Judging Auditor Negligence:
De-biasing Interventions, Outcome Bias, and Reverse Outcome Bias

Jonathan H. Grenier
Doctoral Student
Department of Accountancy
University of Illinois at Urbana-Champaign
jgrenier@uiuc.edu

Mark E. Peecher
Professor of Accountancy
Deloitte & Touche Teaching Fellow
Department of Accountancy
University of Illinois at Urbana-Champaign
peecher@uiuc.edu

M. David Piercey
Assistant Professor of Accounting
Department of Accounting and Information Systems
University of Massachusetts Amherst
piercey@som.umass.edu

September 2007

Comments Appreciated

Data Available Upon Request

We thank Rajib Doogar, Karla Johnstone, Steve Kachelmeier, Marsha Keune, Bill Kinney, Lisa Koonce, Susan Krische, Brian Mayhew, Joel Pike, Steve Smith, Terry Warfield, and participants at workshops at the University of Illinois at Urbana-Champaign, University of Texas-Austin, and University of Wisconsin-Madison for their helpful comments. We also thank University of Illinois doctoral students for their generous help in data collection. Jonathan Grenier acknowledges the support of the Richard and Marie Irwin Foundation.

Judging Auditor Negligence:
De-biasing Interventions, Outcome Bias, and Reverse Outcome Bias

ABSTRACT

Individuals judge audit quality, in part, based on adverse outcome information. Assuming that individuals over-rely on outcomes, prior accounting research attempts to improve their judgments by reducing their reliance on outcome information. Logically, however, individuals could either over-rely on outcomes ("outcome bias") or under-rely on outcomes ("reverse outcome bias"). Peecher and Piercey (2007) provide theory and empirical findings that individuals harshly exhibit outcome bias when the Bayesian probability of negligence is below 40% (e.g., a range that would include frivolous lawsuits), but that individuals also leniently exhibit reverse outcome bias when the Bayesian probability of negligence is above 40% (e.g., above key legal thresholds such as "preponderance of the evidence").

Using Support Theory, we predict and find that, by reducing reliance on outcomes, most interventions from prior literature reduce outcome bias for the lower range of Bayesian probabilities but *exacerbate* reverse outcome bias for the higher range of Bayesian probabilities. Using Cumulative Prospect Theory, we also design a new intervention that, if implemented early during the evaluators' judgment process, successfully reduces both forms of bias. By doing so, we contribute to the accounting literature on de-biasing auditor negligence judgments and to the accounting literature on outcome effects.

Keywords: auditor negligence, de-biasing interventions, outcome effects, outcome bias

Data Availability: available upon request

I. INTRODUCTION

Prior research indicates that evaluators exhibit outcome effects in accounting contexts, including auditing (e.g., Brown and Solomon 1987, 1993; Anderson et al. 1997; Frederickson et al. 1999; Kadous 2001). Specifically, evaluators who receive adverse outcome information judge auditors more harshly than do evaluators who have not yet received outcome information. If systematically too large, outcome effects would bias evaluators against auditors. Assuming that outcome effects tend to be too large, a number of prior accounting studies examine various interventions intended to mitigate or eliminate bias by reducing outcome effects (e.g., Kadous 2001; Clarkson et al. 2002).[1, 2] Recent theory and experimental findings, however, suggest these attempts to *de-bias* may actually increase evaluator bias under commonly occurring conditions (Peecher and Piercey 2007).

Specifically, since individuals generally *should* exhibit outcome effects, Peecher and Piercey (2007) measure whether outcome effects are too harsh (outcome bias), too *lenient* (*reverse* outcome bias), or Bayesian (unbiased). Connecting Tversky and Kahneman's (1992) probability weighting function from Cumulative Prospect Theory to the accounting outcome-effects literature, they predict and find a pattern of outcome bias for Bayesian probabilities of auditor negligence below a threshold that is near 40% and *reverse outcome bias* for Bayesian probabilities above this threshold. If unchecked, this pattern of bias could increase the likelihood of successful frivolous lawsuits but also decrease the likelihood of successful lawsuits for which at least the preponderance of evidence indicates auditor negligence.

[1] While de-biasing is an important motivation in some of these studies, some studies include other motivations as it is psychologically interesting to understand the cognitive, motivational, or contextual factors that moderate outcome effects (cf., Brown and Solomon 1993; Tan and Lipe 1997; Kadous 2001).
[2] Accounting and psychology researchers' attempts to improve evaluators' negligence judgments via de-biasing interventions appear to have had an influence on legal practice and scholarship (e.g., Jolls and Sunstein 2006; Koehler 2006). Attempts to implement de-biasing interventions include modifications of jury instructions and restructuring of defense attorneys' arguments.

Exhibit 1 shows the probability weighting function, which distorts individuals' subjective probabilities, and illustrates related findings in Peecher and Piercey (2007). Because evaluators distort their beliefs in this pattern when relying on outcome information, we examine whether interventions that moderate outcome effects also decrease both outcome bias and reverse outcome bias. Simply lowering outcome effects may reduce outcome bias for the lower range of Bayesian probabilities, but it also may exacerbate reverse outcome bias for the higher range of Bayesian probabilities.

[Insert Exhibit 1 here]

We use *Support Theory* (Tversky and Koehler 1994; Rottenstreich and Tversky 1997; Fox and Tversky 1998; Fox 1999) to predict how interventions designed to lower outcome effects influence outcome bias and reverse outcome bias. We test three interventions from the accounting outcome-effect literature: *counter-explanation* (e.g., Lowe and Reckers 1994; Kennedy 1995; Anderson et al. 1997), *attribution* (Kadous 2001), and *warning* (e.g., Clarkson et al. 2002). And, building upon Peecher and Piercey (2007), we use *Cumulative Prospect Theory* to design and test two *re-weighting* interventions that are new to the accounting outcome-effect literature.

To briefly explain the three traditional interventions, the *counter-explanation* intervention involves listing two ways in which a material misstatement could occur in the absence of auditor negligence. Kadous' (2001) *attribution* intervention is more subtle. Case materials contain language that makes participants realize the role of judging auditor negligence causes them to be anxious and tense. Awareness of their anxiety allows them to understand that alternative causes, besides auditor negligence, plausibly account for their anxiety. Finally, although the *warning* intervention seldom reduces outcome effects, we include it for completeness since it is in both the accounting and psychology literature. In this condition, case materials warned participants against over-relying on outcome information.

As for our new interventions, our *re-weighting* intervention advised participants about how and when evaluators tend to over- or under-react to outcome information. This advice appeared in the materials before the primary facts. The case materials explained that evaluators who should judge the likelihood of auditor negligence to be relatively high (> 40%) tend to judge this likelihood too low. These materials also explained that, on the flip side, evaluators who should judge the likelihood of auditor negligence to be relatively low (< 40%) tend to judge this likelihood too high. The *re-weighting – after* intervention was similar, but it appeared in the materials after participants had reached their judgments of auditor negligence.

Our framework, informed by Support Theory and Cumulative Prospect Theory, allows us to predict that the *re-weighting* intervention will successfully reduce both outcome bias and reverse outcome bias as well as increase participants' conformance with Bayesian belief revision. Moreover, it will improve participants' judgments of auditor negligence on these dimensions to a greater degree than any of the traditional interventions will. According to our framework, these interventions will lack overall effectiveness in reducing evaluator bias and will even exacerbate bias under some conditions. Specifically, we predict the *counter-explanation* and *attribution* interventions to reduce outcome bias but also to exacerbate reverse outcome bias. We do not predict the *warning* or re-*weighting – after* interventions to have a systematic effect on outcome bias or reverse outcome bias.

To test these predictions, we extended the experimentation in Peecher and Piercey (2007). Specifically, participants read case facts about a particular audit adapted from Kadous (2000, 2001). We measured their prior beliefs before providing them with outcome information. This enables us to calculate what each participant's Bayesian judgment of auditor negligence should be when given outcome information. We then presented participants with outcome information, (an SEC investigation alleging material misstatement), and asked them to judge auditor negligence. We measure outcome bias (and reverse outcome bias) by comparing their judgments to the Bayesian benchmark judgments in each of the five intervention conditions and the *control* condition. We also manipulated audit quality at two levels between participants, primarily to ensure sufficient variability in judged auditor negligence to span the entire [0,1] probability interval, but also to explore potential interactions of interventions with audit quality.

Overall, the experimental findings substantially support our predictions, though there are some noteworthy exceptions. As predicted, the *re-weighting* intervention emerges as best, taking into consideration outcome bias, reverse outcome bias, and the extent to which the Bayesian benchmark explains variation in evaluators' judgments. In the *counter-explanation* condition, the predicted reduction in outcome bias relative to the *control* condition occurs, but (unfortunately) so does the predicted exacerbation of reverse outcome bias relative to the *control* condition.

For the *attribution* intervention, we do not observe the predicted reduction in outcome bias relative to the *control* condition. We do, however, observe the predicted worsening of reverse outcome bias compared to the *re-weighting* condition. And, when compared to the *control* condition, we observe that the *attribution* intervention exacerbates reverse outcome bias only in the higher audit quality condition (we had predicted this to occur in both audit quality conditions). Also contrary to predictions, the *attribution* intervention significantly reduces reverse outcome bias relative to the *control* condition given lower audit quality. Thus, while the *attribution* intervention does not fare as well as the *re-weighting* intervention, it generates less reverse outcome bias than does the *counter-explanation* intervention when audit quality is relatively low.

We find that the *re-weighting – after* and *warning* interventions cause significantly greater outcome bias and reverse outcome bias than the *re-weighting* intervention causes, but that they are bias-neutral relative to the *control* condition. Findings nevertheless indicate that these two interventions are undesirable as both decrease the extent to which participants engage in Bayesian belief revision when judging auditor negligence, relative to the *re-weighting* intervention and to the *control* condition.

Overall, our theory and experimental findings contribute to the accounting literature in several ways. First and foremost, we embellish our understanding of where outcome bias and reverse outcome bias fit within existing de-biasing frameworks (e.g., Kennedy 1995). As an example, we show that outcome-effect interventions that the accounting literature currently characterizes as being effective de-biasers fail to decrease outcome bias under some conditions and actually increase reverse outcome bias. This increase in reverse outcome bias, moreover, occurs for relatively high Bayesian posteriors that surpass legal thresholds such as "preponderance of the evidence." As another example, we show that a simple, pre-emptive correction to probability weighting significantly reduces both forms of bias. Second, our empirical evidence suggests a boundary condition of the ability of the *attribution* intervention to reduce outcome effects (cf., Kadous 2001): It appears to require relatively high levels of audit quality. Third, the simple *re-weighting* intervention holds promise for extending the economics and psychology literatures in that distortion of subjective probabilities undermines how well people evaluate risky prospects (in addition to distorting their belief revision as shown herein).

The remainder of this paper unfolds as follows. Section II develops our theory and hypotheses. Section III describes our experimental design. Section IV discusses the results of our experiment. Section V concludes with limitations and offers suggestions for future research.

II. THEORY AND HYPOTHESIS DEVELOPMENT

The accounting outcome-effect literature documents that evaluators judge evaluatees more harshly when evaluators have adverse outcome information (Brown and Solomon 1987; Lipe 1993; Kennedy 1995; Kinney and Nelson 1996; Tan and Lipe 1997; Frederickson et al. 1999; Kadous and Magro 2001). Many accounting outcome-effect studies use audit litigation contexts (Anderson et al. 1993; Lowe and Reckers 1994; Anderson et al. 1997; Kadous 2001; Clarkson et al. 2002; Peecher and Piercey 2007). In these contexts, adverse outcome information is diagnostic to some degree unless an evaluator has certain, accurate, and complete information about the auditor's decision processes (Brown and Solomon 1987; Hershey and Baron 1992, 1995; Tan and Lipe 1997). These pristine conditions rarely occur. As a result, outcome effects do not reveal much at all about whether evaluators are too harsh. People generally should rely (but not over- or under-rely) on outcome information (Baron and Hershey 1992; Peecher and Piercey 2007).[3]

A number of accounting outcome-effect studies argue, imply, or assume that outcome effects tend to be too large, and therefore constitute a harsh outcome bias. Researchers have developed and tested several interventions that reduce outcome effects, with the hope that such reductions also reduce outcome bias.[4] However, none of these studies compare the amount of

[3] Outcome effects differ from hindsight bias. Outcome effects occur when outcome information influences how people judge decision-making quality. Hindsight bias occurs when outcome information causes people to overestimate their prior beliefs about the likelihood of that outcome.

[4] Accounting researchers have attempted to suppress the diagnosticity of outcome information in several ways, with mixed success. One approach, for example, argues that institutional rules proscribe reliance on outcome information or state that reliance would be unfair (Kadous and Magro 2001).

bias present in the outcome effects with and without these interventions (cf., Brown et al. 1999). Thus, the effectiveness of these interventions as a de-biaser remains an untested empirical question. As Peecher and Piercey (2007) point out, to the extent that individuals under-react to outcome information, outcome effects are too small, and interventions that reduce outcome effects could exacerbate reverse outcome bias. As we next discuss, Support Theory provides a basis for predicting that none of the traditional interventions will simultaneously reduce both outcome bias and reverse outcome bias — and that some will exacerbate reverse outcome bias.

We next summarize Support Theory, explain how we lever it to predict conditions under which traditional outcome-effect interventions improve and *impair* evaluators' judgments of auditor negligence, and then how we lever it to design a fruitful intervention that we predict to reduce both outcome bias and reverse outcome bias.

Support Theory's Two Stage Model of Probability Judgment

Support Theory contains a two stage model of subjective probability judgment (see, e.g., Fox and Tversky 1998; Wu and Gonzalez 1999). In stage 1, evaluators amass support for hypotheses and translate support into subjective probabilities. In stage 2, evaluators mis-weight these subjective probabilities via the probability weighting function from Cumulative Prospect Theory. We next describe stage 1 in more detail.

