

Understanding Outcome Bias ^{*}

Andy Brownback[†]
University of Arkansas

Michael A. Kuhn[‡]
University of Oregon

January 31, 2019

Abstract

Disentangling effort and luck is critical when judging performance. In a principal-agent experiment, we demonstrate that principals' judgments of agent effort are biased by luck, despite *perfectly observing the agent's effort*. This bias can erode the power of incentives to stimulate effort when there is noise in the outcome process. As such, we explore two potential solutions to this "outcome bias" – control over irrelevant information about luck, and outsourcing judgment to independent third parties. We find that both are ineffective. When principals have control over information about luck, they do not avoid the information that biases their judgment. In contrast, when agents have control over the principal's information about luck, agents strategically manipulate principals' outcome bias to minimize their punishments. We also find that independent third parties are just as biased as principals. These findings indicate that outcome bias may be more widespread and pernicious than previously understood. They also suggest that outcome bias cannot be driven solely by disappointment nor by distributional preferences. Instead, we hypothesize that luck directly affects beliefs about agents even though effort is perfectly observed. We elicit the beliefs of third parties and principals and find that lucky agents are believed to exert more effort in general than identical, unlucky agents.

JEL Classification: C92, D63, D83

Keywords: Experiment, Reciprocity, Outcome Bias, Attribution Bias, Blame

^{*}We wish to thank Gary Charness, Ryan Oprea, Peter Kuhn, Joshua Miller, Justin Rao, and James Andreoni for valuable feedback.

[†]University of Arkansas, Department of Economics, 220 N. McIlroy Ave, Fayetteville, AR 72701, USA. E-mail: abrownback@walton.uark.edu. Web page: directory.uark.edu/people/apbrownb

[‡]Corresponding author, University of Oregon, Department of Economics, 1285 University of Oregon, Eugene, OR 97403 USA. E-mail: mkuhn@uoregon.edu. Web page: pages.uoregon.edu/mkuhn

1 Introduction

Accurately evaluating effort is essential in any environment with delegated responsibilities where actions are not fully observable. The complexity of using a noisy output—such as profit—to form beliefs about inputs—such as effort or “type”—allows room for systematic stereotypes and biases to impede judgment. For example, Sarsons (2017) finds that primary-care doctors update their beliefs about the skill of specialists in response to a particular outcome differently depending on the specialist’s gender. Our study adds to a literature identifying a consistent inability to ignore luck in favor of strong—even perfect—signals of effort.¹ This empirical phenomenon is referred to as “outcome bias.” Outcome bias is a fundamental misattribution that blames a knowably random event on the actions of an involved party. Other forms of misattribution have been shown, for example, to influence perceptions of taste (Haggag and Pope, 2016) and the supply of labor (Bushong and Gagnon-Bartsch, 2016). In this paper, we explore how far outcome bias reaches, whether sophisticated individuals can avoid or exploit it, and, ultimately, the mechanisms driving this bias.

In many contexts, outcomes are a useful signal of the effort of the actor. A failed project should increase posterior beliefs that a worker exerted low effort and should be punished. Conditioning monetary transfers on outcomes may also serve as a form of risk-sharing when effort is unobservable. Our study abstracts away from any practical way in which outcomes can be instrumental in decision-making. We explore outcome bias in punishment decisions when outcomes possess no signal value. Moreover, our setting involves costly punishment, eliminating the possibility of risk-sharing.

We report the results of experimental treatments that vary the information structure of a principal-agent style interaction. The interaction is designed to make the lack of signal value from outcomes as transparent as possible. Agents exert effort by making costly investments to improve the principal’s chance of winning a prize. Conditional on effort, a random process determines if the principal wins or loses. We refer to this conditional randomness as “luck.” Principals can respond by paying to punish agents. In all treatments, effort is perfectly observed by the principal prior to their punishment decision. However, depending on treatment assignment, principals will observe their luck (and the resulting outcome) either before or after their punishment decision must be made. Since effort is perfectly observed, information about winning or losing is non-diagnostic and carries no added signal value for effort.

Three questions motivated our work. First, how robust and harmful is outcome bias in a transparent

¹See Section 2.1 for a review of this literature.

information environment? Second, are there simple ways of adjusting the environment to eliminate outcome bias? Third, given our answers to 1) and 2), can we pinpoint the behavioral mechanism underlying outcome bias? We derive five key findings:

First, our baseline, Full-information (“*Full*,” hereafter) treatment shows that outcome bias remains strong even in the most transparent of environments. In the *Full* treatment, the principal observes both effort and luck prior to their punishment decision. We find that principals’ punishments are influenced by luck *even though effort is perfectly observable*. Conditional on agent effort, principals punish bad luck the same as they would a 1.4 standard deviation (SD) decrease in effort. Troublingly, this bias crowds out accurate punishment of agents. In the *Hidden* treatment, principals only observe effort prior to the punishment decision. Principals’ punishments are twice as responsive to effort when they cannot observe luck.

Second, we find that control of information by principals is not an effective solution to outcome bias because they fail to avoid information about luck. We implement the “*Commit*” treatment that allows principals to commit against allowing luck to influence their punishment decision by deciding to keep it hidden until after their decision. Principals fail to take advantage of this commitment opportunity. Despite already having perfect information on effort, 30% percent of principals are willing to pay to view their luck prior to making their punishment decision. Nearly 90% will view their luck if it is costless.

Third, in contrast to principals, when agents are given control of information, they leverage it to manipulate outcome bias in principals. We implement a “*Force*” treatment that allows agents to force principals to see their luck prior to the punishment decision. In the *Force* treatment, we find that agents predict and exploit outcome bias better than principals commit against it. Agents manipulate the principal’s bias by strategically revealing information about luck prior to the principal’s punishment choice.² They are nearly twice as likely to reveal good luck than bad luck to the principal and 64% of agents are willing to pay to reveal this information. Agents successfully control information about luck to reduce their punishment. This is intuitive behavior outside of the lab where the impulse to “shoot the messenger” is well understood.

Fourth, using independent third parties does not eliminate outcome bias. We find that uninvolved third party punishers, *who must sacrifice money from their endowment to punish agents*, are also biased by luck. In the third party full-information treatment (“*Full3*,” hereafter), we present third party punishers with complete information about agent effort and luck prior to their punishment decisions.³ Conditional

²Principals will learn about their luck at the end of each round regardless of the agent’s choice.

³We also conduct third party Force and third party Commit treatments (“*Force3*” and “*Commit3*,” hereafter) to test for

on the agent’s effort, third parties punish bad luck the same as they would a 1.4 SD decrease in effort—the same magnitude as with principal punishers.

Finally, we discover that one of the underlying behavioral mechanisms behind outcome bias is that it influences beliefs about agents—lucky agents are thought to be more hardworking than unlucky ones. In the third party “*Guess3*” treatment, we show third parties an agent’s effort and luck from one principal-agent interaction. We then use an incentive-compatible mechanism to elicit beliefs from third parties about the mean effort level of the agent across all of the agent’s interactions, past and future, only one of which is seen by the guesser. We refer to this mean effort as the agent’s “type.” Luck substantially affects the third party guesses about the average effort an agent exerts across all interactions. Bad luck causes third parties to decrease their beliefs about an agent’s type by the same amount as if they observed a 0.25 SD decrease in effort. This finding demonstrates that lucky people are considered to be better intentioned on average than their unlucky counterparts. Thus, while outcome bias itself can potentially be explained by a variety of pathways—disappointment, distributional preferences, mistaken belief updating, to name a few—our set of findings taken together points in the direction of biased beliefs.

This last finding in particular puts additional structure on previous work suggesting mechanisms for outcome bias. Gurdal et al. (2013) outline a model of delegated expertise (Demski and Sappington, 1987) to explain outcome bias. Applied to our study, this model suggests that principals, unfamiliar with an environment in which effort is observable, assume a “salient perturbation” (Myerson, 1991) of their current environment to make their punishment choice. In other words, principals respond to information as if they are operating in a more familiar, hidden-action environment when deciding on punishments. Unobservable effort is presumed to influence luck in the salient perturbation of the interaction because, “In most environments, others do have an influence over outcomes,” (Gurdal et al. (2013), p.1216). Our results are consistent with this model of behavior, but also extend it beyond the concurrent interaction. We show that luck is believed to reveal fixed aspects of the agent’s character.

Relatedly, outcome bias may be a manifestation of attribution bias, which arises when utility derived from one phenomenon is incorrectly attributed to a concurrent phenomenon (Schwarz and Clore (1983); Haggag and Pope (2016); Bushong and Gagnon-Bartsch (2016)).⁴ For example, Weber et al. (2001) find

anticipation of outcome bias amongst third parties, as we did with principals. Similar to *Force*, agents seek to manipulate the outcome bias of third parties in *Force3*. Agents are nearly 75% more likely to reveal winning outcomes to third parties than losing outcomes and 52% are willing to pay to reveal this information. Third parties also show little demand for commitment against information about luck in *Commit3*. 55% of third parties view luck prior to making punishment decisions and 22% are willing to pay for this information.

⁴This is not unlike “projection bias,” where an agent misattributes state-specific utility from this time period onto estimates of utility to be gained in future states of the world (Loewenstein et al. (2003)).

that when leaders are (randomly) assigned more difficult tasks, their group members misattribute worse outcomes to poor leadership. In our study, a principal may misattribute the benefit from a lucky outcome to the agent’s good type and withhold punishment.

Finally, principals may be *attempting* to punish agents based on their efforts, but misinterpreting signals of that effort because of information processing failures such as hindsight bias or correlation neglect. Fischhoff (1975) suggests that a principal’s perception of an agent’s effort may be wrongly colored by the newly arrived information about whether or not that effort was sufficient. That is, agents may be punished as if they knew the necessary effort *ex ante*. Relatedly, when principals observe both outcomes and effort, they should realize that, conditional on effort, outcomes should be disregarded as uncontrollable randomness. However, Enke and Zimmermann (2017) show that individuals tend to treat correlated signals as if they were independent. In our study, an extreme manifestation of correlation neglect would imply that principals may fail to disregard the outcome as a redundant signal of effort.

A potential explanation for a subset of our results is that principals hold distributional preferences and punish agents in order to improve their relative wealth (e.g. Bolton (1991); Fehr and Schmidt (1999); Bolton and Ockenfels (2000); Charness and Rabin (2002)). While this can partially explain the behavior of principals, it is rejected by three aspects of our data. First, we find that winning (and thus richer) principals sometimes punish, indicating a motivation to respond to effort rather than distributional outcomes alone. Second, third parties earn less than half as much money as agents regardless of effort or luck. Nonetheless, third parties condition their punishment on both effort and luck. To be justified by distributional preferences alone, a third party would have to be willing to decrease the payoffs of the poorest player (themselves) in order to reduce differences between the richest two players. Third, distributional preferences offer no explanation for why beliefs about the agent’s type would vary with luck. Beyond our work, Charness and Rabin (2002) find that costly, Pareto-damaging difference aversion is rare. Punishment in our study is costly and Pareto-damaging, yet it is common. Thus, as with Charness and Rabin (2002), we focus on reciprocity issues to explain our data. Moreover, in a closely related setting, Gurdal et al. (2013) explicitly measure the role of distributional preferences on punishment choices and rule it out as a primary motivation of behavior. We discuss this issue further in Section 4.3.⁵

While the applications of these findings are widespread—we discuss the legal context in particular later on—our motivation is rooted in personnel economics. Hölmstrom (1979) shows that, if effort is

⁵Even redistributive motives are not always independent of behavior. For example, Lefgren et al. (2016) find that voters prefer redistributive policies that favor unlucky, hardworking subjects over equally poor, lazy subjects.

observable, contracting based on outcomes leads to sub-optimal risk sharing and weakens incentives for the agent to exert high effort. Thus, when principals rely too heavily on outcomes, they are foregoing the first-best solution and may decrease welfare from principal-agent interactions that benefit from collectively enforced norms of behavior.

