Organizational Behavior and Human Decision Processes 153 (2019) 13–26

Contents lists available at ScienceDirect

Organizational Behavior and Human Decision Processes

journal homepage: www.elsevier.com/locate/obhdp

Relevance insensitivity: A new look at some old biases

Christopher K. Hsee[a], Yang Yang[b], Xilin Li[a]

[a] Booth School of Business, University of Chicago, Chicago, IL 60637, United States
[b] Warrington College of Business, University of Florida, Gainesville, FL 32611, United States

ARTICLE INFO

Keywords:
Behavior decision theory
Heuristics and biases
Integrative theory
Noise model
Overgeneralization
Anchoring
Sunk cost
Regression toward the mean
Base rate neglect
Overconfidence

ABSTRACT

People show systematic biases in judgment and decision making. We propose that many seemingly disparate biases reflect a common underlying mechanism—insensitivity to the relevance of some given information—and that manipulating the relevance of the information can eliminate or even reverse the original bias. We test our theory in four experiments, each focusing on a classic bias—the sunk cost fallacy, non-regressive prediction, anchoring bias, and base rate neglect, and show that people over-rely on a given piece of information when it is irrelevant, thus exhibiting one bias, and under-rely on the same piece of information when it is highly relevant, thus showing a reverse bias. For example, when a past cost is irrecoverable and hence irrelevant to future cost, people over-rely on it when making a decision for the future, thus exhibiting the classic sunk cost fallacy, but when the past cost is fully recoverable and hence highly relevant to future cost, people under-rely on it, thus showing the reverse of the sunk cost fallacy. We also find that when people are made sensitive to the relevance of the information, both the original biases and their reverse biases are attenuated. This research offers a new look at these "old" biases, suggesting that each individual bias is not general because it can be reversed, but collectively, these biases are general because they all reflect relevance insensitivity.

A major contribution of psychological science is the identification of systematic biases in judgment and decision making. A hallmark of many of such biases is that people's responses deviate systematically from a normative benchmark. For example, one of the best-known biases is the sunk cost fallacy (Arkes & Ayton, 1999; Arkes & Blumer, 1985; Thaler, 1999). Normatively, people, when making a decision about an action for which they have already incurred a cost, should ignore the cost if it is irrecoverable, but in actuality, people do not fully ignore irrecoverable past costs. Another well-documented bias is base rate neglect. Normatively, people, when assessing the likelihood of an event, should take into account the event's base rate, but in reality, people often overlook base rates (Tversky & Kahneman, 1971, 1974).

Such biases seem unrelated to each other and traditionally have been attributed to different causes. For example, the sunk cost fallacy has been explained in terms of mental accounting and the use of a "waste-not" heuristic (Arkes & Blumer, 1985; Thaler, 1999), and base rate neglect has been explained in terms of the representativeness heuristic (Tversky & Kahneman, 1971, 1974).

The purpose of this research is not to refute these existing explanations, but to analyze the biases at a different level. Like many other psychological phenomena, biases can be analyzed at different levels. At a very general level, almost all biases can be attributed to bounded rationality (Simon, 1957) and heuristic processing (Kahneman, 2011; Sloman, 1996). At a very specific level, each bias has its own idiosyncratic explanation(s). In this research, we afford a level of analysis that is neither very general nor very specific. It is general enough to integrate multiple seemingly disparate biases, and it is specific enough to render new and falsifiable predictions.

1. Relevance insensitivity

To make good judgments and decisions, people should be sensitive to the relevance of the given information to the judgment or decision they need to make. Normatively, if the information is irrelevant to the focal judgment or decision, people should ignore it; if the information is perfectly relevant, they should fully incorporate it.

Our central tenet is that people are insufficiently sensitive to the relevance of their information. If the information is irrelevant, people will treat it as if it were intermediately relevant, and thus will not completely ignore it. If the information is perfectly relevant, people will also treat it as if it were intermediately relevant, and thus will not fully incorporate it. In other words, people will demonstrate apparently opposite biases depending on the actual relevance of the information. If the information is irrelevant, people will likely "over-rely" on it (i.e., rely on it more than is normatively warranted), committing one form of bias. If the information is perfectly relevant, people will likely "under-

E-mail addresses: Christopher.Hsee@chicagobooth.edu (C.K. Hsee), yang.yang@warrington.ufl.edu (Y. Yang), xilin.li@chicagobooth.edu (X. Li).

https://doi.org/10.1016/j.obhdp.2019.05.002
Received 24 June 2017; Received in revised form 8 April 2019; Accepted 12 May 2019
0749-5978/

rely" on it (i.e., rely on it less than is normatively warranted), exhibiting a reverse bias. We refer to these tendencies as *relevance insensitivity*. Existing studies showing biases often look at only one side of the relevance dimension (e.g., only situations in which the given information is irrelevant) and fail to recognize the other side.

To illustrate the idea of relevance insensitivity, let us consider the sunk cost fallacy. According to our theory, this phenomenon is attributable to decision-makers' insufficient sensitivity to the relevance of the past cost to the future cost; namely, they are insufficiently sensitive to how predictive the past cost is of future cost. If the past cost is irrecoverable and hence not predictive of future cost, people should ignore it, but they do not do so completely, thus exhibiting the typical sunk cost fallacy (i.e., over-reliance on the past cost). On the other hand, if the past cost is perfectly recoverable and hence perfectly predictive of future cost, people should fully incorporate it into their decision, but we predict that they will not fully incorporate it either, thus exhibiting the reverse of the sunk cost bias (i.e., under-reliance on the past cost).

Several caveats: First, relevance insensitivity does not mean insensitivity to the information; rather, it means insensitivity to the *relevance* of the information. For example, in the sunk cost fallacy, relevance insensitivity does not mean insensitivity to the amount of past cost; rather, it means insensitivity to the relevance of past cost to the given situation (i.e., how predictive past cost is of future cost). Second, we use the term "bias" to describe a systematic deviation from the normative benchmark in the given situation and do not claim that a bias is necessarily an act of irrationality. Third, the direction of a "reverse bias" is named relative to the direction of the originally-documented bias. If the originally-documented bias is over-reliance on some information, the reverse bias is under-reliance on the information, and vice versa.

Relevance insensitivity can arise for three reasons: inattention, imperfect knowledge, or imperfect use. Inattention means that people do not pay attention to the information and consequently misunderstand it; for example, in the sunk cost fallacy, people may fail to realize that a past cost is irrecoverable. Imperfect knowledge means that even if people pay attention to the information, they may not know whether the information is normatively relevant or not; for example, in the sunk cost fallacy, people know that a past cost is irrecoverable, they may not know that they should ignore irrecoverable past costs. Imperfect use means that even if people know whether the information is normatively relevant, they may not apply the knowledge to every decision situation, due to insufficient System 2 processing (Kahneman, 2011; Sloman, 1996) or insufficient effort (Bettman, Luce, & Payne, 1998; Fiske & Taylor, 1991); for example, in the sunk cost fallacy, even if people theoretically know that they ought to ignore irrecoverable past costs and incorporate recoverable past costs, they may not apply their knowledge to every decision situation. Of the three reasons, inattention is the least interesting, and in our studies, we use comprehension checks to rule out that reason. We consider both imperfect knowledge and imperfect use to be important aspects of relevance insensitivity, and we do not empirically distinguish between the two in our studies. In the General Discussion section, we will examine the distinction and its implications.

Relevance insensitivity is a form of overgeneralization: people overgeneralize responses they have learned in situations in which the responses are generally optimal to situations in which the responses are not optimal. In their seminal paper introducing the availability, representativeness, and anchoring-and-adjustment heuristics, Tversky and Kahneman (1974) characterize these heuristics as generally useful, but sometimes conducive to serious errors, implying that people over-overgeneralize their habit of using these heuristics from situations in which the heuristics are useful to situations in which they are counterproductive. Other researchers have also evoked the notion of overgeneralization to explain their findings (Amir & Ariely, 2007; Arkes & Ayton, 1999; Baron, 2000; Chinander & Schweitzer, 2003; Haws, Reczek, & Sample, 2017; Hsee, Tu, Lu, & Ruan, 2014; Hsee, Yang, & Ruan, 2015; Peysakhovich & Rand, 2015; Yang, Vosgerau, & Loewenstein, 2013; Yeung & Soman, 2007; Zhu, Bagchi, & Hock, 2019). For example, Baron (2000) argues that people overgeneralize their belief that morality serves self-interest from situations in which the belief reflects reality to situations in which it does not. Peysakhovich and Rand (2015) observe that players of the prisoner's dilemma game overgeneralize the strategy they have learned from one round of the game, in which the strategy is optimal, to a subsequent round, in which the strategy is not optimal. In this research, we argue that people overgeneralize their responses in situations in which a particular type of information is on average intermediately relevant to situations in which that information is irrelevant or highly relevant.

Relevance insensitivity also builds on the idea that people are largely adapted to their environments, but are not perfectly attuned to every situation they encounter (Bar-Anan, Liberman, & Trope, 2006; Hogarth, 1981; Hsee, Zhang, Cai, & Zhang, 2013; Klayman & Brown, 1993; Nisbett & Ross, 1980; Todd & Gigerenzer, 2007). If people are insensitive to a certain situational variable, then they may exhibit one bias when encountering one end of the variable, and a reverse bias when encountering the other end. One prime example of such insensitivity is the hard-easy effect in confidence judgment: when asked to estimate the correctness of their judgment (e.g., answers to general knowledge questions), people overestimate their correctness when the task (e.g., the question) is difficult, and underestimate their correctness when the task is easy, suggesting that people are insensitive to the difficulty of the task (Juslin, Winnan, & Olson, 2000; Larrick, Burson, & Soll, 2007; Lichtenstein & Fischhoff, 1977; see also Kruger & Dunning, 1999, which shows that people are insensitive to their own cognitive ability). As a related example of insensitivity, Gino, Sharek, and Moore (2011) found that people overestimate their level of control in situations in which they have little control, and underestimate their level of control in situations in which they have much control. Similarly, Faro and Rottenstreich (2006) found that predictors are insensitive to others' risk preferences; predictors under-predict both others' risk aversion (Hsee & Weber, 1997) and other's risk-seeking. While the aforementioned research builds on the same idea of situation insensitivity, we are not interested in insensitivity to task difficulty, level of control, or others' risk preferences. Instead, we are interested in insensitivity to information relevance.

