FINDING ERROR

Chris William Sanchirico*

2003 MICH. ST. DCL L. REV. 3

TABLE OF CONTENTS

INTRODUCTION

The view that humans are systematically irrational in ways that matter for evidentiary process has been well represented and well argued at this conference on "Visions of Rationality in Evidence Law."[1] Part of what makes this view so compelling, both here

* Professor of Law, Business & Public Policy. University of Pennsylvania Law School and Wharton School; A.B. 1984, Princeton University; J.D. 1994, Yale Law School; Ph.D. (Economics) 1994, Yale University. Professor Sanchirico's other papers and publications are listed on his website: www.cstone.net/~csanchir.

Professor Sanchirico wishes to thank Jonathon Baron, Jason Johnston, and Jeffrey Rachlinski, for their helpful comments on an earlier draft, Katherine Bierlein for her able research assistance, and Penn Law Librarian Clayton Garthwaite for his help in locating sources.

[1] *See, e.g.*, Erica Beecher-Monas, *Heuristics, Biases, and the Importance of Gatekeeping*, 2003 MICH. ST. DCL L. REV. ___ (arguing in part that heuristics and biases of jury decisions justify the judge's role as gatekeeper for expert testimony) and Michael J. Saks & D. Michael Risinger, *Baserates, The Presumption of Guilt, Admissibility Rulings, and Erroneous Convictions*, 2003 MICH. ST. DCL L. REV. ___ (arguing that "[m]istaken assumptions about baserates, be they low or high, will lead to more erroneous decisions in one

and in general, is the relentless empiricism of its proponents. Every assertion, it often seems, has its counterpart in experimental results. There is the danger, however, that this otherwise admirable emphasis on data will cause one to lose sight of the fact that the empirical investigation of irrationality goes hand in hand with the theoretical investigation of what constitutes rational behavior. Reliably identifying irrational behavior requires first identifying the rational ideal—irrationality being defined by what it is not. Furthermore, because repairing the system requires a clear understanding of what constitutes working order, identifying the rational response is crucial to the task of prescription. Thus, while it is often fair to criticize rational choice theory for being insufficiently grounded in empirical reality,[2] it may in some circumstances also be fair to criticize irrational choice empiricism for being insufficiently grounded in theoretical ideality.

As a case in point, consider applications of "hindsight bias" to legal process. Though it made only a cameo appearance at this conference, this bias plays a leading role in the literature on cognitive psychology and legal process.[3] For the most part, these

---

direction or the other."). But see, Gregory Mitchell, Mapping Evidence Law, 2003 MICH. ST. DCL L. REV. ___ (offering a powerful critique of several well known results in the literature on cognitive error and evidence law).

[2] See, e.g., Ronald J. Allen & Sarah A. Lively, *Burdens of Persuasion in Civil Cases: Algorithms v. Explanations*, 2003 MICH. ST. DCL L. REV. ___ (critiquing two formal models of evidence law for their failure to capture the empirical essence of the problems they purport to address).

[3] A casual count on legal electronic research services shows approximately 120 articles mentioning hindsight bias in connection with jury decision making. A small sample includes: Jonathan D. Casper, Kennette Benedict, and Jo L. Perry, , *Juror Decision Making, Attitudes, and the Hindsight Bias*, 13 L. & HUM. BEHAV. 291 (1989) (finding that whether a police search revealed incriminating evidence affects the jury's view of whether the search was proper); Chris Guthrie, Jeffrey J. Rachlinski, and Andrew J. Wistrich, *Inside the Judicial Mind*, 86 CORNELL L. REV. 777, 800, 802, 827 (2001) (discussing hindsight bias in relation to claims of ineffective assistance from counsel, sanctions under Rule 11, and charges of fraud against corporate officers; finding that judges suffer from hindsight bias); Jon D. Hanson & Douglas A. Kysar, *Taking Behavioralism Seriously: A Response to Market Manipulation*, 6 ROGER WILLIAMS U. L. REV. 259, 384 (2000) (arguing that out of fear of the hindsight bias, manufacturers fail to adopt new safety precautions); Reid Hastie & W. Kip Viscusi, *What Juries Can't Do Well: The Jury's Performance as a Risk Manager*, 40 ARIZ. L. REV. 901, 904-09 (1998) (finding judges less susceptible to the hindsight bias than juries); Kim A. Kamin & Jeffrey J. Rachlinski, *Ex Post ≠ Ex Ante: Determining Liability in Hindsight*, 19

applications concern judgments in hindsight of what *others* knew or should have known in foresight. As I establish within, this form of "across-person" hindsight bias has a rational twin that has not been adequately accounted for in the legal literature.[4] The manifestations of across-person hindsight bias cannot be reliably identified, nor can an

