceptible to contextual cues. While a hypersensitivity to contextual cues may seem contradictory to the notion of an insensitivity to relevance, the two phenomena are actually compatible. For all context effects that involve a hypersensitivity to normatively irrelevant contextual cues, we offer a general mechanism: Many contextual cues are moderately relevant in usually encountered situations, so people over-rely on cues when they are irrelevant (as documented in the context effect literature), and we predict that people will under-rely on cues when they are highly relevant.

Biases are like islands in a sea—some big, some small, some above the water, and some under the water. Our model is far from able to cover all islands in the sea, but it offers a way to map and integrate at least a chunk of the islands that have been considered isolated and disparate, and it has the potential to uncover more islands that remain under the sea.

ORCID

Christopher K. Hsee https://orcid.org/0000-0002-8391-7847

REFERENCES

Ariely, D., Loewenstein, G., & Prelec, D. (2003). "Coherent arbitrariness": Stable demand curves without stable preferences. *The Quarterly Journal of Economics*, *118*(1), 73–106. https://doi.org/10.1162/00335530360535153

Arkes, H. R., & Ayton, P. (1999). The sunk cost and Concorde effects: Are humans less rational than lower animals? *Psychological Bulletin*, *125*(5), 591–600. https://doi.org/10.1037/0033-2909.125.5.591

Bar-Anan, Y., Liberman, N., & Trope, Y. (2006). The association between psychological distance and construal level: Evidence from an implicit association test. *Journal of Experimental Psychology: General*, *135*(4), 609–622. https://doi.org/10.1037/0096-3445.135.4.609

Baumeister, R. F. (1998). "The Self." In D. T. Gilbert, S. T. Fiske & G. Lindzey *The Handbook of Social 90 Psychology* (Vols. 1-2, 4th ed.). McGraw-Hill.

Billeter, D., Kalra, A., & Loewenstein, G. (2010). Underpredicting learning after initial experience with a product. *Journal of Consumer Research*, *37*(5), 723–736. https://doi.org/10.1086/655862

Bolton, L. E., & Alba, J. W. (2012). When less is more: Consumer aversion to unused utility. *Journal of Consumer Psychology*, *22*(3), 369–383. https://doi.org/10.1016/j.jcps.2011.09.002

Buell, R. W., & Norton, M. I. (2011). The labor illusion: How operational transparency increases perceived value. *Management Science*, *57*(9), 1564–1579. https://doi.org/10.1287/mnsc.1110.1376

Campbell, T., O'Brien, E., Van Boven, L., Schwarz, N., & Ubel, P. (2014). Too much experience: A desensitization bias in emotional perspective taking. *Journal of Personality and Social Psychology*, *106*(2), 272–285. https://doi.org/10.1037/a0035148

Chinander, K. R., & Schweitzer, M. E. (2003). The input bias: The misuse of input information in judgments of outcomes. *Organizational Behavior and Human Decision Processes*, *91*(2), 243–253. https://doi.org/10.1016/S0749-5978(03)00025-6

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

Cronley, M. L., Posavac, S. S., Meyer, T., Kardes, F. R., & Kellaris, J. J. (2005). A selective hypothesis testing perspective on price-quality inference and inference-based choice. *Journal of Consumer Psychology*, *15*(2), 159–169. https://doi.org/10.1207/s15327663jcp1502_8

Donnelly, K., Compiani, G., & Evers, E. R. (2022). Time periods feel longer when they span more category boundaries: Evidence from the lab and the field. *Journal of Marketing Research*, *59*(4), 821–839. https://doi.org/10.1177/00222437211073810

Eidelman, S., Pattershall, J., & Crandall, C. S. (2010). Longer is better. *Journal of Experimental Social Psychology*, *46*(6), 993–998. https://doi.org/10.1016/j.jesp.2010.07.008

Epley, N., & Gilovich, T. (2006). The anchoring-and-adjustment heuristic: Why the adjustments are insufficient. *Psychological Science*, *17*(4), 311–318. https://doi.org/10.1111/j.1467-9280.2006.01704.x

Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over-and underconfidence: The role of error in judgment processes. *Psychological Review*, *101*(3), 519–527. https://doi.org/10.1037/0033-295X.101.3.519

