

Intentionally “Biased”: People Purposely Use To-Be-Ignored Information, But Can Be Persuaded Not To

Berkeley J. Dietvorst

The University of Chicago Booth School of Business

Uri Simonsohn

Universitat Ramon Llull, ESADE Business School

Abundant research has shown that people fail to disregard to-be-ignored information (e.g., hindsight bias, curse of knowledge), which has contributed to the popular notion that people are unwillingly and unconsciously affected by information. Here we provide evidence that, instead, people simply do not want to ignore such information. The findings: In Studies 1 and 2, the majority of participants explicitly indicated a desire to use to-be-ignored information in classic paradigms. In Study 3, the effect of receiving to-be-ignored information was driven entirely by the subset of people who wanted to use it. In Study 4, persuading participants to ignore inadmissible evidence in a mock jury trial reduced the impact of such evidence.

Keywords: hindsight bias, curse of knowledge, decision making, judgment, jurors

Supplemental materials: <http://dx.doi.org/10.1037/xge0000541.supp>

Research has repeatedly shown that people make judgments based on information they should ignore: Jurors do not disregard inadmissible evidence (Stebly, Hosch, Culhane, & McWethy, 2006), people exaggerate how predictable outcomes are after they have occurred (Fischhoff, 1975), people use private information when estimating others' valuations (Camerer, Loewenstein, & Weber, 1989), people judge the quality of decisions based on outcomes even when those outcomes are not relevant (Baron & Hershey, 1988), participants continue to believe information provided during a study after it has been discredited (Anderson, Lepper, & Ross, 1980), and so forth. We refer to this behavior as “use of to-be-ignored information.” The demonstrations just described have contributed to the popular notion that once people learn information, they are unable to ignore it.

In contrast to such conclusions, in this article we propose that people largely choose to not ignore such information. Specifically, we found that (a) participants who took part in representative paradigms from the literature intentionally relied on to-be-ignored information, (b) the subset of participants who reported that they did not desire to use the information succeeded in ignoring it, and

(c) providing additional justification to ignore the information (i.e., persuading participants they should ignore it) led more participants to intentionally ignore it.

These findings suggest not only revising the field's understanding of use of to-be-ignored information but also revisiting the implications of this body of research. For example, our findings suggest that people can, at least in the most commonly used paradigms, accurately report whether they are using to-be-ignored information. Moreover, they suggest that efforts to correct for the use of to-be-ignored information could focus on convincing people that they should ignore the information.

Literature on Use of To-Be-Ignored Information

Fischhoff (1975) created the prototypical hindsight bias paradigm in his seminal article “Hindsight ≠ Foresight.” Participants read about historical scenarios and were given multiple-choice questions regarding four possible outcomes. They were then told that one of the outcomes was the actual outcome and were asked the outcome that they would have guessed (his Studies 1 and 2), or that other participants would have guessed (his Study 3), if they had not been told the right answer. Their responses were famously biased toward what they were led to believe was the correct answer. For instance, those told that Option A was the right answer overestimated the probability they themselves or someone else would have guessed A. Since then, over 100 studies have run variations of this basic design (for meta-analyses see Christensen-Szalanski & Willham, 1991; Guilbault, Bryant, Brockway, & Posavac, 2004).

Past research has assumed that participants *wanted* to disregard the to-be-ignored information. This often-implicit assumption was justified because researchers (a) instructed participants to ignore the information (e.g., Fischhoff, 1975), (b) instructed participants to forecast what uninformed others would believe (e.g., Camerer et

Berkeley J. Dietvorst, Department of Marketing, The University of Chicago Booth School of Business; Uri Simonsohn, Department of Operations, Innovation and Data Sciences, ESADE Business School, Universitat Ramon Llull.

This research was presented at the Winter Decision Making Symposium, Miami, Florida (January, 2018), and the Association for Consumer Research Conference, Dallas, Texas (October, 2018).

Correspondence concerning this article should be addressed to Berkeley J. Dietvorst, Department of Marketing, The University of Chicago Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637. E-mail: berkeley.dietvorst@chicagobooth.edu

al., 1989; Fischhoff, 1975), or (c) gave incentives for accuracy of judgments that could be biased if they were influenced by the to-be-ignored information (e.g., Camerer et al., 1989; Hell, Gigerenzer, Gauggel, Mall, & Müller, 1988). The researchers assumed, then, that participants shared their belief that it was wise to disregard the to-be-ignored information.

But this need not be the case. For example, participants may believe (perhaps even correctly under some circumstances) that outcome information makes their forecasts of what uninformed others believe more accurate. Moreover, if participants hold such beliefs, incentivizing accuracy would not only be ineffective but could easily backfire. Thus, without data on participants' intentions to use or ignore the to-be-ignored information, it is uncertain whether participants in these experiments used to-be-ignored information intentionally or despite their best efforts to ignore it.

Researchers have investigated multiple mechanisms behind people's reliance on to-be-ignored information, but these do not speak to whether people share researchers' premise that the to-be-ignored information should be ignored. In particular, researchers have hypothesized that hindsight bias could be a byproduct of a number of possibilities: learning from feedback (e.g., Hawkins & Hastie, 1990; Hoch & Loewenstein, 1989; Hoffrage, Hertwig, & Gigerenzer, 2000; Wasserman, Lempert, & Hastie, 1991), biased reconstruction of judgment (e.g., Fischhoff, 1975; Hawkins & Hastie, 1990; Stahlberg & Maass, 1997), distorted memory (e.g., Fischhoff, 1975; Hawkins & Hastie, 1990), anchoring and adjustment (e.g., Fischhoff, 1975), motivation to seem intelligent¹ (e.g., Campbell & Tesser, 1983; Musch & Wagner, 2007), self-defensive processing (e.g., Louie, 1999), a motivation to make sense of the world (e.g., Pezzo, 2003; Pezzo & Pezzo, 2007), or the identification of causal antecedents of the outcome (e.g., Nestler, Blank, & von Collani, 2008; Wasserman et al., 1991; Yopchick & Kim, 2012). Further, hindsight bias could be multiply determined. Hawkins and Hastie (1990) proposed that hindsight bias could result from some combination of biased reconstruction of a previous judgment, motivation to seem intelligent, biased recall of an original belief, and anchoring on the current belief and adjustment. Roese and Vohs (2012) proposed a model where hindsight bias can result from motivation to have a certain belief, the fluency with which an outcome is understood, and cognitive processes like sensemaking and knowledge updating.