Strategic manipulation of outcome bias has clear practical implications. CEOs may seek control over the timing of financial reports in order to update their stakeholders only when the market conditions are favorable—potentially saving their job despite poor effort. Politicians (and their opponents) may seek to selectively publicize narrow dimensions of the macroeconomy in order to manipulate voters’ biases, thus polarizing information sources. Meanwhile, shareholders or voters may fail to shield themselves from this information believing that it will not influence their perception of effort.⁶

Our third party results indicate that independent arbiters are just as susceptible to outcome bias as affected parties. This is particularly important when considering the implications of our findings on the legal system. The majority of states in the U.S. enforce some version of a legal precedent known as “The Collateral Source Rule.” This rule bars jurors from knowing whether the plaintiff or victim in a case has already been compensated for their loss. Many criminal law statutes present a seemingly contradictory view. Mandatory minimum sentences for attempted or “inchoate” crimes are often statutorily shorter than those for successful crimes.⁷ These inconsistencies in the legal treatment of luck—irrespective of intentions—point to uncertainty in both the normative question of how people *should* respond to luck and the positive question of how people *will* respond when presented with information about it.

2 Literature

There is a considerable experimental literature on outcome bias. Charness and Levine (2007), Cushman et al. (2009), Gurdal et al. (2013), Rubin and Sheremeta (2015) and de Oliveira et al. (2017) all document outcome bias in a variety of real-stakes settings among interested parties. When considering the morality of an action, Gino et al. (2009) demonstrates that people take account of both the intention of the decision-maker and the resulting outcome. Researchers have consistently found that outcomes have a

⁶This asymmetry relates to work from psychology and behavioral economics on “bias blind spots” (Pronin et al. (2002); Ehrlinger et al. (2005); Pronin (2007)). This research often finds that, though people can identify biases in others, they are often unaware of bias in their own decision-making. We are not the first to find a strategic manipulation of others’ biases for personal gain (see, for example, Gneezy and Imas (2014) and Bartling and Fischbacher (2011)). However, we are the first to show a manipulation of outcome bias through information revelation.

⁷For example, the minimum sentence for attempted aggravated murder in Oregon is 10 years (HB 3439), while the minimum sentence for murder is 25 years (Measure 11). In Arkansas, crimes are given rankings based on the seriousness of the offense. Murder is ranked 10 (most serious), while attempted murder is ranked 9.

marked impact on the interpretation of events even when intentions are the primary concern, (Falk et al. (2008); Sutter (2007); Nelson (2002); Landmann and Hess (2016)). But, there are some conflicting results. For example, McCabe et al. (2003) find that intentions alone matter.

The conflict between intentions and luck is particularly stark when considering their impact on reciprocity (Charness and Rabin (2002); Falk and Fischbacher (2006)). Blount (1995) provides an early test of this, showing that outcomes influence reciprocal behavior, but they decrease in importance when the agent’s intentions can be compared to the intentions of others. Charness and Levine (2007), Cushman et al. (2009), Gurdal et al. (2013), and de Oliveira et al. (2017) all study environments with the possibility of costly reciprocity and find strong effects of the outcome on reciprocal behavior. Outcomes are not all that matters, however. Charness (2004) finds that reciprocity does depend critically on the intentions of the first actor and Gurdal et al. (2013) find that the effects of outcomes predictably weaken as the costs of reciprocity increase. Halevy et al. (2009) disentangles the impact of outcomes and effort on reciprocity to show that behavior is most consistent with a model of interdependent preferences—that is, that the desire to improve another person’s outcomes increases as their effort increases. In a principal-agent setting similar to ours, Rubin and Sheremeta (2015) find that any noise weakening the correlation between effort and outcomes lowers reciprocity. Our study expands upon this research area by exploring 1) the sophistication about the impact of random outcomes even when they carry no unique information about effort and 2) the psychological foundations for the influence of these random outcomes.⁸

Several papers explore the role of luck in affecting ex-ante and ex-post social or redistributive preferences (for example, Coate, 1995; Cappelen et al., 2013; Brock et al., 2013; Andreoni et al., 2016). Similarly, Konow (2000) looks at how self-serving biases can influence which allocations people deem to be “fair.” In a related paper, Konow (2005) finds that subjects strategically use information to bias beliefs about fairness and achieve more self-serving outcomes, but that information can also mitigate bias in beliefs about fairness, under the right conditions.

The field of psychology was the first to study outcome bias. In particular, Baron and Hershey (1988) coin the phrase “outcome bias” and identify it in a variety of hypothetical circumstances. They find outcome bias in stated preferences even with uninterested parties and after assuring subjects that outcomes are not potentially valuable signals. Related to our study of third party beliefs, they find that stated beliefs about future competence of doctors are colored by outcome information about patients.

⁸A related literature explores how context and stereotypes impact the attribution of outcomes to luck and skill (Eil and Rao, 2011; Sarsons, 2017; Erkal et al., 2017). Danz et al. (2018) also find failures to process information in an agency setting.

Outside of the lab, outcome bias has been identified using observational data in settings with important policy implications. Competent but unlucky CEOs are more likely to be punished for market shocks than those who are incompetent but lucky (Bertrand and Mullainathan (2001)). Wolfers et al. (2002) and Cole et al. (2012) find similar results with politicians, showing that external, economic factors affect the likelihood of incumbents winning reelection. In professional basketball, Lefgren et al. (2014) show that coaching decisions over-adjust to outcomes relative to other performance information. While observational studies are crucial to ensure the external validity of work on outcome bias, we use the lab to conduct controlled studies of specific, important aspects of the bias, beyond just verifying its existence.

3 Experimental Design

We conducted this study in four waves:

- **Wave 1 (August to November 2016):** Laboratory sessions where principals and agents interact and principals make punishment decisions. Undergraduates at the University of Arkansas and the University of California, San Diego were recruited as subjects.
- **Wave 2 (May 2017):** Incentive-compatible elicitation of third party punishment decisions from subjects recruited from Amazon’s Mechanical Turk website (MTurk). Using a modified strategy method, we presented different potential principal-agent interactions and implemented punishment decisions on the following wave.
- **Wave 3 (July and August 2017):** Laboratory sessions with principals and agents interacting at the University of Arkansas. Punishment of agents was determined using choices from Wave 2.
- **Wave 4 (August 2017):** Incentive-compatible elicitation of beliefs about agent type. Third parties from MTurk were shown data about agents from previous lab sessions.

3.1 Principal-Agent Environments

During Wave 1, we collected responses from 266 subjects in 16 different sessions. All treatment variations were run using zTree software (Fischbacher (2007)) at both campuses. Sessions had a minimum of 12 and a maximum of 18 subjects each. All subjects were randomly assigned to the role of principal or agent at the beginning of the session and maintained that role throughout. We employed neutral language in the study: principals were referred to as “Blue” players and agents were referred to as “Green” players. Each

session consisted of 13 periods. Principals and agents were randomly and anonymously re-matched each period to avoid reputation effects. Periods 1 through 12 featured treatment-specific protocol while period 13 featured a full-information protocol regardless of treatment. One period was randomly selected to determine earnings after all 13 periods were completed. All subjects had to correctly answer all questions of a comprehension quiz about study procedures before beginning period 1. Additionally, we collected an exit survey from all subjects that elicited social preferences and demographic information.⁹

In every treatment, the principal is endowed with \$7. She can win additional money depending on the (simulated) roll of a 6-sided die: if one of her “winning numbers” is rolled, she wins \$6, and if not, she wins nothing. The principal begins with 1 winning number (the number 1). The agent is endowed with \$13 and he can purchase up to 4 additional winning numbers for the principal at a cost of \$0.50 each (he does not get to choose which numbers are purchased). We refer to this investment as agent effort. Thus, the principal always has a chance to win, but cannot be guaranteed to win. The principal always learns of the agent’s effort choice. Depending on the treatment, the principal may or may not learn about the outcome of the die roll as well. The principal then may choose to pay to “reduce the Green player’s earnings” by between \$0 and \$4 in \$1 increments. It costs the principal \$0.25 for each \$1 reduction in the agent’s earnings. After the punishment choice, the outcomes and profits from the interaction are revealed before randomly and anonymously rematching partners.¹⁰

“Full” treatment: The principal is given complete information about the effort investment of the agent and the luck resulting from the die roll. After receiving this information, the principal decides whether to reduce the agent’s payoffs, and if so, by how much.

“Hidden” treatment: The principal is made aware of agent effort, but not the outcome of the die roll until after her punishment decision is made.

“Commit” treatment: The *Commit* treatment introduces a new decision stage for the principal. The principal observes the agent’s effort level, and is given the opportunity to pay p_1 to view the die roll and therefore her luck. The price, p_1 , is drawn from $\{-\$0.25, \$0, \$0.25\}$ with probabilities $\{\frac{1}{4}, \frac{1}{2}, \frac{1}{4}\}$, respectively. That is, the principal may have to pay to see the outcome, may be paid to see the outcome, or may see the outcome for free.

“Force” treatment: The *Force* treatment transfers information control to the agent. After the agent makes his effort decision but before the principal makes her punishment decision, the agent observes

⁹Instructions, comprehension quizzes, and survey questions can be found in Appendix Section A.3.

¹⁰At UCSD, participants were given the same budget and pricing, but were guaranteed a larger payout from participation. Their payout was augmented by a separate study conducted briefly after the conclusion of this study.

the outcome of the die roll and thus whether the principal has won the \$6 prize. The agent then chooses whether to pay p_2 to reveal the luck to the principal. The price, p_2 , is drawn from the set $\{-\$0.25, \$0, \$0.25\}$, with probabilities $\{\frac{1}{4}, \frac{1}{2}, \frac{1}{4}\}$, respectively.

Agents in the *Full*, *Hidden*, and *Commit* treatments are not aware of their treatment status, and therefore the instructions treat them as a pooled group that will randomly encounter one of three treatments. Thus, in the *Commit* treatment, the agent is not made aware of the principal’s control over information. The principal is told that the agent has not been made aware of who controls the information. Agents in the *Force* treatment have a different information set than the other agents because they are aware that they control the principal’s information. Principals in *Force* are informed that they “may or may not observe the dice roll and whether or not [they] won” prior to making their punishment decisions, and the agents know that the principals will not be informed of their control over the information.¹¹

Payoffs: The agent’s payoff in the *Full*, *Hidden*, and *Commit* treatments is given by:

$$\Pi_A = 13 - \frac{x}{2} - y ,$$

where x is the number of sides purchased for the principal and y is the punishment chosen by the principal.

In the *Force* treatment, the agent’s payoff is:

$$\Pi_A = 13 - \frac{x}{2} - y - p_2 \times \text{REVEAL} ,$$

where REVEAL is an indicator variable that equals 1 if the agent reveals the outcome to the principal.

Conditional on outcome, the principal’s payoff in the *Full*, *Hidden* and *Force* treatments is given by:¹²

$$\Pi_P = 7 + 6 \times \text{WIN} - \frac{y}{4} ,$$

where WIN is an indicator variable for whether the die roll shows one of the principal’s winning numbers.

In the *Commit* treatment, the principal’s payoff is:

$$\Pi_P = 7 + 6 \times \text{WIN} - \frac{y}{4} - p_1 \times \text{VIEW} ,$$

where VIEW is an indicator variable for the principal’s choice to view their luck.

3.2 Third Party Punishment Environments

During Wave 2, subjects recruited from MTurk played the roles of third parties tasked with determining the appropriate (costly) punishment for the agents. We again used neutral language and referred to

¹¹While this is no guarantee that the two parties do not develop suspicions about the other party or the other party’s beliefs, this design choice mirrors the ambiguity maintained in the other treatments. Moreover, in Appendix Table A3, we show that principals’ reactions to information does not change over the course of the study.