Relevance insensitivity is also a form of noise or regressiveness in response (Ayton, Hunt, & Wright, 1991; Burson, Larrick, & Klayman, 2006; Costello & Watts, 2014; Erev, Wallsten, & Budescu, 1994; Faro & Rottenstreich, 2006; Fiedler & Unkelbach, 2014; Gino et al., 2011; Hilbert, 2012; Juslin et al., 2000; Kahneman, 2011). If we treat information relevance as the x-axis of a Cartesian coordinate system and response as the y-axis, relevance insensitivity would cause the slope of the actual response to be flatter than the slope of the normative response, thus indicating noise or regressiveness in response. However, our theory is more specific than the general notion of noise or regressiveness, because our theory specifies the x-axis in each bias. In this sense, the contribution of our research is again analogous to the contribution of the aforementioned hard-easy effect in confidence judgment. The hard-easy effect is also an instance of noise or regressiveness in response, but the discovery of that effect is still a contribution, because it specifies the x-axis as the difficulty of the given task, and thereby sheds light on why people are overconfident and predicts when they will become underconfident. Likewise, our theory also specifies the x-axis for each bias (e.g., the relevance of a past cost in the case of the sunk cost fallacy), and thereby sheds light on why people exhibit a given bias (e.g., over-reliance on the past cost when it is irrelevant) and predicts when they will show the reverse bias (e.g., under-reliance on the past cost when it is fully relevant).

If the bias and its reverse bias are due to relevance insensitivity as we have proposed, we should be able to mitigate or eliminate the bias and its reverse bias with an intervention that increases people's

Table 1
Overview of Four Classic Biases and Their Reverse Biases Studied in This Research.

Original bias	Information	Relevance	Bias
Sunk cost fallacy	Past cost	Low (LR), i.e., unrelated to future cost High (HR), i.e., highly related to future cost	Over-reliance (original) Under-reliance (reversed)
Non-regressive prediction	Past outcomes	Low (LR), i.e., due to random factors High (HR), i.e., due to stable factors	Over-reliance (original) Under-reliance (reversed)
Anchoring bias	Anchor number	Low (LR), i.e., unrelated to target judgment High (HR), i.e., highly related to target judgment	Over-reliance (original) Under-reliance (reversed)
Base rate neglect	Base rate	Low (LR), i.e., unrelated to target judgment High (HR), i.e., highly related to target judgment	Over-reliance (reversed) Under-reliance (original)

sensitivity to the relevance of the given information, such as a tip that alerts people to whether the information is relevant and whether they should use it. Consider again the sunk cost fallacy. If the sunk cost fallacy and its reverse bias are due to people's insensitivity to the relevance of the past cost to future decisions, then we should be able to reduce the sunk cost fallacy and its reverse bias by advising people to ignore the past cost if it is irrecoverable and to incorporate the past cost if it is recoverable.

In summary, we propose that *people tend to over-rely on low-relevance information, exhibiting one bias, and under-rely on high-relevance information, committing a reverse bias.* Additionally, we propose that *increasing people's sensitivity to the relevance of the information attenuates or eliminates the bias and its reverse bias.*

As illustrations of our theory, we will analyze four classic biases—the sunk cost fallacy, non-regressive prediction, anchoring bias, and base rate neglect; see Table 1 for an overview. These biases are among the best known and most robust anomalies. Here, we explain all four biases using our relevance-insensitivity theory. We do not claim that our theory is the only explanation for these biases, but hope to convince the reader that our theory offers a uniform account for these seemingly disparate phenomena. Furthermore, we use the same theory to predict the corresponding reverse biases. For each of the four pairs of biases, we have conducted an empirical study to corroborate our analysis. In order not to disrupt the flow of the paper, we summarize the main results of the studies in Fig. 1 and report detailed statistics in Tables 2–5.

A few notes about our studies: While most prior studies on biases used college students as participants, we used MTurkers, because MTurkers are more representative of the general population and are generally reliable (Buhrmester, Kwang, & Gosling, 2011; Peer, Vosgerau, & Acquisti, 2014). Based on previous research showing similar biases, and recognizing that data from MTurkers usually entail more variance than data from college students, we set a target sample size of 320 for studies that included four conditions (Studies 1, 2, and 4) and a target sample size of 640 for the study that included eight conditions (Study 3). To ensure that the participants paid attention to our instructions, we asked a comprehension check question after measuring the dependent variable in each study. Participants who did not pass the comprehension check questions were excluded. This exclusion rule was pre-specified and pre-registered at https://aspredicted.org/ before data collection. Anonymized pre-registration documents can be accessed through the following links for Studies 1–4, respectively: http://aspredicted.org/blind.php?x=aq6666, http://aspredicted.org/blind.php?x=sj3sy2, http://aspredicted.org/blind.php?x=h7ic6n, and http://aspredicted.org/blind.php?x=8uz729.

2. Sunk cost fallacy

2.1. Explanation and prediction

As already mentioned, the sunk cost fallacy (sometimes called the Concorde effect) is a phenomenon in which the more cost a person has already expended on something, the less willing that person is to give it up, even though a better alternative exists (Arkes & Ayton, 1999; Arkes & Blumer, 1985; Thaler, 1999). Traditionally, the sunk cost fallacy has been attributed to mental accounting (Thaler, 1999) and the misuse of the "waste-not" heuristic (Arkes & Ayton, 1999; Arkes & Blumer, 1985).

We offer a relevance-insensitivity account: decision makers are insufficiently sensitive to whether their past cost is relevant to future cost. Normatively, if the past cost is irrecoverable, then it is not relevant to (i.e., not predictive of) future cost, and decision makers should ignore it. If the past cost is recoverable, then it is relevant to (i.e., predictive of) future cost, and decision makers ought to incorporate it into their decisions. Our theory predicts that if the past cost is irrecoverable, decision makers will not fully ignore it and thereby will exhibit the typical sunk cost fallacy (over-reliance on past cost); at the same time, if the past cost is fully recoverable, decision makers still will not fully incorporate it and hence will exhibit a reverse bias (under-reliance on past cost). Finally, we predict that giving participants a tip to alert them to the relevance of the past cost to future cost will attenuate both the sunk cost fallacy and its reverse bias.

2.2. Empirical evidence

Study 1 tested our propositions regarding the sunk cost fallacy. It concerned willingness to pay (WTP) to switch from an old apartment to a new one. We manipulated the relevance of a past cost to future cost (low versus high) and the participants' sensitivity to its relevance (without tip versus with tip). Specifically, to manipulate the relevance, we described the past cost as either a one-time processing fee or a deposit. The processing fee was non-refundable (irrecoverable) regardless of whether one kept or broke the lease, so this past cost was irrelevant to future expenses. In other words, it was a pure sunk cost. By contrast, the deposit would be forfeited if one broke the lease, but would be counted toward future rental expenses if one kept the lease, so this past cost was directly relevant to future expenses, and, hence, it was not a sunk cost. Our theory predicted that participants would not adequately distinguish between the two types of past costs, even if participants knew that one was recoverable and the other was not, unless we gave them an explicit relevance-sensitivity tip.

2.2.1. Method

Three hundred and twenty MTurkers (165 females; mean age = 36.96) participated in the study for a nominal payment. The study constituted a 2 (relevance: LR versus HR) × 2 (tip: without versus with) between-participants design. (We use "LR" to denote low relevance and "HR" to denote high relevance.) We reproduce the stimuli below, italicizing the parts that differed among the conditions:

> You plan to visit Chicago for a month. You have found online an old apartment there and signed a one-month lease. You don't need to pay the rent until the end of your stay. Unless otherwise specified, you may break the lease without penalty. (Assume you have no budgetary constraint, and you value a dollar in the future as much as