---

L. & HUM. BEHAV. 89 (1995) (discussed *infra* note 11); Russell B. Korobkin & Thomas S. Ulen, *Law and Behavioral Science: Removing the Rationality Assumption from Law and Economics*, 88 CAL. L. REV. 1051, 1098 (2000) (suggesting strict liability in cases where the victim was unable to take precautions as a way to avoid hindsight bias); Susan J. LaBine & Gary LaBine, *Determinations of Negligence and the Hindsight Bias*, 20 L. & HUM. BEHAV. 501 (1996) (Presenting experimental results regarding what is herein called "across-person" hindsight bias); Jeffrey J. Rachlinski, *Heuristics and Biases in the Courts: Ignorance or Adaptation?*, 79 OR. L. REV. 61 (2000) [hereinafter Rachlinski, *Heuristics*] (arguing for the existence of hindsight bias and asserting that law has adapted to the bias in some settings but not in others); Jeffrey J. Rachlinski, *A Positive Psychological Theory of Judging in Hindsight*, 65 U. CHI. L. REV. 571 (1998) [hereinafter Rachlinski, *Judging in Hindsight*] (arguing that the law reflects an understanding of hindsight bias); Jeffrey J. Rachlinski, *Regulating in Foresight Versus Judging Liability in Hindsight: The Case of Tobacco* 33 GA. L. REV. 813 (1999) [hereinafter Rachlinski, *Tobacco*] (discussing the impact of hindsight bias on the regulation of potentially dangerous products, including the effect on incentives for precaution and awards of punitive damages); W. Kip Viscusi, *How Do Judges Think About Risk?*, 1 AM. L. & ECON. REV. 26 (1999); Debra L. Worthington, Merrie Jo Stallard, Joseph M. Price, and Peter J. Goss, *Hindsight Bias*, Daubert, *and the Silicone Breast Implant Litigation: Making the Case for Court-Appointed Experts in Complex Medical and Scientific Litigation*, 8 PSYCHOL. PUB. POL'Y & L. 154, 156 (2002) (recommending that courts appoint panels of scientists to educate judges and juries in order to prevent hindsight bias being exacerbated by the fact finders' lack of technical knowledge).

[4] Others scholars have been critical of both the theoretical foundations of hindsight bias and its application to law. Neal R. Feigenson, *The Rhetoric of Torts: How Advocates Help Jurors Think About Causation, Reasonableness, and Responsibility*, 47 HASTINGS L.J. 61, 146 (1995) (explaining how lawyers induce juries not to use hindsight); Keith N. Hylton, *The Theory of Tort Doctrine and the Restatement (Third) of Torts*, 54 VAND. L. REV. 1413, 1430 (2001) (Doubting the irrationality of using outcome information to adjust *own* assessments of the probability of harm); Mark Kelman, David E. Fallas, and Hilary Folger, *Decomposing Hindsight Bias*, 16 J. RISK & UNCERTAINTY 251, 252, 254 (1998) (criticizing experiments on hindsight bias for failing to distinguish among uses and abuses of outcome information); Richard Lempert, *Juries, Hindsight, and Punitive Damage Awards: Failures of a Social Science Case for Change*, 48 DEPAUL L. REV. 867, 870, 874-5 (1999) (critiquing Hastie and Viscusi study, *supra* note 3); Gregory Mitchell, *Taking Behavioralism Too Seriously? The Unwarranted Pessimism of the New Behavioral Analysis of Law*, 43 WM. & MARY L. REV. 1907, 1911 (2002) ("empirical research does not support the dire pronouncements of legal scholars regarding the human capacity for irrational behavior"); Gregory Mitchell, *Why Law and Economics' Perfect Rationality Should Not Be Traded for Behavioral Law and Economics' Equal Incompetence*, 91 GEO. L.J. 67 (2002) [hereinafter Mitchell, *Economics*] (arguing that jurors are no more uniformly irrational than they are uniformly rational and questioning the value of trying to counteract biases that are not uniformly held); Philip G. Peters, Jr., *Hindsight Bias & Tort Liability: Avoiding Premature Conclusions*, 31 ARIZ. ST. L. J. 1277, 1304 (1999) (cautioning against trying to correct for hindsight bias in light of the complexity of trial, the balance between pro-defendant and pro-plaintiff biases, and the inapplicability of existing studies to actual jury trials); C.K. "Pete" Rowland, *Psychological Perspectives on Juror Reactions to the September 11 Events: Terror Management Theory Would Predict More Punitive Jury Reactions, but the More Traditional Social-Cognition Approach May Still Predominate*, 69 DEF. COUNS. J. 180, 184 (2002) (explaining how counterfactual thinking can both weaken and strengthen hindsight bias); Jeanne L. Schroeder, *The Stumbling Block: Freedom, Rationality, and Legal Scholarship*, 44 WM. & MARY L. REV. 263,352-353 (2002) (criticizing suggestions that information be withheld from jurors to avoid hindsight bias).

appropriate response be devised until the bias is clearly distinguished from its rational look-alike.