Fiedler, K., & Unkelbach, C. (2014). Regressive judgment: Implications of a universal property of the empirical world. *Current Directions in Psychological Science*, *23*(5), 361–367. https://doi.org/10.1177/0963721414546330

Frederick, S. W., & Mochon, D. (2012). A scale distortion theory of anchoring. *Journal of Experimental Psychology: General*, *141*(1), 124–133. https://doi.org/10.1037/a0024006

Gilbert, D. T., Gill, M. J., & Wilson, T. D. (2002). The future is now: Temporal correction in affective forecasting. *Organizational Behavior and Human Decision Processes*, *88*(1), 430–444. https://doi.org/10.1006/obhd.2001.2982

Gino, F., Sharek, Z., & Moore, D. A. (2011). Keeping the illusion of control under control: Ceilings, floors, and imperfect calibration. *Organizational Behavior and Human Decision Processes*, *114*(2), 104–114. https://doi.org/10.1016/j.obhdp.2010.10.002

Goodman, J. K., & Irmak, C. (2013). Having versus consuming: Failure to estimate usage frequency makes consumers prefer multifeature products. *Journal of Marketing Research*, *50*(1), 44–54. https://doi.org/10.1509/jmr.10.0396

Haws, K. L., Reczek, R. W., & Sample, K. L. (2016). Healthy diets make empty wallets: The healthy= expensive intuition. *Journal of Consumer Research*, *43*(6), 992–1007.

Hilbert, M. (2012). Toward a synthesis of cognitive biases: How noisy information processing can bias human decision making. *Psychological Bulletin*, *138*(2), 211–237. https://doi.org/10.1037/a0025940

Hsee, C. K., Yang, Y., Gu, Y., & Chen, J. (2009). Specification seeking: How product specifications influence consumer preference. *Journal of Consumer Research*, *35*(6), 952–966. https://doi.org/10.1086/593947

Hsee, C. K., Yang, Y., & Li, X. (2019). Relevance insensitivity: A new look at some old biases. *Organizational Behavior and Human Decision Processes*, *153*, 13–26. https://doi.org/10.1016/j.obhdp.2019.05.002

Hsee, C. K., Yang, Y., & Ruan, B. (2015). The mere-reaction effect: Even nonpositive and noninformative reactions can reinforce actions. *Journal of Consumer Research*, *42*(3), 420–434. https://doi.org/10.1093/jcr/ucv022

Hsee, C. K., Zhang, J., Cai, C. F., & Zhang, S. (2013). Overearning. *Psychological Science*, *24*(6), 852–859. https://doi.org/10.1177/0956797612464785

Irmak, C., Naylor, R. W., & Bearden, W. O. (2011). The out-of-region bias: Distance estimations based on geographic category membership. *Marketing Letters*, *22*(2), 181–196. https://doi.org/10.1007/s11002-010-9120-3

Jung, M. H., Perfecto, H., & Nelson, L. D. (2016). Anchoring in payment: Evaluating a judgmental heuristic in field experimental settings. *Journal of Marketing Research*, *53*(3), 354–368. https://doi.org/10.1509/jmr.14.0238

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, *80*(4), 237–251. https://doi.org/10.1037/h0034747

Kahneman, D., & Tversky, A. (1984). Choices, values and frames. *American Psychologist*, *39*, 341–350. https://doi.org/10.1037/0003-066X.39.4.341

Kardas, M., & O'Brien, E. (2018). Easier seen than done: Merely watching others perform can foster an illusion of skill acquisition. *Psychological Science*, *29*(4), 521–536. https://doi.org/10.1177/0956797617740646

Kardes, F. R., Cronley, M. L., Kellaris, J. J., & Posavac, S. S. (2004). The role of selective information processing in price-quality inference. *Journal of Consumer Research*, *31*(2), 368–374. https://doi.org/10.1086/422115

Kim, S., & Labroo, A. A. (2011). From inherent value to incentive value: When and why pointless effort enhances consumer preference. *Journal of Consumer Research*, *38*(4), 712–742. https://doi.org/10.1086/660806