In stage 1, individuals assign subjective probabilities to *descriptions* of events, called hypotheses. The subjective probability they assign to a focal hypothesis, A, depends on the amount of descriptive support ($S(\cdot)$) that they amass for A relative to the support for its mutually exclusive residual, A^C:

$$p(A) = S(A)/(S(A) + S(A^C)). \quad (1)$$

A focal hypothesis could be “the auditor is negligent” with the residual being “the auditor is not negligent”.

Different descriptions of the focal or residual hypothesis can significantly distort individuals’ subjective probabilities. These distortions sometimes result in violations of laws of probability. For example, an individual may assign a probability of 0.05 to the hypothesis “car will not start” but assign a much larger probability, say 0.35, to the same event when it is “unpacked” into various components (e.g., car will not start because the battery is dead, car will not start because it is still in “drive”, etc.). As such, interventions that influence the extent to which evaluators unpack events such as “the auditor is negligent” or otherwise change how evaluators describe such events have the potential to distort individuals’ subjective probabilities.

In stage 2, individuals further distort probabilities assigned to hypotheses (i.e., event descriptions) by applying a probability weighting function (see Exhibit 1). We submit that evaluators’ probability weights in judging the likelihood of auditor negligence, and the extent to which such weights are Bayesian-inconsistent, likely are sensitive to attributes of interventions that moderate outcome effects. Hereafter, we analyze several interventions, with respect to this two-stage model, in order to predict their effects on outcome bias and reverse outcome bias. We begin with the traditional interventions that have been used to reduce outcome effects and follow up with the *re-weighting* interventions that we introduce to the accounting literature.

Traditional Interventions

Counter-explanation Intervention

Counter-explanation is a common intervention in the accounting outcome-effect literature (Lowe and Reckers 1994; Kennedy 1995; Anderson et al. 1997).[5] Typically, because individuals’

[5] We note that Kennedy (1995) is not an outcome effect study but rather a hindsight bias study.

tend to not spontaneously counter-explain (Koonce 1992), researchers provide participants with alternative explanations for the adverse outcome information besides evaluatee negligence. The participants then rate the likelihood of these explanations and develop their own alternative explanation. Designed to weaken causal linkages automatically formed in evaluators' mental representations between the adverse outcome information and auditor negligence (i.e., Fischhoff's (1975) *creeping determinism*), counter-explanation interventions typically reduce outcome effects.

Several accounting studies provide evidence consistent with Fischhoff's theoretical basis for the success of counter-explanation interventions in reducing outcome effects (Lowe and Reckers 1994; Kennedy 1995; Anderson et al. 1997). Recent work, however, suggests that certain factors moderate the extent to which the counter-explanation intervention reduces outcome effects. Findings from Kadous et al. (2006), for example, suggests that the *availability* of counter-explanations — or the cognitive ease with which counter-explanations come to mind — would moderate the intervention's reduction in outcome effects.

Consistent with the moderating effect of availability, we use stage 1 of Support Theory (Tversky and Koehler 1994) to predict how counter-explanation moderates outcome bias and reverse outcome bias for relatively low and high Bayesian posteriors of auditor negligence, respectively. Specifically, counter-explanation unpacks, and thus increases the salience of, how material misstatement could occur *without* auditor negligence (i.e., $MM|AN^C$).[6] Unpacking an event such as "material misstatement given that the auditor is not negligent" increases support for that event and so decreases judged support for auditor negligence (Tversky and Koehler

[6] One also could argue that counter-explanations unpack the event of the auditor not being negligent given material misstatement ($AN^C|MM$) by describing scenarios under which the auditor conformed to standards of care in the presence of a material misstatement. This variant of unpacking also decreases the observed posterior of $p(AN|MM)$. Thus, the underlying logic behind our forthcoming hypotheses is unchanged.

1994; Rottenstreich and Tversky 1997). In terms of our experimental conditions, unpacking the event of material misstatement given that the auditor was not negligent ($MM|AN^C$) would result in the inequality:

$$S_{counter\text{-}explanation}(MM|AN^C) > S_{control}(MM|AN^C). \tag{2}$$

As theory holds that support gathered increases subjective probabilities, an implication of Equation (2) is that counter-explanation likely decreases the judged posterior probability of auditor negligence.

In addition, the counter-explanation-induced unpacking of ($MM|AN^C$) likely also causes evaluators to incrementally distort the weight they give to the probability of material misstatement conditional on the absence of auditor negligence when evaluators assess auditor negligence. This incremental distortion would reduce outcome bias for relatively low Bayesian posteriors of auditor negligence but also exacerbate reverse outcome bias for relatively high Bayesian posteriors of auditor negligence. Formally, relative to the *control* condition, the implications of the *counter-explanation* condition are:

H1: The *counter-explanation* intervention will decrease outcome bias for low Bayesian posteriors of auditor negligence (< 40%) relative to the *control* condition, but will exacerbate reverse outcome bias for high Bayesian posteriors of auditor negligence (≥ 40%) relative to the *control* condition.

Attribution Intervention

Kadous (2001) introduces the *attribution* intervention to the accounting literature by levering the affect-as-information literature in psychology (Schwarz and Clore 1983). Because individuals prefer to see the world as fair and just (Lerner and Simmons 1966), they experience negative affect after receiving information about negative outcomes that others experience, which motivates them to lay blame (Walster 1966). So, upon learning of adverse audit outcomes, including investor losses, evaluators of auditors experience negative affect and treat it as a cue suggestive of auditor negligence.

Anticipating that negative affect and increased motivation to lay blame moderate outcome effects, Kadous (2001) employs an intervention that weakens their influence on evaluators' judgments of auditor negligence. The intervention reminds participants that their tension and anxiety could be attributable to their difficult role as an evaluator and so redirects the influence of negative affect on the evaluative process (Kadous 2001, 430). As predicted, Kadous (2001) observes smaller outcome effects in the *attribution* intervention condition.

We posit that the *attribution* intervention also will change the extent to which evaluators' posterior judgments of auditor negligence depart from the Bayesian benchmark posterior via the support accumulation stage and probability weighting stages of probability judgment. In Support Theory terms, evaluators' realization that their negative affect plausibly does not signify auditor negligence increases the support of the residual hypothesis (e.g., something besides auditor negligence caused the undetected misstatement), and so decreases the net support of the focal hypothesis (e.g., auditor negligence caused the undetected misstatement). This shift in relative support reduces the assigned subjective probability of auditor negligence. In addition, however, this realization likely further distorts the weight afforded to the probability of material misstatement in the absence of auditor negligence during evaluators' belief revision. One consequence of this incremental distortion is a reduction in outcome bias for relatively low Bayesian posteriors of auditor negligence. A second consequence, however, is exacerbation of reverse outcome bias for relatively high Bayesian posteriors of auditor negligence. Thus, the directional predictions of the effects of the *attribution* intervention are the same as the predictions for the *counter-explanation* intervention:

H2: The *attribution* intervention will decrease outcome bias for low Bayesian posteriors of auditor negligence (< 40%) relative to the *control* condition, but it also will exacerbate reverse outcome bias for high Bayesian posteriors (≥40%) relative to the *control* condition.

Warning Intervention

Building on interventions employed in the psychology literature (see Fischhoff 1982 for a review), some interventions in the accounting outcome-effect literature advise participants not to over-rely (or not to rely *at all*) on outcome information when judging auditor negligence (e.g., Clarkson et al. 2002). This *warning* intervention empirically causes limited, if any, reduction in outcome effects (Fischhoff 1982). Relative to counter-explanation, warning-like interventions have not been extensively examined in the audit litigation literature. Clarkson et al. (2002) employ two warning interventions termed *weak* and *awareness*. In the *weak* condition, participants receive advice about the ostensible non-normativeness of using outcome information to judge auditor negligence and encouragement to ignore outcome information completely when judging auditor negligence. We note that this warning advocates a biased decision rule and asks participants to forgo natural and reasonable reliance on diagnostic outcome information. Consistent with prior research, the researchers did not observe a reduction in outcome effects in this condition.

Clarkson et al. (2002) do report a reduction in outcome effects in the *awareness* condition. In this condition, the researchers provide additional information to participants ("…deepened recession, nonavailability of new financing due to 'tight money' conditions, unexpected litigation, etc." p. 18). Provision of this additional information, however, creates a hybrid intervention that mixes *warning* and *counter-explanation*. That is, the additional information is alternative explanations for the adverse outcome that do not implicate auditor negligence.

Herein, we employed a variation of the *weak* intervention employed by Clarkson et al. (2002). Because asking participants to ignore outcome information would be non-normative, our language encourages participants to avoid over-relying on outcome information (see Appendix 1 Panel C). In terms of outcome bias and reverse outcome bias, our framework does not warrant a hypothesis about how the *warning* intervention will affect either form of bias relative to the *control* condition. Yet because of the *warning* intervention's prevalence in the literature, we pose a research question and explore whether it affects outcome and/or reverse outcome bias.

> **RQ**: Does the *warning* intervention affect either form of bias relative to the *control* condition?

New Interventions

Re-weighting Interventions

One likely can reduce both outcome bias and reverse outcome bias by attacking the distortion to subjective probabilities caused by evaluators' reliance on their probability weighting functions. Cumulative Prospect Theory (Tversky and Kahneman 1992) holds that people over-weight low probabilities (typically < 40%), the range in which outcome bias manifests itself in Peecher and Piercey (2007). It also holds that people under-weight relatively high probabilities (> 40%), the range in which reverse outcome bias manifests itself in Peecher and Piercey (2007).

Our new intervention essentially explains for participants how people tend to distort probabilities via the probability weighting function (see Appendix 1 Panel D) and admonishes participants to avoid doing so. Specifically, it directs participants to consider decreasing (increasing) their judgment of auditor negligence if the case suggest that it is relatively unlikely (likely) that the auditor is negligent.

We devote two conditions to *re-weighting*, and these conditions differ in multiple psychologically important respects. One difference is timing. The *re-weighting* intervention occurs before participants read the main case, but the *re-weighting – after* intervention occurs after participants finish the case and reach their posterior judgments of auditor negligence. We walk these latter participants through a manual correction process whereby they are informed of probability mis-weighting and given an opportunity to revise their preliminary negligence judgment. This type of manual revision has the benefit of being a particularly deliberate version of the re-weighting intervention. However, it naturally must occur after individuals have already formed their judgments. While individuals commonly think they can counteract biases after forming judgments (Gilbert 1991), ex post interventions tend to have little influence on judgment (see, e.g., Lerner and Tetlock 1999). Thus, whether our *re-weighting – after* intervention helps individuals overcome this inability to correct for biases ex post is an empirical question. Another difference is that the *re-weighting – after* intervention emphasizes how the probability distortion occurs "on average." Since people are overconfident about their susceptibility to judgment bias and since ex post interventions tend to be less influential than ex ante interventions in altering judgment (Kruger and Dunning 1999), we predict the *re-weighting – after* intervention to be less effective in promoting Bayesian-consistent belief revision or in reducing outcome bias or reverse outcome bias, relative to the *re-weighting* intervention. Whether the *re-weighting – after* intervention improves bias relative to the *control* condition or to other interventions is also an empirical question.[7]

[7] We made the design choice to have our re-weighting interventions differ in multiple respects within one experiment to explore the robustness of the *re-weighting* intervention. An alternative design choice would entail tweaking one facet of the re-weighting at time. Our findings, as discussed later, enable us to learn that, while successful, the *re-weighting* intervention is sensitive to timing, language that stresses "on average" biases, or both of these. Isolating the full set of boundary conditions for the *re-weighting* intervention represents a fruitful area of future research.

As such, our next hypothesis focuses on how the *re-weighting* intervention will improve participants' judgments of auditor negligence, relative to the *control* condition and to the *re-weighting – after* intervention.

H3a: The *re-weighting* intervention will reduce outcome bias for relatively low Bayesian-revised probabilities of auditor negligence (< 40%) and reduce reverse outcome bias for relatively high Bayesian-revised probabilities of auditor negligence (≥ 40%), relative to the *control* condition and to the *re-weighting – after* intervention.

Comparing the *re-weighting* intervention to the traditional interventions, our theory predicts that the *re-weighting* will out-perform the *warning* intervention with respect to outcome bias. With respect to reverse outcome bias, our theory predicts the *re-weighting* intervention will result in less reverse outcome bias compared to the *counter-explanation*, *attribution,* and *warning* interventions.

H3b: The *re-weighting* intervention will reduce outcome bias for relatively low Bayesian-revised probabilities of auditor negligence (< 40%), relative to the *warning* intervention, and reduce reverse outcome bias for relatively high Bayesian-revised probabilities of auditor negligence (≥ 40%), relative to the *counter-explanation, attribution, and warning* interventions.

Summary of H1 through H3

Figure 1 summarizes the H1-H3 predictions in linear terms.[8] As illustrated, we predict outcome bias to be lower in the *counter-explanation, attribution*, and *re-weighting* conditions

[8] In cubic polynomial regressions, Peecher and Piercey (2007) observe that linear, squared, and cubed forms of the Bayesian judgment of auditor negligence (*Bayes*) each explain a statistically significant portion of the variation in participants' judgments of auditor negligence. Yet, the linear *Bayes* term accounts for over 97% of the variation in participants' judgment that is explained by their estimation models. And, both $Bayes^2$ and $Bayes^3$ provide very little incremental adjusted-R^2. Accordingly, we focus on linear *Bayes* term herein to simplify exposition. We note, however, that nonlinear forms of *Bayes* attain statistical significance only in a minority of our experimental conditions. This is not surprising for a few reasons. One, researchers have yet to address factors that moderate conditions under which the probability weighting function is predominantly linear. Two, our experimental design affords a low-power test for quadratic and cubic forms of nonlinearity. Three, the extent to which individuals' probability weighting functions are linear in *Bayes* is not readily discernable from related psychology research, as the estimation models used therein generally do not allow a linear fit to the data. We leave it to future research to examine interesting questions about what individual or contextual factors influence the shape of individuals' probability weighting functions, including its quadratic and cubic elements.

than in the *control*, *re-weighting -after*, and *warning* conditions. We also predict reverse outcome bias to be lowest in the *re-weighting* condition, next lowest in the *control*, *re-weighting – after*, and *warning* conditions, and highest in the *attribution* and *counter-explanation* conditions.