I. HINDSIGHT BIAS: SAME-PERSON VERSUS ACROSS-PERSON

Baruch Fischhoff, who is credited with identifying hindsight bias, describes it this way:

> In hindsight, people consistently exaggerate what could have been anticipated in foresight. They not only tend to view what has happened as having been inevitable but also to view it as having appeared "relatively inevitable" before it happened. People believe that others should have been able to anticipate events much better than was actually the case.[5]

As so laid out, and later investigated, hindsight bias may be decomposed into two distinct components.[6] First, "[people] not only tend to view what has happened as having been inevitable, but also view it as having appeared 'relatively inevitable' before it happened."[7] This first phenomenon has to do with an individual's ability to understand what her *own* assessment of an uncertainty would have been in the counterfactual that she did not receive additional information on how the uncertainty was resolved. In the canonical experiment confirming this phenomenon, the experimenter gives the control group no information about whether event *E* has occurred and asks them to judge the likelihood of *E*'s occurrence. At the same time, the experimenter tells the experimental group that *E* did in fact occur and asks them to imagine what they *would* have thought *E*'s chance of occurring to be before learning that it actually did. Because the numerous subjects are randomly divided, the two

---

[5] Baruch Fischhoff, *For Those Condemned to Study the Past: Heuristics and Biases in Hindsight*, *in* JUDGMENT UNDER UNCERTAINTY: HEURISTICS AND BIASES 335, 341 (Daniel Kahneman et al. eds., 1982) [hereinafter Fischhoff, *Condemned*].

[6] *C.f.*, Kelman et al., *supra* note 4, at 252-54 (1998) (decomposing hindsight bias into primary, secondary, and tertiary components).

[7] Fischhoft, *Condemned*, *supra* note 5, at 341.

groups are unlikely to differ systematically in composition. Yet, the outcome-informed group's assessment of *E*'s ex ante likelihood shows up higher in experiment after experiment.[8]

The second component of hindsight bias has to do with outcome-informed assessments of what *other* people did understand or should have understood about the likelihood of an event. Thus, Fischhoff states, "[p]eople believe that *others* should have been able to anticipate events much better than was actually the case."[9] It is this second component that is perhaps most relevant for the legal process, and specifically fact finder decision making, for it implies that "the [hindsight] bias …causes people to hold decisionmakers legally liable for outcomes that they could not have predicted."[10] Like same-person hindsight bias, across-person hindsight bias has been the subject of experimental inquiry.[11]

---

[8] *See* Rachlinski, *Heuristics*, *supra* note 3, at 67 ("Psychologists have conducted nearly 130 experiments demonstrating the existence of the hindsight bias using a variety of different methods, materials, and subjects." (citation omitted)).

In Fischhoff's seminal experiment, for example, subjects were presented with background information on the Nineteenth Century conflict between the British and the Nepalese Gurkhas as well as four possible outcomes for the conflict (e.g., British win, Gurkha's win, etc). The subjects were divided into five groups. One group was given no information on the outcome of the conflict and asked to assess the likelihood of each of the four outcomes. Each of the four remaining groups was told that a different outcome had in fact occurred and was asked what they would have thought the likelihood of each outcome to be had they not been given this additional information. For each outcome, the group that was told that a given outcome had occurred tended to rate that outcome as having been more probable ex ante than did the other groups. *See* Baruch Fischhoff, *Hindsight ≠ Foresight: The Effect of Outcome Knowledge on Judgment Under Uncertainty*, 1 J. EXPERIMENTAL PSYCHOL. 288, 289-93 (1975) [hereinafter Fischhoff, *Hindsight ≠ Foresight*].

See Rachlinski, *Heuristics*, *supra* note 3, at 67, and Rachlinski, *Judging in Hindsight*, *supra* note 3, at 576-79 for cogent commentary on this and other experiments confirming hindsight bias.

[9] Fischhoff, *Hindsight ≠ Foresight*, *supra* note 8, at 289.

[10] Rachlinski, *Judging in Hindsight*, *supra* note 3, at 588. On the importance for law of this form of hindsight bias see also Kelman et al., *supra* note 4, at 254 ("[Across person] bias has proven especially interesting to those studying the impact of bias on legal judgment.") and Jeffrey J. Rachlinski, *Heuristics*, *supra* note 3, at 69 ("Whenever a court must determine what a party 'should have known,' it is susceptible to the influence of the hindsight bias.")

[11] *See, e.g.*, Kamin & Rachlinski, *supra* note 3, at 101 (discussed below); Casper et al., *supra* note 3 (subjects judge whether police search was proper); Kelman et al., *supra* note 4 (subjects judge quality of others' gambles).

The first, *same*-person component of the hindsight bias seems unassailable, not just empirically, but also theoretically. It is indeed irrational to use the information that an event has occurred in answering the question of what you yourself might have thought the event's likelihood to be if you could *not* use the information that *E* did occur.

The second, across-person version of the hindsight bias is, however, more problematic. Is it wrong to use knowledge that an event *E* has occurred in judging whether *others* should have or did in fact know that it was likely to occur ex ante? I argue in the next section that it is not wrong in the case that one believes that these others: (1) may have been in a position to know about the chance of events like *E;* and 2) may now be unable or unwilling to fully reveal what information that position provided.