¹²Prior to receiving information about the outcome, expected payoff is a more appropriate measure for the principal. This is characterized by: $E(\Pi_P) = 7 + 6 \times \frac{1+x}{6} - \frac{y}{4}$.

these third parties as “Orange” players. Third parties viewed 16 different possible combinations of agent effort and principal luck and determined the amount of punishment they wished to impose in each case. These combinations appeared alongside randomly generated rolls of the die.¹³ Since the principal-agent interactions had not yet taken place, this can be thought of as a form of the strategy method.

In every treatment, third parties begin with \$3.50 and can punish the agent between \$0 and \$4 in \$1 increments. It costs the third party \$0.10 for each \$1 reduction in the agent’s earnings. These values were chosen so that \$1 in punishment cost third parties a similar fraction of their wealth as it cost principals in the previous treatments.¹⁴

In Wave 3, we recruited new principals and agents to interact in the laboratory. Principals and agents experienced the same interaction as described above in the *Force* treatment, except the principal was a passive partner with no option to punish. Instead, we probabilistically applied the punishment decisions from the third parties to the agents based on their effort level and luck.

“Full3” treatment: In the full information treatment, the third party is given information about the agent’s effort choice and the luck of the principal before making a punishment decision. With positive probability, this choice is matched to a lab interaction featuring the same effort, die roll and luck.

“Commit3” treatment: Prior to punishment, third parties have the opportunity to pay p_3 to view the principal’s luck. We draw prices, $p_3 \in \{-.10, 0, .10\}$, with probabilities $\{\frac{5}{16}, \frac{6}{16}, \frac{5}{16}\}$, respectively. With positive probability, we apply the punishment choice to a lab interaction with the same effort and luck.

“Force3” treatment: Agents in Wave 3—the subsequent, in-lab interactions—are subject to the punishment decisions of the third parties. These agents know their punishment will be determined by third parties and are given the opportunity to determine the information that these third parties had about the principal’s luck. Just like the principal-agent environment, the agent can pay p_2 to reveal the principal’s outcome to the third party. p_2 is again drawn from $\{-\$0.25, \$0, \$0.25\}$ with probabilities $\{\frac{1}{4}, \frac{1}{2}, \frac{1}{4}\}$, respectively. If the agent chooses to reveal the outcome to the third party, they are matched to a punishment decision featuring the same effort and luck from a third party who observed the outcome of the interaction. If the agent chooses not to reveal the third party the outcome, they are matched to a punishment decision featuring the same effort and luck from a third party who did not observe the outcome of the interaction.¹⁵

¹³The scenarios reflected by the dice were presented with similar frequency to their empirical likelihood. For example, if the principal had only 1 winning side, we presented more losing scenarios with high dice rolls than winning ones.

¹⁴Principals were given \$7 in endowment and received approximately \$3 in expected earnings from the roll of the dice (mean effort was just under 2 sides, yielding a 50% chance of winning \$6). Thus, \$1 in punishment cost 2.5% of a principal’s wealth. Similarly, \$1 in punishment costs 2.85% of a third party’s wealth.

¹⁵A third party who observed the outcome could either be from the *Full* treatment or could be a third party who chose to

Payoffs: The payoff for a third party in the *Full3* treatment is:

$$\Pi_T = 3.5 - \frac{y}{10},$$

where y again represents the punishment choice.

A third party’s payoff in the *Commit3* treatment is:

$$\Pi_T = 3.5 - \frac{y}{10} - p_3 \times \text{VIEW}.$$

where VIEW is, again, an indicator variable for the choice to view the principal’s luck.

Similar to the previous *Force* treatment, the agent’s payoff in the *Force3* treatment is:

$$\Pi_A = 13 - \frac{x}{2} - y - p_2 \times \text{REVEAL},$$

where REVEAL is an indicator variable that equals 1 if the agent shows the principal’s luck to the third party. Principals in this environment were mostly passive observers of outcomes.¹⁶

3.3 Eliciting Beliefs about Agents

“Guess3” treatment: Wave 4 was designed to evaluate the mechanism behind outcome bias in our study. This environment has only one treatment. We again presented third parties with a series of 16 combinations of agent effort and principal luck. These scenarios were drawn from Wave 1. For each agent in each interaction, we calculated the mean effort investment they made across their 13 periods of the study. The third parties are endowed with \$2.50 and can earn an additional \$1 if they guess the mean effort choice of the agent within 0.10. For one randomly selected round, we compare the third party’s guess with the actual mean effort for the agent they observed in that round to determine their payments.

Payoffs: The third party’s payoff in this environment is:

$$\Pi_T = 2.50 + 1 \times \text{CORRECT GUESS},$$

where CORRECT GUESS indicates if the third party’s guess was within 0.10 of the agent’s mean effort.

4 Theory and Hypotheses

The subgame perfect Nash equilibrium in our principal-agent interaction involves no effort or punishment.¹⁷ However, we observe substantial amounts of both, consistent with results on reciprocity such as

reveal the outcome in the *Commit* treatment. A third party who did not observe the outcome would come from the *Commit* treatment after not choosing to reveal the outcome.

¹⁶As before, they earned a \$7 endowment plus a \$6 prize if a winning number of their was rolled in the round randomly selected for payment. At the end of each round, we elicited their beliefs about the mean effort of the agent they were paired with in that round. If their guess was within 0.10 of the correct answer in the round randomly chosen for payment, they received a \$5 bonus. We analyze these data in the same way as the third party guesses.

¹⁷See Appendix Section A.1 for more detail.

the “gift exchange” (Akerlof, 1982). This can be explained by a model of reciprocity, altruistic enforcement of efficient effort as a social norm, or distributional preferences. However, as we will discuss in detail below, no standard model of distributional preferences can explain our third party results. Therefore, at minimum, a large fraction of the punishment behavior must be driven by norm-enforcement, reciprocity, or other social-efficiency concerns.

Punishment as a form of norm-enforcement has been well documented in economics (Rabin, 1993; Fehr and Gächter, 2000; Fehr and Gächter, 2002; Fehr and Fischbacher, 2004). Our results from both interested parties (principals) and disinterested third parties are consistent with this motivation. In this section, we explore a simple model of belief updating that can explain the role of luck in norm-enforcement.

4.1 Belief Updating: An Example

Suppose principals and third parties desire to use punishment to enforce norms. Therefore, they must estimate an agent’s mean effort level (his “type”) in order to determine how deserving of punishment the agent is. This type can be inferred based on an agent’s effort choice, $x \in \{0, 1, 2, 3, 4\}$. Let \bar{x} represent the agent’s type and \hat{x} be the principal’s belief about \bar{x} with $\hat{x}, \bar{x} \in [0, 4]$. Since principals and agents are randomly and anonymously rematched, inference about an agent’s type will be restricted to information gained from their one-shot interaction.¹⁸

A principal who forms beliefs according to Bayes’ rule only takes into account x , the perfect information about the agent’s effort choice. Information about the outcome—a noisy signal of the agent’s effort choice—is made redundant by the perfect signal of x .

To illustrate, let us consider a simple example with the following behavior:

- The principal holds uniform priors about \bar{x} . That is, $\hat{x} \sim U[0, 4]$.
- \bar{x} can hold non-integer values, that is, \bar{x} does not perfectly determine behavior. Suppose agents randomize effort choices between the two nearest integer values with probabilities such that their average effort equals \bar{x} .¹⁹ For example, an agent with $\bar{x} = 2.25$ will exert effort x according to the probabilities, $\Pr(x = 2) = 0.75$ and $\Pr(x = 3) = 0.25$.
- Thus, conditional on observing x , the principal knows that \bar{x} is strictly less than 1 unit away.

¹⁸The population distribution of x could also be informative for the principal. However, just like the current observed x , the population distribution of x is conditionally uncorrelated with luck in the current period, which is our explanatory variable of interest.

¹⁹If behavior were deterministic, then the first news about x would perfectly determine the posterior distribution, leaving no uncertainty to resolve with further signals. This is inconsistent with our experimental data where the vast majority of subjects changed effort choices throughout the experiment.

Figure 1 plots the posterior beliefs generated after observing each possible value of x . Censoring at the upper and lower bounds creates taller peaks of the probability distributions for $x = 0$ and $x = 4$. In this way, the principal updates her beliefs about \bar{x} based only on x . With a higher posterior belief about the agent's type, principals gain less utility from punishing and decrease punishment accordingly.

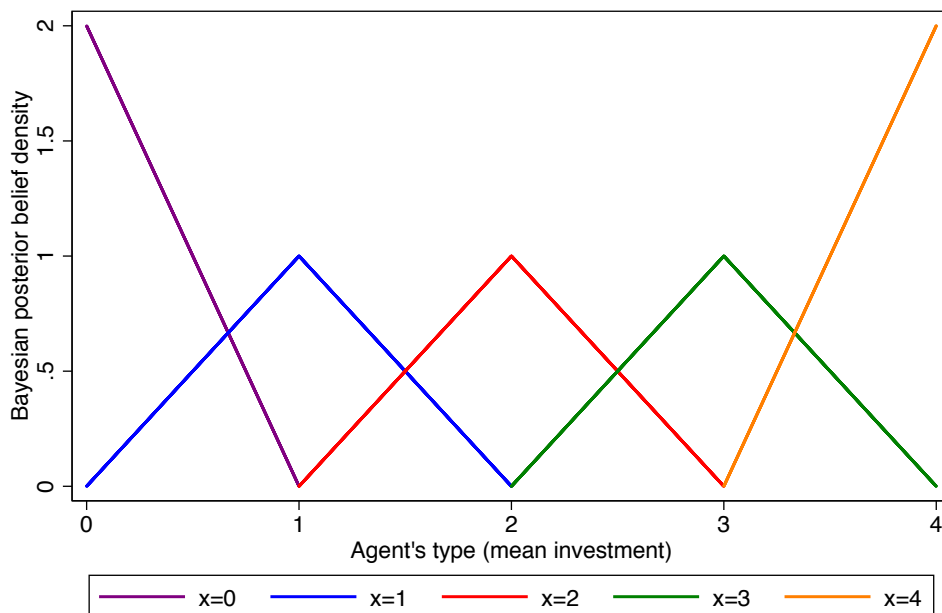


Figure 1: Posterior belief distribution, \hat{x} , conditional on observed effort, x .

Rational belief updating not only eliminates any punishment based on luck, it also eliminates any benefit to control of information about luck. Principals and third parties should be unwilling to pay money to avoid outcome information and agents should similarly place no value on forcing outcome information on principals or third parties.

Prior research shows that people often neglect the correlation between news sources (Enke and Zimmermann, 2017), perceive hidden information within outcomes (Gurdal et al., 2013), and misattribute the effect of state variables (like luck) to choice variables (like effort) (Haggag and Pope, 2016). We incorporate these possibilities by deriving \hat{x} for principals with extreme correlation neglect, believing that the outcome and x represent independent signals of \bar{x} . First, principals update \hat{x} based on their luck. This generates a non-uniform distribution over type. Then, the principal observes x and updates \hat{x} again. The now non-uniform distribution of prior beliefs skews the posterior belief distributions associated with each value of x . This skewness can be seen in Figure 2, which plots the posterior belief distributions for each possible value of x for a principal that observed good luck. Each distribution is now skewed to the

right reflecting the higher beliefs after observing good luck.²⁰ Appendix Section A.2 offers more details on Bayesian and biased belief updating in our study.