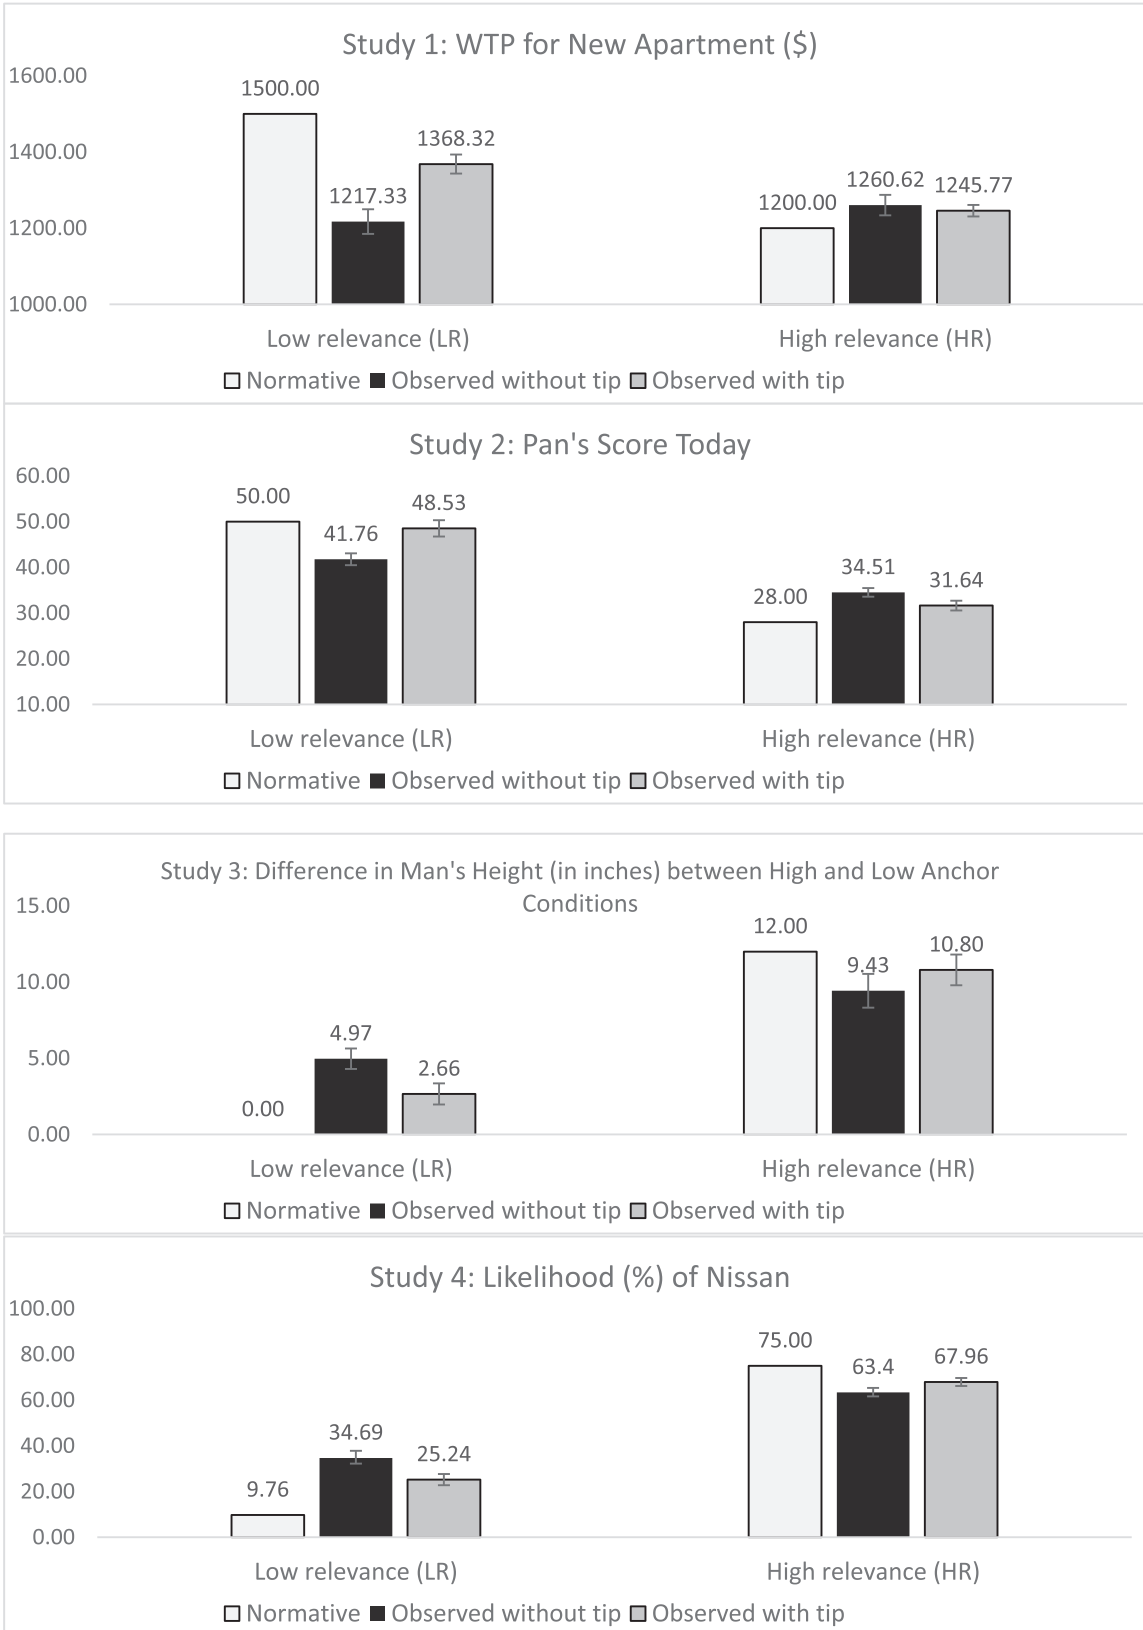

Note: The error bars represent one ±SE.

Fig. 1. A glance at the results of all studies (see Tables 2–5 for details). Note: The error bars represent one ± SE.

you value a dollar now.)

Right after signing the lease for the old apartment, you find a new apartment in the same region. The new apartment requires no deposit or processing fee. You are debating whether to keep the lease and rent the old apartment, or to break the lease for the old apartment and rent the new apartment instead.

If you had not signed the lease for the old apartment and had not paid anything for the old apartment, the maximum you are willing to pay to rent the new apartment would be $1,500.

[LR] However, you have already paid a $300 fee for the old apartment. It is a one-time processing fee and won't be refunded regardless of whether you break the lease for the old apartment or not.

[HR] However, you have already paid a $300 fee for the old apartment. It is a one-time deposit; it will be forfeited if you break the lease for the old apartment, or will be counted toward your rent if you keep the lease and rent the old apartment.

Given that, what's the maximum you are willing to pay for the new apartment now?

[Without Tip] (No tip was provided.)

[With Tip/LR] Before you make your prediction, read the following tip carefully. Because the $300 you have already paid for the old apartment is non-refundable regardless of whether you break the lease or not, it is irrelevant to how much you should be willing to pay for the new apartment. In other words, you should ignore the $300 when answering the question.

[With Tip/HR] Before you make your prediction, read the following tip carefully. Because the $300 you have already paid for the old apartment would be counted toward your rent if you did not break the lease for the old apartment, it is highly relevant to how much you should be willing to pay for the new apartment. In other words, you should fully incorporate the $300 when answering the question.

Note that the main dependent variable (WTP for the new apartment) was continuous rather than discrete (e.g., whether to switch to the new apartment). We chose a continuous variable because it allowed us to quantify the biases (i.e., the differences between people's actual responses and the normative benchmarks) in both the LR and the HR conditions.

After answering the above question, participants were directed to the next page and asked a comprehension check question: according to the instructions, the $300 fee paid for the old apartment was (a) a one-time processing fee, which would not be refunded whether they broke the lease or not, or (b) a one-time deposit fee, which would be forfeited if they broke the lease or would be counted as a part of their rent if they kept the lease. Tables 2A and 2B report the number of participants who passed this question. The results below are based on those participants. At the end of the study, participants also answered a relevance-knowledge question, which we will explain in the General Discussion section.

Table 2A

Study 1 Low Relevance Condition Results: Participants Over-Relied on the Past Cost to Determine Willingness to Pay (WTP) for the New Apartment.

	Without tip	With tip
N of Ps assigned to the condition	83	76
N of Ps who passed the comprehension check question	75	74
M (SD) of observed WTP	1217.33 (281.71)	1368.32 (214.60)
Normative WTP	1500.00	1500.00
N of Ps who gave normative WTP	18	42
Bias (normative WTP minus observed WTP)	+282.67	+131.68
Significance of bias	$t(74) = 8.69$, $p < .001$	$t(73) = 5.28$, $p < .001$
Significance of tip	$t(147) = 3.68$, $p < .001$	

Table 2B

Study 1 High Relevance Condition Results: Participants Under-Relied on the Past Cost to Determine Willingness to Pay (WTP) for the New Apartment.

	Without tip	With tip
N of Ps assigned to the condition	71	90
N of Ps who passed the comprehension check question	47	71
M (SD) of observed WTP	1260.62 (184.72)	1245.77 (127.26)
Normative WTP	1200.00	1200.00
N of Ps who gave normative WTP	25	49
Bias (normative WTP minus observed WTP)	−60.62	−45.77
Significance of bias	$t(46) = 2.25$, $p = .029$	$t(70) = 3.03$, $p = .003$
Significance of tip	$t(116) = 0.52$, $p = .606$	

2.2.2. Results

A 2 (relevance: LR vs. HR) × 2 (tip: without vs. with) ANOVA on participants' WTP for the new apartment revealed a relevance × tip interaction effect, $F(1, 263) = 9.74$, $p = .002$. Below, we elaborate on the results, first examining the LR condition and then examining the HR condition.

In the LR condition, the $300 (processing fee) was irrecoverable and irrelevant to future cost. Therefore, normatively, participants' WTP for the new apartment should not be influenced by the $300; it should still be $1500—the same as if they had not signed the lease for the old apartment. But as Fig. 1 and Table 2A show, in the without-tip condition, the observed WTP was significantly lower than the normative WTP of $1500 ($t(74) = 8.69$, $p < .001$), indicating that the participants failed to ignore the sunk cost and committed the classic sunk cost fallacy. In the with-tip condition, even though the observed WTP was still lower than $1500, the bias was significantly less severe than in the with-tip condition ($t(147) = 3.68$, $p < .001$).

In the HR condition, the $300 (deposit) was fully refundable and highly relevant to future cost. Therefore, normatively, participants' WTP should be $1200 (i.e., $300 less than the original $1500). But as Fig. 1 and Table 2B show, the actual WTP was significantly higher than $1200 ($t(46) = 2.25$, $p = .029$), suggesting that participants under-relied on the past cost and showed the reverse of the sunk cost effect. The presence of the tip did not significantly reduce this effect ($t(116) = 0.52$, $p = .606$).

2.2.3. Discussion

Prior research demonstrating the sunk cost fallacy focuses on situations in which the past cost is irrecoverable, and hence not predictive of future cost. Study 1 examined two situations, one with an irrecoverable past cost and one with a recoverable past cost, and found the typical sunk cost bias in the former situation and a reverse bias in the latter situation. This result suggests that the typical sunk cost fallacy occurs not because people universally over-rely on past cost, but rather because they are insensitive to whether the past cost is predictive of future cost.

3. Non-regressive prediction

3.1. Explanation and prediction

Statistically, extreme observations often regress towards the mean over time (Campbell & Kenny, 1999; Galton, 1886). Non-regressive prediction refers to the phenomenon that when making predictions, people under-predict the temporal regression of extreme observations (Kahneman & Tversky, 1973; Stigler, 1997).