The rational across-person use of outcome information seems far more important to law than its almost complete absence from the literature applying hindsight bias to

---

For example, in the Kamin and Rachlinski experiment, whose design was inspired by an actual legal case, the uncertainty in question was the possibility of upstream flood damage precipitated by river flood conditions combined with the risk of debris blockage by a new city-owned draw bridge. The control group was given background information and asked whether the city should hire an operator during the winter months who could continuously monitor the build up of debris under the bridge and raise it should that become imperative to prevent blockage and upstream flooding. *See* Kamin & Rachlinski, *supra* note 3, at 97. The control group was specifically asked to imagine itself as an auxiliary member of an urban planning committee charged with making this decision in foresight before any outcome became known. *See id*. at 94. After learning the problem, the control group was asked whether the city should hire the winter operator, and the background information was such that they were induced to answer "yes" if they thought that the chance of a flood damage was greater than ten percent. The experimental group, on the other hand, was told that the city had not hired an operator, that flooding had occurred due to build up of debris, and that an upstream bakery business had suffered economic damage as a result. *See id*. at 94-95. In addition to this outcome information, the experimental group was given the same background information as the control group, but all this in the context of a jury trial, rather than an urban planning meeting. Members of the experimental group were asked whether they believed that the city "was responsible for the flood damage." *See id*. at 97. Lengthy jury instructions provided to the experimental group defined "responsibility." Subjects were essentially encouraged to find the city responsible if the ex ante chance of flooding had been more than ten percent. The researchers found markedly different assessments across the control and experimental groups. Less than a quarter of the control group thought the likelihood of flood exceeded ten percent, as compared to significantly more than half of the experiment group. *See id*. at 99-101.

legal problems would indicate.[12] The two conditions for this rational use—that others were then informed and now unforthcoming—will often obtain when a legal fact finder must decide what the defendant knew or should have known ex ante about the chance of an uncertain outcome, such as an accident. Such an inquiry would be relevant to determining whether due care was exercised under a negligence standard, whether warnings were adequate in product liability, or whether an accident was reasonably foreseeable under strict liability.[13] Likely, wherever across-person hindsight bias is an issue for law, the rational across-person use of outcome information will also be present.

---

[12] I thank Jonathan Baron for pointing out to me that the experimental psychology literature does explicitly acknowledge, and attempt to control for, the rational across-person use of outcome information. *See, e.g.*, Jonathan Baron & John C. Hershey, *Outcome Bias in Decision Evaluation*, 54 J. OF PERSONALITY & SOC. PSYCHOL. 569, 569-570 (1988) (in an introduction that discusses both outcome and hindsight bias: "Although outcome information plays no direct role in the evaluation of decisions, it may play an appropriate indirect role. In particular, it may affect a judge's beliefs about actor information [i.e., information known only to the decision maker at the time the decision is made]. A judge who does not know the decision-maker's probabilities may assume that the probability was higher for an outcome that occurred than for the same outcome, had it not occurred."); John C. Hershey & Jonathan Baron, *Judgment by Outcomes: When is it Justified?*, 53 ORGANIZATIONAL BEHAV. & HUM. DECISION PROCESSES 89 (1992) ("whenever the decision maker has, or ought to have, more information than the judge, outcome can be a valid, albeit imperfect, indicator of decision quality."); John C. Hershey & Jonathan Baron, *Judgment by Outcomes: When is it Warranted?*, 62 ORGANIZATIONAL BEHAV. & HUM. DECISION PROCESSES 127 (1995) (same).

In contrast this author's search of the large literature applying hindsight bias to law revealed only two references, both in footnotes. Kelman et al., *supra* note 4, at 261, 267 n.10 ("It is possible as well that finding contraband [i.e., outcome information] *is* probative of the reasonableness of suspicion [the other person's probability belief at the time of decision]; it may be impossible to specify in a short written problem, or even real courtroom testimony, all of the "soft" signs that those conducting searches [the others] use to justify their suspicions, but factfinders may believe that the soft signs were more likely present if inculpatory material was actually found"); *C.f.*, *id*. at 252, 260-1, 267-68 n.11 (criticizing experiments for conflating cross person bias with a *different* dynamic—namely, subjects rejection of "the normative significance of the ex ante perspective."). *See also* Mitchell, *Economics*, *supra* note 4, at 127 n.182.

This deficit in the legal literature on hindsight bias is specially perplexing for two reasons. First, as argued in the text, the conditions giving rise to the rational use of outcome information are pervasive in real legal problems. Second, the legal literature does often cite to Baron & Hershey, *supra*, for the *irrational* use of outcome information without mentioning its rational use, even though, as noted, Baron & Hershey, *supra*, is clear about the presence of the potentially confounding rational use of such information. *See, e.g.*, Janice Nadler & Mary R. Rose, *Victim Impact Testimony and the Psychology of Punishment*, 88 CORNELL L. REV. 419, 442 n.126, 443 (2003); Rachlinski, *Judging in Hindsight*, *supra* note 3, at 581, 581 n.35, 591, 591 n.91 (same); Rachlinski, *Heuristics*, *supra* note 3, at 78 n.89 (same).