Klayman, J., & Brown, K. (1993). Debias the environment instead of the judge: An alternative approach to reducing error in diagnostic (and other) judgment. *Cognition*, *49*(1–2), 97–122, 122. https://doi.org/10.1016/0010-0277(93)90037-V

Kruger, J., Wirtz, D., Van Boven, L., & Altermatt, T. W. (2004). The effort heuristic. *Journal of Experimental Social Psychology*, *40*(1), 91–98. https://doi.org/10.1016/S0022-1031(03)00065-9

Kupor, D., Jia, J., & Tormala, Z. (2021). Change appeals: How referencing change boosts curiosity and promotes persuasion. *Personality and Social Psychology Bulletin*, *47*(5), 691–704. https://doi.org/10.1177/0146167220941294

Langer, E. J. (1975). The illusion of control. *Journal of Personality and Social Psychology*, *32*(2), 311–328. https://doi.org/10.1037/0022-3514.32.2.311

Lee, C.-Y., & Morewedge, C. (2022). Noise increases anchoring effects. *Psychological Science*, *33*(1), 60–75. https://doi.org/10.1177/09567976211024254

Li, X., & Hsee, C. K. (2021). The psychology of marginal utility. *Journal of Consumer Research*, *48*(1), 169–188. https://doi.org/10.1093/jcr/ucaa064

Li, X., Hsee, C. K., & O'Brien, E. (forthcoming). "It could be better" can make it worse: When and why people mistakenly communicate upward counterfactual information. *Journal of Marketing Research*.

Loewenstein, G. (2005). Hot-cold empathy gaps and medical decision making. *Health Psychology*, *24*(4S), S49–S56. https://doi.org/10.1037/0278-6133.24.4.S49

Loewenstein, G., O'Donoghue, T., & Rabin, M. (2003). Projection bias in predicting future utility. *The Quarterly Journal of Economics*, *118*(4), 1209–1248. https://doi.org/10.1162/003355303322552784

Longoni, C., Cian, L., & Kyung, E. J. (forthcoming). AI in the government: Responses to failures. *Journal of Marketing Research*.

Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, *37*(11), 2098–2109. https://doi.org/10.1037/0022-3514.37.11.2098

Matarese, J. (2018). Before you go gambling: The best and worst casino game odds. News 5 Cleveland. https://www.news5cleveland.com/before-you-go-gambling-the-best-and-worst-casino-game-odds

Mittelman, M., Gonçalves, D., & Andrade, E. B. (2020). Out of sight, out of mind: Usage frequency considerations in purchase decisions. *Journal of Consumer Psychology*, *30*(4), 652–659. https://doi.org/10.1002/jcpy.1155

Norman, D. (2013). *The design of everyday things* (Revised and expanded ed.). Basic Books.

Park, S. K., Yang, Y., & Zhang, S. (2022). Mitigating inequalities caused by awareness of algorithmic bias. *Working paper*.

Pelham, B. W., Sumarta, T. T., & Myaskovsky, L. (1994). The easy path from many to much: The numerosity heuristic. *Cognitive Psychology*, *26*(2), 103–133. https://doi.org/10.1006/cogp.1994.1004

Price, P. C., & Stone, E. R. (2004). Intuitive evaluation of likelihood judgment producers: Evidence for a confidence heuristic. *Journal of Behavioral Decision Making*, *17*(1), 39–57. https://doi.org/10.1002/bdm.460

Raghunathan, R., Naylor, R. W., & Hoyer, W. D. (2006). The unhealthy = tasty intuition and its effects on taste inferences, enjoyment, and choice of food products. *Journal of Marketing*, *70*(4), 170–184. https://doi.org/10.1509/jmkg.70.4.170

Read, D., & Loewenstein, G. (1995). Diversification bias: Explaining the discrepancy in variety seeking between combined and separated choices. *Journal of Experimental Psychology: Applied*, *1*(1), 34–49.

Read, D., & van Leeuwen, B. (1998). Predicting hunger: The effects of appetite and delay on choice. *Organizational Behavior and Human Decision Processes*, *76*(2), 189-205.