[Insert Figure 1 here]

The Explanatory Power of Bayesian Revision

We posit one additional hypothesis. The *re-weighting* intervention likely reduces distortions to participants' probabilities and so likely improves the extent to which Bayesian benchmark judgments explain variation in participants' actual judgments of auditor negligence. In contrast, the *counter-explanation*, *attribution*, and *warning* interventions likely incrementally distort participants' probability weighting and so undermine the extent to which the Bayesian benchmark explains variation in participants' actual judgments of auditor negligence. This potential inhibition of Bayesian reasoning is likely due to reducing individuals' reasonable and natural reliance on diagnostic outcome information (i.e., *counter-explanation* and *attribution*) or by advocating a biased decision rule (i.e., *warning*).

H4: Relative to the *control* condition, the *re-weighting* intervention improves the explanatory power of Bayesian revision for participants' actual auditor negligence probability judgments.

Relative to the *control* condition and to the *re-weighting* intervention, the *counter-explanation*, *attribution*, and *warning* interventions decrease the explanatory power of Bayesian revision for participants' actual auditor negligence probability judgments.

III. EXPERIMENTAL DESIGN

Participants

Like most accounting outcome-effect studies, we use non-professional evaluators.[9] Specifically, we employ 1,746 undergraduates enrolled in introductory accounting courses at a large Midwest university.[10] We exclude 136 participants for nonsensical, illegible, or incomplete responses resulting in a final sample of 1,610 participants.[11] Participants volunteered for modest extra-credit and had completed 1.54 years of post-high-school education (s = 1.1), 0.56 accounting courses (s = 0.9), and 4.37 business, accounting, and economics courses (s = 2.9). On a 9-point Likert scale, centered at 0, participants were neither sympathetic nor unsympathetic to auditors (mean = 0.05, s = 1.5, t = 1.39, $p_{\text{two-tailed}} = 0.16$).

Task

Participants completed experimental materials adapted from experiment 2 in Peecher and Piercey (2007).[12] See Exhibit 2 for an overview of the experimental procedures. The average time to complete the task was 22.61 minutes (s = 5.0). Participants began by receiving a packet that included an introduction and the case materials. The introduction explained basic concepts such as material misstatements, unqualified audit opinions, reasonable assurance, auditor negligence, and due professional care. The case materials included review questions to reinforce and aid in comprehension of these concepts. As in Peecher and Piercey (2007), we wanted to

---

[9] Many non-professional and professional evaluators assess the quality of auditors' judgments and decisions in many contexts (e.g., business school students, readers of the popular business press, jurors, mediators, judges, lawyers, regulators, etc.). Since we are aware of no theory or evidence that students revise their beliefs differently than potential jurors or other evaluators of auditors (see, e.g., Bornstein and Rajki 1994), our use of students as participants is justified (Libby, Bloomfield, and Nelson, 2002; Solomon and Peecher, 2001).

[10] A large number of participants helps ensure reasonably powerful tests of our hypotheses and judged probabilities of auditor negligence that reasonably span the [0,1] interval within each experimental condition. As explained later, we effectively have 24 experimental conditions with a 2 × 6 between-subjects manipulated design coupled with a measured binary variable (2 × 6 × 2 = 24).

[11] An example of a nonsensical response would be a Bayesian posterior greater than 100%.

[12] Peecher and Piercey (2007) adapt their materials from the Big Time Gravel case in Kadous (2000, 2001).

provide sufficient opportunity for evaluators' assessments to be overly harsh to demonstrate the robust nature of reverse outcome bias (i.e., over-leniency). Thus, the introduction included language adapted from Kadous' (2000) *severe consequences* condition. Our case emphasizes the financial and emotional losses associated with undetected material misstatements, including the severe consequences associated with the Enron and WorldCom alleged audit failures.

[Insert Exhibit 2 here]

After the introduction and review questions, the materials presented the main case. The first part of the case presented information about the audit of Big Time Gravel, conducted by the audit firm Jones & Company. This part also included the audit quality manipulation and some of the intervention manipulations, both of which we explain hereafter in the independent variable section of this paper. After this part, the materials provided audit quality manipulation check questions as well as questions that allowed participants to convey their prior beliefs using natural frequencies.[13] Specifically, the participants reported their prior beliefs about the base rates of material misstatements (*p(MM)*), auditor negligence (*p*(*AN*)), and material misstatements given auditor negligence (*p*(*MM*|*AN*)). We used these measures to help compute the participants' Bayesian posterior (*Bayes*; see equation (3) in Footnote 14).

The next part of the materials featured the outcome information: The Securities and Exchange Commission (SEC) had investigated Big Time Gravel's financial statements and concluded they were materially misstated. The materials included a question asking participants about the probabilistic nature of this outcome information. That is, it asked participants to report the probability that the SEC is correct in concluding that financial statements are materially

[13] As in Peecher and Piercey (2007), we solicited participants' prior beliefs using natural frequencies as opposed to probabilities as individuals come closer to being Bayesians when using the former compared to the latter (Gigerenzer and Hoffrage 1995, 1999; Gigerenzer 2000). Thus, using frequencies biases us against observing our predicted form of non-Bayesian belief revision.

misstated (P^*).[14] To complete the main case, participants reported their revised beliefs about auditor negligence conditional on a probabilistic outcome ($p(AN|MM)$; the observed posterior). After completion, participants turned in the main case to the administrator and completed a post-experimental questionnaire. This questionnaire included demographic questions and the intervention manipulation checks.

Design and Independent Variables

The experiment features a 2 × 6 × 2 factorial design. One factor manipulates the quality of audit performed. We included this factor primarily to ensure participants' auditor negligence judgments reasonably spanned the [0,1] probability interval but also to allow exploration of whether our interventions interact with audit quality. The second manipulated factor is the one of primary theoretical interest and includes the five interventions that we predict to influence the extent and nature of bias as well as a control condition. The third factor is a measured variable that facilitates observation of whether outcome bias and reverse outcome bias obtain in the predicted ranges of Bayesian judgments of auditor negligence. This measured variable categorizes participants according to whether their Bayesian judgments of auditor negligence fall below or above 40%, an empirical threshold associated with Cumulative Prospect Theory's probability weighting function. We next explain each factor in turn.

We manipulated audit quality between-participants at two levels: lower and higher.[15] In the higher audit quality condition, an audit senior with considerable industry experience and a

[14] If a misstatement outcome, MM, has a probability P^*, the *Bayesian* posterior of auditor negligence is:

$$Bayes = P^* \times \left(p(AN) \times \left(\frac{p(MM \mid AN)}{p(MM)} \right) \right) + (1 - P^*) \times \left(p(AN) \times \left(\frac{p(\overline{MM} \mid AN)}{p(\overline{MM})} \right) \right), \tag{3}$$

where $\overline{MM}$ is a no-misstatement outcome. Note that, when $P^* = 1$, equation (3) reduces to $Bayes = p(AN|MM)$, and when $P^* = 0$, equation (3) reduces to $Bayes = p(AN|\overline{MM})$.

[15] These conditions only represent relative levels of audit quality. In neither audit quality condition do we conclude whether the auditor did or did not conform to standards of due professional care.

second year staff auditor with some industry experience visited 5 of 11 inventory locations while informing management of the specific locations two weeks in advance. In the lower audit quality condition, two second-year staff auditors, each with some industry experience, visited 4 of 11 inventory locations while informing management of the specific locations one month in advance.

We manipulated interventions between-participants at six levels. In the *control* condition, participants completed the experimental case with no additional instructions. This condition serves as a baseline for several of our planned comparisons and helps establish the generalizability of theory presented in Peecher and Piercey (2007).

In the *counter-explanation* condition and prior to their revised auditor negligence probability judgment, participants rated the likelihood of two alternative explanations under which a material misstatement could have occurred without Jones and Company's negligence on a 9-point Likert scale.[16] The first alternative explanation described a scenario where the material misstatement was isolated to one audit location reasonably not sampled for a visit by the auditor. The second alternative explanation described a scenario where the auditor judged a market price decline to be temporary where, in hindsight, the SEC judged it to be permanent. Consistent with prior research, participants encounter the *counter-explanation* intervention after the outcome information but before their final auditor negligence judgments since the intervention is designed to break down the causal linkages automatically created by the outcome information. See Appendix 1 Panel A for the intervention used in the experiment.

In the *attribution* condition, we used language similar to that in Kadous (2001). Participants read a paragraph describing how real evaluators often feel anxious or tense when making decisions. We then asked participants to report any feelings of anxiety or tension, and

[16] We did not ask participants to provide their own alternative explanation(s) because we wanted to ensure the salience of at least some counter-explanations. Kadous et al. (2006) provide evidence that salience of counter-explanations would help the counter-explanation intervention reduce outcome effects.

whether these feelings were potentially attributable to the participant's role as an evaluator. Consistent with Kadous (2001), participants encountered this manipulation at the beginning of the experimental case. See Appendix 1 Panel B for the experimental materials related to the *attribution* manipulation.

In the *warning* condition, we adapted the *weak* manipulation used in Clarkson et al. (2002) to our experimental case. After receiving the outcome information but before making their final auditor negligence probability judgment, participants read a paragraph describing outcome bias and how it relates to their upcoming judgments. The intervention's language encouraged participants not to over-rely on outcome information when judging auditor negligence. Participants then answered a manipulation check question that also served to reinforce the manipulation. See Appendix 1 Panel C for the related experimental materials.

We devoted two conditions to *re-weighting*. In both *re-weighting* conditions, we informed the participants of how research has found that evaluators tend to over-react to outcome information when the mathematical (i.e., Bayesian) probability of auditor negligence is low, but tend to under-react to outcome information when this probability is high. Also, both conditions included a question to reinforce the manipulation. In the *re-weighting* condition, the manipulation came at the beginning of the experimental case. In the *re-weighting – after* condition, the manipulation appeared *after* participants already had made their posterior auditor negligence probability judgment. It asked them if they wanted to revise their judgment. See Appendix 1 Panel D for the experimental materials for both *re-weighting* interventions.

Our final independent variable, *Hilo*40, classifies participants' Bayesian posteriors as either < 40% or ≥ 40%. Consistent with prior research on the probability weighting function, Peecher and Piercey (2007) observe outcome bias for Bayesian posteriors below 40% and reverse outcome bias for Bayesian posteriors above 40%. Thus, the 40% threshold for this dichotomous variable is theory consistent and allows us to separately test for the influence of de-biasing interventions on outcome bias and reverse outcome bias over ranges in which we predict each form of bias most likely to exist.

Dependent Variables

The primary dependent variables concern participants' judgments of auditor negligence given adverse outcome information ($p(AN|MM)$) and how these judgments differ from Bayesian posterior judgments of auditor negligence. Secondary dependent variables include the participants' prior beliefs about the frequency of material misstatements ($p(MM)$), auditor negligence ($p(AN)$), and material misstatements given auditor negligence ($p(MM|AN)$). Using these prior beliefs, coupled with outcome information, we computed the participants' Bayesian posterior (*Bayes*). See equation (3) in footnote 14 for the calculation of *Bayes*. After elicitation of their prior beliefs and presentation of outcome information, we elicited their revised belief or conditional auditor negligence probability judgment ($p(AN|MM)$). From these dependent variables, we calculated outcome bias, reverse outcome bias, and outcome effects. *Bias* equals the difference between the participant's Bayesian posterior (*Bayes*) and observed posterior ($p(AN|MM)$), and represents either outcome bias or reverse outcome bias. If $Bias = p(AN|MM) - Bayes$ is positive, the difference is outcome bias. If $Bias = p(AN|MM) - Bayes$ is negative, the difference is reverse outcome bias. We measure outcome effects by the extent to which the participants' auditor negligence judgments changed after consideration of the outcome information (i.e., outcome effects = $p(AN|MM) - p(AN)$).

Manipulation Checks

Three audit quality manipulation check questions were asked immediately after the participants read the information about the audit. We chose this timing in order to increase the salience of the manipulation and the likelihood that the audit quality information would affect, but not polarize, their judgments. See Appendix 2 Panel A for the audit quality manipulation check questions.[17] On the post experimental questionnaire, we also asked a manipulation check question related to the intervention manipulation. Without being able to look back the case materials, the answers to this question provide evidence that the participant actually processed the manipulation.[18] See this manipulation check question in Appendix 2 Panel B.

IV. RESULTS

Manipulation Checks

For the audit quality manipulation, 98.7% of participants answered the three manipulation check questions successfully. In addition, untabulated results indicate a significant main effect of audit quality on both the observed posteriors ($p_{\text{one-tailed}} = 0.003$) and Bayesian posteriors ($p_{\text{one-tailed}} = 0.001$). Thus, the audit quality manipulation was effective. For the intervention manipulation check, 70.7% of participants answered the question on the post-experimental questionnaire correctly. Excluding the participants who did answer the manipulation check questions correctly does not qualitatively influence our results. Overall, it appears that the participants successfully attended to the experimental interventions.