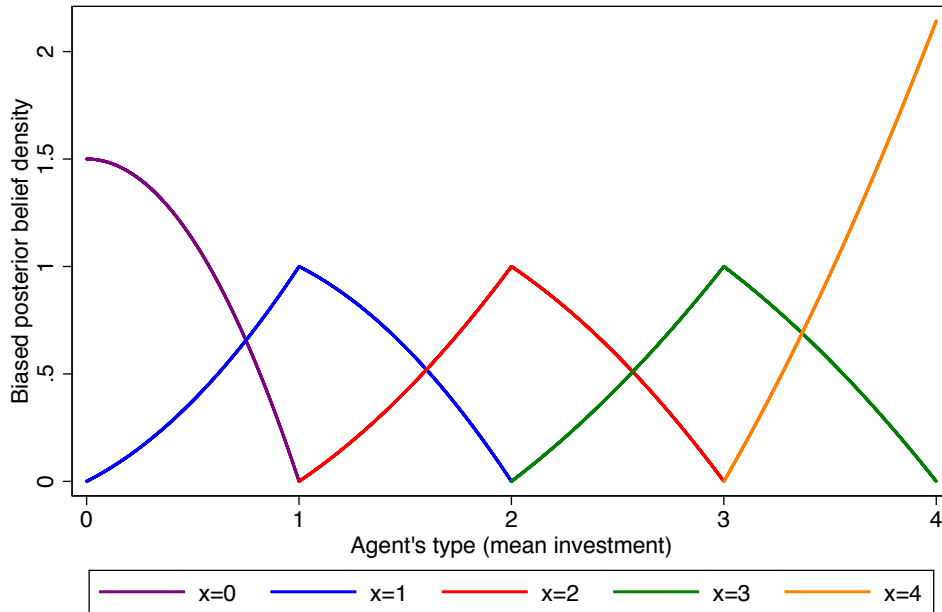


Figure 2: Posterior belief distribution conditional on x after observing good luck.

4.2 Hypotheses

Regardless of the information structure, after conditioning on x , any attribution of luck to the agent's type produces a bias in \hat{x} . This leads to subsequent biases in the nature of punishment by principals and third parties who punish as a form of norm-enforcement. From this model, we derive three primary hypotheses about our environment.

Hypothesis 1 - Outcome biased punishments: Biases in belief updating will cause both principals and third parties to punish unlucky agents more than lucky ones, weakening the relationship between effort and punishment.

Hypothesis 2 - Demand for information control: Agents should be willing to pay to reveal good luck and to keep bad luck hidden. Punishers who anticipate their biases will pay to *not* view information about luck, while naive punishers may seek out this information. That is, under biased belief updating, agents unambiguously benefit from control over what information is revealed to principals and third parties.

²⁰The ordering of the signals is unimportant for generating biased posteriors. If x is observed first, the ex-interim belief distribution will, again, be captured by Figure 1. Observing the outcome will then bias those posteriors as captured in Figure 2. Also, the principal need not believe these signals to be entirely independent. Believing that the outcome possesses *any* information independent of x will lead to biased posteriors.

But, information control has an ambiguous welfare effect on principals and third parties depending on whether or not they anticipate their own biases and desire to enforce norms efficiently.

Hypotheses 1 and 2 are consistent with any form of outcome bias, but we will be able to distinguish biased beliefs from other possible mechanisms like disappointment or distributional preferences by eliciting beliefs, \hat{x} . The most direct result from our model suggests that beliefs, as elicited from third parties, should reveal outcome bias in exactly the same way that punishment decisions do.

Hypothesis 3 - Outcome biased beliefs: Beliefs about an agent's type should be lower (higher) after observing bad (good) luck associated with that agent.

4.3 Distributional Preferences

Inequity aversion and distributional preferences in general (e.g. Bolton (1991); Fehr and Schmidt (1999); Bolton and Ockenfels (2000)) are important models to consider for punishment in our study because knowledge of relative wealth may be a key input to preferences. When the principal learns that she has won the prize, this means that she will be wealthier than the agent. When the principal learns that she has not won the prize, this means that she will be poorer than the agent. In this way, punishment that is motivated by inequity will respond to the revelation of luck.

However, there are a variety of reasons why our data cannot be fully characterized by distributional preferences. First, an inequity-averse principal should never punish when they win, even after the agent exerts the minimum effort. This is not true in our data. In the Full treatment, punishment by winning principals is \$0.40 (S.E. = 0.09) on average. Second, third parties are always poorer than both principals and agents. Thus, the wealth of third parties relative to agents is unaffected by luck, yet, we observe that third parties condition their punishment on both effort and luck, neither of which influence the third parties' relative wealth ranking. Punishment by third parties would therefore have to be concerned with the relative wealth only between the two richer players rather than the relative wealth between all players, which would increase when the principal wins. Finally, distributional preferences offer no pathway through which luck should affect *beliefs* about agent behavior in other interactions. For these reasons, we choose to focus on the model of biased belief updating.

5 Principal Punishment Results

We begin our results with the principal punishment environment. First, we explore the *Full* and *Hidden* treatments to replicate prior results on outcome bias and blame. We then use the *Commit* and *Force* treatments to test for commitment against and manipulation of this bias among principals and agents, respectively.

Before analyzing treatment effects, however, we want to verify an aspect of our study that is important to our interpretation of results: does outcome bias diminish the effectiveness of punishment? That is, is accurate punishment effective at increasing effort from agents? For outcome bias to have welfare consequences, it must be the case that it stifles effective norm enforcement. We obtain the effect of lagged punishment by regressing agent effort this period on their experienced punishment in the prior period, controlling for effort in the prior period. This recovers a causal estimate of punishment because, conditional on agent effort, the punishment is determined by the principals, who are randomly matched with agents. We pool data across the *Full*, *Hidden*, and *Commit* treatments for these estimates because agents in these treatments face identical information feedback. Because maximum effort (four sides purchased) is socially optimal, punishment at all non-maximum effort levels is, in a social sense, accurate. Appendix Table A1 presents results from linear random-effects regressions of effort on lagged punishment with standard errors clustered by agent, separately for non-maximum and maximum effort. Punishment of non-maximum effort is effective: a \$1 increase in punishment raises subsequent effort by 0.12 sides invested ($p < 0.01$). On the other hand, punishment of maximum effort is counterproductive: a \$1 increase in punishment decreases subsequent effort by 0.19 sides invested ($p = 0.02$).

5.1 *Full* and *Hidden* Treatments

The *Full* treatment presents principals with all available information prior to their punishment decisions. Figure 3 shows agent effort and principal punishment—conditional on good or bad luck—in the *Full* treatment. Panel A shows that mean effort is stable throughout the study at roughly two sides purchased. Panel B shows that two is also the modal choice of effort with roughly one-third of effort choices. The remaining effort choices are evenly spread across zero, one, three, and four sides purchased. Panels A and B both show that our data feature substantial outcome bias. With bad luck, mean punishment is around \$1.50, and with good luck, mean punishment is around \$0.50. Outcome bias occurs at every level

of effort, but is larger when effort is less than three sides purchased.²¹

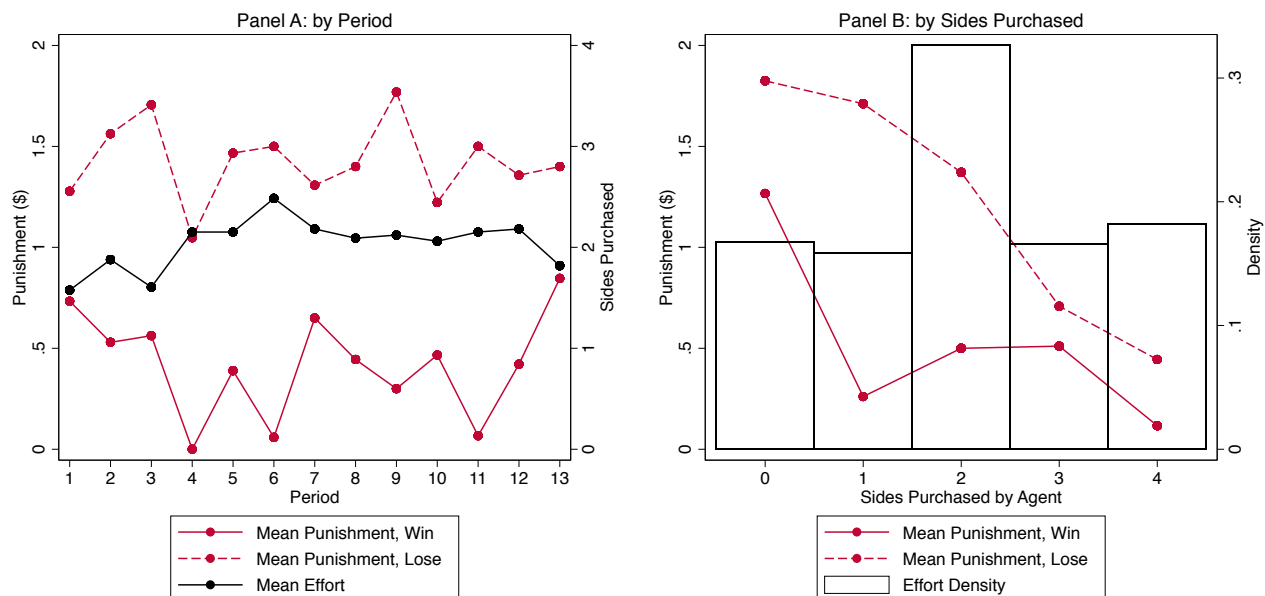


Figure 3: Effort and punishment in the *Full* treatment.

Table 1 formally tests for outcome bias in the *Full* treatment. We regress punishment on an indicator variable for good luck (“Win”), and control for effort to access the conditionally random variation in luck. In column (1), the control for effort is linear (Figure 3, Panel B suggests this should fit the data well), and in column (2), we use effort fixed effects as a non-parametric control. Standard errors are clustered at the principal level. We find very large and statistically significant estimates of the Win coefficient, indicating the presence of outcome bias. Good luck reduces punishments by roughly 33%. Holding luck fixed, the linear effort specification indicates that higher effort is associated with a significant decrease in punishment. Thus, punishment does reflect reciprocity along with outcome bias. However, good luck diminishes punishment as much as a 1.4 SD increase in effort. These effects are not driven by outliers; Appendix Figure A1 shows individual-specific estimates of outcome bias and reciprocity. Both are widespread.

The *Hidden* treatment represents the counterfactual to *Full* in which principals do not have access to information about the bias-generating random outcomes until after their punishment decisions are made. Absent information about luck on which to condition punishment, we find that punishment is more responsive to agent effort. That is, outcome bias weakens responsiveness to effort.

In columns (3) and (4) of Table 1, we compare punishment behavior between *Full* and *Hidden*. Using

²¹This is roughly consistent with the delegated-expertise model of Gurdal et al. (2013), which predicts that the impact of luck is decreasing in the agent’s effort level.

Table 1: Impact of Luck and Effort on Punishment

	(1)	(2)	(3)	(4)
Constant (Full, Effort = 0, Win = 0)	1.99 (0.29)	2.08 (0.33)	1.99 (0.28)	2.08 (0.33)
Win	-0.67*** (0.17)	-0.66*** (0.18)	-0.67*** (0.17)	-0.66*** (0.17)
Effort	-0.36*** (0.08)		-0.36*** (0.07)	
Hidden			0.11 (0.39)	0.52 (0.45)
Win×Hidden			0.73*** (0.19)	0.69*** (0.19)
Effort×Hidden			-0.21* (0.12)	
Effort-level FEs	N	Y	N	Y†
Clusters (Principals)	33	33	66	66
Observations	429	429	825	825

*** $\Rightarrow p < 0.01$, * $\Rightarrow p < 0.10$, † \Rightarrow Effort-level fixed effects are included alone and interacted with the Hidden variable. Estimates are from linear random-effects models.

the same linear random-effects model, as in columns (1) and (2), we include the data from the *Hidden* treatments, and allow for the responses to effort and luck to differ by treatment. Reassuringly, the coefficient on the interaction between Win and *Hidden* perfectly counteracts the Win coefficient; luck cannot affect principals' decisions in *Hidden*. In column (3), we find that the relationship between effort and punishment is 58% stronger in the *Hidden* treatment. An extra side of the die purchased reduces punishment by an average of \$0.36 in *Full*, and \$0.57 in *Hidden*. This difference is marginally statistically significant ($p = 0.08$).²²

5.2 *Commit* Treatment

In the *Commit* treatment, principals learn of the agent's effort level and then are given the opportunity to avoid viewing their luck prior to their punishment decision. Recall that all principals learn the outcomes at the end of each round regardless of their choice. Thus, keeping this information temporarily hidden can be interpreted as an opportunity for principals to commit against outcome bias.