Some of these mechanisms seem more consistent with people having insight into their use of to-be-ignored information (e.g., learning from feedback) than do others (e.g., biased reconstruction of judgment). But none of them preclude or necessitate intentionality. For example, self-defensive processing, learning, or the desire to seem more intelligent could be intentional or unintentional. Thus, in our view, prior work on the mechanism behind reliance on to-be-ignored information is not necessarily informative about people's desire to use to-be-ignored information.

Past research has also explored moderators of hindsight bias, including asking participants to consider why alternative outcomes may have occurred (Arkes, Faust, Guilmette, & Hart, 1988), discrediting feedback (Hasher, Attig, & Alba, 1981), manipulating how surprising the outcome is (Ofir, Mazursky, 1997; Pezzo, 2003), and attributing the outcome to chance (Wasserman et al., 1991). This research demonstrates that people's use of to-be-ignored information is malleable, but like research on mechanisms, it also does not speak to whether use of to-be-ignored information

is intentional. For example, showing that people can ignore information after it is discredited does not imply that they could have ignored that same information if it had not been discredited. People can blink intentionally, but this does not mean that all blinking is intentional. Further, this research does not demonstrate whether this malleability in use of to-be-ignored information is mediated by changes in the desire to use information. Most important, it is not clear whether these manipulations interfere with the impact of to-be-ignored information or whether they introduce separate main effects that reduce the net overall effect. For instance, it seems likely that asking people to consider why Option B is true should reduce their tendency to choose Option A in general, whether they were previously told A was the right answer or not. Considering other outcomes may introduce a bias in the opposite direction, rather than debiasing the effect.

Researchers have demonstrated people's failure to disregard to-be-ignored information in many other contexts beyond research on hindsight bias. The possibility that participants intentionally pay attention to to-be-ignored information applies to them as well. For example, in their seminal article on the curse of knowledge, Camerer et al. (1989) incentivized participants who had private information to trade assets whose values were determined by past participants who did not have the private information. They found that participants failed to ignore their private information when trading. However, they did not ask participants whether they intended to ignore the private information. Thus, participants in this experiment may have used their private information intentionally. The results of our Study 2 suggest that they did.

Many studies have demonstrated that jurors fail to disregard inadmissible evidence after they are instructed to do so (for a meta-analysis see Steblay et al., 2006). These studies did not address the question of whether participants intended to ignore the inadmissible evidence.² Further, researchers have found factors that moderate jurors' use of inadmissible evidence, including suspicion about why the information was introduced (Fein, McCloskey, & Tomlinson, 1997; Fein, Morgan, Norton, & Sommers, 1997), the reason given for disregarding the evidence (Kassin & Sommers, 1997; Pickel, 1995), and appeals explaining that outcome information can make reasonable judgments look bad in hindsight (Stallard & Worthington, 1998). But again, the fact that jurors are better at disregarding inadmissible evidence under some conditions does not imply that their use of inadmissible evidence is intentional in other conditions.

Overall, our read of the literature is that when people have been shown to rely on to-be-ignored information, it is not clear whether they used that information intentionally or used it despite their intentions to ignore it. In this article, we seek to tease apart these two alternative possibilities. We conclude that people often use to-be-ignored information intentionally.

¹ Some research has failed to find that motivation to seem intelligent drives hindsight bias (Connolly & Bukszar, 1990; Leary, 1981, 1982; for a review see Pezzo, 2011).

² Kassin and Sukel (1997) proposed that mock jurors in their studies used inadmissible evidence even "when jurors said it had no influence" (p. 27). However, the article does not report results for jurors who indicated that the to-be-ignored information did not influence their judgment. Further, the studies did not show significant differences in guilty verdicts across the conditions of interest.

Why People May Want to Use To-Be-Ignored Information

There are many reasons why people may use to-be-ignored information intentionally. First, people may believe that doing so improves their accuracy, and in some cases they are probably right. For example, imagine a participant who reads a multiple-choice question that she does not know the answer to, learns the correct answer, and is then asked how others would answer that question. If she believes that some people know the correct answer to the question (presumably some people do), she could shade her answer toward the correct answer and possibly increase her accuracy (and note that all that is needed to produce an effect is for some participants believe that this is true; as participants, both authors would).³ Additionally, people may use to-be-ignored information in these ways: in an attempt to “appear intelligent, knowledgeable, or perspicacious” (see Hawkins & Hastie, 1990, p. 316; see also Campbell & Tesser, 1983), as a shortcut if they believe that they would have figured out the correct answer with enough effort, because relying on the to-be-ignored information requires lower effort than does generating an independent answer, because they misunderstood or chose not to follow the instructions, or for other reasons that we have not considered. Thus, we conjecture that people often use to-be-ignored information intentionally.

Predictions From the Hypothesis That People Intentionally Use To-Be-Ignored Information

Our hypothesis that people intentionally use to-be-ignored information led us to several predictions that we tested in studies reported here. First, we predicted that when asked whether they wanted to ignore to-be-ignored information, a substantial share of people would indicate they did not want to ignore it. As predicted, across Studies 1–3, only 34% of participants in classic paradigms from past research agreed that to-be-ignored information should actually be ignored. Second, we predicted that people who desire to ignore the information do so and that the previously documented biases are driven by the (majority of) participants who want to use the information. As predicted, in Studies 3 and 4a we did not find that obtaining to-be-ignored information had a significant effect (either practically or statistically significant) on participants’ estimates, *after accounting for their plan to use that information* (the other studies do not speak to this hypothesis). Third, we predicted that giving people persuasive arguments about the importance of disregarding information would increase the degree to which they ignored such information. As predicted, in Study 4 we found that such arguments reduced people’s plan to ignore the to-be-ignored information and their use of that information as a result.

We conclude that when people use to-be-ignored information, they may not be acting rationally but they are often acting intentionally.

General Method

Overview of Studies

In Study 1, we used materials from the classic Fischhoff (1975) hindsight bias paradigm and found that the majority of participants desired to use the to-be-ignored information (his Experiment 3). In Study 2 we replicated this finding for the curse-of-knowledge paradigm by Camerer et al. (1989). In Study 3 we used a novel

paradigm and found that people who desire to ignore to-be-ignored information do indeed ignore it. Study 4 employed a mock-jury paradigm to show that providing different reasons for ignoring inadmissible evidence changes people’s reliance on that evidence by changing their intention to use it, and demonstrates evidence for our propositions in a practically important domain.