[17] We wanted the audit quality manipulation to be transparent though not extreme. So, we structured the alternative answers of the multiple-choice manipulation check question such that, in the higher audit quality condition, there was only one answer that featured a better audit procedure than that in the correct answer. In the lower audit quality condition, two alternatives featured better audit procedures than that in the correct answer. As such, the manipulation check question not only verified attention to the manipulation, but also served to strengthen the manipulation.

[18] Based on responses from the first 24 participants, we learned that this question was vaguely worded. Thus, we clarified it for remaining participants. The results are qualitatively unchanged if we omit these 24 participants.

Overall Findings for Outcome Bias and Reverse Outcome Bias

In Figure 2, we present fitted linear regression lines of participants' judged probability of auditor negligence. These graphs visually depict how the interventions moderated participants' outcome bias and reverse outcome bias. Observe that, relative to other graphed conditions, the regression line in the *re-weighting* condition most closely approximates the 45-degree angle (the benchmark Bayesian line) and suggests the least amount of outcome bias and reverse outcome bias. Observe also how the regression lines in the other intervention conditions tend to be flatter and suggest greater bias. We next turn to inferential tests of our hypotheses.[19]

[Insert Figure 2 here]

Before running planned contrasts to test our specific hypotheses about outcome bias and reverse outcome bias, we ran an ANOVA model with *bias* as the dependent variable and with *Intervention*, *Audit Quality*, and *Hilo*40 as independent factors. (We also ran a general linear model with a continuous specification of *Bayes* instead of the *Hilo*40 binary variable and obtained similar results.)

The *Bias* dependent variable is positive for outcome bias (expected when *Bayes* < 40%) and negative for reverse outcome bias (expected when *Bayes* ≥ 40%). Panel A of Table 1 provides summary statistics based on this model, Panel B provides the corresponding ANOVA table, and Figure 3 plots least-square means. Because the three-way interaction between *Intervention*, *Audit Quality*, and *Hilo*40 is significant at $p_{\text{two-tailed}} = 0.036$, we provide descriptive statistics for each of the 2 × 6 × 2 experimental conditions. As it turns out (and as explained

[19] In untabulated results, we examine whether the interventions decreased *outcome effects* relative to the *control* condition to examine consistency with prior research. We observe that, relative to the *control* condition (+0.184), the *counter-explanation* (+0.114) intervention decreases outcome effects (p = 0.002), consistent with prior research. We also observe that whether the *attribution* intervention decreases outcome effects depend on the level of audit quality. We elaborate on this latter effect later, when discussing post hoc tests related to the *attribution* intervention in particular. The *warning* intervention does not decrease outcome effects compared to the *control* condition (p = 0.27), consistent with prior research.

later), the *attribution* condition drives this three-way interaction because its results unexpectedly differ depending on *Audit Quality*. For now, we turn to our overall and within-condition tests for outcome bias and reverse outcome bias.

[Insert Table 1 and Figure 3 here]

Recall that we predict outcome bias (positive *Bias*) to be lower in the *attribution*, *counter-explanation*, and *re-weighting* conditions than in the *control*, *re-weighting – after*, and *warning* conditions. For *Bayes* < 40% (Hilo40 = 0), the planned contrast of the average of *attribution* (*Bias* = +0.132), *counter-explanation* (*Bias* = +0.094), and *re-weighting* (*Bias* = +0.077) against the average of *control* (*Bias* = +0.144), *re-weighting – after* (*Bias* = +0.176), and *warning* (*Bias* = +0.135) conditions is statistically significant ($F_{1,\ 1586} = 8.94$, $p < 0.01$). A departure is that outcome bias in *attribution* condition is higher than expected.

With respect to reverse outcome bias (negative *Bias*), recall that we predict it to be closest to zero in the *re-weighting* condition, next lowest in the *control*, *re-weighting – after*, and *warning* conditions, and highest in the *attribution* and *counter-explanation* conditions. For *Bayes* ≥ 40% (Hilo40 = 1), a planned contrast with weights of -3 for *re-weighting* (*Bias* = -0.051), -1 for *control* (*Bias* = -0.126), *re-weighting – after* (*Bias* = -0.125), and *warning* (*Bias* = -0.165), and +3 for *attribution* (*Bias* = -0.133) and *counter-explanation* (*Bias* = -0.211) is statistically significant ($F_{1,\ 1586} = 8.05$, $p < 0.01$). The lone departure from the predicted order occurs in the *attribution* condition. Instead of being significantly worse, its magnitude of reverse outcome bias unexpectedly approximates that of the *control* condition. We address this departure hereafter discussing the aforementioned three-way interaction among *Intervention*, *Audit Quality*, and *Hilo*40. We next consider findings within each intervention condition.

Control Condition – Replication of Peecher and Piercey (2007)

Following Peecher and Piercey (2007), we expected a pattern of outcome bias and reverse outcome bias for relatively low and relatively high Bayesian posterior probabilities of auditor negligence in the *control* condition, respectively. For outcome bias, Table 1 Panel A shows evidence of significant outcome bias and reverse outcome bias in the *control* condition. For participants with Bayesian posteriors less than 40%, their observed posterior was on average 14.4 percentage points higher than their Bayesian posterior. This outcome bias is significantly different from zero at the $p_{one\text{-}tailed} < 0.01$ level. For participants with Bayesian posteriors greater than 40%, their observed posterior was on average 12.6 percentage points lower than their Bayesian posterior. This reverse outcome bias is significantly different from zero at the $p_{one\text{-}tailed} < 0.01$ level. Combined, these results demonstrate a successful replication of Peecher and Piercey (2007) at the linear level and facilitate examination of the incremental effects of the interventions on outcome bias and reverse outcome bias relative to the *control* condition (see Table 1 Panel C).

Counter-explanation Intervention

Linearly transforming the Likert scale for each counter-explanation to a 0-1 scale, participants rated the likelihood of the two counter-explanations as 0.638 and 0.456 indicating that they felt the explanations were plausible. In H1, we predict the *counter-explanation* intervention to decrease outcome bias relative to the *control* condition, but to exacerbate reverse outcome bias. For Bayesian posteriors less than 40%, the *counter-explanation* intervention reduces outcome bias compared to the *control* condition (+0.094 versus +0.144; $F = 2.98$, $p_{one\text{-}tailed} = 0.041$). For Bayesian posteriors greater than 40%, the *counter-explanation* intervention exacerbates reverse outcome bias relative to the *control* condition (-0.211 versus -0.126; $F = 3.95$, $p_{one\text{-}tailed} = 0.024$). H1 is supported.

Since we observe that *counter-explanation* exacerbates reverse outcome bias, we conducted supplemental analyses to investigate further. We calculated how many participants exhibited reverse outcome bias in the *counter-explanation* and *control* conditions when *Bayes* ≥ 40%. The number of such participants in *counter-explanation* (66 out of 91, 72.5%) is significantly higher than the number of participants in the *control* condition (41 out of 81, 50.6%; $\chi^2 = 17.48$; $p < 0.001$).

Attribution Intervention

Participants' average score on the question asking for the extent of their anxiety or tension was 1.52 (See Appendix 1 Panel B), while the average score for the question asking whether this anxiety or tension was attributable to their role as evaluator was 1.66.[20] Overall, 244 out of 266 (91.7%) participants in this condition reported feeling some anxiety or tension.

In H2, we predict the *attribution* intervention to decrease outcome bias relative to the *control* condition, but to exacerbate reverse outcome bias. For Bayesian posteriors less than 40%, the *attribution* intervention did not significantly reduce the extent of outcome bias compared to the *control* condition (+0.144 versus +0.132; $F = 0.19$, $p_{\text{one-tailed}} = 0.330$). This result is inconsistent with H2. For Bayesian posteriors greater than 40%, the *attribution* intervention does not increase the extent of reverse outcome bias relative to the *control* condition (-0.133 versus -0.126; $F = 0.03$, $p_{\text{one-tailed}} = 0.483$).

Our outcome bias ANOVA revealed a significant three-way interaction of *Intervention, Audit Quality,* and *Hilo*40 that is driven by the *attribution* intervention. In post hoc analyses, we observe that when audit quality is relatively high, the *attribution* intervention increases reverse outcome bias relative to the *control* condition ($F = 4.45$, $p_{\text{two-tailed}} = 0.035$). When audit quality is

[20] We coded the responses to the *attribution* questions as 0, 1, 2, or 3, with the score increasing in anxiety or tension experienced and attribution to their role as evaluator.

relatively low, the *attribution* intervention decreases reverse outcome bias relative to the *control* condition ($F = 4.75$, $p_{\text{two-tailed}} = 0.029$). A reasonable post hoc explanation of this pattern of results is that intervention fails to redirect attribution of experienced anxiety and tension from the auditor when audit quality is relatively low. That is, if the audit is perceived to be of relatively low quality, participants may find it implausible that their anxiety and tension primarily is due to the tough role as an evaluator and so continue to lay blame on the auditor.

This post-hoc explanation becomes more plausible when one considers the pattern of *outcome effects* (i.e., $p(AN|MM) - p(AN)$) that participants exhibit. In untabulated results, we observe that when audit quality is lower, outcome effects findings are inconsistent with Kadous (2001): Outcome effects are larger in the *attribution* condition than in the *control* condition (difference = 0.057; $F = 2.78$, $p_{\text{two-tailed}} = 0.096$). When audit quality is higher, however, outcome effect findings are consistent with Kadous (2001): Smaller outcome effects occur in the *attribution* condition than in the *control* condition (difference = -0.087; $F = 5.54$, $p_{\text{two-tailed}} = 0.019$).

To explore the plausibility of our post-hoc explanation further, we examine the responses to the second *attribution*-specific question (see Appendix 1 Panel B). This question measures the extent to which participants attributed their anxiety and tension to their role as an evaluator. According to our post-hoc explanation, when audit quality is relatively low, the intervention would be less successful in redirecting participants' attribution of their anxiety to their tough role as an evaluator. Participants' responses to this question significantly differ with the level of audit quality ($\chi^2 = 10.56$; $p = 0.014$), and the pattern of responses support our post-hoc proposition.

Specifically, if audit quality is lower, participants are inclined reject the idea that their anxiety and tension stem from their role as an evaluator and so continue to blame the auditor.[21]

Warning Intervention

98.1% of participants answered the *warning*-specific question correctly. We explore whether this traditional intervention increases or decreases either outcome bias or reverse outcome bias relative to the *control* condition. For Bayesian posteriors less than 40%, the *warning* intervention does not influence the extent of outcome bias compared to the *control* condition (+0.135 versus +0.144; $F = 0.11$, $p_{\text{two-tailed}} = 0.74$). Likewise, for Bayesian posteriors greater than 40%, the *warning* intervention does not influence the extent of reverse outcome bias relative to the *control* condition (-0.165 versus -0.126; $F = 0.84$, $p_{\text{two-tailed}} = 0.36$).

Re-Weighting Interventions

95.2% answered the *re-weighting*-specific question correctly compared to 88.2% in the *re-weighting – after* condition. In H3a, we predict the *re-weighting* intervention to simultaneously reduce both outcome bias and reverse outcome bias relative to the *control* condition. The *re-weighting* intervention reduces outcome bias for low Bayesian posteriors (+0.077 versus +0.144; $F = 5.19$, $p_{\text{one-tailed}} = 0.011$). For reverse outcome bias, the results are particularly strong. In Table 1 Panel A, we show that the *re-weighting* intervention eliminates reverse outcome bias from statistical significance when audit quality is relatively high. The overall reduction in reverse outcome bias in the *re-weighting* condition compared to the *control* condition is significant (-0.051 versus -0.126; $F = 3.18$, $p_{\text{one-tailed}} = 0.037$). H3a is supported.

[21] Of course, we cannot rule out the possibility of other explanations in this single study. Other factors potentially causing the difference in our results versus results in Kadous (2001) include the more extensive legal context in her study, a higher-yet level of audit quality in her study (e.g., use of an independent inventory specialist), and the different control group design. We leave examination of these and other possibilities to future research.

As expected, compared to the *re-weighting* intervention, the *re-weighting – after* intervention has greater outcome bias (+0.176 versus +0.077; F = 11.04, $p_{one\text{-}tailed} < 0.001$) and greater reverse outcome bias (-0.125 versus -0.051; F = 3.10, $p_{one\text{-}tailed} = 0.039$). It also features similar levels of outcome bias and reverse outcome bias as the *control* condition (see Panel C of Table 1).

In H3b, we predict that the *re-weighting* intervention to reduce outcome bias compared to the *warning* intervention and reduce reverse outcome bias relative to the *counter-explanation, attribution,* and *warning* interventions. Relative to the *warning* intervention, the *re-weighting* intervention decreases outcome bias (+0.135 versus +0.077; F = 3.79, $p_{one\text{-}tailed} = 0.026$). The *re-weighting* intervention also decreases reverse outcome bias relative to the *counter-explanation* (-0.051 versus -0.211; F = 15.37, $p_{one\text{-}tailed} < 0.001$), *attribution* (-0.051 versus -0.133; F = 3.81, $p_{one\text{-}tailed} = 0.026$), and *warning* interventions (-0.051 versus -0.165; F = 7.60, $p_{one\text{-}tailed} = 0.003$). H3b is supported.

Explanatory Power of Bayesian Belief Revision (H4)

In H4, we predict that, relative to the *control* condition, the *re-weighting* intervention will improve the Bayesian consistency of participants' judgments of auditor negligence. And, we predict that *attribution*, *counter-explanation*, *warning* interventions will erode the relationship between the Bayesian revision and participants' judgments of auditor negligence. To test this hypothesis, we regress the observed posterior ($p(AN|MM)$) on *Bayes* and *Intervention* condition. We tabulate this regression in Table 2 Panel A, contrasts in Table 2 Panel B, and descriptive statistics regarding adjusted R^2 in Table 2 Panel C. Turning to the latter panel first, observe that descriptively, the adjusted R^2 statistics generally support this pattern. The largest adjusted R^2

occurs in the *re-weighting* condition (32.6%), followed by the *control* and other intervention conditions (ranging from 8.8% to 18.8%).