[13] Regarding negligence, one treatise author writes:

> The idea of [unreasonable] risk in this context necessarily involves a recognizable danger, *based upon some knowledge of the existing facts*, and some reasonable belief that

To be sure, identifying a rational doppelganger for across-person hindsight bias does not negate the existence of the bias itself. However, the coexistence of a rational deduction ought to give legal scholars serious pause in judging the relevance of experimental results in this area. Firstly, not all experiment designs are entirely clear about whether the subjects are meant to judge their own or another's ex ante assessment.[14] Secondly, not all experimental designs that purport to measure the unwarranted outcome adjustment of across-person hindsight bias adequately control for the presence of the correlate rational adjustment. Such results may thus overstate the magnitude of the irrational adjustment.[15] Thirdly, even if an experiment fully

---

> harm may possibly follow. Risk, for this purpose, may then be defined as a danger which is apparent or should be apparent, *to one in the position of the actor*.

PROSSER AND KEETON ON THE LAW OF TORTS § 31 (W. Page Keeton et al. eds., 5th ed. 1984) (emphasis added). Regarding the materiality of what the defendant did and should have know in other tort settings, see, e.g., MCCORMICK ON EVIDENCE § 200 (John W. Stronged., 5th ed. 1999) (discussing the admissibility of similar happenings evidence for this purpose). *C.f.*, Rachlinski, *Tobacco*, *supra* note 3, at 840-841 (discussed infra note 3, focusing narrowly on a recklessness determination for punitive damages.)

[14] For example, reading the jury instructions in Kamin and Rachlinski's drawbridge experiment described above in note 111, one might well conclude that the jurors were meant to determine what they themselves would have thought the chance of flood damage to be to be before learning that it had occurred. *See* Kamin & Rachlinski, *supra* note 3, at 96-97. Putting aside the law of negligence, this would seem to be testing the first single-person component of hindsight bias. Yet in later writing by one of the authors of this study, the flood study is offered to illustrate that "[t]he research on the hindsight bias shows that people blame others for failing to have predicted adverse outcomes that could not have been predicted." Rachlinski, *Judging in Hindsight*, *supra* note 3, at 588. *See also* Ward Edwards & Detlof von Winterfeldt, *Cognitive Illusions and Their Implications for the Law*, 59 S. CAL. L. REV. 225, 243-4 (1986) (offering a similar critique).

But see, Rachlinski, *Tobacco*, *supra* note 3, at 840-841 (acknowledging this distinction in the context of recklessness and punitive damages). *Accord*, Stephen M. Bainbridge & G. Mitu Gulati, *How Do Judges Maximize? (The Same Way Everybody Else Does–Boundedly): Rules of Thumb in Securities Fraud Opinions*, 51 EMORY L.J. 83, 128-9 (2002) (contending that the "no fraud by hindsight" rule in securities disclosure class actions is not meant to reduce hindsight bias, but rather aims to simplify the judge's analysis in complex cases).

[15] Consider, for example, Kamin & Rachlinski's drawbridge experiment described in *supra* note 111. Apropos of *supra* note 14, let us assume for purposes of argument that subjects were meant to judge what the *city* believed or should have believed at the time of decision making. A subject in that experiment might well implicitly conclude that the city—which had built the bridge and had been operating it for several months—was in a better position to know about the dangers of winter flooding. To be sure, whether the city knew more about the likelihood of river thaw in the winter might be open to question. But from its prior experience operating the bridge, the city quite plausibly knew more about what might get caught under the bridge, what effect this would have upstream, and how closely one would have to monitor the bridge to prevent damage. Given that subjects believed that the city had special knowledge, it would have been rational to apply outcome information to judge the city's culpability.

controls for the rational adjustment, its implications for policy are far from clear. The irrational adjustment may be compensating for individuals' failure to adequately make the rational adjustment. If this is so, policies that prevent the irrational adjustment will have the perverse result of increasing decision-making error. Fourthly, and more generally, the rational adjustment—which is, at best, controlled for in existing experimental designs—needs to be explicitly studied, both in its own right and in interaction with its irrational cousin.[16]

II. THE RATIONAL ACROSS-PERSON USE OF OUTCOME INFORMATION

Why might it be rational to use outcome information to adjust upward one's assessment of the assessment of another who one believes was in a position to know? The easiest case in point is a fanciful extreme. Suppose that one is certain that the

---

Baron & Hershey, *supra* note 12, at 569-570, do attempt to explicitly control for the rational use of outcome information in their five experiments testing for "outcome bias," a phenomenon related to hindsight bias. But it is not clear that they are entirely successful in doing so. In experiments 1, 2, and 4, subjects judge the decision of a doctor regarding patient treatment or testing. *Id*. 571-76. While subjects are given the *statistical* information available to the doctor, they are not given other information that they might have imagined was available to the doctor at the time of the decisions, including specific information about the patient (as opposed to statistical information about the "average" patient) and information gleaned from medical training. In experiment 5, *see id*. at 577-78, involving drawing cards from one of two decks of cards, it is not clear whether subjects understood the probability theory behind the choices made by the decision maker—it appears from the debriefing questionnaire that many did not. Subjects might well have imagined, therefore, that the decision maker had, at the time of decision, additional "information" that they did not have in the form of knowledge about probability theory. To this extent, they were justified in assessing the decision maker's choice according to outcome. (Note here the irony that their "irrationality" viz. probability theory generates a rational use of outcome information.) Experiment 4 does appear to have adequate controls, but it is not clear in that experiment whether, when subjects said a decision was "good" or "bad," they meant to indicate, at least in part, whether it had *turned out* well or poorly, respectively. To be sure, many of these qualifications are made clear in Baron and Hershey's article.