²²Because these regressions feature across-treatment variation that is randomized at the session level, we also run OLS version of these regressions that allow for wild-boostapped standard errors at the session level (Cameron et al., 2008). In column (3), the Win, Effort and Win×Hidden coefficients remain statistically significant with $p = 0.04$, 0.04 , and 0.02 , respectively. The Effort×Hidden coefficient becomes marginally insignificant with $p = 0.13$. In column (4), both the Win and Win×Hidden coefficients remain very precise, with $p < 0.01$ in both cases.

Figure 4, Panel A shows the rates of information revelation by the agent’s effort level. We find high rates of principals viewing their luck. Principals choose to view the outcome 71% of the time, on average. If principals wanted to avoid large errors in judgment—punishing maximum effort, or failing to punish minimum effort—we should see the lowest levels of revelation at those effort levels. Instead, we observe the opposite. Principals are 15% more likely to view the outcome if effort is at the maximum or minimum level than at any of the moderate effort levels ($p = 0.06$).²³ This is also inconsistent with a model of curiosity: uncertainty is lowest at maximum and minimum effort. Principals appear to be searching out outcome information to ensure that selection of low or high punishment is “justified” by the outcome.

Figure 4, Panel B plots the principal demand function for information. Principals who anticipate their outcome bias and wish to avoid it should never view the outcome when it is costly or free. However, 30% of principals pay \$0.25 to see the outcome ($p < 0.01$, tested against zero). A striking 90% of principals choose to view the outcome when avoiding it is free ($p < 0.01$, tested against zero). A notable feature of Panel B is that demand is non-monotonic in the price: when principals are paid to see the outcome, the rate declines to 70% ($p = 0.03$, tested against demand at price of \$0), indicating that principals’ preferences about information revelation lack sophistication in multiple dimensions.




Figure 4: Principal information revelation in the *Commit* treatment.

While principals may not avoid information to reduce outcome bias, it could still be that the act of thinking about whether to view the information helps them to better condition their punishment on the

²³Linear probability model, standard errors clustered by principal.

behavior of agents. If this is the case, we would expect the relationship between effort and punishment in the *Commit* treatment to look more like the relationship in *Hidden* than in *Full*. From column (3) of Table 1, we found that an additional side purchased by the agent reduced punishment by \$0.36 in *Full*, and \$0.57 in *Hidden*. We estimate the same linear random effects model (standard errors clustered by principal), with the *Commit* treatment included. We find that the relationship between effort and punishment is significantly weaker in *Commit* than in *either Hidden or Full*: an additional side purchased by the agent reduces punishment by \$0.07 ($p = 0.24$ tested against zero, $p < 0.01$ both tested against \$0.57 in *Hidden* and \$0.36 in *Full*). Rather than improving the responsiveness of punishment to luck, principal information control has the opposite effect. Punishment in *Commit* remains responsive to luck—good luck decreases punishment by a marginally significant \$0.29 ($p = 0.06$)—but it may be less responsive than in *Full* (difference: $p = 0.10$).²⁴

To assess how actively principals seek to commit against outcome bias, we asked their motivations behind avoiding the outcome prior to making their punishment decisions. 32% indicated that they did so to avoid punishing an agent who did not deserve to be punished. Of those 32%, only 12% indicated that they did so to make sure they punished deserving agents. This suggests that relatively few principals understood, or cared, that luck would influence their choice.²⁵

5.3 Force Treatment

Figure 5, Panel A shows the rates of information revelation by the agent’s effort level. For each level of effort, agents are more likely to reveal good luck than bad luck, with good luck being revealed 85% of the time, relative to 46% of the time for bad luck. Panel B plots the agent demand function for showing the outcome to the principal, separated by good and bad luck. Agents are willing to be paid or costlessly show good luck to principals the vast majority of the time. Agents that perceive sufficient outcome bias amongst principals may be willing to pay to show the principal good luck. We find that 64% of agents are willing to pay to show good luck ($p < 0.01$, tested against zero). On the other hand, we expect agents who anticipate this bias to be willing to forgo \$0.25 to hide bad luck from outcome biased principals. We find that 40% of principals are willing to forgo payment to hide bad luck ($p < 0.01$, tested against zero).²⁶

In Appendix Table A2, we formally estimate the impact of luck on agent information revelation. We employ both random effect linear probability models and random effect probit models. Good luck has a

²⁴Using the wild-bootstrap procedure for this regression, we still find that punishment in *Commit* is significantly less responsive to effort than in *Hidden* ($p < 0.01$) and *Full* ($p = 0.02$). Within *Commit*, luck has a marginally insignificant

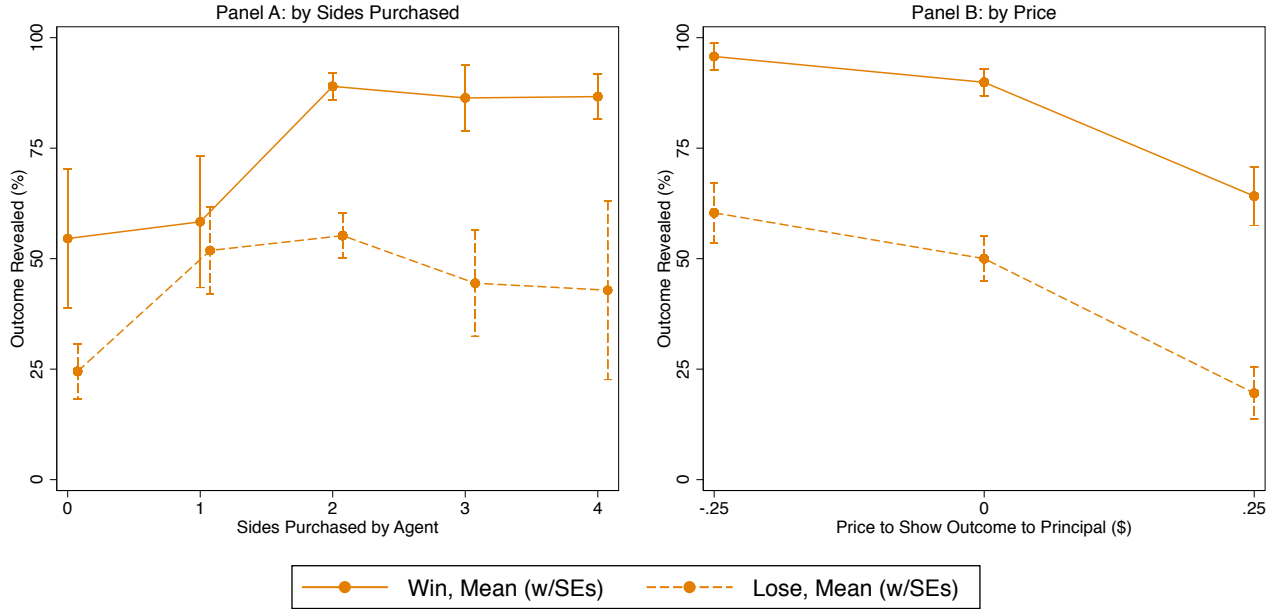


Figure 5: Agent information revelation in the *Force* treatment.

large, positive, and precisely estimated impact—between 33 and 39 percentage points—on the likelihood of revealing the outcome to the principal.²⁷

In a world of rational principals, positive expenditure on information control is strictly welfare decreasing. Thus, the observed expenditures demonstrate the perception amongst agents that there is value in manipulating the principal’s information. This perception is supported by the different responses to effort and luck between the *Hidden* and *Full* treatments. However, principals are randomly assigned *between* the *Full* and *Hidden* treatments, while in the *Force* treatment, there is within-principal variation in whether the outcome is revealed. Thus, principals may treat the information differently even though they do not know why the outcome is revealed selectively.

Table 2 presents a test of the ability of agents to successfully manipulate principals using outcome information. We regress principal punishment on 1) a revealed winning outcome and 2) a revealed

effect $p = 0.14$, which may be less responsive than in *Full* ($p = 0.10$)

²⁵44% never avoided the outcome, 3% failed to respond.

²⁶The distribution of individual-level mean behavior after winning stochastically dominates the distribution after losing (Kolmogorov-Smirnov test, $p < 0.01$), though both distributions span from 0 – 1. 48% of agents always show good luck and 21% never show bad luck.

²⁷In columns (1) and (4) of Appendix Table A2, holding luck fixed, effort is positively correlated with the likelihood of revealing the outcome. This is the opposite of what we’d expect; showing good luck should be more valuable when effort is low. This relationship could be obscured by the fact that 75% of the time, it is free, or agents are paid to show the outcome. Indeed, when we interact the price with agent effort in columns (3) and (6), we find that a higher price predicts a lesser positive relationship between effort and revealing the outcome, although this interaction is not statistically significant ($p = 0.30$ in column (3) and $p = 0.17$ in column (6)). To actually observe a negative relationship between effort and information revelation, these estimates suggest we would need a price of at least \$0.50.

losing outcome. While the choice to reveal outcomes is endogenous to the agent, the randomly assigned pairing of principals and agents ensures that the revelation is conditionally random to the principal after controlling for agent effort and luck.²⁸ In each specification, we control for outcome. Column (1) includes a linear effort control, while Column (2) includes effort fixed effects. We use the same linear random effects models as in Table 1. Standard errors are clustered at the principal level. We find that a revealed win has a large negative impact on Punishment in the *Force* treatment.²⁹ In Appendix Table A3, we show that this effect does not diminish over time. Showing a loss has no substantial positive impact.³⁰

Table 2: Effect of Information on Punishment in the *Force* Treatment

	(1)	(2)
Constant (Outcome Hidden, Lose, Effort = 0)	1.50 (0.27)	1.70 (0.32)
Outcome Shown×Win	-0.44*** (0.17)	-0.34** (0.17)
Outcome Shown×Lose	0.10 (0.17)	0.18 (0.16)
Win	0.22 (0.18)	0.18 (0.18)
Effort	-0.35*** (0.08)	
Effort-level FEs	N	Y
Clusters (Principals)	33	33
Observations	396	396

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$. Standard errors clustered by individual. Estimates are from linear random-effects models.

Though these results demonstrate an awareness of outcome bias on the part of the agents, they do not reveal whether these are considered decisions or mere instinct. To address this, we surveyed all agents at the end of the experiment to ask about their decision-making process. When asked what motivated them to reveal the outcome of the dice, 70% indicated that they did so because they thought good luck

²⁸Alternatively, we can use the price of showing the outcome as an instrumental variable for information revelation. This yields a similar qualitative pattern of results but with larger magnitudes of the effects of information. Because the instrumental variables approach yields a local average treatment effect at the margin where revelation is influenced by its price, we prefer the estimates in Table 2.

²⁹Considering costly punishment for low effort a sort of public good, this is consistent with existing results showing that costly pro-social behavior is often avoided under the type of uncertainty that would be present without knowledge of the outcome (Dana et al. (2007), Andreoni and Bernheim (2009), Andreoni and Sanchez (2014), Exley (2016)).

³⁰This suggests that principals in the *Force* treatment associate not seeing the outcome with losing, providing an explanation for why agent willingness to pay to reveal winning outcomes is higher than agent willingness to forgo payment to hide losing outcomes.

would decrease the principal’s selected punishment.³¹ When asked what motivated them to hide the outcome of the dice, 58% indicated that they did so because they thought bad luck would increase the principal’s selected punishment.³² These responses suggest that most agents have a keen awareness of the disproportionate impact of outcomes even when their effort levels are observed.