To illustrate, consider a game with a score that ranges from 20 (lowest) to 80 (highest), with all numbers in the range equally likely. Assume that this game is purely luck-based, namely, that scores in the game depend purely on random factors. If a particular player scored low (say, 28) the last time she played the game, what will she score the next time she plays the game? Normatively, the best prediction is 50, because scores in the game depend purely on random factors and therefore will fully regress toward the mean (50). But extant research (e.g., Harrison & Bazerman, 1995; Kahneman & Tversky, 1973) finds that people under-predict such temporal regression, which, in this example, means that people will under-predict her score. Traditionally, non-regressive prediction has been attributed to the use of the representativeness heuristic and the tendency to invent spurious explanations (Kahneman & Tversky, 1973; Kahneman, 2011).

This bias can also be interpreted as a manifestation of relevance insensitivity, namely, insufficient sensitivity to the relevance of past outcomes to future outcomes. Consider two alternative scenarios: low relevance and high relevance. We already described the low-relevance scenario, in which scores in the game depend purely on random factors. In this case, the score in one round is not predictive of the score in the next, and therefore, the best normative prediction is always the average of all possible scores: 50. Now, consider an alternative high-relevance scenario, in which the game is purely skill-based; namely, scores in the game depend purely on stable skills. In this case, the last score is perfectly predictive of the next score, and, therefore, if a person scored 28 the last time, she will also score 28 the next time.

Our theory posits that forecasters are insufficiently sensitive to the relevance (predictiveness) of past outcomes to future outcomes. Consequently, they will under-predict regression in the purely-random-cause (i.e., low-relevance) condition but will not under-predict, and may even over-predict, regression in the purely-stable-skill (i.e., high-relevance) condition. Our theory also predicts that giving participants a tip to highlight the relevance of past outcomes to future outcomes will reduce both of these biases.

3.2. Empirical evidence

Study 2 tested our propositions regarding non-regressive prediction. In the study, we informed participants of a player's score from the day before and asked participants to predict the score the player will get today. We manipulated the relevance of yesterday's score (low versus high) and the sensitivity to the relevance of the score (without tip versus with tip). In the low-relevance (LR) condition, we told participants that scores in the game depended solely on random factors (luck), so the player's score yesterday was not predictive of her score today. In the high-relevance (HR) condition, we told participants that scores in the game depended solely on stable (unchanging) skills, so the player's score yesterday was highly predictive of her score today. Our theory predicted that participants would not adequately distinguish between the two types of games, even if participants knew whether the scores depended solely on random factors or unchanging skills, unless we gave them a relevance-sensitivity tip.

3.2.1. Method

Three hundred and twenty-one MTurkers (155 females; mean age = 36.43) participated in the study for a nominal payment. The study constituted a 2 (relevance: LR versus HR) × 2 (tip: without versus with) between-participants design. We reproduce the stimuli below, italicizing the parts that differed among the conditions:

> Game RS is a newly invented two-player game.
> *[LR] Scores in the game depend solely on random factors (luck).*
> *[HR] Scores in the game depend solely on stable (unchanging) skills.*
> Of all the players who have played the game so far, the scores are distributed equally within its possible range, which is between 20 (lowest) to 80 (highest). Yesterday, Pan played this game with Qin, and scored 28. Pan is going to play this game with Qin again today. Take a guess at Pan's score today.
> *[Without Tip] (No tip was provided.)*
> *[With Tip/LR] Before you make your prediction, read the following tip carefully. Because scores in the game depend solely on random factors (luck), the information about Pan's score yesterday is not relevant to your prediction of her score today. You should ignore that information when making your prediction.*
> *[With Tip/HR] Before you make your prediction, read the following tip carefully. Because scores in the game depend solely on stable (unchanging) skills, the information about Pan's score yesterday is highly relevant to your prediction of her score today. You should make your prediction based on that information.*

Table 3A
Study 2 Low Relevance Condition Results: Participants Over-Relied on Pan's Score Yesterday to Predict Pan's Score Today.

	Without tip	With tip
N of Ps assigned to the condition	85	77
N of Ps who passed the comprehension check question	80	72
M (SD) of observed score	41.76 (11.45)	48.53 (15.14)
Normative score	50.00	50.00
N of Ps who gave normative score	10	14
Bias (normative score minus observed score)	+8.24	+1.47
Significance of bias	$t(79) = 6.43$, $p < .001$	$t(71) = 0.83$, $p = .412$
Significance of tip	$t(150) = 3.13$, $p = .002$	

After answering the above question, participants were directed to the next page to answer a comprehension check question, which asked whether, according to the instructions, the scores depended solely on random factors or (unchanging) skills. Tables 3A and 3B report the number of participants who passed the question. The results below are based on those participants. At the end of the study, participants also answered a relevance-knowledge question, which we will explain in the General Discussion section.

3.2.2. Results

A 2 (relevance: LR vs. HR) × 2 (tip: without vs. with) ANOVA on participants' predictions of Pan's score today revealed a relevance × tip interaction effect, $F(1, 288) = 13.22$, $p < .001$. Below, we report more detailed results.

In the LR condition, scores depended purely on random factors, so the best normative prediction of Pan's score today was complete regression, namely, the average of all possible scores (i.e., 50). But as Fig. 1 and Table 3A show, in the without-tip condition, the observed prediction was significantly lower than 50 ($t(79) = 6.43$, $p < .001$), indicating that participants over-relied on Pan's past score. This result replicated the classic non-regressive prediction bias. In the with-tip condition, even though the observed prediction was still lower than 50, it was closer to 50 than in the without-tip condition ($t(150) = 3.13$, $p = .002$).

In the HR condition, scores depended solely on stable (unchanging)

Table 3B
Study 2 High Relevance Condition Results: Participants Under-Relied on Pan's Score Yesterday to Predict Pan's Score Today.

	Without tip	With tip
N of Ps assigned to the condition	74	85
N of Ps who passed the comprehension check question	67	73
M (SD) of observed score	34.51 (7.82)	31.64 (9.16)
Normative score	28.00	28.00
N of Ps who gave normative score	8	25
Bias (normative score minus observed score)	−6.51	−3.64
Significance of bias	$t(66) = 6.81$, $p < .001$	$t(72) = 3.40$, $p = .001$
Significance of tip	$t(138) = 1.98$, $p = .05$	

skills, so the normative prediction of Pan's score today was no regression, namely, the same as her score yesterday (i.e., 28). But as Fig. 1 and Table 3B show, in the without-tip condition, the observed prediction was significantly higher than 28 ($t(66) = 6.81$, $p < .001$), suggesting that participants under-relied on the past score. This result is the reverse of the classic non-regressive prediction bias. In the with-tip condition, even though the observed prediction was still higher than 28, it was closer to 28 than in the without-tip condition ($t(138) = 1.98$, $p = .05$).

3.2.3. Discussion

In this study, we examined a purely-random-cause condition and a purely-stable-skill condition because these conditions entailed clear normative predictions. We replicated the classic non-regressive prediction bias in the purely-random-cause condition, and showed the opposite over-regressive prediction bias in the purely-stable-skill condition. In reality, the cause of most outcomes is neither purely random nor purely stable. We suspect that the more purely random the outcomes of a particular event are, the more likely people are to under-predict temporal regression, and vice versa.

4. Anchoring bias

4.1. Explanation and prediction

The anchoring bias is one of the most celebrated psychological phenomena, and it takes on many different forms (Epley & Gilovich, 2006; Janiszewski & Uy, 2008; Mussweiler & Strack, 2000; Simmons, LeBoeuf, & Nelson, 2010; Tversky & Kahneman, 1974). It is beyond the scope of this research to conduct a review of the different forms of the anchoring bias and the numerous theories that have been advanced to explain them. Here, we focus on one, and probably the best-known, form of anchoring bias, in which one's judgment of a numerical value is influenced by an externally-given and obviously-irrelevant number—the so-called anchor. In a classic demonstration of the effect, Tversky and Kahneman (1974) first showed participants a roulette wheel that stopped at either 10 or 65, and asked them whether the number of African nations in the United Nations was greater or less than 10 (or 65), and then asked the participants to estimate the number. Participants whose wheel stopped at 10 made a lower estimate than those whose wheel stopped at 65, even though participants knew (or ought to have known) that the externally-given roulette-wheel number was irrelevant. This effect has been replicated in a wide range of contexts (Northcraft & Neale, 1987; Stewart, 2009; Wansink, Kent, & Hoch, 1998) and has been explained in terms of the anchoring-and-adjustment process and selective accessibility (Epley & Gilovich, 2006; Mussweiler & Strack, 2000; Simmons et al., 2010; Tversky & Kahneman, 1974).

We offer a relevance-insensitivity account for this effect: decision makers succumb to the anchoring bias because they are insufficiently sensitive to the relevance of the anchor to the target number (i.e., the number to be judged). This account makes three specific predictions. First, it predicts that in situations in which an external anchor is irrelevant to the target number, people will still consider the anchor and rely on it more than is normatively warranted. This is the typical anchoring bias. In most prior studies demonstrating the anchoring bias, the anchor is irrelevant to the target number; for example, the number on a roulette wheel is completely irrelevant to the number of African nations in the United Nations. In such situations, people should ignore the anchor, but they do not; in other words, people over-rely on the anchor. Second, our theory predicts that in situations in which an externally-given anchor is highly relevant to the target number, people will not base their judgment fully on the number (unless it is obvious that the target number is the same as the anchor number); in other words, people will under-rely on the anchor. Finally, we predict that a relevance-sensitivity tip will make people less likely to over-rely on the irrelevant anchor and under-rely on the highly relevant anchor, mitigating both the anchoring bias and its reverse bias.

4.2. Empirical evidence

Study 3 tested our propositions regarding the anchoring bias. We told participants the height of a woman (the anchor) at a dance party and asked them to predict the height of the man (the target) with whom she was paired. We manipulated the anchor (low versus high), its relevance to the target (low versus high), and sensitivity to the relevance of the anchor (without tip versus with tip). In the low-relevance (LR) condition, we told participants that men and women were randomly paired at the party, so there was no relationship between the height of the man and the height of the woman within a pair. In other words, the anchor (the height of a woman in a pairing) was not diagnostic of the target (the height of the man in the same pairing). In the high-relevance (HR) condition, we told participants that men and women were paired according to their height such that the tallest man was paired with the tallest woman, the second-tallest man was paired with the second-tallest woman, and so on. In this case, the anchor (the height of the woman) was highly predictive of the target (the height of the man). Our theory predicted that participants would not adequately distinguish between the two ways in which men and women were paired, even if participants knew exactly how they were paired, unless we sensitized them to the relevance of the anchor by giving a tip.