[Insert Table 2 here]

Table 2, Panel B contrasts show the *re-weighting* intervention improves the explanatory power of Bayesian revision relative to the *control* condition ($p_{\text{one-tailed}} = 0.031$). Although the *counter-explanation* and *attribution* interventions do not decrease the explanatory power of Bayesian revision relative to the *control* condition ($p_{\text{one-tailed}} = 0.20$ and $p_{\text{one-tailed}} = 0.17$, respectively), both of these interventions decrease the explanatory of Bayesian revision relative to the *re-weighting* intervention (all p's$_{\text{one-tailed}} < 0.01$). The *warning* intervention decreases the explanatory power of Bayesian revision relative to the *control* condition ($p_{\text{one-tailed}} = 0.023$) and relative to the *re-weighting* condition ($p_{\text{one-tailed}} < 0.001$). Overall, these results suggest that our new, theory-based *re-weighting* intervention improves the explanatory ability of Bayesian reasoning while several of the traditional interventions used in accounting outcome-effect studies incrementally inhibit Bayesian revision. H4 is supported.

V. CONCLUSION

Our theory and experimental findings provide evidence on the effects of *de-biasing* interventions on outcome bias and reverse outcome bias. We predict and test how two categories of interventions affect the decision processes participants use to gather support and/or weight subjective probabilities. One category is a new *re-weighting* intervention designed to reduce both outcome bias and reverse outcome bias, and the other category features traditional interventions designed to reduce reliance on outcome information and so reduce outcome effects (with the untested hope of also reducing bias). We find that the latter variety of interventions sometimes exacerbate reverse outcome bias and only sometimes reduce outcome bias (e.g., *counter-*

explanation and *attribution* interventions). Meanwhile, our *re-weighting* intervention simultaneously reduces both of these cognitive biases. It also out-performs all other interventions and the *control* condition with respect to outcome bias, reverse outcome bias, and conformance with Bayesian belief revision.

Like any experimental study, the inferences one can reach in this paper are subject to several limitations. First, we have examined the de-biasing effects of five interventions in isolation. Perhaps future work will show that the best simultaneous de-biaser of outcome bias and reverse outcome bias is a combination of these or other interventions. Second, our theory is not sufficiently rich to predict effect sizes of interventions in real-world contexts, such as jury trials. Third, our experimental materials did not separately ask participants to assess the strength of evidence, standards of due professional care, support of hypotheses, as well as subjective probabilities (cf., Fox 1999). Instead, we collected only the latter. One consequence is that we cannot tease out the extent to which the interventions distorted participants' support amassing activities as opposed to their probability mis-weighting activities.

Despite these limitations, our theory and experimental findings embellish existing de-biasing frameworks in the accounting literature (e.g., Kennedy 1995), demonstrate how traditional interventions that decrease outcome effects significantly exacerbate reverse outcome bias, and demonstrate how a re-weighting intervention can reduce both outcome bias and reverse outcome bias.

REFERENCES

Anderson, J. C., M. M. Jennings, D. J. Lowe and P. M. J. Reckers. 1997. The mitigation of hindsight bias in judges' evaluation of auditor decisions. *Auditing: A Journal of Practice and Theory* 16(2): 20-39.

Anderson, J. C., D. J. Lowe and P. M. J. Reckers. 1993. Evaluation of auditor decisions: Hindsight bias effects and the expectation gap. *Journal of Economic Psychology* 14(4): 711-737.

Bornstein, B. H., and M. Rajki. 1994. Extra-legal factors and product liability: The influence of mock jurors' demographic characteristics and intuitions about the cause of an injury. *Behavioral Sciences and the Law* 12: 127-147.

Brown, C. E., M. E. Peecher and I. Solomon. 1999. Auditors' hypothesis testing in diagnostic inference tasks. *Journal of Accounting Research* 37(1): 1-26.

Brown, C. E. and I. Solomon. 1987. Effects of outcome information on evaluations of managerial decisions. *The Accounting Review* 62(3): 564 - 577.

Brown, C. E. and I. Solomon. 1993. An experimental investigation of explanations for outcome effects on appraisals of capital-budgeting decisions. *Contemporary Accounting Research* 10(1): 83-111.

Clarkson, P. M., C. Emby and V. W. S. Watt. 2002. Debiasing the outcome effect: The role of instructions in an audit litigation setting. *Auditing: A Journal of Practice and Theory* 21(2): 1-20.

Fischhoff, B. 1975. Hindsight does not equal foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance* 1: 288 - 299.

Fischhoff, B. 1982. Debiasing. *Judgment under uncertainty: Heuristics and biases.* D. Kahneman, A. Tversky and P. Slovic. New York, Cambridge University Press.

Fox, C. R. 1999. Strength of evidence, judged probability, and choice under uncertainty *Cognitive Psychology* 38: 167 - 189.

Fox, C. R. and A. Tversky. 1998. A belief-based account of decision under uncertainty. *Management Science* 44(7): 879 - 895.

Frederickson, J. R., S. A. Peffer and J. Pratt. 1999. Performance evaluation judgments: Effects of prior experience under different performance evaluation schemes and feedback frequencies. *Journal of Accounting Research* 37(1): 151 - 165.

Gigerenzer, G. 2000. *Adaptive Thinking: Rationality in the Real World*. Oxford, UK, Oxford University Press.

Gigerenzer, G. and U. Hoffrage. 1995. How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review* 102: 684 - 704.

Gigerenzer, G. and U. Hoffrage. 1999. Overcoming difficulties in Bayesian reasoning: A reply to Lewis & Keren and Mellers & McGraw. *Psychological Review* 106: 425 - 430.

Gilbert, D. T. 1991. How mental systems believe. *American Psychologist* 46: 107-119.

Hershey, J. and J. Baron. 1992. Judgments by outcomes: When is it justified? *Organizational Behavior and Human Decision Processes* 53: 89 - 93.

Hershey, J. C. and J. Baron. 1995. Judgment by outcomes: When is it warranted? *Organizational Behavior and Human Decision Processes* 62(1): 127.

Jolls, C. and C. R. Sunstein. 2006. Debiasing through law. *Journal of Legal Studies* 35: 199 - 241.

Kadous, K. 2000. The effects of audit quality and consequence severity on juror evaluation of auditor responsibility for plaintiff losses. *The Accounting Review* 75(3): 327-341.

Kadous, K. 2001. Improving jurors' evaluations of auditors in negligence cases. *Contemporary Accounting Research* 18(3): 425-444.

Kadous, K., S. D. Krische and L. M. Sedor. 2006. Using counter-explanation to limit analysts' forecast optimism. *The Accounting Review* 81: 377 - 397.

Kadous, K. and A. M. Magro. 2001. The effects of exposure to practice risk on tax professionals' judgments and recommendations. *Contemporary Accounting Research* 18(3): 451 - 475.

Kennedy, J. 1995. Debiasing the curse of knowledge in audit judgment. *The Accounting Review* 70(2): 249 - 273.

Kinney, W. R. and M. W. Nelson. 1996. Outcome information and the "expectation gap": The case of loss contingencies. *Journal of Accounting Research* 34(2): 281 - 299.

Koehler, J. J. 2006. Train our jurors. *Heuristics and the Law*. G. Gigerenzer and C. Engel. Cambridge, MA, The MIT Press.

Koonce, L. 1992. Explanation and counterexplanation in audit analytical review. *The Accounting Review*, 67 (1): 59-76.

Kruger, J. and D. Dunning. 1999. Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology* 77(6): 1121 - 1134.

Lerner, J. S. and P. E. Tetlock. 1999. Accounting for the effects of accountability. *Psychological Bulletin* 125(2): 255 - 275.

Lerner, M. J. and C. H. Simmons. 1966. Observer's reaction to the "innocent victim": Compassion or rejection. *Journal of Personality and Social Psychology* 4(2): 203 - 210.

Libby, R., R. J. Bloomfield and M. W. Nelson. 2002. Experimental Research in Financial Accounting. *Accounting, Organizations and Society* 27: 775 - 810.

Lipe, M. G. 1993. Analyzing the variance investigation decision: The effects of outcomes, mental accounting, and framing. *The Accounting Review* 68(4): 748 - 764.

Lowe, D. J. and P. M. J. Reckers. 1994. The effects of hindsight bias on jurors' evaluations of auditor decisions. *Decision Sciences* 25(3): 401-420.

Peecher, M. E. and M. D. Piercey. 2007. Judging audit quality in light of adverse outcomes: Evidence of outcome bias and reverse outcome bias. *Contemporary Accounting Research* (forthcoming).

Rottenstreich, Y. and A. Tversky. 1997. Unpacking, repacking, and anchoring: Advances in support theory. *Psychological Review* 104(2): 406 - 415.

Schwarz, N. and G. L. Clore. 1983. Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states. *Journal of Personality and Social Psychology* 45(3): 513 - 523.

Solomon, I. and M. E. Peecher. 2001. Theory and Experimentation in Studies of Audit Judgments and Decisions: Avoiding Common Research Traps. *International Journal of Auditing* 5(3): 193 - 203.

Tan, H.-T. and M. G. Lipe. 1997. Outcome effects: The impact of decision process and outcome controllability. *Journal of Behavioral Decision Making* 10: 315 - 325.

Tversky, A. and D. Kahneman. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5: 297 - 323.

Tversky, A. and D. J. Koehler. 1994. Support theory: A nonextensional representation of subjective probability. *Psychological Review* 101: 547 - 567.

Walster, E. 1966. Assignment of responsibility for an accident. *Journal of Personality and Social Psychology* 3(1): 73 - 79.

Wu, G. and R. Gonzalez. 1999. Nonlinear decision weights in choice under uncertainty. *Management Science* 45(1): 74 - 85.

Exhibit 1
Probability Weighting

A Hypothetical Probability Weighting Function
(Tversky and Kahneman 1992)

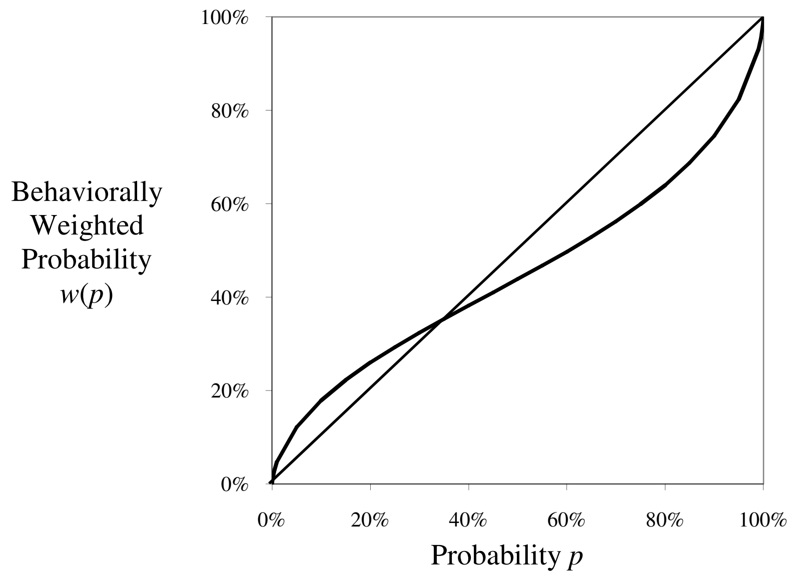

Observed Outcome Bias (Above Diagonal) and Reverse Outcome Bias (Below Diagonal)
(Peecher and Piercey 2007, Experiment 2)