[16] For example, in the portion of Kelman, Fallas, and Folger's experiment in which outcome-informed subjects judge what others should have known, whether these others were in a position to know is not manipulated. In the relevant experimental conditions, the subjects are effectively told that these others (gamblers betting on rolls of a die) had no way of knowing that the die was loaded. *See* Kelman et al., *supra* note 4, 265-266.

defendant is omniscient. Then, if one learns that an accident did occur, one also learns that the defendant knew it was going to occur. What actually happened is completely identified with what the defendant knew was going to happen. In this limiting case, then, information about what actually happened is enormously relevant—indeed conclusive—in determining what the defendant knew.

It is not that great a leap from this hyper-hypothetical to a case commonly encountered in legal process, wherein the fact finder has reason to believe, based on other evidence, that the defendant, though hardly omniscient, was in a position to know about the likely results of her actions. Take, for example, the classic product liability suit. In such a case, the defendant designed or manufactured the product and so is likely to have understood something about the product's dangers. It is certainly true that learning that the accident did occur should not properly cause the fact finder to increase *her* current guess about the probability *she* would have placed on the event if she did not know of its occurrence. But why should it not influence her assessment of whether the defendant knew ex ante? In fact, it should.

A detailed justification for this assertion begins with the recognition that the problem of rationally judging what the defendant knew ex ante may be decomposed into two related questions. First, based on the state of the world at the time of the event, what was the true likelihood that the accident would occur—i.e., what was the "ex ante objective" probability[17] of the event? Second, what was the defendant's subjective assessment of this ex ante objective probability? Given these two questions, the argument may be decomposed into two steps. The first step links the

[17] As is common, I use the terms "probability," "likelihood," and "(probabilitistic) assessment" interchangeably.

fact finder's knowledge of the actual outcome to her answer to the first question, concerning the ex ante objective likelihood of the event. The second step links her answer to the first question to her answer to the second, concerning the defendant's subjective assessment of the ex ante objective probability. Transitivity across these two steps establishes the proper use of outcome information in judging the defendant's subjective assessment.

A. Outcomes to Own Estimates

The first step is uncontroversial. No one could doubt that the fact finder ought to use its knowledge of the outcome in assessing the ex ante objective probability that the event would occur, the first probability mentioned in the prior paragraph. After hearing that the event did occur, the fact finder ought to increase her assessment of the accident's ex ante objective probability. This is as natural as increasing one's assessment that a horse is fast after hearing that it won a race, or that a die is loaded after seeing it land on six ten times in a row. Updating one's beliefs about the uncertain parameters of the probabilistic process generating outcomes is well accepted to be a rational response to observing those outcomes.[18]

B. Own Estimates to Others' Estimates

[18] See, e.g., Rachlinski, *Judging in Hindsight*, *supra* note 3, at 576 ("It is important to distinguish between the [own person] hindsight bias and the more ordinary process of learning from experience."); Hylton, *supra* note 4, at 1430 ("Bayesians will update their predictions of the probability an event will happen after it happens. This much of hindsight judgment is entirely rational and should not be treated as a type of cognitive bias.")

The second step of the argument links the fact finder's assessment of the ex ante objective probability of the accident to its assessment of the ex ante subjective probability of a defendant who is believed to have been in a position to know. Consider the product liability example. Not knowing, and unable to ever know with certainty what the product liability defendant was or should have been thinking at the time of product design, the fact finder is nevertheless justified in imagining that the defendant's assessment is positively correlated with the true chance of the accident. In other words, if independent of the current case, one told the fact finder that the event had been objectively very likely to occur, the fact finder, believing that the defendant was in a position to know this, would rationally increase its assessment of the defendant's subjective probability that the event would occur. Similarly, if one told the fact finder that the objective probability of the event was quite low, then, imagining that the defendant was in a position to know *this*, the rational response for the trier of fact would be to adjust its assessment of the defendant's subjective probability of the accident downward. This is part of what it means to think that someone else knows something that one does not. Thus, one's assessment of a manufacturer's confidence in the safety of her product ought to be correlated with what one learns of the product's actual objective safety.