6 Third Party Punishment Results

We now examine outcome bias among third parties. We use the *Full3* treatment to estimate the magnitude of outcome bias among third parties. The *Force3* treatment then measures agents’ anticipation of outcome bias among third parties. Finally, the *Commit3* treatment tests for demand among third parties for commitment against their own outcome bias.

Punishment decisions by third parties in the *Full3* treatment are shown in Figure 6. Despite not being affected by the agent’s effort choices, third parties regularly pay to punish them. Punishment declines with effort, indicating that third parties are likely trying to enforce the same norms of effort provision as principals. As with the principals in the *Full* treatment, they exhibit clear outcome bias.

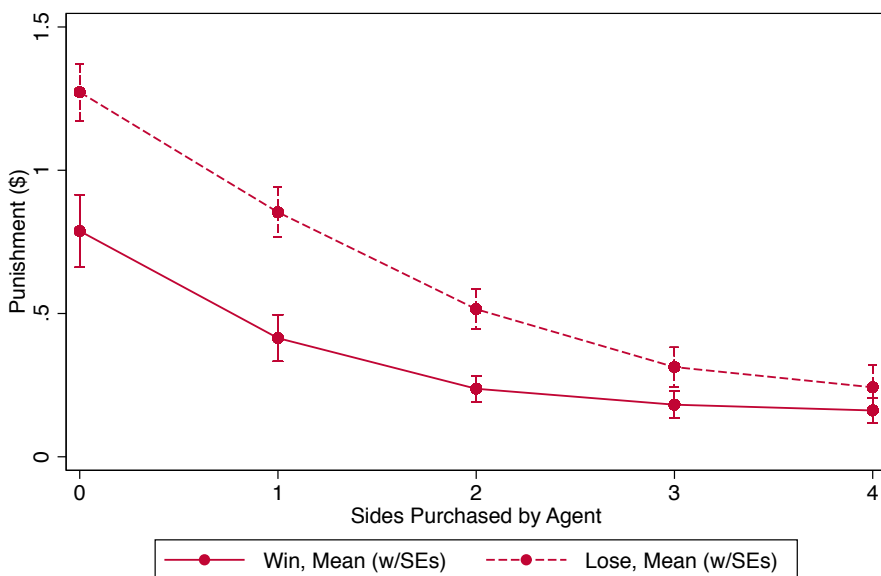


Figure 6: Punishment in the *Full3* treatment.

In Table 3, we formally estimate the magnitude outcome bias using a linear random effects model

³¹These individuals either responded that they revealed the dice when they had invested a large amount and wanted the principal to know they had won, or that they revealed the outcome when they had invested little and wanted to show the principal that they had gotten lucky (or both). 3% never revealed the outcome, and 9% failed to respond.

³²These individuals either responded that they hid the outcomes when they had exerted little effort and did not want the principal to know they had lost, or that it was to hide bad luck when they exerted substantial effort (or both). 12% of agents never hid the outcome, and 9% failed to respond.

with standard errors clustered at the third-party level. Good luck reduces punishment from third parties by about 26% (compared to 33% in the *Full* treatment). However, relative to the impact of effort on punishment, the magnitude of outcome bias for third parties is *identical* to that of principals: good luck reduces punishment in the same way as a 1.4 SD increase in effort.

Table 3: Impact of Luck and Effort on Third Party Punishment in the *Full3* Treatment

	(1)	(2)
Constant (Win = 0, Effort = 0)	1.06 (0.13)	1.21 (0.14)
Win	-0.28*** (0.06)	-0.28*** (0.06)
Effort	-0.20*** (0.03)	
Effort-level FEs	N	Y
Clusters (Third Parties)	99	99
Observations	1584	1584

*** $\Rightarrow p < 0.01$. Estimates are from linear random-effects models.

Sensitivity to luck among third parties who have lower wealth levels implies that it cannot be inequity aversion alone driving the outcome bias. Moreover, these results cannot be explained by the emotion associated with winning or losing money, since the third parties have no monetary stake in the interaction.

6.1 *Force3* Treatment

We use the *Force3* treatment to test for demand among agents in the laboratory for control over the information of the third parties. Figure 7, Panel A shows that, regardless of the agent's effort, they are more likely to reveal good luck to the third parties than bad luck. Controlling for effort, good luck increases the probability of revealing the outcome to the third party by 24 percentage points ($p = 0.01$).³³ Panel B shows the demand functions for revealing the outcome to the third party separated by good and bad luck. We find that 52% of agents are willing to pay to show good luck to third parties ($p < 0.01$, tested against zero), and 46% of principals are willing to forgo payment to hide bad luck ($p < 0.01$, tested against zero). Agents' positive willingness to pay for control over information and their choice to condition revelation on luck both indicate that agents correctly predict outcome bias in third parties.

³³Random effects linear probability model, standard errors clustered by agent.




Figure 7: Agent information revelation in the *Force3* Treatment.

6.2 *Commit3* Treatment

Similar to the *Commit* treatment the third parties in *Commit3* were presented with the choice whether to see information about the principal’s luck prior to making their punishment choices. Figure 8, Panel A shows the rates of viewing information by the agent’s effort level. Third parties in *Commit3* view the outcome 55% of the time, slightly less often than principals in *Commit*, who view it 71% of the time. Unlike principals, who sought information most strongly at maximum or minimum effort, the rate of revelation by third parties is decreasing in agent effort, although the magnitude of this relationship is very small.³⁴ However, the information demand function for third parties in *Commit3* shown in Figure 8, Panel B corresponds closely to that of principals in *Commit*. 22% of third parties pay to view luck ($p < 0.01$, tested against zero), 72% choose to view the principal’s luck for free ($p < 0.01$, tested against zero), and demand *decreases* to 58% when the price decreases from \$0 to -\$0.10 ($p = 0.10$).

³⁴From a baseline of 58%, each additional side purchased reduces revelation by 1.5 percentage points ($p = 0.02$). This estimate is from a random effect linear probability model, standard errors clustered by third party.

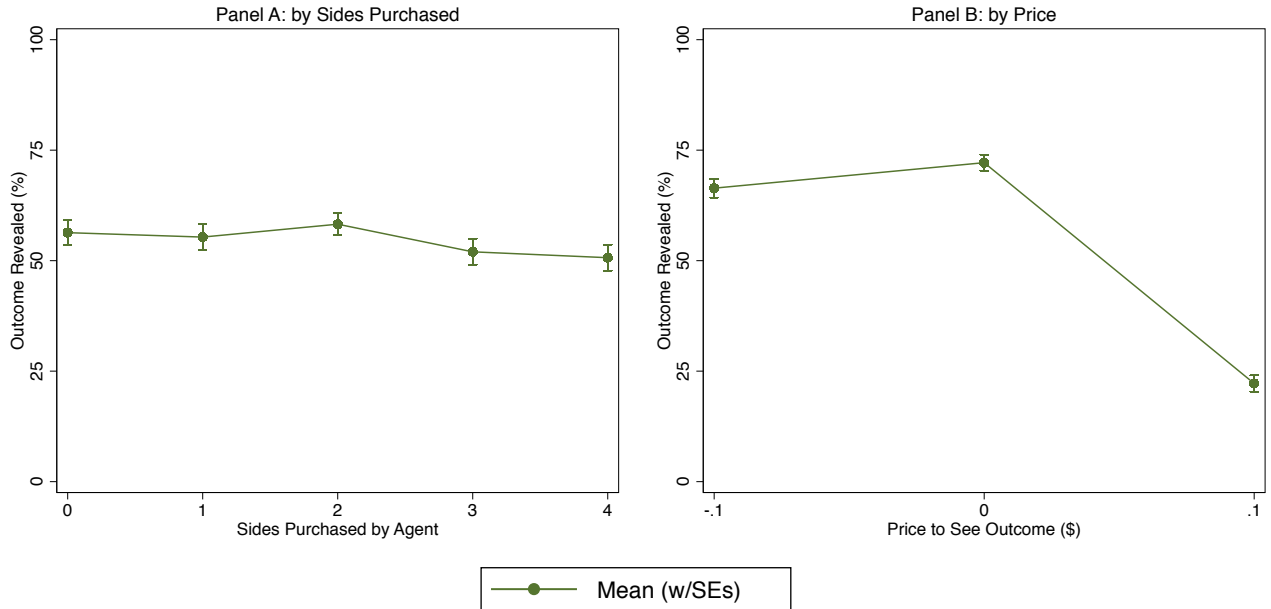


Figure 8: Third party information revelation in the *Commit3* treatment.

7 Belief Measurement Results

In this section, we measure the impact of luck on the beliefs that third parties hold about the agents' types. We presented third parties in the *Guess3* treatment with effort and outcomes for 16 different principal-agent interactions. For each interaction, we elicited the third party's belief about the *average effort choice* of the agent currently being observed. This average includes the agent's choice in the interaction observed by the third party and twelve choices unobserved by the third party.

Information about luck in the current interaction should be useless in the belief updating process.³⁵ Nonetheless, we observe a strong influence of luck on the beliefs elicited about agents' types. Figure 9 shows the relationship between effort, luck, and third party guesses. Panel A shows that observed effort is the first-order determinant of guesses, but, for every effort level, the guess is higher with good luck than bad luck. In Panel B, we plot the same relationship using the predicted residual after regressing guesses on effort-level fixed effects. The effect of good luck on guessed average agent effort ranges from roughly 0.1 sides to 0.3 sides, with the biggest effects at maximum and minimum effort.

In Table 4, we estimate the impact of luck on beliefs about the agent's type using a linear random effects model with standard errors clustered by third party. Good luck increases estimates of the agent's type by 16%. This effect is precisely estimated. Despite observing the agent's effort, good luck is

³⁵There are no dynamic effects of luck on subsequent effort choices of agents.

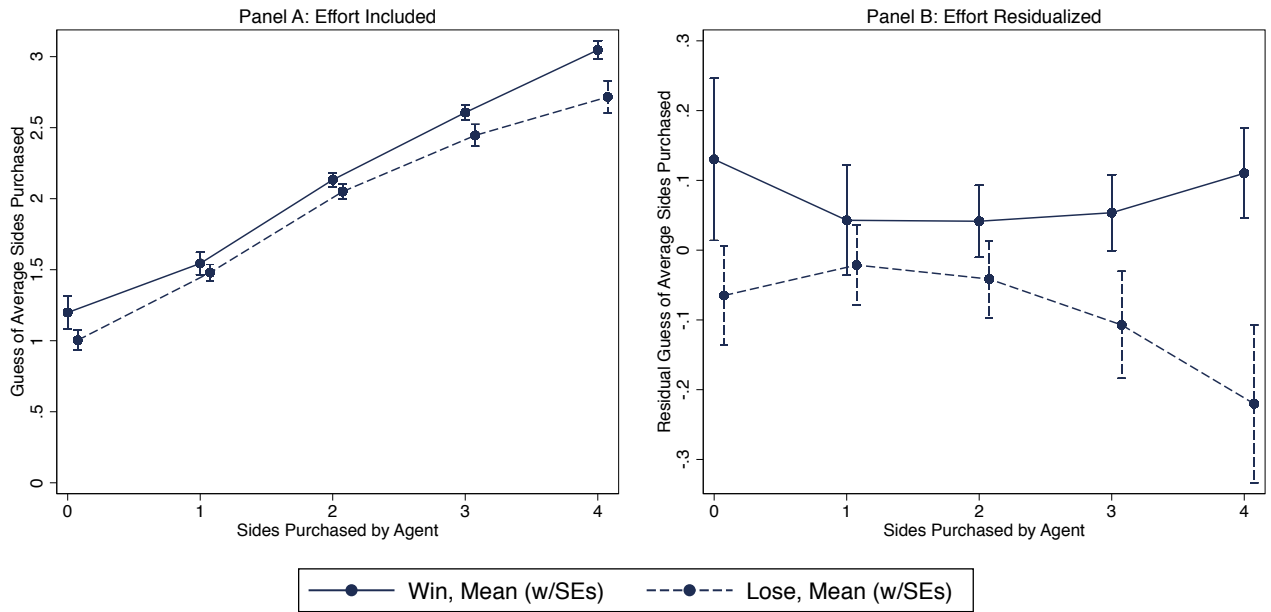


Figure 9: Third party guesses about average agent effort in the *Guess3* treatment.