Transparent Reporting

For each study, we set the sample size before any data were collected and collected over 50 observations per condition (see Simmons, Nelson, & Simonsohn, 2013). We report all exclusions (if any), all manipulations, and all measures in the article and online supplemental materials. We also include additional details for each study and report results from additional studies that conceptually replicate our findings in the online supplemental materials. We preregistered Studies 2, 3, and 4a in AsPredicted.org. On this project’s OSF page (<https://osf.io/uvqhs/>), we posted raw data, code for all analyses, original study materials, and preregistrations. We ran all studies on MTurk using institutional review board–approved procedures. Workers were required to be in the United States and have an approval rating of at least 90%. We included several attention and comprehension checks based on which we carried out exclusions. All exclusions took place before random assignment, and we detail them in Supplement 1 in the online supplemental materials. In all studies with comprehension checks, participants who failed any one check were not allowed to continue in the survey, were not assigned to a condition, and did not complete the dependent measure. The sample sizes presented in our studies are for reflect the net number of participants who completed the dependent measure, after exclusions.

Study 1—Intentional Hindsight Bias

Method

We ran three similar studies designed to directly test our hypothesis that participants in hindsight bias experiments intentionally rely on the to-be-ignored information. These studies were based on the paradigm created by Fischhoff (1975). We focused on the key elements of the study designs here and described them in full detail in Supplements 2–4 in the online supplemental materials. We avoided paradigms where participants reported their own likelihood of selecting the correct answer to a question after learning to-be-ignored information in order to reduce concerns that participants’ responses were affected by a desire to appear intelligent or demonstrate that they would have selected the correct answer without learning the to-be-ignored information (see Hawkins & Hastie, 1990).

In the three studies, participants read a historical passage used by Fischhoff (1975) describing a conflict between the British and the Gurka in Nepal.⁴ Then they were told an outcome for the

³ This argument is analogous to the rationalization of the false-consensus effect (Dawes, 1990).

⁴ This passage is reproduced in Supplement 2 in the online supplemental materials.

conflict: For example, the Gurka won.⁵ Then they were told about a multiple-choice question that other people, who were not told the outcome, had been asked. The possible answers to this question were the same ones used by Fischhoff: (a) *The Gurka won*, (b) *The British won*, (c) *Military stalemate with no peace settlement*, and (d) *Military stalemate with a peace settlement*. Finally, in all three studies the dependent variable involved asking participants whether using versus ignoring the outcome information would increase accuracy in estimating the percentage of respondents choosing a, b, c, or d.⁶

We ran three similar studies with different dependent measures to assure that our results did not depend on how we chose to ask the question of whether using outcome information improves accuracy. Moreover, each study had detailed instructions and presented information sequentially in separate screens, and Studies 1b and 1c included four comprehension check questions. The original materials are available in the online supplemental materials.

Results

Study 1a. In Study 1a, participants ($N = 393$, $M_{\text{age}} = 35$; 46% female) were asked “If your task were to estimate the shares of the [respondents] who gave an answer of ‘a’, ‘b’, ‘c’, and ‘d’ to the multiple choice question, do you think that using the ‘additional information’ shown above (The Gurka won this conflict) would make you more or less accurate?” Their options were “I think my answer would be more accurate if I used the ‘additional information’ to some degree” and “I think my answer would be more accurate if I completely ignored the ‘additional information’” (emphasis as presented). The majority of participants (73%; 287/393) indicated that using the outcome information would increase accuracy.

Study 1b. In Study 1b, participants ($N = 295$, $M_{\text{age}} = 37$; 55% female) learned that a university student, John, had estimated the proportion of other university students giving each answer both before and after being told the actual outcome, and that the estimates were different. Participants were not permitted to complete the dependent measure if they failed any one of four comprehension check questions asking the following: how many estimates John made (2), what John’s occupation was (university student), who the Gurka were in conflict with (the British), and where the conflict took place (Nepal). Participants then chose between these two options: “I think that *estimate #1* (John’s estimate before he was told the outcome of the conflict) was more accurate” and “I think that *estimate #2* (John’s estimate after he was told the outcome of the conflict) was more accurate”⁷ (emphasis as presented). Once again, a majority of participants (67%; 198/295) indicated that the outcome information would increase accuracy.

Study 1c. In this study, participants ($N = 274$, $M_{\text{age}} = 34$; 47% female) were told that two different university students, John and Robert,⁸ had made predictions with versus without the outcome information. Participants were not permitted to complete the dependent measure if they failed any one of four comprehension check questions asking the following: what the subject of the multiple-choice question was (history), what John and Robert’s occupations were (university students), who the Gurka were in conflict with (the British), and where the conflict took place (Nepal). Participants chose between these two counterbalanced options: “I think that John [Robert], who was *not told the outcome of the conflict*, was more accurate” versus “I think that Robert [John], who was *told the outcome of the conflict*, was more

accurate” (emphasis as presented). Once again, a majority of participants (62%; 171/274) indicated that the outcome information would increase accuracy.

Discussion

Whereas past research has assumed that people who are making estimates of uninformed others’ judgments should ignore private outcome information that they have learned, the results from Study 1 suggest that many people intentionally attend to this information. Thus, although the results from past research do demonstrate that outcome information can bias forecasts, it does not demonstrate this occurs because people are unable to ignore to-be-ignored information. Study 2 explores whether people intentionally used to-be-ignored information in another classic bias: the *curse of knowledge* by Camerer et al. (1989), where MBA students with outcome information estimated how uninformed MBA students valued assets.

Study 2—Intentional Bias in Curse of Knowledge

We asked respondents to indicate whether they believed participants in the Camerer et al. (1989) curse-of-knowledge study would be more accurate if they did versus did not ignore the to-be-ignored information.

Method

Respondents ($N = 113$, $M_{\text{age}} = 39$; 52% female) read that they would learn about estimates that two finance students made during an exercise that contained two stages. In particular, respondents learned that in Stage 1, a group of finance students were told a company’s earnings for the previous 10 years, and then they predicted its earnings for the following quarter. In Stage 2, several months later, we asked two new students, Participant A and Participant B, to make estimates about what their fellow students had predicted earlier in Stage 1. Respondents were not permitted to continue if they failed any one of four comprehension check questions asking the following: how many estimates they would answer questions about (two), what Participant A and B studied (finance), what the large group of finance students predicted (a company’s earnings), and what A and B estimated (the large group of finance students’ predictions about the company’s earnings).