4.2.1. Method

Six hundred and forty-four MTurkers (297 females; mean age = 36.25) participated in the study for a nominal payment. While the other studies in this paper employed a 2 (relevance) × 2 (tip) design, this study adopted a 2 (relevance: LR versus HR) × 2 (anchor: low versus high) × 2 (tip: without versus with) design, because it also manipulated the anchor. We reproduce the stimuli below, italicizing the parts that differed among the conditions:

> One hundred men and one hundred women are attending a dance party. These are typical adult men and typical adult women in the US. Some men are taller than other men, some women are taller than other women, and, on average, men are taller than women.
> *[LR] The hostess randomly pairs one man with one woman. Thus, the height of the man and the height of the woman are unrelated.*
> *[HR] The hostess pairs one man with one woman according to their height. Specifically, the hostess pairs the tallest man with the tallest woman, the second-tallest man with the second-tallest woman, and so on and so forth.*
> *[Low Anchor] In a certain pairing, the height of the woman is 4 feet 11 in.*
> *[High Anchor] In a certain pairing, the height of the woman is 5 feet 11 in..*
> Predict the height of the man in the pairing.
> *[Without Tip] (No tip was provided.)*
> *[With Tip/LR] Before you make your prediction, read the following tip carefully. Because the height of the man and the height of the woman are unrelated, the information about the height of the woman is not relevant to your prediction of the height of the man in the pairing. You should ignore that information when making your prediction.*
> *[With Tip/HR] Before you make your prediction, read the following tip carefully. Because the hostess pairs the tallest man with the tallest woman, the second-tallest man with the second-tallest woman, and so on and so forth, the information about the height of the woman is highly relevant to your prediction of the height of the man in the pairing. You should fully incorporate it into your prediction.*

After answering the above question, participants were directed to the next page to answer additional questions. One was a comprehension check, which asked whether, according to the instructions, the hostess (a) randomly paired one man with one woman so that the height of the man and the height of the woman were unrelated, or (b) paired men and women according to their height such that she paired the tallest man with the tallest woman, the second-tallest man with the second-tallest woman, and so on.

The second question checked whether participants knew that the variation in height among men was at least as large as that among women. We will explain later why this question was critical in establishing normative estimates. Specifically, it asked,

> According to your common sense, which of the following do you think is true about the men and the women in the party?

(a) There is more variation in height among the men than among the women. Thus, the difference in height between the tallest and the shortest man is greater than that between the tallest and the shortest woman.
(b) There is less variation in height among the men than among the women. Thus, the difference in height between the tallest and the shortest man is smaller than that between the tallest and the shortest woman.
(c) The men and the women have the same variation in height. Thus, the difference in height between the tallest and the shortest man is the same as that between the tallest and the shortest woman.

In reality, the correct answer is (a), but to make our test more conservative, we considered participants to have passed this question if they chose either (a) or (c). Tables 4A and 4B tally the number of participants who passed this question, as well as the number of participants who passed the aforementioned comprehension question. The results below are based on participants who passed both questions. At the end of the study, participants also answered a relevance-knowledge question, which we will explain in the General Discussion section.

4.2.2. Results

A 3-way ANOVA on participants' estimates of the man's height revealed a relevance × tip × anchor interaction effect, $F(1, 499) = 4.27$, $p = .039$. Below, we elaborate on the results. In this study, we were not interested in the absolute height estimate made by participants in the low-anchor and the high-anchor conditions; rather, we were interested in the signed *difference* in their estimates of the man's height between the low-anchor and the high-anchor conditions (i.e., the estimate in the high-anchor condition minus the estimate in the low-anchor condition).

In the LR condition, the man's height was unrelated to the woman's height (anchor). Therefore, normatively, the difference in estimates should have been 0. However, as Fig. 1 and Table 4A show, in the without-tip condition, the observed difference was significantly greater than 0 ($t(113) = 7.42$, $p < .001$), replicating the classic anchoring bias. In the with-tip condition, even though the observed difference was still greater than 0, it was much closer to 0 ($F(1, 254) = 5.58$, $p = .019$).

In the HR condition, the man's height was perfectly correlated with the woman's, and the participants included in the analysis understood that the variation in height among men was at least as large as that among women. Therefore, the normative difference in the estimate of the man's height between the high-anchor and the low-anchor conditions should have been at least as large as the difference in the woman's height between the high anchor and the low anchor conditions, namely, 5′11″ − 4′11″ = 12″. But as Fig. 1 and Table 4B show, in the without-tip condition, the observed difference was significantly smaller than 12″($t(122) = 2.31$, $p = .023$), indicating that the participants under-relied on the anchors when they were highly relevant. In the with-tip condition, the observed difference was closer to 12″ than in the without-tip condition, although the difference was not statistically significant ($p > .10$).

4.2.3. Discussion

To summarize, the study found that when predicting the height of the man, participants over-relied on the height of the woman when her height was not diagnostic of the man's height, exhibiting the typical anchoring bias, and under-relied on the height of the woman when her height was highly diagnostic of the man's height, committing a reverse bias. Reminding participants of the relevance of the anchor was effective at reducing the biases.

The extant literature on anchoring effects suggests that decision makers are influenced "too much" by anchors. That is probably because the existing studies adopt irrelevant anchors. Our research examines both irrelevant anchors and highly-relevant anchors, and finds that decision makers are not universally influenced "too much" by anchors. Rather, they are insufficiently sensitive to whether or not an anchor is relevant.

5. Base rate neglect

5.1. Explanation and prediction

Base rate neglect refers to the tendency to judge the likelihood of a situation erroneously by neglecting the prior probability. In a classic illustration of this fallacy, Kahneman and Tversky (1973) showed participants personality sketches of several individuals, randomly drawn from a group of 100 people comprising either 70 engineers and 30 lawyers in one condition, or 30 engineers and 70 lawyers in another condition. Then, participants indicated the probability that each sketch belonged to an engineer. Though the odds should be 5.44 times higher in the former condition than in the latter, participants' estimated likelihood did not differ between conditions, violating Bayes' rule. More interestingly, base rate neglect occurred not only when the personality description matched the stereotypes of engineers or lawyers, but also when it was completely uninformative. Base rate neglect has been attributed to the representativeness heuristic (Kahneman & Tversky, 1973; Tversky & Kahneman, 1974).

Table 4A
Study 3 Low Relevance Condition Results: Participants Over-Relied on the Woman's Height to Estimate the Man's Height.

	Without tip	With tip
N of Ps assigned to the condition	159 (72 in the low-anchor condition and 87 in the high-anchor condition)	164 (77 in the low-anchor condition and 87 in the high-anchor condition)
N of Ps who passed the comprehension check question	142	160
N of Ps who passed the question about height variation	129	145
N of Ps who passed both questions	115	143
M (SD) of estimate in high-anchor condition	72.64 (4.02)	71.18 (4.71)
M (SD) of estimate in low-anchor condition	67.68 (3.07)	68.52 (3.29)
Observed difference (estimate in high-anchor condition minus estimate in low-anchor condition)	4.97	2.66
Normative difference	0.00	0.00
Bias (normative difference minus observed difference)	−4.97	−2.66
Significance of bias	$t(113) = 7.42$, $p < .001$	$t(141) = 3.83$, $p < .001$
Significance of tip	$F(1, 254) = 5.58$, $p = .019$	

Table 4B
Study 3 High Relevance Condition Results: Participants Under-Relied on the Woman's Height to Estimate the Man's Height.

	Without tip	With tip
N of Ps assigned to the condition	163 (84 in the low-anchor condition and 79 in the high-anchor condition)	158 (92 in the low-anchor condition and 66 in the high-anchor condition)
N of Ps who passed the comprehension check question	143	142
N of Ps who passed the question about height variation	136	133
N of Ps who passed both questions	124	125
M (SD) of estimate in high-anchor condition	72.57 (8.37)	74.25 (2.50)
M (SD) of estimate in low-anchor condition	63.14 (3.21)	63.46 (6.87)
Observed difference (estimate in high-anchor condition minus estimate in low-anchor condition)	9.43	10.80
Normative difference	12.00	12.00
Bias (normative difference minus observed difference)	+2.57	+1.20
Significance of bias	$t(122) = 2.31$, $p = .023$	$t(123) = 1.20$, $p = .233$
Significance of tip	$F(1, 245) = 0.83$, $p = .365$	

According to our relevance-insensitivity account, decision makers neglect base rate because they are insufficiently sensitive to the relevance of the base rate to the target probability (i.e., the probability to be judged). Normatively, if the base rate is not relevant to (not predictive of) the target probability, then decision makers should ignore it. If the base rate is relevant to (predictive of) the target probability, then decision makers should incorporate it into their prediction. Our theory predicts that if the base rate is highly relevant, decision makers will not fully incorporate it, and hence will exhibit the typical base rate neglect (under-reliance on the base rate). Our theory also predicts that if the base rate is irrelevant, decision makers will not fully ignore it, and thereby will exhibit the reverse of base rate neglect (over-reliance on the base rate).

Of note, although base rate neglect, like the three aforementioned biases, can be explained by relevance insensitivity, it differs from the other three biases in one aspect: whereas the other three (original) biases happen when the information is irrelevant (i.e., over-reliance on irrelevant information), the (original) base rate neglect occurs when the information (the base rate) is highly relevant (i.e., under-reliance on highly relevant base rate). Accordingly, the reverse bias of the original base rate neglect fallacy is over-reliance on the base rate when the base rate is irrelevant. Finally, our theory also predicts that giving participants a tip to increase their sensitivity to the relevance of the base rate will attenuate both base rate neglect and its reverse bias.