Risen, J. L., Gilovich, T., & Dunning, D. (2007). One-shot illusory correlations and stereotype formation. *Personality and Social Psychology Bulletin*, *33*(11), 1492–1502. https://doi.org/10.1177/0146167207305862

Sayette, M. A., Loewenstein, G., Griffin, K. M., & Black, J. J. (2008). Exploring the cold-to-hot empathy gap in smokers. *Psychological Science*, *19*(9), 926–932. https://doi.org/10.1111/j.1467-9280.2008.02178.x

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

Schwarz, N., & Clore, G. L. (1983). Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states. *Journal of Personality and Social Psychology*, *45*(3), 513–523. https://doi.org/10.1037/0022-3514.45.3.513

Shin, J., & Ariely, D. (2004). Keeping doors open: The effect of unavailability on incentives to keep options viable. *Management Science*, *50*(5), 575–586. https://doi.org/10.1287/mnsc.1030.0148

Simonson, I., & Drolet, A. (2004). Anchoring effects on consumers' willingness-to-pay and willingness-to-accept. *Journal of Consumer Research*, *31*(3), 681–690. https://doi.org/10.1086/425103

Skinner, B. F. (1953). *Science and Human Behavior*. New York: 45 Macmillan.

Sokolova, T., Seenivasan, S., & Thomas, M. (2020). The left-digit bias: When and why are consumers penny wise and pound foolish? *Journal of Marketing Research*, *57*(4), 771–788. https://doi.org/10.1177/0022243720932532

Staw, B. M. (1976). Knee-deep in the big muddy: A study of escalating commitment to a chosen course of action. *Organizational Behavior and Human Performance*, *16*, 27–44. https://doi.org/10.1016/0030-5073(76)90005-2

Stewart, N. (2009). The cost of anchoring on credit-card minimum repayments. *Psychological Science*, *20*(1), 39–41. https://doi.org/10.1111/j.1467-9280.2008.02255.x

Strack, F., & Mussweiler, T. (1997). Explaining the enigmatic anchoring effect: Mechanisms of selective accessibility. *Journal of Personality and Social Psychology*, *73*(3), 437–446. https://doi.org/10.1037/0022-3514.73.3.437

Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral Decision Making*, *12*(3), 183–206. https://doi.org/10.1002/(SICI)1099-0771(199909)12:3<183::AID-BDM318>3.0.CO;2-F

Thomas, M., & Morwitz, V. (2005). Penny wise and pound foolish: The left-digit effect in price cognition. *Journal of Consumer Research*, *32*(1), 54–64. https://doi.org/10.1086/429600

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

VanEpps, E. M., Downs, J. S., & Loewenstein, G. (2016). Advance ordering for healthier eating? Field experiments on the relationship between the meal order–consumption time delay and meal content. *Journal of Marketing Research*, *53*(3), 369–380. https://doi.org/10.1509/jmr.14.0234

Yang, Y., Hsee, C. K., & Li, X. (2021). Prediction biases: An integrative review. *Current Directions in Psychological Science*, *30*(3), 195–201. https://doi.org/10.1177/0963721421990341

Yang, Y., Vosgerau, J., & Loewenstein, G. (2013). Framing influences willingness to pay but not willingness to accept. *Journal of Marketing Research*, *50*(6), 725–738. https://doi.org/10.1509/jmr.12.0430

Yeung, C. W., & Soman, D. (2007). The duration heuristic. *Journal of Consumer Research*, *34*(3), 315–326. https://doi.org/10.1086/519500

Zhang, S., & Yang, Y. (2022). The unintended consequences of raising awareness: Knowing about the existence of algorithmic racial bias worsens racial inequality. *Working paper*.

Zhu, M., Bagchi, R., & Hock, S. J. (2019). The mere deadline effect: Why more time might sabotage goal pursuit. *Journal of Consumer Research*, *45*(5), 1068–1084. https://doi.org/10.1093/jcr/ucy030

Zhu, M., Yang, Y., & Hsee, C. K. (2018). The mere urgency effect. *Journal of Consumer Research*, *45*(3), 673–690. https://doi.org/10.1093/jcr/ucy008

How to cite this article: Yang, Y., Li, X., & Hsee, C. K. (2023). Relevance insensitivity: A framework of psychological biases in consumer behavior and beyond. *Consumer Psychology Review*, *6*(1), 121–132. https://doi.org/10.1002/arcp.1082