Participant A (B in the counterbalanced version) learned the company’s actual earnings before making their estimate, and Participant B (A) did not. Respondents read that A and B gave different answers. Our respondents then indicated whether they

⁵ In Study 1a, all participants learned that the Gurka won the conflict. In Studies 1b and 1c, the “correct answer” was randomly selected.

⁶ In Study 1a, we manipulated whether the other people asked the question were MTurk participants versus college students. Contrary to our expectations, there was no effect of this manipulation: The shares of people wishing to use the information were quite similar ($M_{\text{MTurk}} = 73.74\%$, $M_{\text{Students}} = 72.31\%$), $\chi^2(1, N = 393) = .10, p = .750$. This led us to revise our hypothesizing for future studies, and we did not rely on such a manipulation again.

⁷ We counterbalanced both the order of these options and which was labeled #1 and #2.

⁸ We randomized which name appeared first in the instructions and orthogonally randomized which name was later associated with learning the outcome information.

thought A or B would be more accurate in estimating the group of finance students’ estimates, by choosing between these two counterbalanced options: “I think that Participant A [B] (who was *not told the company’s earnings*) was more accurate” and “I think that Participant B [A] (who was *told the company’s earnings*) was more accurate” (emphasis as presented).

Results

Consistent with Study 1, and with the notion that people use to-be-ignored information intentionally, we found that 81% (91/113) of participants thought that the participant who had learned the company’s actual earnings would be more accurate in estimating what uninformed students predicted. In our next study, we explore whether use of to-be-ignored information is driven by the majority of participants who see value in using to-be-ignored information.

Study 3—Participants Show Hindsight Bias Only if They Want To

In a two-cell between-subjects design, half of participants were assigned to learn to-be-ignored information (the answer to a multiple-choice question) and the other half of participants were assigned not to. Before exposure to the to-be-ignored information, all participants reported whether they would use the to-be-ignored information if they learned it. We tested whether the overall reliance on to-be-ignored information was driven by the subset of participants who wanted to use it, by testing for hindsight bias within participants who indicated they wanted versus did not want to use such information.

Method

Participants ($N = 511$, $M_{\text{age}} = 34$; 49% female) answered 10 multiple-choice trivia questions, drawn at random out of a pool of 11 questions, and learned how many they answered correctly.⁹ Next they were asked to predict how another participant, who answered the same number of questions correctly, would answer an 11th trivia question. Participants then read the 11th trivia question and answered these two questions regarding how they would use outcome information: “Do you think that you could more accurately predict the past participant’s answer if you learned the correct answer to this question?” and “If you learned the correct answer to this question, would you use that information to predict the past participant’s answer?” Next, participants were randomly assigned to learn versus not learn the correct answer. Then they completed the study’s dependent measure: estimating the past participant’s probability of choosing each of the four possible answers to the 11th trivia question.

Results

First, aggregating across all participants, we replicated the hindsight bias: Participants who learned the right answer gave a higher estimate of the popularity of that answer ($M = 47.3\%$) than did those who did not learn the right answer ($M = 37.7\%$), $t(509) = -4.08$, $p < .001$, $d = .36$.¹⁰ Note that if participants thought all four answers were equally probable, then this mean would be 25%.

Second, we conceptually replicated Studies 1 and 2; a majority of participants (64%; 328/511) answered “yes” to at least one of the questions inquiring about a desire to use the to-be-ignored informa-

tion. In particular, 45% (232/511) indicated they believed they could more accurately predict the past participant’s answer if they learned the correct answer, and 59% (300/511) reported that they would use the correct answer to predict the past participant’s answer if they learned it.

Last, and most important for the purposes of this study, we found that the hindsight bias was only observed among participants who indicated that they desired to use the to-be-ignored information (see Figure 1). In particular, participants who answered “no” to both of the intention-to-use-information questions made a similar guess regarding the popularity of the right answer when told the right answer ($M = 42.6\%$) versus not ($M = 43.2\%$), $t(181) = .17$, $p = .867$, $d = .03$. In contrast, participants who answered “yes” to at least one question gave a higher estimate when told the right answer ($M = 49.9\%$) versus not ($M = 34.6\%$), $t(326) = -5.23$, $p < .001$, $d = .58$. The interaction was statistically significant ($b = 15.98$), $t(507) = 3.28$, $p = .001$. If people are classified as wishing to use the information based on answering “yes” to just the first or second question, rather than to both, statistically significant interactions in the predicted direction are also obtained ($p = .001$ and $p = .019$, respectively).¹¹

A challenge to the interpretation we are giving to the results from Study 3 is the possibility that participants who indicated they did not plan to use the information about the correct answer did so because they already knew the correct answer. In other words, they were not affected by the information given, because they already had that information. We do not believe in this interpretation of the results for three reasons. First, it predicts that, because participants who had more correct answers among the first 10 trivia questions would be more likely to know this answer, they would be more likely to indicate they

⁹ See Supplement 7 in the online supplemental materials for a list of all questions.

¹⁰ We used the formula $d = 2 \times t/\sqrt{df}$ to calculate the effect sizes for t tests.

¹¹ We designate the predicted and actual share of people guessing the right answer as p and s , respectively. Our preregistration indicated that p would be our dependent variable (DV) but that we would also report the results using $(p - s)/s$ (standardizing across the 11 different questions by dividing by their actual share of right answers). With this alternative DV we also replicated the hindsight bias aggregating across all participants, $t(509) = -3.94$, $p = .001$, $d = .35$. This was again driven by the subset of participants planning to use the information, $t(326) = -4.14$, $p < .001$, $d = .46$, and there was no significant hindsight bias among those who didn’t plan to do so, $t(181) = -.938$, $p = .350$, $d = .14$. But the interaction was statistically much weaker, $t(507) = 1.84$, $p = .066$, $b = .27$. This could occur for a few reasons. First, the dependent variable could be noticeably noisier because it incorporates s —if participants don’t predict s very well, taking s into account adds rather than reduces noise. Second, the possibility we worried about, it could indicate that our preferred specification is spurious, somehow capitalizing on differences across the 11 questions. For robustness, therefore, we ran a nonpreregistered regression with p as the dependent variable, as in the main text, but controlling for s rather than building it into the dependent variable. This managed to take out systematic variation without amplifying random variation. Using this specification, the interaction was again statistically strong ($b = 12.75$), $t(506) = 2.74$, $p = .006$.