Table 4: Impact of Luck and Effort on Third Party Guesses about Agent's Mean Effort in the *Guess3* Treatment

	(1)	(2)
Constant (Win = 0, Effort = 0)	1.03 (0.08)	1.02 (0.09)
Win	0.16*** (0.06)	0.16*** (0.06)
Effort	0.46*** (0.03)	
Effort-level FEs	N	Y
Clusters (Third Parties)	100	100
Observations	1600	1600

*** $\Rightarrow p < 0.01$. Estimates are from linear random-effects models.

misattributed to effort and has the same effect as if the agent’s effort increased by 0.25 SD.³⁶

The third parties are outside observers whose payoffs are unaffected by the principal-agent interaction. Thus, we do not expect them to be making punishment decisions based on emotional responses to principal outcomes. The influence of luck on the beliefs that third parties hold about an agent’s type—and, thus, how deserving agents are of punishment—is strong evidence that biased beliefs are a driver of our results.

8 Conclusion

This paper studies three previously-unexplored aspects of outcome bias: 1) anticipation and manipulation of outcome bias, 2) outcome bias among third parties, 3) outcome-biased belief updating.

We find that agents strategically manipulate principals, but principals do not leverage commitment opportunities. If susceptibility to the bias is heterogeneous in the population, this could be a mechanism through which rents can be extracted from susceptible individuals by politicians, financial advisors, employees, attorneys or others with an informational advantage over their audience. The threat of this manipulation should cause policy makers to think carefully about who controls information revelation and if this information control can have a negative welfare impact on society. For example, a nefarious financial advisor with control over information revelation could selectively reveal portfolio performance when market returns are high and would be perceived more positively than a financial advisor required to reveal performance at predetermined intervals.³⁷

We also extend the study of outcome bias to third parties. We find that, just like interested parties, they exhibit strong outcome bias. They can also be manipulated by agents with control over information. This finding has broad consequences, since it eliminates the possibility that outcome bias can be “solved” by simply using third party arbiters. It also rules out two origins for outcome bias: 1) the disappointment or elation associated with bad or good luck, respectively and 2) the inequalities created by good and bad luck. For example, not only might the plaintiff judge a defendant more harshly because of a unlucky, negative outcome, but juries and judges may do so as well. We do, however, find that third parties avoid basing their punishments on luck when they have control over their information set.

Finally, we construct a test of whether biased belief updating can explain our data: are beliefs that

³⁶We used passive principals in the *Force3* treatment to perform a similar test, and find a similar estimate—good luck has the same effect on the principal’s guess as a 0.28 SD increase in effort—but the test of this effect is underpowered ($p = 0.22$). Results are in Appendix Table A4.

³⁷The efficiency of legislation on fiduciary responsibility suffers in the face of outcome bias. Advisors are only likely to face litigation after bad luck, and therefore excessively invest in low-risk investments.

third parties hold about the effort level of agents *in general*, sensitive to luck *in the specific, observed interaction*? We find that bad luck causes third parties to decrease their beliefs about an agent’s type by the same amount as a 0.25 SD decrease in effort.

These results are consistent with a model of biased updating where observers of interactions fail to ignore outcomes—a noisy signal—despite observing effort—a perfect signal. This model can be thought of as a combination of extreme correlation neglect (Enke and Zimmermann, 2017) and attribution bias (Haggag and Pope, 2016). When individuals fail to recognize the redundancy between signals, they end up attributing more of the outcome to effort than they should. The fact that we observe beliefs respond to luck is direct support for this behavioral pathway.

Using outcomes to infer the effort, intentions, or skill of others is a common feature of many important decisions, especially in personnel and legal settings. The degree of bias exhibited in our experimental study may be a lower-bound due to the simple and salient information structure. Also, considering that we clearly document outcome bias among third parties, this phenomenon may be quite widespread indeed. We believe that the workplace and legal contexts are ripe for practical extensions of this work. For example, performance assessments are almost always made based on outcomes that depend on both effort and luck. Incentives for effort—for example, through compensation and promotion policies—depend on the predictive validity of performance assessments, which can be limited by outcome bias. Moreover, given that we find that manipulation is effective, employees may be rewarded for their differential willingness to manipulate their employer’s information. These mistaken evaluations may be persistent: good luck may be misattributed to an employee’s work-ethic and bias future assessments as well. Given these long-run implications, whether principals learn over a longer horizon, whether role reversal leads to more symmetric behavior, and whether experienced evaluators exhibit such biases will be important avenues for future research.

References

- Akerlof, George.** (1982). ‘Labor contracts as partial gift exchange’, *Quarterly Journal of Economics* 97, 543–569.
- Andreoni, James, Aydin, Deniz, Barton, Blake, Bernheim, B Douglas and Naecker, Jeffrey.** (2016). ‘When fair isn’t fair: Sophisticated time inconsistency in social preferences’.
- Andreoni, James and Bernheim, B Douglas.** (2009). ‘Social image and the 50–50 norm: A theoretical and experimental analysis of audience effects’, *Econometrica* 77(5), 1607–1636.
- Andreoni, James and Sanchez, Alison.** (2014), Do beliefs justify actions or do actions justify beliefs? An experiment on stated beliefs, revealed beliefs, and social-image manipulation, Technical report, National Bureau of Economic Research.
- Baron, Jonathan and Hershey, John C.** (1988). ‘Outcome bias in decision evaluation.’, *Journal of personality and social psychology* 54(4), 569.
- Bartling, Björn and Fischbacher, Urs.** (2011). ‘Shifting the blame: On delegation and responsibility’, *The Review of Economic Studies* 79(1), 67–87.
- Bertrand, Marianne and Mullainathan, Sendhil.** (2001). ‘Are CEOs rewarded for luck? The ones without principals are’, *The Quarterly Journal of Economics* 116(3), 901–932.
- Blount, Sally.** (1995). ‘When social outcomes aren’t fair: The effect of causal attributions on preferences’, *Organizational behavior and human decision processes* 63(2), 131–144.
- Bolton, Gary E.** (1991). ‘A Comparative Model of Bargaining: Theory and Evidence’, *American Economic Review* 81(5), 1096–1136.
- Bolton, Gary E. and Ockenfels, Axel.** (2000). ‘ERC: A Theory of Equity, Reciprocity, and Competition’, *American Economic Review* 90(1), 166–193.
- Brock, J. Michelle, Lange, Andreas and Ozbay, Erkut Y.** (2013). ‘Dictating the Risk: Experimental Evidence on Giving in Risky Environments’, *American Economic Review* 103(1), 415–37.
URL: <http://www.aeaweb.org/articles?id=10.1257/aer.103.1.415>
- Bushong, Benjamin and Gagnon-Bartsch, Tristan.** (2016), Misattribution of Reference Dependence: Evidence from Real-Effort Experiments, Technical report, Working Paper.
- Cameron, A Colin, Gelbach, Jonah B and Miller, Douglas L.** (2008). ‘Bootstrap-based improvements for inference with clustered errors’, *The Review of Economics and Statistics* 90(3), 414–427.
- Cappelen, Alexander W, Konow, James, Sørensen, Erik Ø and Tungodden, Bertil.** (2013). ‘Just luck: An experimental study of risk-taking and fairness’, *American Economic Review* 103(4), 1398–1413.
- Charness, Gary.** (2004). ‘Attribution and reciprocity in an experimental labor market’, *Journal of labor Economics* 22(3), 665–688.
- Charness, Gary and Levine, David I.** (2007). ‘Intention and stochastic outcomes: An experimental study’, *The Economic Journal* 117(522), 1051–1072.
- Charness, Gary and Rabin, Matthew.** (2002). ‘Understanding social preferences with simple tests’, *The Quarterly Journal of Economics* 117(3), 817–869.

- Coate, Stephen.** (1995). ‘Altruism, the Samaritan’s Dilemma, and Government Transfer Policy’, *American Economic Review* 85(1), 46–57.
- Cole, Shawn, Healy, Andrew and Werker, Eric.** (2012). ‘Do voters demand responsive governments? Evidence from Indian disaster relief’, *Journal of Development Economics* 97(2), 167–181.
- Cushman, Fiery, Dreber, Anna, Wang, Ying and Costa, Jay.** (2009). ‘Accidental outcomes guide punishment in a “trembling hand” game’, *PloS one* 4(8), e6699.
- Dana, Jason, Weber, Roberto A and Kuang, Jason Xi.** (2007). ‘Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness’, *Economic Theory* 33(1), 67–80.
- Danz, David, Madarász, Kristóf and Wang, Stephanie.** (2018). ‘The Biases of Others: Projection Equilibrium in an Agency Setting’.
- Demski, Joel S. and Sappington, David E.M.** (1987). ‘Delegated Expertise’, *Journal of Accounting Research* 25(1), 68–89.
- de Oliveira, Angela C.M., Smith, Alexander and Spraggon, John.** (2017). ‘Reward the Lucky? An experimental investigation of the impact of agency on luck and bonuses’, *Journal of Economic Psychology* 62, 87–97.
- Ehrlinger, Joyce, Gilovich, Thomas and Ross, Lee.** (2005). ‘Peering into the bias blind spot: People’s assessments of bias in themselves and others’, *Personality and Social Psychology Bulletin* 31(5), 680–692.
- Eil, David and Rao, Justin M.** (2011). ‘The good news-bad news effect: asymmetric procession of objection information about yourself’, *American Economic Journal: Microeconomics* 3(2), 114–138.
- Enke, Benjamin and Zimmermann, Florian.** (2017). ‘Correlation Neglect in Belief Formation’, *The Review of Economic Studies* p. rdx081.
- Erkal, Nisvan, Gangadharan, Lata and Koh, Boon Han.** (2017). Attribution biases in leadership: is it effort or luck? Working Paper.
- Exley, Christine L.** (2016). ‘Excusing Selfishness in Charitable Giving: The Role of Risk’, *Review of Economic Studies* 83(2), 587–628.
- Falk, Armin, Fehr, Ernst and Fischbacher, Urs.** (2008). ‘Testing theories of fairness—Intentions matter’, *Games and Economic Behavior* 62(1), 287–303.
- Falk, Armin and Fischbacher, Urs.** (2006). ‘A theory of reciprocity’, *Games and economic behavior* 54(2), 293–315.
- Fehr, Ernst and Fischbacher, Urs.** (2004). ‘Third-party punishment and social norms’, *Evolution and human behavior* 25(2), 63–87.
- Fehr, Ernst and Gächter, Simon.** (2000). ‘Cooperation and punishment in public goods experiments’, *American Economic Review* 90(4), 980–994.
- Fehr, Ernst and Gächter, Simon.** (2002). ‘Altruistic punishment in humans’, *Nature* 415(6868), 137.
- Fehr, Ernst and Schmidt, Klaus M.** (1999). ‘A theory of fairness, competition, and cooperation’, *The quarterly journal of economics* 114(3), 817–868.
- Fischbacher, Urs.** (2007). ‘z-Tree: Zurich toolbox for ready-made economic experiments’, *Experimental economics* 10(2), 171–178.