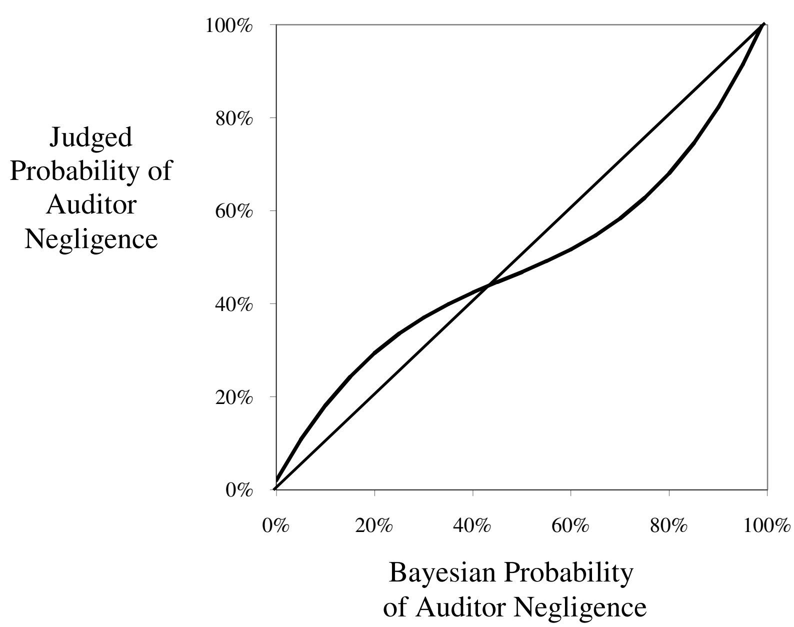

Exhibit 2
Experimental Procedures

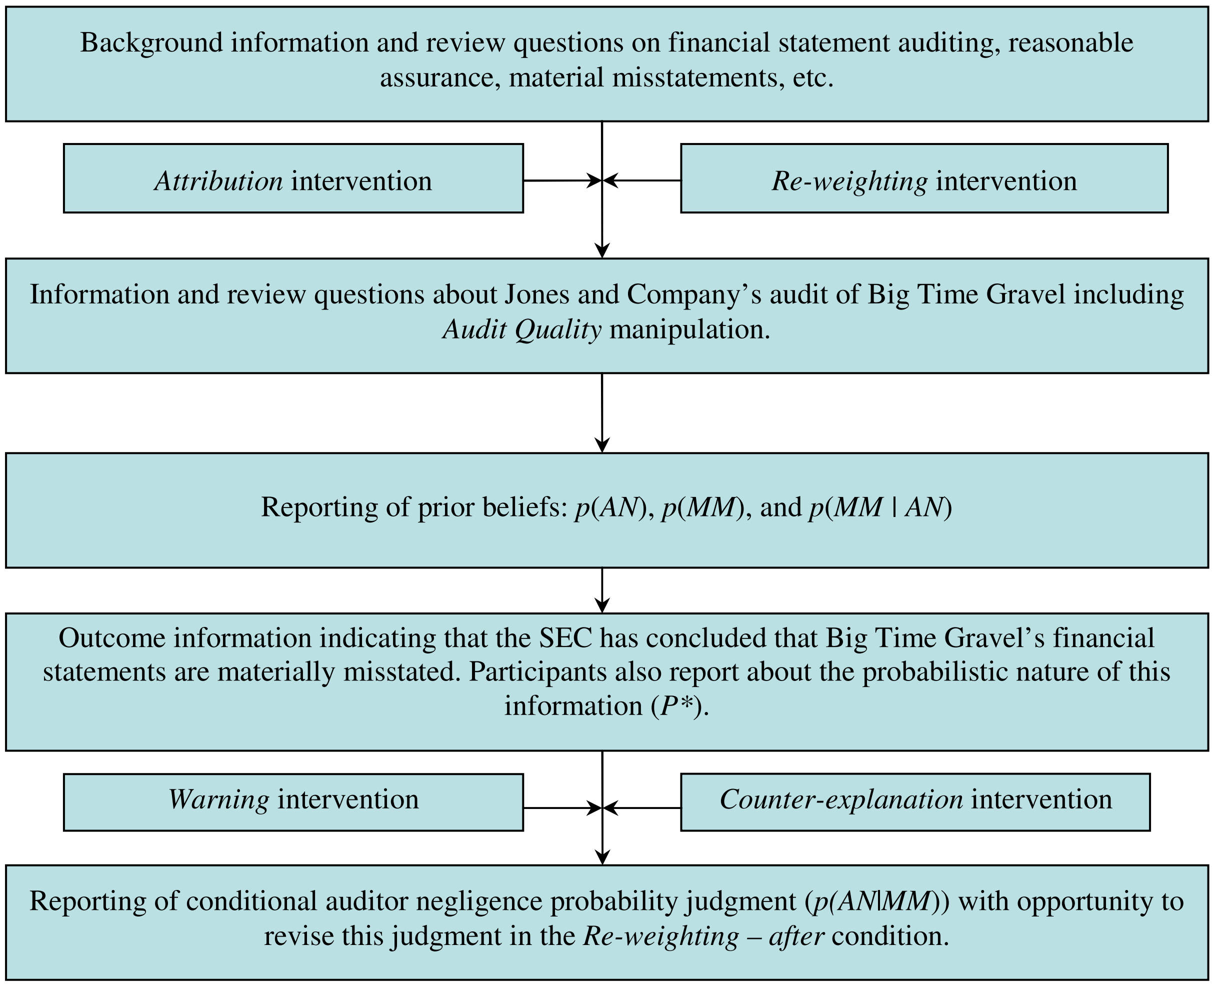

Figure 1
Graphical Representation of Hypotheses H1-H3

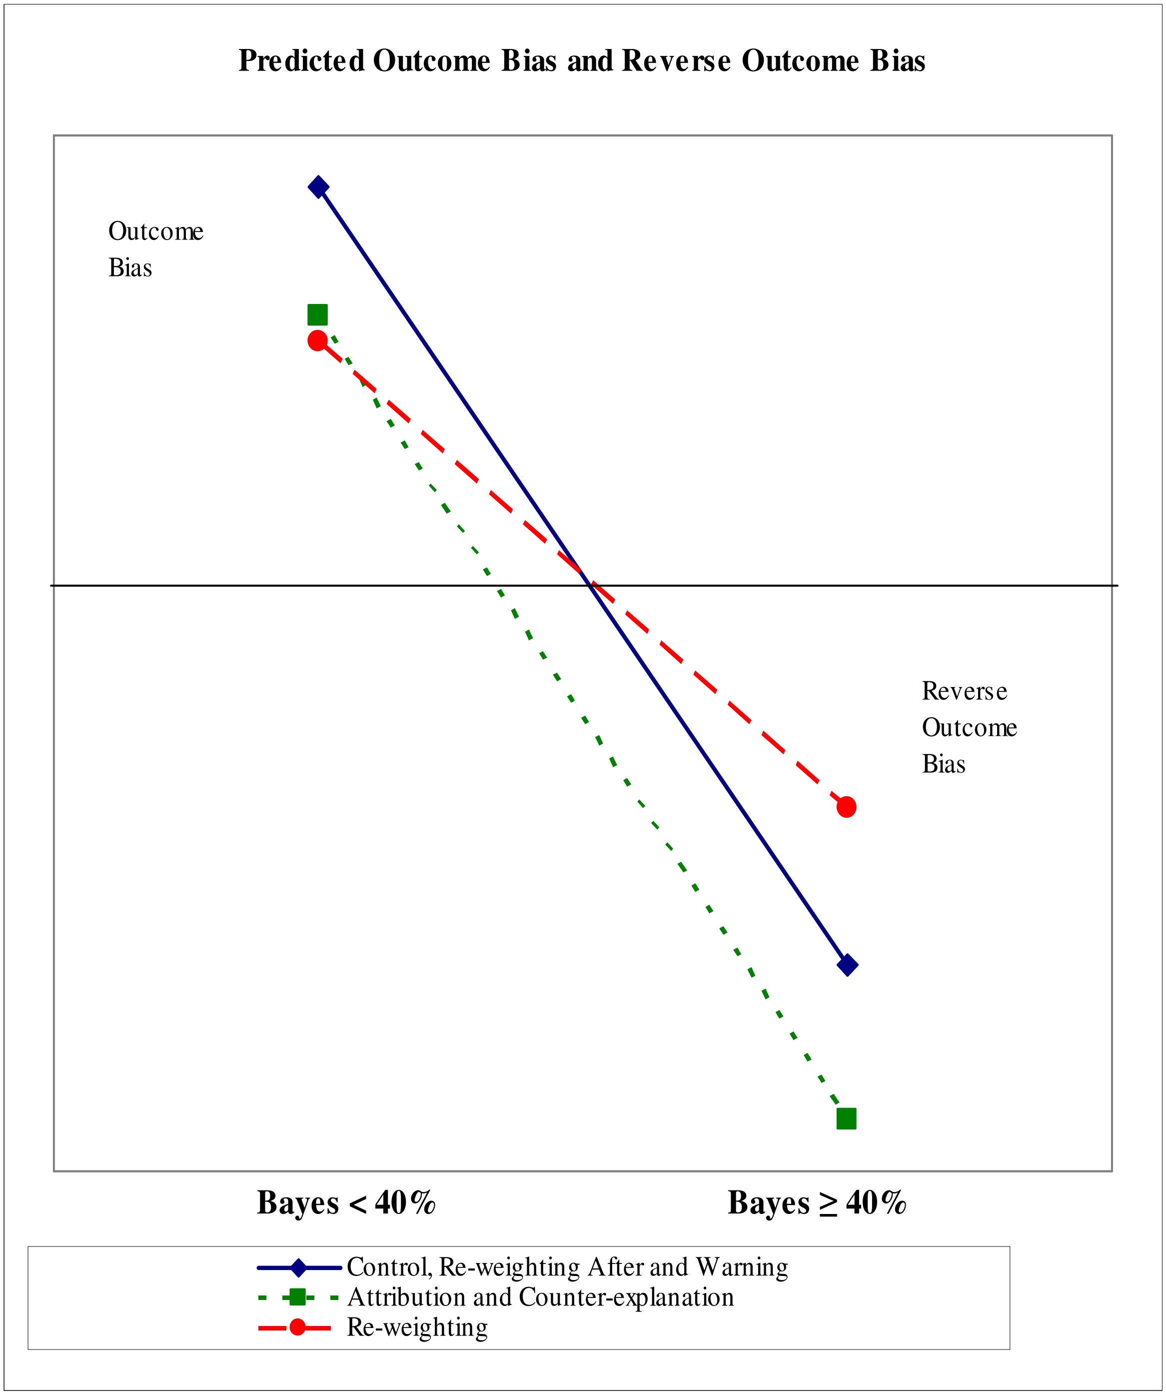

Figure 2
Graphical Depiction of Participants' Judgments of Auditor Negligence

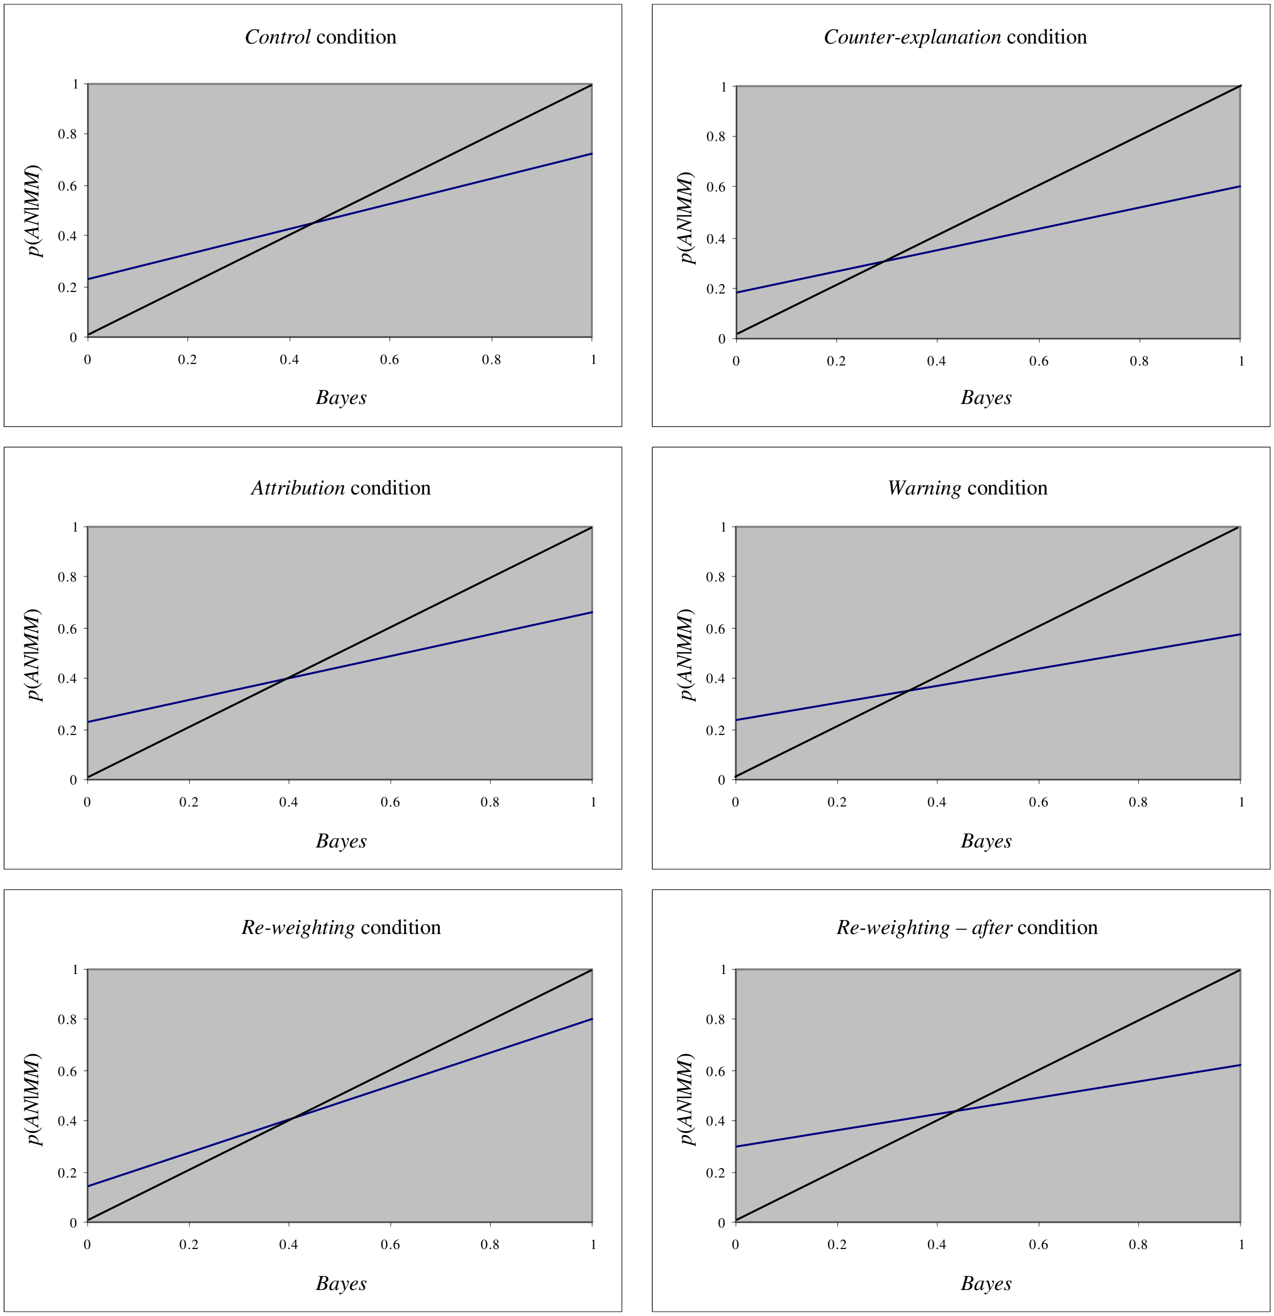

These graphs compare participants' judged probabilities of auditor negligence (vertical-axis) against the Bayesian probability of auditor negligence (horizontal-axis). The 45-degree line emanating at the origin is a benchmark line that would obtain if participants were perfect Bayesians in judging auditor negligence. The other lines are fitted by-condition regression lines (See Table 2 for the regression coefficients). Regression lines above the benchmark line suggest overly harsh outcome bias, whereas regression lines below the benchmark line suggest overly lenient reverse outcome bias. Inspection of the graphs suggests the least outcome bias and reverse outcome bias occur in the *re-weighting* condition and that several interventions increased reverse outcome bias relative to *control* condition.