Following the suggestion of a reviewer, we report the pairwise differences across bars in Figure 1: The first versus third and first versus fourth bars are $t(258) = 2.57$, $p = .011$, and $t(255) = -2.42$, $p = .016$, respectively. The second versus third and second versus fourth are $t(252) = -1.87$, $p = .063$, and $t(249) = -2.08$, $p = .039$, respectively.

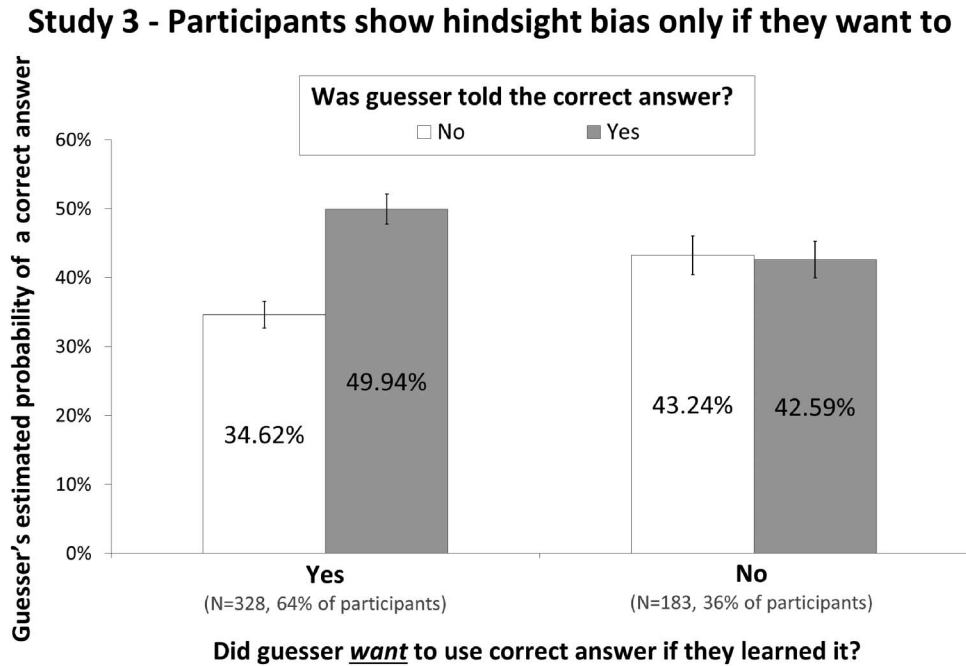


Figure 1. Participants guessed the probability that a past participant would choose each of four answers to a multiple-choice trivia question. Before that, they indicated whether they would like to take into account what the right answer was when making their guesses. People exhibit hindsight bias when they give the right answer a higher probability after learning the answer. Only participants who wanted to use such information exhibited the hindsight bias. Error bars show 1 standard error above or below sample means.

did not plan to use the information. But, despite ample variation in correct answers across participants ($SD = 1.83$), the correlation between participants' number of correct answers and their indication that they would use the information is negligible, $r(N = 511) = .016, p = .711$. Second, and in a similar vein, this explanation predicts that people who are being asked an easier 11th question will be more likely to say they do not plan to use the answer (recall that we randomized the order of the questions across participants). Despite ample variation in the percentage of correct answers across questions ($SD = 16.56\%$), the correlation between the 11th question difficulty and plans to use the answer is negligible, $r(N = 511) = -.047, p = .288$. Last, because these questions were difficult (50% of people answered them correctly on average), among those who indicated planning not to use the information, there must have been a substantial share who had wrong beliefs about the right answer, and thus, they must have been receiving new information when we revealed it. Thus, we would expect them to have exhibited a larger effect from having obtained the information.

Study 4 Overview—Persuading Jurors to Ignore Inadmissible Evidence

In our last study, we manipulated, rather than measured, intention to use to-be-ignored information. Study 4 is based on research by Kassin and Sommers (1997), who showed that mock jurors are more likely to disobey a judge's instructions to ignore inadmissible evidence when the reason is "procedural" (e.g., an audiotape was obtained without a warrant) than when the reason is "causal" (e.g.,

an audiotape is difficult to understand, and the transcript lacks credibility).

The notion that people do not ignore information when they do not want to, and that people might not want to ignore information for procedural reasons, provides a parsimonious explanation for the aforementioned result.¹² Moreover, this explanation leads to two predictions that we tested in Studies 4a and 4b, respectively. In Study 4a, we replicated past studies adding a question about intention to use the to-be-ignored information. We tested whether the desire to ignore the to-be-ignored information mediates the previously documented pattern that causal reasons to ignore information are more effective than are procedural reasons. In Study 4b, we tested whether providing a more persuasive justification for the procedural instruction to ignore information reduces reliance on such information.

Study 4a—Jurors Ignore Inadmissible Evidence if They Want To

Method

Participants ($N = 480, M_{\text{age}} = 36; 61\%$ female) read about a case involving a man charged with murdering his estranged wife

¹² See Study S4a in Supplement 11 of the online supplemental materials, in which we hypothesized and found that participants' preference for a juror who uses the inadmissible evidence (vs. ignores it) followed the same pattern that Kassin and Sommers (1997) found for guilty verdicts.

and male neighbor. They were then asked to answer questions as if they were a juror in the trial. The prosecutor claimed that the defendant killed the victims in a fit of jealous rage. The defendant said that he found the bodies when he returned to his former home to retrieve personal papers. Participants had to correctly answer the following four multiple-choice comprehension questions, on the first attempt, in order to proceed: how many victims there were (two), what the charge was in the trial (murder), how the defendant was related to the first victim (estranged wife), and how the defendant was related to the second victim (neighbor).

Next, participants read a police officer’s testimony that consisted of reading off a transcript made from a wiretapped telephone conversation in which the defendant appeared to confess to a friend after fleeing the scene. As in Kassin and Sommers (1997), the wording of the transcript was:

“I killed Mary and some bastard she was with. God, I don’t . . . yeah, I ditched the blade.”

Then participants were randomly assigned to one of three conditions, differing in what they were told to do with this evidence. In particular, they read this (with wording not used by Kassin and Sommers [1997]):

At this point the defense lawyer objects.

[“Admissible” condition]

The judge overrules the objection and admits the evidence as valid.

[Inadmissible/“Due-process” condition]

The judge sustains the objection and admonishes the jury to disregard the evidence because the phone was intercepted without a warrant, deeming it an illegal search.

[Inadmissible/“Unreliable” condition]

The judge sustains the objection and admonishes the jury to disregard the evidence because the original audio is of poor quality (it is barely understandable) and hence it is difficult to determine what was actually said during the call. The transcript that was read is not trustworthy.