5.2. Empirical evidence

Study 4 tested our propositions regarding base rate neglect. We told participants that in one region of a foreign country, 75% of the cars were Nissans (i.e., the base rate). Participants then learned that a car was involved in a hit-and-run accident and were asked to estimate the likelihood that the car was a Nissan. We manipulated the relevance of the base rate (low versus high) and sensitivity to the relevance of the base rate (without tip versus with tip). In the low-relevance (LR) condition, we told participants that the accident occurred in the region where the participant was actually living (rather than in the region where 75% of the cars were Nissans). Because the market share of Nissans differed across regions, the number "75%" was irrelevant to the target probability, and participants should have based their prediction on their impression of the local market share. In the high-relevance (HR) condition, we told participants that the accident happened in the region where 75% of the cars were Nissans, so the "75%" base rate was highly relevant to their prediction. Our theory predicted that participants would not sufficiently distinguish between the two regions, even if participants knew exactly where the accident took place, unless they were given a relevance-sensitivity tip.

5.2.1. Method

Three hundred and nineteen MTurkers (177 females; mean age = 36.00) participated in the study for a nominal payment. The study constituted a 2 (relevance: LR versus HR) × 2 (tip: without versus with) between-participants design. We reproduce the stimuli below, italicizing the parts that differed among the conditions:

> Nissans are more popular in some regions than in others. For example, in one region in a foreign country, 75% of the cars are Nissans.
> *[LR] One night in the region where you are actually living now, a car was involved in a hit-and-run accident.*
> *[HR] One night in that region, a car was involved in a hit-and-run accident.*
> The police asked a witness whether the car was a Nissan and he replied yes. But after a careful investigation, the police found that the witness had answered the question randomly by flipping a coin. Given all the information above, please estimate the likelihood that the car involved in the accident was a Nissan.
> *[Without Tip] (No tip was provided.)*
> *[With Tip/LR] Before you make your prediction, read the following tip carefully. Because the accident happened in the region where you are actually living now, and not in the region in a foreign country where 75% of the cars are Nissans, the number "75%" is not relevant to your estimate. You should ignore that number when answering the question.*
> *[With Tip/HR] Before you make your prediction, read the following tip carefully. Because the accident happened in the region in a foreign country where 75% of the cars are Nissans, the number "75%" is highly relevant to your estimate. You should use that number when answering the question.*

After answering the above question, participants were directed to the next page to answer additional questions. One was a comprehension check question, which asked where, according to the instructions, the hit-and-run accident occurred. Tables 5A and 5B report the number of participants who passed the comprehension check. The results below are based on those participants.

The second question asked participants to estimate the shares of various car brands, including Nissans, in the participants' own region. The answer to this question served as a normative benchmark in the LR condition. Specifically, it asked,

> *In the region where you are actually living now, what percentage of the cars is of each of the following brands? Please note: the sum of the percentages must add up to 100%.*
> *Audi:* _____
> *BMW:* _____
> *Chrysler:* _____
> *Chevrolet:* _____
> *Ford:* _____
> *Honda:* _____
> *Hyundai:* _____
> *Jeep:* _____

Table 5A
Study 4 Low Relevance Condition Results: Participants Over-Relied on the Given Base Rate to Estimate the Likelihood that the Involved Car Was a Nissan.

	Without tip	With tip
N of Ps assigned to the condition	76	83
N of Ps who passed the comprehension check question	55	68
M (SD) of observed likelihood	34.69 (23.16)	25.24 (20.10)
Normative likelihood	10.33	9.19
N of Ps who gave normative likelihood	8	20
Bias (normative likelihood minus observed likelihood)	−24.36	−16.05
Significance of bias	$t(54) = 7.81$, $p < .001$	$t(67) = 6.01$, $p < .001$
Significance of tip	$t(121) = 2.42$, $p = .017$	

Mazda: _____
Mercedes-Benz: _____
Nissan: _____
Toyota: _____
Volkswagen: _____
Other: _____
Total: _____

At the end of the study, participants also answered a relevance-knowledge question, which we will explain in the General Discussion section.

5.2.2. Results

A 2 (relevance: LR vs. HR) × 2 (tip: without vs. with) ANOVA on participants' likelihood judgment yielded a relevance × tip interaction effect, $F(1, 270) = 9.47$, $p = .002$. Below, we elaborate on the results.

In the LR condition, because the accident happened in the region where participants were actually living, rather than in the region where 75% of the cars were Nissans, the 75% base rate was irrelevant. Therefore, the normative likelihood that the car involved in the accident was a Nissan was the local market share of Nissans, which participants estimated at 10.33%. However, as Fig. 1 and Tables 5A show, in the without-tip condition, the observed likelihood judgment was significantly higher than 10.33% ($t(54) = 7.81$, $p < .001$), suggesting that the participants over-relied on the 75% base rate. In the with-tip condition, although the observed likelihood judgment was still higher than 10.33%, it was closer to 10.33% than in the without-tip condition ($t(121) = 2.42$, $p = .017$).

We now turn to the HR condition. Here, the accident happened in the region where 75% of the cars were Nissans. Therefore, the normative likelihood that the car involved in the accident was a Nissan was 75%. But as Fig. 1 and Table 5B show, in the without-tip condition, the observed likelihood judgment was significantly lower than 75% ($t(72) = 5.97$, $p < .001$), suggesting that participants under-relied on the 75% base rate and displayed the classic base rate neglect. In the with-tip condition, although the likelihood judgment was still lower than 75%, it was closer to 75% than in the without-tip condition ($t(149) = 1.75$, $p = .083$).

5.2.3. Discussion

Prior research demonstrating base rate neglect focuses on situations in which the base rate is highly relevant. Study 4 examined two types of situations, one in which the base rate was highly relevant and one in which the base rate was irrelevant, and found the typical base rate neglect in the former situation and the reverse of base rate neglect n the latter situation. These results suggest that people do not universally ignore the base rate when estimating the target probability. Rather, they are generally insensitive to whether the base rate is predictive of the target probability.

6. General Discussion

In making judgments and decisions, people often encounter a wealth of information. Normatively, they should ignore irrelevant information and incorporate relevant information. In reality, people are not perfectly attuned to the relevance of the information they have, so they may over-rely on irrelevant information and under-rely on highly relevant information. In this research, we illustrate relevance insensitivity with four seemingly unrelated classic biases: sunk cost fallacy, non-regressive prediction, anchoring bias, and base rate neglect. Existing research on these biases has focused primarily on one end of the relevance continuum, creating the impression that each of these biases occurs in only one direction (Fiedler, 2000, 2011; Wells & Windschitl, 1999). We examine both ends of the relevance continuum and show that a manipulation of relevance can turn these biases into their reverse biases. Moreover, alerting people to the relevance of the information reduces both the original and the reverse biases.

6.1. Open questions

Our research is far from conclusive. Below, we raise some open questions and offer our tentative answers.

6.1.1. Is relevance insensitivity due to imperfect knowledge or imperfect use?

In the Introduction, we distinguished between three possible causes of relevance insensitivity: inattention (failure to understand the instructions), imperfect knowledge (failure to know whether the given

Table 5B
Study 4 High Relevance Condition Results: Participants Under-Relied on the Given Base Rate to Estimate the Likelihood that the Involved Car Was a Nissan.

	Without tip	With tip
N of Ps assigned to the condition	78	82
N of Ps who passed the comprehension check question	73	78
M (SD) of observed likelihood	63.40 (16.59)	67.96 (15.41)
Normative likelihood	75.00	75.00
N of Ps who gave normative likelihood	40	55
Bias (normative likelihood minus observed likelihood)	+11.60	+7.04
Significance of bias	$t(72) = 5.97$, $p < .001$	$t(77) = 4.04$, $p < .001$
Significance of tip	$t(149) = 1.75$, $p = .083$	

Table 6
Correlations between responses to the relevance-knowledge questions and degrees of biases. The results indicate that the less correctly participants answered the questions, the greater their biases.

Study	Relevance condition	Overall correlation
1	LR	$r(149) = -0.31$, $p < .001$
	HR	$r(118) = -0.21$, $p = .024$
2	LR	$r(152) = -0.44$, $p < .001$
	HR	$r(140) = -0.28$, $p = .001$
4	LR	$r(123) = 0.38$, $p < .001$
	HR	$r(151) = 0.35$, $p < .001$

Note: We could not calculate correlations for Study 3 because the observed difference in height was derived from responses in two between-subjects conditions (i.e., the estimate in the high-anchor condition minus the estimate in the low-anchor condition).

information is normatively relevant), and imperfect use (failure to apply the relevance knowledge). As noted, all participants included in our analyses passed a comprehension check question about the key manipulated information, and therefore, it is unlikely that our results were due to inattention.

To elucidate whether the biases were due to imperfect knowledge or imperfect use, we included, at the end of each study, a relevance-knowledge question, which asked participants to assess the relevance of the key information in the study (the past cost in Study 1, the past score in Study 2, the anchor in Study 3, and the base rate in Study 4). For example, in Study 1 (sunk cost fallacy), we asked participants whether the $300 they had already paid for the old apartment was relevant to the maximum they should be willing to pay for the new apartment. In each study, participants answered the relevance-knowledge question by choosing among three options: 1 = highly relevant, 2 = somewhat relevant, 3 = not relevant.

Table 6 presents the correlations between participants' responses to the relevance-knowledge question and the degree of bias they exhibited, and it shows that the more incorrectly participants answered the relevance-knowledge question, the greater their biases. These results reinforce the statement we made earlier: that imperfect knowledge played a role in producing the biases. These results also corroborate the recent finding by Dietvorst and Simonsohn (2019) that many people who incorporate normatively-irrelevant information in their decisions simply do not realize that they should ignore the information.

Table 7 focuses on the without-tip condition of each study and compares the degree of bias among participants who answered the relevance-knowledge question correctly (i.e., chose the "not relevant" answer in the LR condition and the "highly relevant" answer in the HR condition) with the degree of bias among participants who answered the relevance-knowledge question incorrectly (i.e., chose other

Table 8
Responses to the relevance-knowledge question by participants in the without-tip and the with-tip conditions. The results reveal that the tip manipulation was effective at improving responses to the relevance-knowledge question.