Figure 3
Outcome Bias and Reverse Outcome Bias Findings

Panel A: Collapsed Across Audit Quality

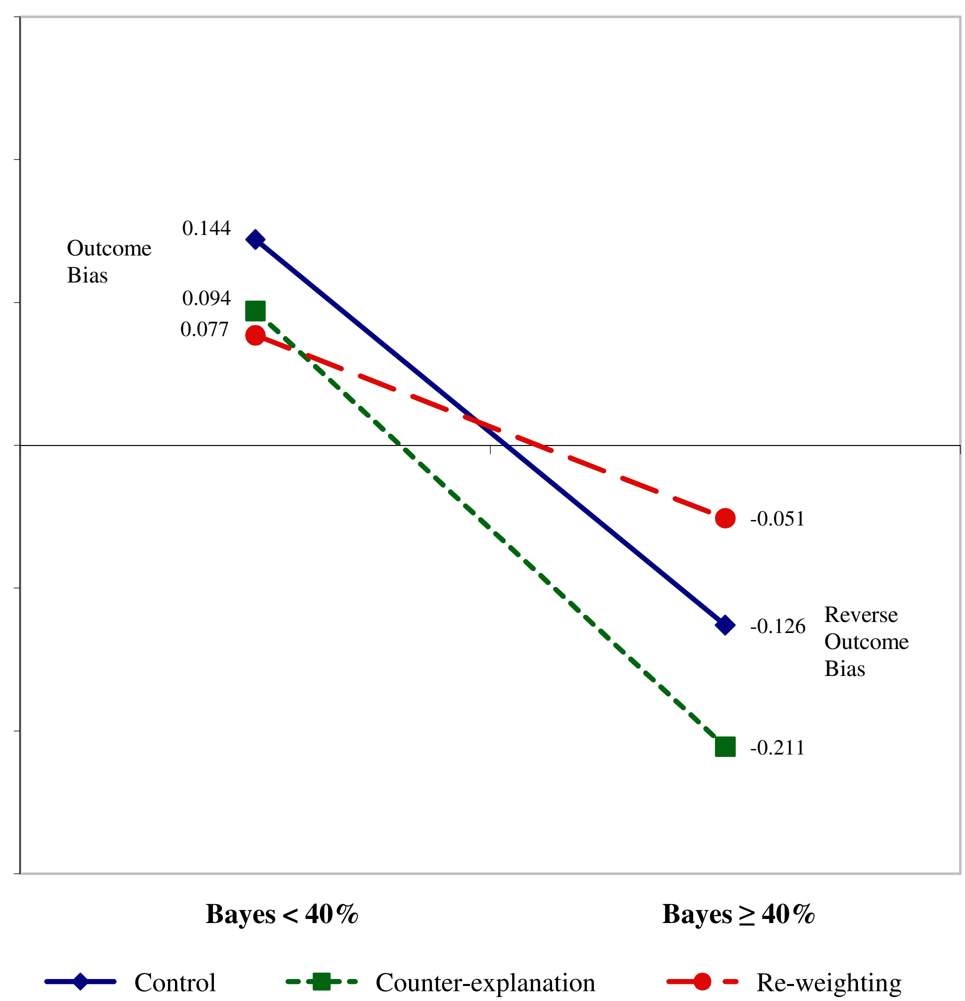

Figure 3 (continued)

Panel B: Separated by Audit Quality

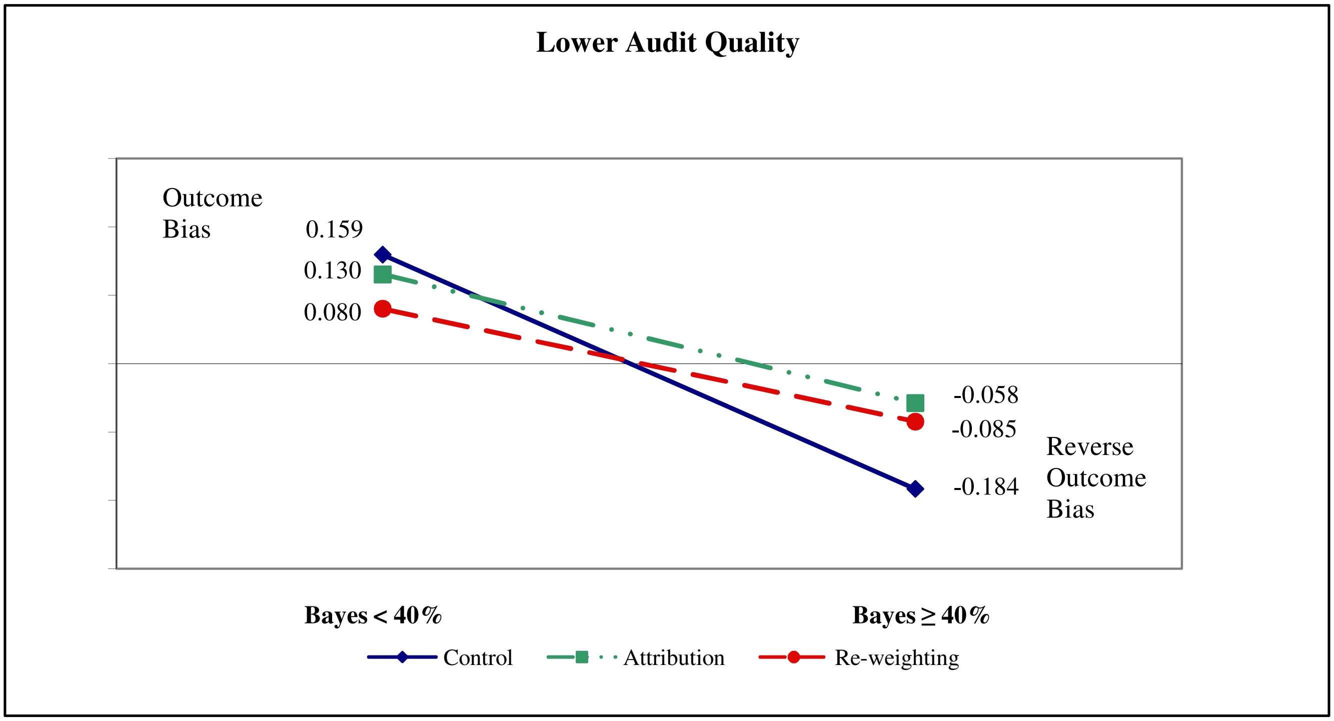

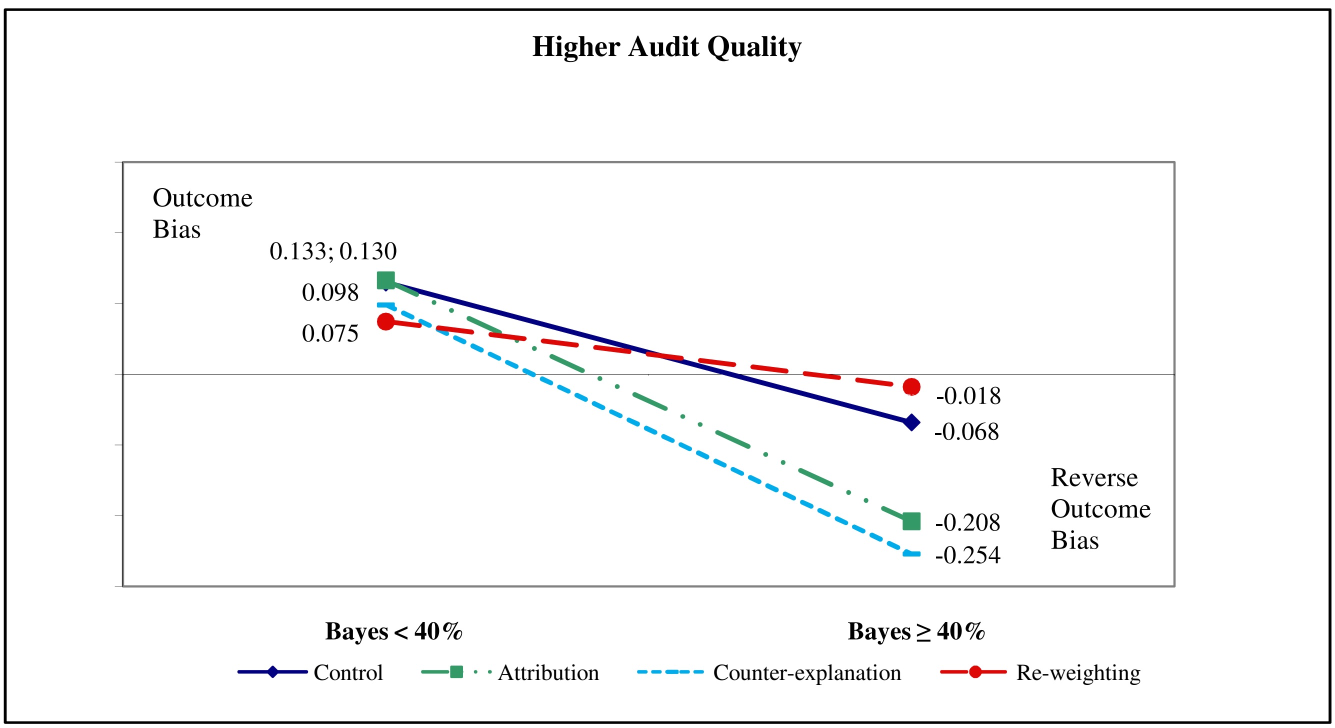

These graphs show results for those conditions that differ significantly from the *control* condition.

Table 1
Outcome Bias (+), Reverse Outcome Bias (-), and Total Bias
Bias = *p*(*AN*|*MM*) – *Bayes* = *f*(*Intervention, Audit Quality,* and *Hilo*40)

Panel A: Descriptive Statistics: Least-Square Means, (Standard Errors), [n]

	Overall			Lower *Audit Quality*			Higher *Audit Quality*		
	***Bayes* < 40%**	***Bayes* ≥ 40%**	**Total Bias**	***Bayes* < 40%**	***Bayes* ≥ 40%**	**Total Bias**	***Bayes* < 40%**	***Bayes* ≥ 40%**	**Total Bias**
Control	0.144***	-0.126***	0.270	0.159***	-0.184***	0.342	0.130***	-0.068*	0.198
	(0.020)	(0.031)	(0.019)	(0.028)	(0.043)	(0.026)	(0.029)	(0.045)	(0.027)
	[191]	[81]	[272]	[101]	[42]	[143]	[90]	[39]	[129]
Attribution	0.132***	-0.133***	0.265	0.130***	-0.058*	0.188	0.133***	-0.208***	0.341
	(0.021)	(0.031)	(0.019)	(0.031)	(0.038)	(0.025)	(0.028)	(0.049)	(0.028)
	[181]	[85]	[266]	[79]	[53]	[132]	[102]	[32]	[134]
Counter-explanation	0.094***	-0.211***	0.305	0.090***	-0.169***	0.259	0.098***	-0.254***	0.352
	(0.021)	(0.030)	(0.018)	(0.031)	(0.038)	(0.027)	(0.028)	(0.047)	(0.024)
	[176]	[91]	[267]	[80]	[55]	[135]	[96]	[36]	[132]
Re-weighting	0.077***	-0.051**	0.128	0.080***	-0.085***	0.165	0.075***	-0.018	0.093
	(0.021)	(0.028)	(0.018)	(0.031)	(0.037)	(0.026)	(0.030)	(0.042)	(0.024)
	[169]	[103]	[272]	[83]	[58]	[141]	[86]	[45]	[131]
Re-weighting – After	0.176***	-0.125***	0.301	0.202***	-0.123***	0.325	0.150***	-0.126***	0.276
	(0.021)	(0.031)	(0.019)	(0.029)	(0.045)	(0.026)	(0.030)	(0.045)	(0.027)
	[182]	[81]	[263]	[94]	[39]	[133]	[88]	[42]	[130]
Warning	0.135***	-0.165***	0.300	0.138***	-0.205***	0.343	0.132***	-0.126***	0.258
	(0.020)	(0.031)	(0.018)	(0.029)	(0.042)	(0.027)	(0.029)	(0.045)	(0.025)
	[187]	[83]	[270]	[94]	[44]	[138]	[93]	[39]	[132]
Total	0.126***	-0.135***	0.261	0.133***	-0.137***	0.269	0.120***	-0.133***	0.253
	(0.008)	(0.012)	(0.008)	(0.012)	(0.017)	(0.010)	(0.012)	(0.018)	(0.011)
	[1,086]	[524]	[1,610]	[531]	[291]	[822]	[555]	[233]	[788]

*, **, *** Significantly different from zero at the 0.10, 0.05, and 0.01 levels. All t-tests are one-tailed.