Next, participants answered two questions, which we used as the dependent measure and the mediator, in counterbalanced order. The dependent measure was a binary choice: “I would vote not guilty” versus “I would vote guilty.” The mediator was also binary: “My plan is [was] to completely disregard the police officer’s testimony when deciding on my verdict” versus “My plan is [was] to take the police officer’s testimony into account when deciding on my verdict.”

Results

The results for the judgment of guilt (see Figure 2) were qualitatively consistent with those reported by Kassin and Sommers




Figure 2. Participants gave guilty or not-guilty votes on a hypothetical murder case that included an audio recording of an apparent confession. Across conditions, the confession was deemed admissible, inadmissible because it was obtained without a warrant, or inadmissible because the tape was not clear. Results are shown controlling (direct effect) and not controlling (total effect) for self-reported intention to use the tape when voting (mediator). Error bars show 1 standard error above or below coefficient estimates.

(1997), who found that 79%, 55%, and 24% of participants voted guilty in the three admissible, no-warrant, and invalid-information conditions, respectively ($p = .1049$). We preregistered that we would test to see whether these differences in guilty verdicts were explained by participants' desire to ignore the information from the wiretap. This conjecture was supported by the black dotted line in Figure 2, which represents the conditional mean controlling for the mediator (i.e., the direct effect). The flatter pattern indicates that the differences in guilty verdicts between conditions are almost entirely accounted for by participants' desire to ignore the information. The omnibus test dropped from $F(2, 477) = 50.61, p < .001$, to $F(2, 476) = 2.26, p = .105$, after accounting for participants' plan to either ignore or attend to the wiretap. Bootstrapped confidence intervals (CIs) for the indirect effects (mediation) do not include 0 for either the omnibus test [33.24, 68.21] or for each of the three pairwise comparisons (Admissible = Unreliable [65.34, 139.48], Due-process = Admissible [18.72, 59.54], Unreliable = Due-process [1.28, 29.26]). Thus, the results of this study are consistent with the notion that the evidence affected participants' judgments only when they intended to be affected by it.

Study 4b—Persuading to Ignore Evidence Obtained Without a Warrant

Thus far, our studies have provided evidence consistent with the ideas that (a) people use to-be-ignored information intentionally, (b) people who intend to disregard to-be-ignored information can ignore it, and (c) people's use of to-be-ignored information can be manipulated. These findings suggest a novel intervention to get jurors to ignore evidence for procedural reasons: persuade them that the procedural reasons for ignoring evidence have merit. In Study 4b, we investigated this intervention by modifying the paradigm from Study 4a to manipulate the strength of the procedural arguments given for ignoring evidence. In other words, we tested whether one can convince people to ignore information without discrediting that information—by “simply” providing stronger arguments for ignoring it.

Method

A total of 571 MTurk participants passed the attention and comprehension checks and completed the dependent measure ($M_{\text{age}} = 37$; 58% female). In relation to Study 4a, we collected (only) the measure of guilt vote, and we added three conditions: a control condition where participants did not learn about the wiretapped phone call and two conditions where the judge gave either strong or very strong procedural reasons to ignore evidence obtained without a warrant. The latter two read:

Strong Procedural Reason

The judge sustains the objection and admonishes the jury to disregard the evidence because the phone was intercepted without a warrant, deeming it an illegal search. The judge explains that it is very important not to consider evidence that was gathered with an illegal search. In fact, illegal searches violate constitutional rights and threaten your own individual privacy.

Very Strong Procedural Reason

[Strong reason, followed by] If we allowed such evidence, the police will conduct searches proactively, routinely violating our rights and

probably especially those of minorities. For example, the police could pull you over for a routine traffic stop and search your glove compartment, phone, trunk, and so forth without probable cause, and anything they found could be used against you.

Results

As predicted, the three warrantless search conditions with different justifications to ignore the evidence did exhibit significantly different guilty vote rates. We found significant differences across the three no-warrant conditions ($N = 285$) in a probit regression, with a predictor taking values of 1, 2, and 3 for the three levels of strength of argument to ignore evidence ($b = -.253, z = -2.69, p = .007$; with ordinary least squares, also $p = .007$).¹³ Thus, providing stronger reasons to ignore evidence collected without a warrant did get participants to ignore the evidence more often, even though these arguments did not discredit the validity of the inadmissible evidence.

Turning our attention to comparing the condition with the very strong argument to ignore evidence obtained without a warrant and the condition where the audio-recording was deemed unreliable (last two bars in Figure 3), we found a small (53% vs. 46%) and not statistically significant, $\chi^2(1, N = 191) = .88, p = .347$, difference. This suggests that procedural reasons to ignore evidence may have the potential to be as effective as causal reasons to ignore evidence if participants find the procedural reasons to be persuasive. However, the confidence interval around the observed 6.8% difference [-7% , 21%] includes large effects, so we cannot say that the difference was completely eliminated by the very strong argument that we chose for this study.

After running this study, we worried that giving stronger arguments to ignore warrantless searches may have created a demand effect where participants felt pressured to vote not guilty. Thus, we ran a new study (see Study S4b in Supplement 12 in the online supplemental materials), where we presented participants with two pieces of evidence, one trivial and one consequential. We manipulated whether the strongest argument to ignore the evidence applied to the trivial evidence, the consequential evidence, both, or neither. We found that when only the trivial evidence was discredited with the strongest argument, participants voted guilty (94%) at virtually the same rate as when neither piece of evidence was discredited with the strongest argument (95%). Thus, we believe it is unlikely that the results of Study 4b were due to experimenter demand to vote not guilty.

General Discussion

In Studies 1 and 2, we found that participants exposed to classic paradigms disagreed with the premise that the to-be-ignored information should in fact be ignored. In Study 3, we documented that only participants who indicated that they would like to use to-be-ignored information were affected by it. In Study 4, we manipulated participants' desire to ignore information by providing stronger reasons to do so, finding that this indeed reduced their intentions to use, and reliance on, to-be-ignored information. Taken together, we interpret our findings as indicating that use of

¹³ The lower powered nondirectional omnibus test is, $\chi^2(2, N = 285) = 7.30, p = .026$.

Study 4b - Impact of audio confession on guilty vote




Figure 3. Participants ($N = 571$) read about a scenario involving a murder case. In all conditions except the control, they were told of a transcript of a confession made on a recorded phone call. The defense lawyer objects, and the conditions differ on the judge's response. Error bars show 1 standard error above or below the mean.

to-be-ignored information is often intentional, arising from people's desire to use information that decision-making researchers believe they should not rely on.