Study	Relevance condition	Without-tip condition		With-tip condition		Difference between the without-tip and the with-tip conditions
		Mean	SD	Mean	SD	
1	LR	1.89	0.83	2.70	0.68	$p < .001$
	HR	1.87	0.85	1.20	0.55	$p < .001$
2	LR	2.46	0.71	2.92	0.28	$p < .001$
	HR	1.40	0.60	1.08	0.32	$p < .001$
3	LR	2.63	0.66	2.87	0.43	$p < .001$
	HR	1.31	0.60	1.06	0.25	$p < .001$
4	LR	2.58	0.71	2.82	0.46	$p = .024$
	HR	1.34	0.56	1.15	0.43	$p = .021$

Note: Responses on the relevance-knowledge question were made on a 1–3 scale, with greater numbers indicating less relevance. Therefore, in the LR condition, the correct response was 3, and in the HR condition, the correct response was 1.

answers). The results reveal two general patterns. First, even participants who answered the relevance-knowledge question correctly showed biases, suggesting that imperfect use played a role in producing the biases. Second, participants who answered the relevance-knowledge question incorrectly showed even greater biases, suggesting that imperfect knowledge also played a role.

Table 8 reports participants' answers to the relevance-knowledge question in both the without-tip and the with-tip conditions. As the data show, more participants in the with-tip conditions answered the question correctly than in the without-tip condition, suggesting that the tip manipulation was effective at improving participants' relevance sensitivity. Specifically, the tip improved both participants' knowledge of the relevance of the information and their use of the information; it improved knowledge because the tip told participants whether the given information was relevant or not, and it improved use because the tip told participants whether they should incorporate or ignore the information.

6.1.2. What contributes to imperfect knowledge and imperfect use?

As the previous section indicates, both imperfect knowledge and imperfect use contribute to relevance insensitivity. Then, what contributes to imperfect knowledge and imperfect use? We speculate that it is a combination of imperfect feedback, different past experiences, and insufficient reflection. For example, after a person decides to continue a course of action for which she has already incurred a cost, she rarely

Table 7
Degree of bias shown by participants in the without-tip condition who answered the relevance-knowledge question correctly and participants in the without-tip condition who answered it incorrectly. The results reveal that (a) even those who answered correctly showed biases, and (b) those who answered incorrectly showed greater biases.

Study	Relevance condition	Degree of bias among participants who answered the relevance-knowledge question correctly		Degree of bias among participants who answered the relevance-knowledge question incorrectly		Difference in degree of bias between those who answered correctly and those who answered incorrectly
		Mean	Significance	Mean	Significance	
1	LR	168.16	$p = .002$	330.19	$p < .001$	$p = .022$
	HR	−50.00	$p = .204$	−68.48	$p = .083$	$p = .739$
2	LR	3.70	$p = .036$	14.70	$p < .001$	$p < .001$
	HR	−5.61	$p < .001$	−8.22	$p = .001$	$p = .198$
3	LR	−2.90	$p < .001$	−10.45	$p < .001$	$p < .001$
	HR	3.01	$p = .036$	2.02	$p = .122$	$p = .715$
4	LR	−18.31	$p < .001$	−39.13	$p < .001$	$p = .002$
	HR	7.64	$p < .001$	20.77	$p < .001$	$p = .001$

finds out what would have happened if she had aborted the action instead. This lack of feedback makes it difficult for her to learn whether she should or should not have incorporated the past cost in her decisions; namely, it contributes to imperfect knowledge. Moreover, even if she had perfect knowledge, she may not apply it to every decision. In one's learning history, past costs are sometimes irrelevant (e.g., a deposit for a lease is non-refundable if the lease is nullified), sometimes perfectly relevant (e.g., a deposit is fully refundable), and sometimes intermediately relevant (e.g., a deposit is partially refundable). Thus, on average, past costs are intermediately relevant. Without sufficient effort and deliberation (System 2 processing), people may overgeneralize past experiences—in which the information was, on average, intermediately relevant—to situations in which the information is irrelevant or perfectly relevant. This leads to imperfect use of the information.

6.1.3. What other biases can or cannot be explained by relevance insensitivity?

As noted earlier, the four sets of biases studied in this research are only examples, and are not the only biases that our theory can explain. We suggest that a decision maker is likely to show a bias induced by relevance-insensitivity if (a) there is a normative (i.e., correct) response, and (b) the decision maker faces some information that is either of lower relevance or of higher relevance to the normative response in the current decision context than is typical for this type of information in his/her learning history. Under these conditions, the decision maker will likely over-rely on the information when it is of low relevance and under-rely on the information when it is of high relevance. All the biases and their reverse biases studied in this research meet the above criteria.

These criteria also allow us to apply our theory to other phenomena. For example, Hsee, Yang, Gu, and Chen (2009) found that consumers tend to base their purchase decisions on the specifications of a product (e.g., the resolution count of a screen) even if they can sample the product (e.g., view the screen) and if the specifications carry no additional information. This bias meets the criteria above: (a) there is a normative response—to base one's purchase decision on the sampling experience—and (b) specifications are generally informative, but in the specific situation studied in the research, the specification information is uninformative. Therefore, people show an over-reliance bias in the situation. The same analysis applies to Yeung and Soman (2007) finding that consumers over-rely on irrelevant duration information to assess the quality of a service, Chinander and Schweitzer (2003) finding that employers over-rely on irrelevant input information (e.g., time spent in office) when assessing employee performance, and Hsee, Yu, Zhang, and Zhang (2003) finding that workers over-rely on immediate payoff rather than the final outcome when deciding how much effort to expend. In all of these cases, people over-rely on information that is irrelevant, but we predict that if the given information were instead highly relevant (e.g., if the specification information were highly diagnostic of a product's quality, the duration information were highly diagnostic of a service's quality, the input information were highly diagnostic of an employee's performance, and the immediate payoff were highly diagnostic of the final outcome), then people would under-rely on the information.

While the above biases are due to over-reliance on low-relevance information, other biases are due to under-reliance on high-relevance information. One such example is the egocentric bias. For example, when asked to judge something, such as the length of the Nile River, research participants tend to rely more on their own opinions than on other participants' opinions, even when the other participants' opinions are as valid (i.e., as relevant) as their own (Soll & Larrick, 2009). We suggest that this bias also fits our criteria: (a) there exists a normative benchmark, namely, that each participant should weigh the opinion of another participant as much as his/her own, and (b) in real life, if a person is asked about something, it implies that the person has more knowledge on the issue than others and the person should normatively give less weight to others' opinions than to his/her own opinion, but in the context that produces the egocentric bias, the participant has no privileged knowledge. Hence, the participant under-relies on the other participants' opinions because he/she is insensitive to the relevance of the others' opinions. On the other hand, we predict that if the others' options are totally irrelevant (e.g., the others' options are randomly generated), people will not completely ignore those opinions either, thus showing the reverse of the egocentric bias, namely, over-reliance on the others' (irrelevant) opinions.

Of course, many biases do not meet the criteria postulated above and cannot be explained by our theory. For example, relevance insensitivity cannot explain many forms of preference reversals (e.g., Hsee, 1996; Hsee & Zhang, 2010; Lichenstein & Slovic, 1971; Nowlis & Simonson, 1997), because there is no clear normative benchmark. Nor can relevance insensitivity explain many forms of prediction biases (e.g., Campbell, O'Brien, Van Boven, Schwarz, & Ubel, 2014; Cooney, Gilbert, & Wilson, 2017; Hsee & Weber, 1997; Hsee & Zhang, 2004; Loewenstein, O'Donoghue, & Rabin, 2003; O'Brien, 2019; Ruan, Hsee, & Lu, 2018; Wilson & Gilbert, 2005). For instance, people under-predict how much they will continue to enjoy an event after experiencing it multiple times (O'Brien, 2019), and people over-predict how excited others will be to hear about their novel experiences (Cooney et al., 2017). Even though each of these biases has a normative benchmark (i.e., the correct prediction), and even though these biases may be attributed to insensitivity to situational variables (e.g., the actual adaptation rate of an event) and fit the general framework of situation insensitivity, people showing these biases do not face information that is more or less relevant than usual. Thus, these biases cannot be ascribed to relevance insensitivity.

6.1.4. Does each bias always have a corresponding reverse bias?

A bias does not necessarily have a reverse bias. People's tendency to over-rely or under-rely on a piece of information depends on the average relevance of that type of information in their learning history. Most people correctly perceive the average relevance of a certain type of information, but they are not attuned to how the relevance may deviate from average in specific situations. Therefore, if a certain type of information is, on average, of high relevance, then people are likely to exhibit the over-reliance bias (over-relying on a specific piece of information when it is actually irrelevant) but unlikely to exhibit the under-reliance bias (under-relying on a specific piece of information when it is highly relevant). If a certain type of information is, on average, of low relevance, then people are likely to exhibit the under-reliance bias but unlikely to exhibit the over-reliance bias. If a certain type of information is, on average, of intermediate relevance, then people may exhibit both the over-reliance bias and the under-reliance bias in different situations. The types of information investigated in this research—past cost, past outcomes, externally-given anchor, and base rates—are presumably of moderate relevance on average. Therefore, for these types of information, we are able to show both the over-reliance bias and the under-reliance bias. These analyses are only our speculations. Nevertheless, they are consistent with the idea, reviewed in the Introduction of this article, that people are generally well-adapted to their environments but are not finely-tuned to each specific situation they encounter (Bar-Anan et al., 2006; Hogarth, 1981; Hsee et al., 2013; Klayman & Brown, 1993; Nisbett & Ross, 1980; Todd & Gigerenzer, 2007). Relevance insensitivity is a manifestation of situation insensitivity.

6.2. Conclusion

This research offers a new look at various "old" biases. We show that, individually, each of these biases is not general because manipulating relevance can reverse the bias, but, collectively, these biases share a general underlying mechanism—relevance insensitivity.