Table 1 (continued)

Panel B: Analysis of Variance

Source	SS	df	MS	F	p
Intervention	1.03	5	0.21	2.65	0.022
Audit Quality	0.01	1	0.01	0.09	0.760
*Hilo*40	23.68	1	23.68	304.25	0.000
Intervention × *Audit Quality*	0.67	5	0.13	1.73	0.124
Intervention × *Hilo*40	1.41	5	0.28	3.61	0.003
Audit Quality × *Hilo*40	0.03	1	0.03	0.35	0.555
Intervention × *Audit Quality* x *Hilo*40	0.93	5	0.19	2.39	0.036
Error	123.42	1586	0.08		

Panel C: Planned Comparisons of Interventions Relative to the Control Condition

	***Bayes* < 40%** **- Reduction in outcome bias** **+ Increase in outcome bias**			***Bayes* ≥ 40%** **- Reduction in reverse outcome bias** **+ Increase in reverse outcome bias**		
	Contrast	*F*	*p*	*Contrast*	*F*	*p*
Control vs. *Attribution*[1]	-0.012	0.19	0.330	-0.007	0.03	0.483
Control vs. *Counter-explanation*[1]	-0.050	2.98	0.041	-0.085	3.95	0.024
Control vs. *Warning*[3]	-0.009	0.11	0.740	-0.039	0.84	0.358
Control vs. *Re-weighting*[1]	-0.067	5.19	0.011	0.075	3.18	0.037
Control vs. *Re-weighting – after*[2]	0.032	1.20	0.273	0.001	0.00	0.984
Re-weighting vs. *Attribution*[2]	0.055	3.32	0.070	-0.082	3.81	0.051
Re-weighting vs. *Counter-explanation*[2]	0.017	0.31	0.579	-0.160	15.37	<0.001
Re-weighting vs. *Warning*[1]	0.058	3.79	0.026	-0.114	7.60	0.003
Re-weighting vs. *Re-weighting – after*[1]	0.099	11.04	0.001	-0.074	3.10	0.039

Post Hoc Comparisons at Lower and Higher *Audit Quality*

	***Bayes* < 40%** **- Reduction in outcome bias** **+ Increase in outcome bias**			***Bayes* ≥ 40%** **- Reduction in reverse outcome bias** **+ Increase in reverse outcome bias**		
	Contrast	*F*	*p*	*Contrast*	*F*	*p*
Control vs. *Attribution* at lower *Audit Quality*[2]	-0.029	0.45	0.502	0.126	4.75	0.029
Control vs. *Attribution* at higher *Audit Quality*[2]	0.003	0.00	0.949	-0.140	4.45	0.035

[1] One-tailed, as findings are consistent with predictions.
[2] Two-tailed, as at least one finding is not consistent with predictions.
[3] Two-tailed, as no directional prediction is made.

Table 2
Linear Regressions

Panel A: Linear Regressions of *p*(*AN*|*MM*) on *Bayes* and Intervention Condition

Dependent Variable**:** *p*(*AN*|*MM*) is the participants' observed posterior belief of auditor negligence.

Independent Variables [a, b]	**Regression Coefficient**	**t**	**p**
Intercept	0.225	9.24	<0.001
Attribution	0.006	0.17	0.866
Counter-explanation	-0.046	-1.31	0.185
Warning	0.015	0.44	0.668
Re-weighting	-0.082	-2.27	0.023
Re-weighting – after	0.073	2.08	0.038
Control × *Bayes*	0.502	8.90	<0.001
Attribution × *Bayes*	0.431	6.93	<0.001
Counter-explanation × *Bayes*	0.427	7.33	<0.001
Warning × *Bayes*	0.334	5.40	<0.001
Re-weighting × *Bayes*	0.659	10.59	<0.001
Re-weighting – after × *Bayes*	0.323	5.21	<0.001
Model F	33.35		<0.001
R^2	18.7%		
R^2_{adj}	18.1%		
N	1610		

[a] *Bayes* = participant's Bayesian posterior of auditor negligence given material misstatement; see equation (3) for calculation. *Control* = 1 if participant was in *control* condition, 0 otherwise. *Attribution* = 1 if participant was in *attribution* condition, 0 otherwise. *Counter-explanation* = 1 if participant was in *counter-explanation* condition, 0 otherwise. *Warning* = 1 if participant was in *warning* condition, 0 otherwise. *Re-weighting* = 1 if participant was in *re-weighting* condition, 0 otherwise. *Re-weighting – after* = 1 if participant was in *re-weighting – after* condition, 0 otherwise.

[b] We also ran a model that includes the *Audit Quality* manipulation and interactions of *Audit Quality* with all of the independent variables. This more complex model does not improve the explanatory power and does not change any of our inferences, with one exception. The exception stems from a significant *Intervention* × *Bayes* x *Audit Quality* interaction. The *attribution* intervention deteriorates Bayesian reasoning compared to the *control* condition, but only when *Audit Quality* is relatively high.

Table 2 (continued)

Panel B: Comparison of Coefficients

Coefficient Comparisons	**Pred. Sign**	**Contrast**	**F**	**p**
Control × *Bayes* = *Attribution* × *Bayes*	-	-0.071	0.73	0.394
Control × *Bayes* = *Counter-explanation* × *Bayes*	-	-0.075	0.89	0.347
Control × *Bayes* = *Re-weighting* × *Bayes*	+	0.157	3.48	0.062
Control × *Bayes* = *Re-weighting – after* × *Bayes*	?	-0.179	4.59	0.032
Control × *Bayes* = *Warning* × *Bayes*	-	-0.167	4.02	0.045
Re-weighting × *Bayes* = *Attribution* × *Bayes*	-	-0.228	6.74	0.010
Re-weighting × *Bayes* = *Counter-explanation* × *Bayes*	-	-0.232	7.48	0.006
Re-weighting × *Bayes* = *Re-weighting – after* × *Bayes*	-	-0.336	14.67	<0.001
Re-weighting × *Bayes* = *Warning* × *Bayes*	-	-0.325	13.68	<0.001

Comparisons of the coefficients on the *Intervention* × *Bayes* term represent the average difference in the relationship between *Bayes* and the observed posterior between intervention conditions. The contreasts equal the difference between the coefficients on the *Intervention* × *Bayes* terms.

Panel C: Explanatory Power of Regression Models by Intervention Condition

$p(AN|MM) = \beta_0 + \beta_1 \, Bayes + \varepsilon$	R^2_{adj}	**F**	**n**
Re-weighting	32.6%	131.88	272
Counter-explanation	18.8%	62.55	267
Control	18.5%	62.50	272
Attribution	13.1%	40.99	266
Re-weighting – after	11.6%	35.53	263
Warning	8.8%	26.97	270

Appendix 1: Intervention Manipulations

Panel A: Counter-Explanation Intervention

14. While the SEC investigation concluded that material misstatement due to inventory overstatement did occur, this may not necessarily mean auditor negligence. Listed below are two other feasible explanations about how the SEC might have reached this conclusion without the auditors being negligent.

Please examine each and indicate on the scales below how likely you believe it is that the audit of Big Time Gravel could have included the following explanations:

Explanation A: The SEC said that the alleged overstatement occurred at one inventory location. This location was not selected by the auditor. However, the auditor had no reason to believe that this location was any more likely to contain material misstatements than the locations selected. As such, the auditor exercised due professional care by selecting a representative sample of locations.

This explanation is **extremely <u>un</u>likely.**	This explanation is **extremely likely.**

Explanation B: At the time of the audit, the market price of gravel had dropped significantly. This price decline was reasonably judged as temporary by both the auditors and management. In accordance with professional standards, Jones & Company did not require management to write down the inventory balance as the decline was deemed temporary. However, after the completion of the audit, news surfaced that the market price decline was, in fact, permanent. Thus, with the benefit of 20/20 hindsight, the SEC concluded that the inventory balance was materially overstated.

This explanation is **extremely <u>un</u>likely.**	This explanation is **extremely likely.**

Panel B: Attribution Intervention

IMPORTANT! You are about to evaluate an auditor's behavior performed in connection with an audit of a gravel company. Evaluating auditors' behavior is a serious responsibility. Based on your evaluations, professional reputations and large amounts of money are at stake. As a result, evaluators have to make very difficult decisions. Real evaluators feel **anxious and tense** when faced with making these decisions. Your goal is to get completely into the evaluator role as possible. To the extent that you are successful in taking on this role, you may feel at least a little tense, anxious or otherwise uncomfortable. Please try hard to take on this role.

Case Questions

7. Are you able to detect any feelings of tension, anxiety, or discomfort?

- ❑ Yes, quite a lot.
- ❑ Yes, some.
- ❑ Yes, at least a small amount.
- ❑ No, none at all.

8. To the extent that you are able to detect at least a small amount of tension, anxiety, or discomfort, could those feelings arise partly from your role as evaluator?

- ❑ Yes, definitely.
- ❑ Yes, probably.
- ❑ Yes, possibly.
- ❑ No, definitely not.

Panel C: Warning Intervention

IMPORTANT! Next we will ask you to judge the probability that Jones and Company was negligent. **Outcome bias** is said to exist when individuals overestimate the extent to which an outcome could have been anticipated prior to its occurrence. This bias occurs when people unknowingly over-rely on outcome information in evaluating decisions made by others. In this case, outcome bias would exist if an evaluator over-relied on conclusions of the SEC investigation in evaluating auditor negligence.

Imagine that you are asked to evaluate the work of the auditors in this case. Please beware of outcome bias and do not over-rely on outcome information.

14. In the current case, outcome bias would exist if I:

- ❑ Over-relied on outcome information when evaluating the auditors.
- ❑ Under-relied on outcome information when evaluating the auditors.

Panel D: Re-weighting Interventions

Both *Re-weighting* and *Re-weighting – after* Interventions

Researchers have found something interesting about the way people judge auditor negligence after they find out about a bad outcome (e.g., a lawsuit accusing an auditor of negligence). When their own beliefs about various likelihoods of material misstatement and of auditor negligence mathematically indicate that the chance of auditor negligence is relatively high (more than a 40% chance), people tend to **under-react to the bad outcome** and conclude the chance of auditor negligence **to be too low**. On the flip side, when their own beliefs about various likelihoods of material misstatements and of auditor negligence indicate that the chance of auditor negligence is relatively low (less than a 40% chance), people tend to **over-react to the bad-outcome** and conclude the chance of auditor negligence **to be too high**. The 2 x 2 box below may help you make better judgments of the probability of auditor negligence for this case.

If you conclude that:	Then consider:
It is relatively <u>likely</u> that the auditor is negligent (i.e., above 40%)	You may well be too lenient on the auditor. <u>Consider</u> raising your auditor negligence probability judgment.
It is relatively <u>unlikely</u> that the auditor is negligent (i.e., below 40%)	You may be too harsh on the auditor. <u>Consider</u> lowering your auditor negligence probability judgment

(7.) 15. If I conclude that it is relatively <u>unlikely</u> that the auditor is negligent, I should consider:

- ❑ That I may well be too lenient on the auditor and consider raising my auditor negligence probability judgment.
- ❑ That I may well be too harsh on the auditor and consider lowering my auditor negligence probability judgment.

***Re-weighting – after* Intervention Only**

16. Now that you are aware that individuals' judgments exhibit this pattern, you have the opportunity to reconsider your judgment in #14. It is important to remember that this pattern occurs, **on average** across many individuals, but your probability judgment may not exhibit this pattern and reflect your true beliefs. However, to the extent that you feel that this pattern applies to your probability judgment in #14, please revise your probability judgment accordingly. Please realize that not revising your judgment is perfectly acceptable.

<u>Initial Judgment from #14</u>:

A: I believe that the auditors would have been **negligent** on ______ out of every 100 audits like this one.

<u>Revised Judgment</u>:

A: I believe that the auditors would have been **negligent** on ______ out of every 100 audits like this one.

Appendix 2: Manipulation Check Questions

Panel A: *Audit Quality* Manipulation Check Questions

7. To how many inventory sites did Jones & Company send its auditors?

- ❑ 6 out of 11
- ❑ 5 out of 11 (higher *Audit Quality* condition)
- ❑ 4 out of 11 (lower *Audit Quality* condition)
- ❑ 3 out of 11

8. How did Jones & Company inform Big Time Gravel about which inventory sites it would be visiting:

- ❑ Completely unannounced, surprise basis
- ❑ Two weeks in advance (higher *Audit Quality* condition)
- ❑ One month in advance (lower *Audit Quality* condition)
- ❑ Two months in advance

9. Which auditors conducted the inventory observation?

- ❑ Two audit seniors, each with considerable industry experience
- ❑ An audit senior with considerable industry experience and a second year staff with some industry experience (higher *Audit Quality* condition)
- ❑ Two second year staff auditors, each with some industry experience (lower *Audit Quality* condition)
- ❑ A second year staff with some industry experience and an intern with no industry experience

Panel B: *Intervention* Manipulation Check Question

9. Which of the following statements is true? (Check only one.)

- ☐ The case encouraged me to be aware that I might feel anxious or tense. It asked me to answer a question about the extent to which I felt anxious or tense.
- ☐ The case included a 2 x 2 table to discuss a research finding that people tend to be too lenient on auditors if their judged probability of auditor negligence is relatively high (i.e., above 40%) but too harsh on auditors if their judged probability of auditor negligence is relatively low (i.e., below 40%).
- ☐ The case discussed the concept of outcome bias as being bias that occurs when people overestimate the extent to which an outcome could have been anticipated prior to its occurrence. It advised me to beware of outcome bias and to not to over-rely on outcome information.
- ☐ The case provided and asked me to rate the likelihood of two alternative explanations under which a material misstatement could have occurred without Jones & Company being negligent. One of these alternative explanations had to do with a temporary versus permanent decline in market value and the other had to do with the misstatement being isolated to one audit location.
- ☐ None of the above happened in the case.