C. Transitivity

We may now combine the foregoing steps. If there is any correlation in the fact finder's mind between the defendant's assessment of the chance of an accident and the actual chance of an accident, then on learning that the accident occurred, the fact finder should not only increase its assessment of the objective chance of the accident, but also increase its assessment that the defendant knew the accident would occur. Thus the rational fact finder would look at the fact that the accident occurred, look at the fact that the defendant was in a position to know whether it would occur, and properly increase its assessment of the chance the defendant did in fact know. It is, therefore, not irrational for the fact finder to change its view of the defendant's state of mind with regard to an uncertain event on hearing about the outcome. The prior state of the defendant's mind and the current state of the defendant's product are not probabilistically independent variables in a rational fact finder's mind.

Below, in proposing an experiment, I make this same point in more formal terms, but before turning to that presentation, some additional discussion of the argument is in order.

D. Discussion

First, it is of course relevant whether the defendant was in fact in a position to know. Evidence to this effect would include the defendant's admission that she designed, manufactured, or operated the product (or drawbridge) in question. If the defendant is not in a position to know, then the second link disappears and there is no

appropriate connection made between the event's objective probability and the defendant's subjective assessment thereof.

Second, another important ingredient for the rational use of outcome information is the suspicion that the defendant may not be inclined to, or even able to relate to the fact finder all the relevant information that it had in its possession prior to the event's occurrence. Outcome information is in some sense a proxy for what the defendant knew. If we could fully determine the defendant's information set by other means, then that proxy would have no incremental informational value.

Third, to reiterate and expand upon a point made earlier, the fact that *some* upward adjustment in the fact finder's assessment of the defendant's subjective belief about the chance of accident is rational, does not mean that the *full* adjustment actually made by fact finders is warranted. It could well be that these two effects—the rational component and the irrational—are piled on top of each other so that across-person hindsight bias is still causing fact finders to find defendants negligent where they could not have predicted an accident ex ante.

Yet, for at least two reasons this possibility does not detract from the importance of experimental redesign to account for the rational component of outcome adjustment. Suppose, in the first place, that juries do rationally incorporate outcome information into their assessments of what better informed others knew ex ante. Then experiments failing to account for this rational adjustment—and attributing the entire upward adjustment to hindsight bias—will exaggerate the importance of the hindsight bias. Proposed fixes, which to be appropriate will have to be properly calibrated, will then overshoot and may end up doing more harm than good.

Second, suppose alternatively that juries do not rationally incorporate outcome information into their assessment of well-informed others' ex ante beliefs. Suppose, in particular, that they do this too little. Then the hindsight bias would be a good thing because it would supply the missing upward adjustment, and, for the compensating role that it plays, eliminating the hindsight bias would be harmful, not helpful.

III. SKETCH OF AN EXPERIMENT AND FORMAL ARGUMENT

How could one experimentally separate hindsight bias from the proper form of upward outcome adjustment I have just described? The following sketch of an experiment takes a step toward answering that question. The description also serves as a more formal explanation of the rational use of outcome information in judging others' ex ante assessments, as described informally above. I will describe the experiment in the form of instructions to the experimenter:

> First, provide all subjects with the following background information. The day before, you procured two "urns" and filled each with 100 ping pong balls. Of the balls in the first urn, you painted 25 black and 75 red. For the second urn, you painted the balls in reverse proportions: 75 black, 25 red. You then randomly chose one of the urns by flipping a coin. This is the urn they now see before them but cannot see into.
>
> Next, divide the subjects into two groups. Send Group 2 out of the room, leaving Group 1. Group 1 will become the group "in a position to know."[19] Draw a ball from the urn showing its color to Group 1.[20] Replace the ball in the urn and stir. This part of the experiment gives Group 1 private information with respect to which urn is actually placed before them. Have all members of Group 1 write down their probability assessment of the chance that the urn before them is the mostly red urn.
>
> Next have Group 2 return to the room. Explain to Group 2 that Group 1 has just seen a draw from the urn. Do not tell Group 2 anything

[19] The following is just one means of placing Group 1 in a position to know. It has the convenience of easy calculation. Other settings may more readily correspond to legal contexts.

[20] The experiment can be done with any number of draws. The math would be only slightly more complicated.

about the actual color of the ball that Group 1 observed. Make sure that there is neither intra- nor inter- group communication. The announcement that Group 1 has seen a draw from the urn instills in Group 2 the *belief* that Group 1 is in a position to know more about the composition of the urn than they. Ask Group 2, which has not seen any draws thus far, let alone those just observed by Group 1, what they think Group 1 now believes about the composition of the urn based on the draw that it saw in private.[21]

Next, draw another ball in front of both groups. This part of the experiment corresponds to the universal provision of outcome information. Now ask Group 2 what they think Group 1 believed about the probability that the urn was mostly red before Group 1 saw the second, public draw and after Group 1 observed the first, private draw. (Use the same probability "bins" as before.)

The experiment is quite removed from the legal context, to be sure. But it has the virtue that one can calculate the purely rational responses for Group 2 and examine how these compared to actual responses. Making this calculation requires working through three simple preliminary calculations.