Acknowledgement

The authors thank the following individuals (in alphabetical order of last names) for their helpful comments during different stages of this project: Joe Alba, Dan Bartels, Lyle Brenner, Danny Kahneman, Josh Klayman, Jingyi Lu, Sang Kyu Park, Bowen Ruan, Aner Sela, Anuj Shah, David Tannenbaum, Richard Thaler, Oleg Urminsky, Minzhe Xu, and Ying Zeng. We particularly wish to thank editor Rick Larrick for his guidance during the review process, especially his suggestion for us to include the tip condition in the studies.

Appendix A. Supplementary material

Supplementary data to this article can be found online at https://doi.org/10.1016/j.obhdp.2019.05.002.

References

Amir, O., & Ariely, D. (2007). Decisions by rules: The case of unwillingness to pay for beneficial delays. *Journal of Marketing Research, 44*, 142–152.
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.
Arkes, H. R., & Blumer, C. (1985). The psychology of sunk cost. *Organizational Behavior and Human Decision Processes, 35*(1), 124–140.
Ayton, P., Hunt, A. J., & Wright, G. (1991). Randomness and reality. *Journal of Behavioral Decision Making, 4*(3), 222–226.
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.
Baron, J. (2000). *The effects of overgeneralization on public policy*Working paper. University of Pennsylvania.
Bettman, J., Luce, M., & Payne, J. (1998). Constructive consumer choice processes. *Journal of Consumer Research, 25*, 187–217.
Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science, 6*(1), 3–5.
Burson, K. A., Larrick, R. P., & Klayman, J. (2006). Skilled or unskilled, but still unaware of it: How perceptions of difficulty drive miscalibration in relative comparisons. *Journal of Personality and Social Psychology, 90*(1), 60–77.
Campbell, D. T., & Kenny, D. A. (1999). *A primer on regression artifacts.* New York, NY: Guilford Press.
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.
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.
Cooney, G., Gilbert, D. T., & Wilson, T. D. (2017). The novelty penalty: Why do people like talking about new experiences but hearing about old ones? *Psychological Science, 28*(3), 380–394.
Costello, F., & Watts, P. (2014). Surprisingly rational: Probability theory plus noise explains biases in judgment. *Psychological Review, 121*(3), 463–480.
Dietvorst, B., & Simonsohn, U. (2019). Intentionally "biased": People purposely use to-be-ignored information, but can be persuaded not to. *Journal of Experimental Psychology: General.*.
Epley, N., & Gilovich, T. (2006). The anchoring-and-adjustment heuristic: Why the adjustments are insufficient. *Psychological Science, 17*(4), 311–318.
Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over-and under-confidence: The role of error in judgment processes. *Psychological Review, 101*(3), 519–527.
Faro, D., & Rottenstreich, Y. (2006). Affect, empathy, and regressive mispredictions of others' preferences under risk. *Management Science, 52*(4), 529–541.
Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review, 107*(4), 659.
Fiedler, K. (2011). Voodoo correlations are everywhere—Not only in neuroscience. *Perspectives on Psychological Science, 6*(2), 163–171.
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.
Fiske, S., & Taylor, S. (1991). *Social cognition* (2nd ed.). New York: McGraw-Hill.
Galton, F. (1886). Regression towards mediocrity in hereditary stature. *Journal of the Anthropological Institute of Great Britain and Ireland, 15*, 246–263.
Gino, F., Sharek, Z., & Moore, D. (2011). Keeping the illusion of control under control: Ceilings, floors, and imperfect calibration. *Organizational Behavior and Human Decision Processes, 114*, 104–114.
Harrison, J. R., & Bazerman, M. H. (1995). *Regression to the mean, expectation inflation, and the winner's curse in organizational contexts.* Los Angeles, CA: Sage Publications Inc.
Haws, K. L., Reczek, R. W., & Sample, K. L. (2017). Healthy diets make empty wallets: The healthy = expensive intuition. *Journal of Consumer Research, 43*, 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.
Hogarth, R. M. (1981). Beyond discrete biases: Functional and dysfunctional aspects of judgmental heuristics. *Psychological Bulletin, 90*(2), 197–217.
Hsee, C. K. (1996). The evaluability hypothesis: An explanation for preference reversals between joint and separate evaluations. *Organizational Behavior and Human Decision Processes, 67*(3), 247–257.
Hsee, C. K., Tu, Y., Lu, Y., & Ruan, B. (2014). Approach aversion: Negative hedonic reactions towards approaching stimuli. *Journal of Personality and Social Psychology, 106*(5), 699–712.
Hsee, C. K., & Weber, E. U. (1997). A fundamental prediction error: Self-other discrepancies in risk preference. *Journal of Experimental Psychology: General, 126*(1), 45–53.
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.
Hsee, C. K., Yang, Y., & Ruan, B. (2015). The mere reaction effect: Even non-positive and non-informative reactions can reinforce actions. *Journal of Consumer Research, 42*(3), 420–434.
Hsee, C. K., Yu, F., Zhang, J., & Zhang, Y. (2003). Medium maximization. *Journal of Consumer Research, 30*(1), 1–14.
Hsee, C. K., & Zhang, J. (2004). Distinction bias: Misprediction and mischoice due to joint evaluation. *Journal of Personality and Social Psychology, 86*(5), 680–695.
Hsee, C. K., & Zhang, J. (2010). General evaluability theory. *Perspectives on Psychological Science, 5*(4), 343–355.
Hsee, C. K., Zhang, J., Cai, C. F., & Zhang, S. (2013). Overearning. *Psychological Science, 24*(6), 852–859.
Janiszewski, C., & Uy, D. (2008). Precision of the anchor influences the amount of adjustment. *Psychological Science, 19*(2), 121–127.
Juslin, P., Winman, A., & Olsson, H. (2000). Naive empiricism and dogmatism in confidence research: A critical examination of the hard–easy effect. *Psychological Review, 107*(2), 384–396.
Kahneman, D. (2011). *Thinking, fast and slow.* New York, NY: Macmillan.
Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review, 80*(4), 237–251.
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), 97–122.
Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology. 77*(6), 1121–1134.
Larrick, R. P., Burson, K. A., & Soll, J. B. (2007). Social comparison and confidence: When thinking you're better than average predicts overconfidence (and when it does not). *Organizational Behavior and Human Decision Processes, 102*(1), 76–94.
Lichenstein, S., & Slovic, P. (1971). Reversals of preference between bids and choices in gambling decisions. *Journal of Experimental Psychology, 89*(1), 46–55.
Lichtenstein, S., & Fischhoff, B. (1977). Do those who know more also know more about how much they know? *Organizational Behavior and Human Performance, 20*(2), 159–183.
Loewenstein, G., O'Donoghue, T., & Rabin, M. (2003). Projection bias in predicting future utility. *The Quarterly Journal of Economics, 118*(4), 1209–1248.
Mussweiler, T., & Strack, F. (2000). The use of category and exemplar knowledge in the solution of anchoring tasks. *Journal of Personality and Social Psychology, 78*(6), 1038–1052.
Nisbett, R. E., & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment.* Englewood Cliffs, NJ: Prentice-Hall.
Nowlis, S., & Simonson, I. (1997). Attribute-task compatibility as a determinant of consumer preference reversal. *Journal of Marketing Research, 34*, 205–218.
Northcraft, G. B., & Neale, M. A. (1987). Experts, amateurs, and real estate: An anchoring-and-adjustment perspective on property pricing decisions. *Organizational Behavior and Human Decision Processes, 39*(1), 84–97.
O'Brien, E. (2019). Enjoy it again: Repeat experiences are less repetitive than people think. *Journal of Personality and Social Psychology, 116*(4), 519–540.
Peer, E., Vosgerau, J., & Acquisti, A. (2014). Reputation as a sufficient condition for data quality on Amazon Mechanical Turk. *Behavior Research Methods, 46*(4), 1023–1031.
Peysakhovich, A., & Rand, D. G. (2015). Habits of virtue: Creating norms of cooperation and defection in the laboratory. *Management Science, 62*(3), 631–647.
Ruan, B., Hsee, C. K., & Lu, Y. (2018). The teasing effect: An underappreciated benefit of creating and resolving an uncertainty. *Journal of Marketing Research*.
Simmons, J. P., LeBoeuf, R. A., & Nelson, L. D. (2010). The effect of accuracy motivation on anchoring and adjustment: Do people adjust from provided anchors? *Journal of Personality and Social Psychology, 99*(6), 917–932.
Simon, H. A. (1957). *Models of man: Social and rational.* Oxford, England: Wiley.
Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119*(1), 3–22.
Soll, J. B., & Larrick, R. P. (2009). Strategies for revising judgment: How (and how well) people use others' opinions. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35*(3), 780–805.
Stewart, N. (2009). The cost of anchoring on credit-card minimum repayments. *Psychological Science, 20*(1), 39–41.
Stigler, S. M. (1997). Regression towards the mean, historically considered. *Statistical Methods in Medical Research, 6*(2), 103–114.
Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral Decision Making, 12*(3), 183–206.
Todd, P. M., & Gigerenzer, G. (2007). Environments that make us smart: Ecological rationality. *Current Directions in Psychological Science, 16*(3), 167–171.
Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological*

Bulletin, 76(2), 105–110.
Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*(4157), 1124–1131.
Wansink, B., Kent, R. J., & Hoch, S. J. (1998). An anchoring and adjustment model of purchase quantity decisions. *Journal of Marketing Research, 35*(1), 71–81.
Wells, G. L., & Windschitl, P. D. (1999). Stimulus sampling and social psychological experimentation. *Personality and Social Psychology Bulletin, 25*(9), 1115–1125.
Wilson, T. D., & Gilbert, D. T. (2005). Affective forecasting: Knowing what to want. *Current Directions in Psychological Science, 14*(3), 131–134.
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.
Yeung, C. W., & Soman, D. (2007). The duration heuristic. *Journal of Consumer Research, 34*(3), 315–326.
Zhu, M., Bagchi, R., & Hock, S. J. (2019). The mere deadline effect: Why more time might sabotage goal pursuit. *Journal of Consumer Research, 45*, 1068–1084.