It is important to note that we do not interpret our findings as showing that desiring to ignore information is a sufficient condition for successfully doing so. We believe there may be cases where, even if people wanted to ignore information, they may be unable to, such as when it affects perception and thus high-order cognition cannot easily intervene—for example in studies on auditory hindsight bias (Higham, Neil, & Bernstein, 2017), visual hindsight bias (Harley, Carlsen, & Loftus, 2004), and change-blindness blindness (Levin, Momen, Drivdahl, & Simons, 2000). Although not a *sufficient* condition, our findings suggest that the desire to ignore information is a *necessary* condition. Further, the studies presented in this article demonstrate that previous research on use of to-be-ignored information did not meet this necessary condition, and suggest that participants could have ignored the to-be-ignored information if this necessary condition had been satisfied. Our findings may also explain why Simonsohn (2011) found that when Consumer Reports retracted a car seat safety ranking, people were able to successfully ignore the retracted information: The instruction to ignore the information was persuasive, perhaps because it came from the original source.

Although we cannot say that meeting this condition in any context would be enough to eliminate use of to-be-ignored information, we provided evidence suggesting that it is probably enough to eliminate it in some of the famous contexts that researchers have studied. These findings should fundamentally change the way that one thinks about people's ability to disregard to-be-ignored information.

Taking this necessary condition into account led us to reinterpret the large literature on the use of to-be-ignored information. First, it suggests refocusing the attention away from possible unconscious influences on people's behavior until one has accounted for conscious influences on people's behavior. Second, by identifying a different psychological mechanism for the impact of to-be-ignored information, our results have novel practical implications that provide concrete debiasing tools to avoid use of to-be-ignored information.

For example, our findings suggest that one can sort people based on their desire to use to-be-ignored information. As an illustration, a judge could ask jurors whether they are attending to inadmissible evidence, trust that jurors have the ability to accurately report the truth, and dismiss jurors who report that they are using such information. Our findings also suggest methods for debiasing individual decision makers: persuade decision makers to ignore information that needs to be ignored. For example, teachers who struggle to teach to the level of knowledge of their students may avoid this curse of knowledge by simply giving a diagnostic test to the students before the course begins, so that they know what information they can and cannot take as known by everyone; our results suggest that teachers could teach without relying on concepts that are unknown to their students if they simply knew which concepts the students had yet to learn.

It is important to clarify what we do and do not believe we are documenting. First, we have not studied situations where a person who wants to use to-be-ignored information attempts to take the perspective of someone who doesn't want to use the information. Second, we did not investigate the reasons why people desire to use to-be-ignored information. Third, we do not take a stance on the issue of whether the use of to-be-ignored information is nor-

mative, is rational, or boosts accuracy. To be clear, we are not proposing that people use to-be-ignored information normatively; we are proposing they often use it intentionally.