First, we need to determine the rational responses for Group 1 to the first, private draw. Group 1 begins with a prior belief that there is a fifty percent chance that the urn is mostly red, the urns having been chosen by the flip of a fair coin. Then, according to Bayes' rule, on seeing a red ball on the first draw, Group 1 ought to believe that the chance that the urn is mostly red is seventy-five percent.[22] On seeing a black ball, twenty-five percent.[23]

Second, we calculate Group 2's assessment of the probability that Group 1 saw a red ball. The chance that Group 1 saw a red ball depends on which urn was being

[21] We are here asking Group 2 for a probability distribution rather than a point. We might then ask Group 2 to divide 100 percentage points between four or five different ranges (e.g., 0-20%, 21%-40%, 41-60%, 61-80%, 81-100%) for Group 1's current probability assessment that the urn before them is the first, mostly red urn. If, for example, a member of Group 2 believes it equally likely that the assessment of a member of Group 1 falls in any of the ranges in the above parenthetical, then the Group 2 member would place twenty percentage points in each of the five bins.

[22] P(red urn| red ball) = P(red ball| red urn)P(red urn)/P(red ball) = (.75) (.5)/((.75) (.5)+(.25)(.5)) = .75.

[23] The calculation is similar to that in *supra* note 22.

drawn from. There is a fifty percent chance that the urn was mostly red, in which case, the chance of drawing a red ball was .75. There is also a fifty percent chance that the urn was mostly black, in which case the chance of drawing a black ball was .25. Therefore, the chance that Group 1 saw a red ball was .5(.75) + .5(.25) = .5. Complementarily, the chance that Group 1 saw a black ball was also .5.

Third, we calculate Group 2's rational assessment, in the period after Group 1 has seen its private draw and before both groups have seen the public draw, of Group 1's beliefs about whether the urn is mostly red. Group 2 knows that, to the event that the urn is mostly red, Group 1 assigns either probability 75 (if Group 1 privately saw a red ball) or probability 25 (if Group 1 saw a black ball). As just calculated in the prior paragraph, Group 2 rationally believes that there was a fifty percent chance that Group 1 saw a red ball, and the same chance that Group 1 saw a black ball. Therefore, Group 2 should believe that there is a fifty percent chance that Group 1 believes that there is a seventy-five percent chance that the urn is mostly red and a fifty percent chance that Group 1 believes that there is a twenty-five percent chance that the urn is mostly black. Thus, if a mostly red urn corresponds to a dangerous condition, Group 2 thinks it equally likely that Group 1 rationally thought the condition likely to occur (seventy-five percent) as that Group 1 thought the event unlikely (twenty five percent).

We may now calculate how, after seeing the public draw, Group 2 should change it's beliefs about what Group 1 believed in the period between the private and public draws. One possibility is that Group 2 sees a red ball—and it suffices for our purposes to work through this possibility alone. In this case, just as with Group 1 in the first step above, Group 2 no longer believes that the two urns are equally likely.

Rather, it believes that there is a seventy-five percent chance that the mostly red urn is the one being drawn from. If we then ask Group 2 again what it thinks Group 1 believed in the period between the private and public draws, Group 2 must redo its calculation because its beliefs about Group 1's past beliefs will need to be informed by its beliefs about what Group 1 saw in private, which in turn depends on its beliefs about which urn is being drawn from. This is the crux of the matter.

In more detail, given that the urn is seventy-five percent likely to be mostly red, the chance that Group 1 saw a red ball is .75(.75) + .25(.25) = .5625 + .0625 = .625. Notice how this number is higher than before. When it had no outcome information Group 2 thought it equally likely that Group 1 privately saw a red ball as a black ball. Now having itself seen a red ball, Group 2 thinks that the chance Group 1 also saw a red ball in private is.625.

As before, if Group 1 saw red ball in private, its belief immediately thereafter that the urn was mostly red would have been .75. If it saw a black ball, .25. Therefore, Group 2, thinking the chance that Group 1 saw a red ball to be .625, now believes that there was a 62.5% chance that Group 1 thought the urn was 75% likely to mostly red and a 37.5% chance that Group 1 thought the urn was mostly black.

Therefore, outcome information—the public draw—causes Group 2 to decide that it was more likely than not (62.5%) that Group 1 believed that the urn was very likely (75%) mostly red, whereas prior to the public draw, Group 2 had been in equipoise regarding whether Group 1 thought the urn very likely to be mostly red or very likely to be mostly black. Analogously, the fact that the microwave did blow up in the plaintiff's apartment (i.e., the public draw was red) ought to inform the fact

finder's (i.e., Group 2's) assessment of how often the microwave blew up in the defendant's (i.e., Group 1's) product testing laboratory.

CONCLUSION

I have argued that the experimental identification of hindsight bias in the context of judging what others were in a position to know is subject to a confounding factor that is consistent with rational decision making. The argument hardly negates the value of existing research applying hindsight bias to legal process. Nor does it reduce the necessity for vigilance regarding whether this bias is at work in yet untested applications. It does, however, serve to illustrate the more general point that the empirics of irrationality and the theory of rationality are inextricably linked. In experiments designed to identify departures from ideal rationality, the ideal human of rational choice theory is as important a participant as real human subjects.