References

- Anderson, C. A., Lepper, M. R., & Ross, L. (1980). Perseverance of social theories: The role of explanation in the persistence of discredited information. *Journal of Personality and Social Psychology*, *39*, 1037–1049. <http://dx.doi.org/10.1037/h0077720>
- Arkes, H. R., Faust, D., Guilmette, T. J., & Hart, K. (1988). Eliminating the hindsight bias. *Journal of Applied Psychology*, *73*, 305–307. <http://dx.doi.org/10.1037/0021-9010.73.2.305>
- Baron, J., & Hershey, J. C. (1988). Outcome bias in decision evaluation. *Journal of Personality and Social Psychology*, *54*, 569–579. <http://dx.doi.org/10.1037/0022-3514.54.4.569>
- Camerer, C., Loewenstein, G., & Weber, M. (1989). The curse of knowledge in economic settings: An experimental analysis. *Journal of Political Economy*, *97*, 1232–1254. <http://dx.doi.org/10.1086/261651>
- Campbell, J. D., & Tesser, A. (1983). Motivational interpretations of hindsight bias: An individual difference analysis. *Journal of Personality*, *51*, 605–620. <http://dx.doi.org/10.1111/j.1467-6494.1983.tb00868.x>
- Christensen-Szalanski, J. J., & Willham, C. F. (1991). The hindsight bias: A meta-analysis. *Organizational Behavior and Human Decision Processes*, *48*, 147–168. [http://dx.doi.org/10.1016/0749-5978\(91\)90010-Q](http://dx.doi.org/10.1016/0749-5978(91)90010-Q)
- Connolly, T., & Bukszar, E. W. (1990). Hindsight bias: Self-flattery or cognitive error? *Journal of Behavioral Decision Making*, *3*, 205–211. <http://dx.doi.org/10.1002/bdm.3960030305>
- Dawes, R. M. (1990). The potential nonfalsity of the false consensus effect. In R. M. Hogarth (Ed.), *Insights in decision making: A tribute to Hillel J. Einhorn* (pp. 179–199). Illinois: University of Chicago Press.
- Fein, S., McCloskey, A. L., & Tomlinson, T. M. (1997). Can the jury disregard that information? The use of suspicion to reduce the prejudicial effects of pretrial publicity and inadmissible testimony. *Personality and Social Psychology Bulletin*, *23*, 1215–1226. <http://dx.doi.org/10.1177/01461672972311008>
- Fein, S., Morgan, S. J., Norton, M. I., & Sommers, S. R. (1997). Hype and suspicion: The effects of pretrial publicity, race, and suspicion on jurors' verdicts. *Journal of Social Issues*, *53*, 487–502. <http://dx.doi.org/10.1111/j.1540-4560.1997.tb02124.x>
- Fischhoff, B. (1975). Hindsight is not equal to foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, *1*, 288–299. <http://dx.doi.org/10.1037/0096-1523.1.3.288>
- Guilbault, R. L., Bryant, F. B., Brockway, J. H., & Posavac, E. J. (2004). A meta-analysis of research on hindsight bias. *Basic and Applied Social Psychology*, *26*, 103–117. <http://dx.doi.org/10.1080/01973533.2004.9646399>
- Harley, E. M., Carlsen, K. A., & Loftus, G. R. (2004). The “saw-it-all-along” effect: Demonstrations of visual hindsight bias. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 960–968. <http://dx.doi.org/10.1037/0278-7393.30.5.960>
- Hasher, L., Attig, M. S., & Alba, J. W. (1981). I knew it all along: Or, did I. *Journal of Verbal Learning and Verbal Behavior*, *20*, 86–96. [http://dx.doi.org/10.1016/S0022-5371\(81\)90323-6](http://dx.doi.org/10.1016/S0022-5371(81)90323-6)
- Hawkins, S. A., & Hastie, R. (1990). Hindsight: Biased judgments of past events after the outcomes are known. *Psychological Bulletin*, *107*, 311–327. <http://dx.doi.org/10.1037/0033-2909.107.3.311>
- Hell, W., Gigerenzer, G., Gauggel, S., Mall, M., & Müller, M. (1988). Hindsight bias: An interaction of automatic and motivational factors? *Memory & Cognition*, *16*, 533–538. <http://dx.doi.org/10.3758/BF03197054>
- Higham, P. A., Neil, G. J., & Bernstein, D. M. (2017). Auditory hindsight bias: Fluency misattribution versus memory reconstruction. *Journal of Experimental Psychology: Human Perception and Performance*, *43*, 1144–1159. <http://dx.doi.org/10.1037/xhp0000405>
- Hoch, S. J., & Loewenstein, G. F. (1989). Outcome feedback: Hindsight and information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 605–619. <http://dx.doi.org/10.1037/0278-7393.15.4.605>
- Hoffrage, U., Hertwig, R., & Gigerenzer, G. (2000). Hindsight bias: A by-product of knowledge updating? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 566–581. <http://dx.doi.org/10.1037/0278-7393.26.3.566>
- Kassin, S. M., & Sommers, S. R. (1997). Inadmissible testimony, instructions to disregard, and the jury: Substantive versus procedural considerations. *Personality and Social Psychology Bulletin*, *23*, 1046–1054. <http://dx.doi.org/10.1177/01461672972310005>
- Kassin, S. M., & Sukel, H. (1997). Coerced confessions and the jury: An experimental test of the “harmless error” rule. *Law and Human Behavior*, *21*, 27–46. <http://dx.doi.org/10.1023/A:1024814009769>
- Leary, M. R. (1981). The distorted nature of hindsight. *Journal of Social Psychology*, *115*, 25–29. <http://dx.doi.org/10.1080/00224545.1981.9711984>
- Leary, M. R. (1982). Hindsight distortion and the 1980 presidential election. *Personality and Social Psychology Bulletin*, *8*, 257–263. <http://dx.doi.org/10.1177/0146167282082012>
- Levin, D. T., Momen, N., Drivdahl, S. B., IV, & Simons, D. J. (2000). Change blindness blindness: The metacognitive error of overestimating change-detection ability. *Visual Cognition*, *7*, 397–412. <http://dx.doi.org/10.1080/135062800394865>
- Louie, T. A. (1999). Decision makers' hindsight bias after receiving favorable and unfavorable feedback. *Journal of Applied Psychology*, *84*, 29–41. <http://dx.doi.org/10.1037/0021-9010.84.1.29>
- Musch, J., & Wagner, T. (2007). Did everybody know it all along? A review of individual differences in hindsight bias. *Social Cognition*, *25*, 64–82. <http://dx.doi.org/10.1521/soco.2007.25.1.64>
- Nestler, S., Blank, H., & von Collani, G. (2008). Hindsight bias doesn't always come easy: Causal models, cognitive effort, and creeping determinism. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*, 1043–1054. <http://dx.doi.org/10.1037/0278-7393.34.5.1043>
- Ofir, C., & Mazursky, D. (1997). Does a surprising outcome reinforce or reverse the hindsight bias? *Organizational Behavior and Human Decision Processes*, *69*, 51–57. <http://dx.doi.org/10.1006/obhd.1996.2671>
- Pezzo, M. V. (2003). Surprise, defence, or making sense: What removes hindsight bias? *Memory*, *11*, 421–441. <http://dx.doi.org/10.1080/09658210244000603>
- Pezzo, M. V. (2011). Hindsight bias: A primer for motivational researchers. *Social and Personality Psychology Compass*, *5*, 665–678. <http://dx.doi.org/10.1111/j.1751-9004.2011.00381.x>
- Pezzo, M. V., & Pezzo, S. P. (2007). Making sense of failure: A motivated model of hindsight bias. *Social Cognition*, *25*, 147–164. <http://dx.doi.org/10.1521/soco.2007.25.1.147>
- Pickel, K. L. (1995). Inducing jurors to disregard inadmissible evidence: A legal explanation does not help. *Law and Human Behavior*, *19*, 407–424. <http://dx.doi.org/10.1007/BF01499140>
- Roese, N. J., & Vohs, K. D. (2012). Hindsight bias. *Perspectives on Psychological Science*, *7*, 411–426. <http://dx.doi.org/10.1177/1745691612454303>
- Simmons, J., Nelson, L., & Simonsohn, U. (2013). Life after P-Hacking [Abstract]. *SSRN Electronic Journal*. Advance online publication. <http://dx.doi.org/10.2139/ssrn.2205186>
- Simonsohn, U. (2011). Lessons from an “oops” at Consumer Reports: Consumers follow experts and ignore invalid information. *Journal of Marketing Research*, *48*, 1–12. <http://dx.doi.org/10.1509/jmkr.48.1.1>

- Stahlberg, D., & Maass, A. (1997). Hindsight bias: Impaired memory or biased reconstruction? *European Review of Social Psychology*, 8, 105–132. <http://dx.doi.org/10.1080/14792779643000092>
- Stallard, M. J., & Worthington, D. L. (1998). Reducing the hindsight bias utilizing attorney closing arguments. *Law and Human Behavior*, 22, 671–683. <http://dx.doi.org/10.1023/A:1025706823554>
- Stebly, N., Hosch, H. M., Culhane, S. E., & McWethy, A. (2006). The impact on juror verdicts of judicial instruction to disregard inadmissible evidence: A meta-analysis. *Law and Human Behavior*, 30, 469–492. <http://dx.doi.org/10.1007/s10979-006-9039-7>
- Wasserman, D., Lempert, R. O., & Hastie, R. (1991). Hindsight and causality. *Personality and Social Psychology Bulletin*, 17, 30–35. <http://dx.doi.org/10.1177/0146167291171005>
- Yopchick, J. E., & Kim, N. S. (2012). Hindsight bias and causal reasoning: A minimalist approach. *Cognitive Processing*, 13, 63–72. <http://dx.doi.org/10.1007/s10339-011-0414-z>

Received January 30, 2018

Revision received October 14, 2018

Accepted October 22, 2018 ■