- Fischhoff, Baruch.** (1975). ‘Hindsight is not equal to foresight: The effect of outcome knowledge on judgment under uncertainty.’, *Journal of Experimental Psychology: Human perception and performance* 1(3), 288.
- Gino, Francesca, Moore, Don A and Bazerman, Max H.** (2009). ‘No harm, no foul: The outcome bias in ethical judgments’.
- Gneezy, Uri and Imas, Alex.** (2014). ‘Materazzi effect and the strategic use of anger in competitive interactions’, *Proceedings of the National Academy of Sciences* 111(4), 1334–1337.
- Gurdal, Mehmet Y, Miller, Joshua B and Rustichini, Aldo.** (2013). ‘Why blame?’’, *Journal of Political Economy* 121(6), 1205–1247.
- Haggag, Kareem and Pope, Devin.** (2016), Attribution bias in consumer choice. Working Paper.
- Halevy, Yoram, Peters, Michael et al.** (2009), Other-Regarding Preferences: Outcomes, Intentions or Interdependence?, Technical report, Vancouver School of Economics.
- Hölmstrom, Bengt.** (1979). ‘Moral hazard and observability’, *The Bell journal of economics* pp. 74–91.
- Konow, James.** (2000). ‘Fair shares: Accountability and cognitive dissonance in allocation decisions’, *American economic review* 90(4), 1072–1091.
- Konow, James.** (2005). ‘Blind spots: The effects of information and stakes on fairness bias and dispersion’, *Social Justice Research* 18(4), 349–390.
- Landmann, Helen and Hess, Ursula.** (2016). ‘What elicits third-party anger? The effects of moral violation and others’ outcome on anger and compassion’, *Cognition and emotion* pp. 1–15.
- Lefgren, Lars J, Sims, David P and Stoddard, Olga B.** (2016). ‘Effort, luck, and voting for redistribution’, *Journal of Public Economics* 143, 89–97.
- Lefgren, Lars, Platt, Brennan and Price, Joseph.** (2014). ‘Sticking with what (barely) worked: A test of outcome bias’, *Management Science* 61(5), 1121–1136.
- Loewenstein, George, O’Donoghue, Ted and Rabin, Matthew.** (2003). ‘Projection bias in predicting future utility’, *The Quarterly Journal of Economics* 118(4), 1209–1248.
- McCabe, Kevin A, Rigdon, Mary L and Smith, Vernon L.** (2003). ‘Positive reciprocity and intentions in trust games’, *Journal of Economic Behavior & Organization* 52(2), 267–275.
- Myerson, Roger B.** (1991), *Game Theory: Analysis of Conflict*, Harvard University Press.
- Nelson, William Robert.** (2002). ‘Equity or intention: it is the thought that counts’, *Journal of Economic Behavior & Organization* 48(4), 423–430.
- Pronin, Emily.** (2007). ‘Perception and misperception of bias in human judgment’, *Trends in cognitive sciences* 11(1), 37–43.
- Pronin, Emily, Lin, Daniel Y and Ross, Lee.** (2002). ‘The bias blind spot: Perceptions of bias in self versus others’, *Personality and Social Psychology Bulletin* 28(3), 369–381.
- Rabin, Matthew.** (1993). ‘Incorporating fairness into game theory and economics’, *The American economic review* pp. 1281–1302.
- Rubin, Jared and Sheremeta, Roman.** (2015). ‘Principal–agent settings with random shocks’, *Management Science* 62(4), 985–999.

- Sarsons, Heather.** (2017), Interpreting Signals: Evidence from Doctor Referrals. Working Paper.
- Schwarz, Norbert and Clore, Gerald L.** (1983). ‘Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states.’, *Journal of personality and social psychology* 45(3), 513.
- Sutter, Matthias.** (2007). ‘Outcomes versus intentions: On the nature of fair behavior and its development with age’, *Journal of Economic Psychology* 28(1), 69–78.
- Weber, Roberto, Camerer, Colin, Rottenstreich, Yuval and Knez, Marc.** (2001). ‘The illusion of leadership: Misattribution of cause in coordination games’, *Organization science* 12(5), 582–598.
- Wolfers, Justin et al.** (2002), *Are voters rational?: Evidence from gubernatorial elections*, Graduate School of Business, Stanford University.

A Appendix for Online Publication

A.1 Subgame Perfect Nash Equilibrium

The subgame perfect solution to our principal-agent interaction involves no effort or punishment. Because subjects are randomly and anonymously rematched, we can solve each interaction independently. Using backwards induction, we begin at the final subgame with the punishment decision of the principal:

$$\Pi_P = 7 + 6 \times \frac{1+x}{6} - \frac{y}{4}$$

For any value of x , and for any positive, monotonic utility function $u(\Pi_P)$, $y^*(x) = 0$. That is, the optimal punishment is zero. Note that this is independent of any beliefs that the principal may hold about the agent's mean effort level (his "type").

Moving backwards to the effort decision, the agent forecasts that the optimal punishment for the principal is $y^*(x) = 0 \forall x$. Therefore, the agent's payoff function is

$$\Pi_A = 13 - \frac{x}{2} - 0.$$

This is solved with $x^* = 0$ for any positive, monotonic $u(\Pi_a)$. Thus, the SPNE involves no effort and no punishment. This framework does not allow principals or agents to place positive value on information. Likewise, principals and agents should not be willing to pay to avoid this information.

A.2 Example of Bayesian Updating

Suppose types, θ , are believed to be uniformly distributed. That is, $\theta \sim U[0, 4]$ and the PDF of θ is $f(\theta) = 1/4 \forall \theta \in [0, 4]$. Each agent chooses a number of sides, x , in which to invest. The probability distribution associated with each value of x is:

$$\Pr(X = x) = \begin{cases} 1 - |\theta - x| & \text{IF } |\theta - x| \leq 1 \\ 0 & \text{OTHERWISE} \end{cases}$$

We know that the density function for winning conditional on x is given by $f(\text{WIN}|x) = \frac{1+x}{6}$. Now, given the mapping from θ to x , we have that the density function for winning conditional on θ is:

$$f(\text{WIN}|\theta) = \frac{1+x}{6} \times \Pr(x|\theta) + \frac{1+(x+1)}{6} \times \Pr(x+1|\theta) \quad (1)$$

$$f(\text{WIN}|\theta) = \frac{1+x}{6} \times (1-\theta-x) + \frac{1+(x+1)}{6} \times (\theta-x) \quad (2)$$

$$f(\text{WIN}|\theta) = \frac{1+\theta}{6} \quad (3)$$

The unconditional probability of winning is determined by the integral of $f(\text{WIN}|\theta)f(\theta)$ over all possible values of θ :

$$g(\text{WIN}) = \int_0^4 1/4 \times \frac{1+\theta}{6} = 1/2$$

By Bayes Rule, we can calculate the posterior distribution of types conditional on observing $\text{WIN} = 1$ or $\text{WIN} = 0$. These conditional probability distributions are:

$$f(\theta|\text{WIN} = 1) = \frac{f(\text{WIN}|\theta)f(\theta)}{g(\text{WIN})} = \frac{(1+\theta)/24}{1/2} = \frac{1+\theta}{12} \quad (4)$$

$$f(\theta|\text{WIN} = 0) = \frac{(5-\theta)/24}{1/2} = \frac{5-\theta}{12} \quad (5)$$

These distributions can be seen in Figure ??.

A.2.1 Bayesian Inference

Conditional on observing x , Bayesian inference requires that you supplant any information gathered from observing WIN . Thus, the conditional probability distributions become:

$$f(\theta|x) = \frac{f(x|\theta)f(\theta)}{g(x)}$$

For $x \in \{1, 2, 3\}$, $g(x) = 1/4$ and these are piecewise functions. For $x \in \{0, 4\}$, $g(x) = 1/8$ and these are simple linear probability density functions. All Bayesian conditional probability density functions are plotted in Figure 1.

A.2.2 Biased Bayesian Inference

Treating news about x and news about WIN as independent signals leads to a biased conditional probability distribution, since this unnecessarily incorporates luck into the inference about effort. To illustrate, consider our conditional probability distributions after observing x but in the place of the uniform priors

associated with $f(\theta)$ we incorporate the posterior distribution of θ given each outcome:

$$f(\theta|x) = \frac{f(x|\theta)f(\theta|W_{IN} = 1)}{g(x)} \tag{6}$$

$$f(\theta|x) = \frac{f(x|\theta)f(\theta|W_{IN} = 0)}{g(x)} \tag{7}$$

We incorporate the conditional probability distribution, $f(\theta|W_{IN} = 1)$, into the updating process to arrive at the biased posterior distribution in Figure 2. To arrive at the opposite bias, we can similarly incorporate $f(\theta|W_{IN} = 0)$ and find distributions slanted to the left instead of the right.

A.3 Supplementary Tables and Figures

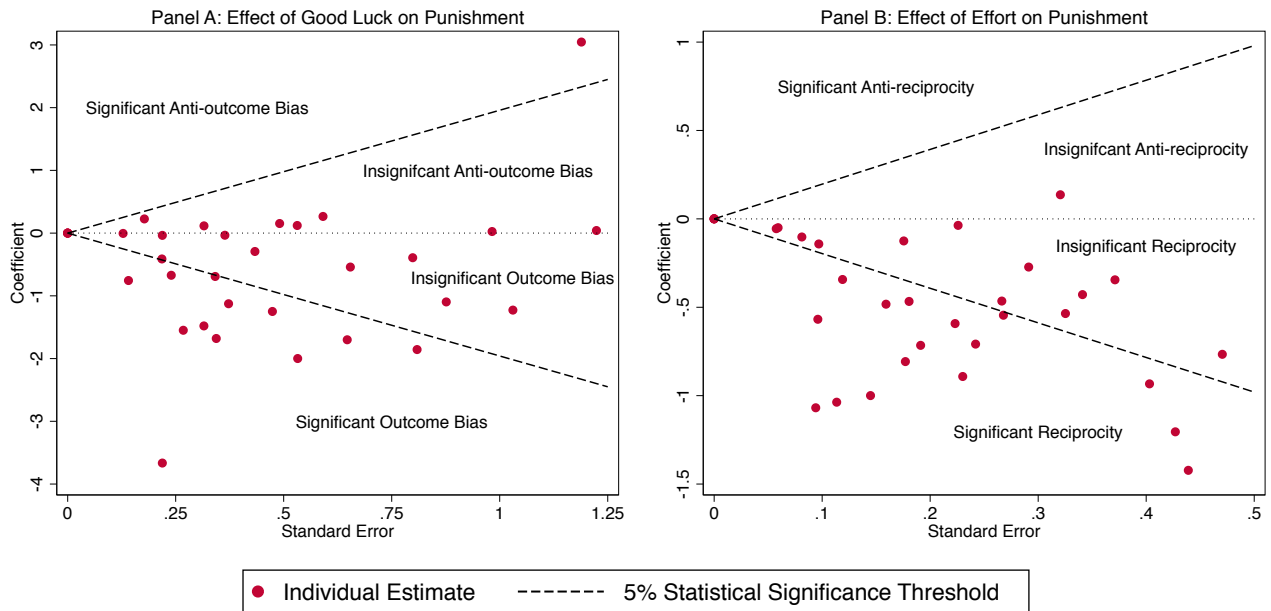


Figure A1: Individual estimates of outcome bias and reciprocity in the *Full* treatment. Estimates obtained from linear regression of punishment on an indicator for good luck and a linear effort variable, separately for each principal in the *Full* treatment. Heteroskedasticity-robust standard errors.

Table A1: Impact of Punishment on Subsequent Effort in the *Full*, *Hidden*, and *Commit* Treatments

Lagged Effort:	Less than Maximum		Maximum
	(1)	(2)	(3)
Constant (Lagged Punishment = 0)	0.54 [†] (0.12)	0.46 [†] (0.13)	2.88 [†] (0.16)
Lagged Punishment	0.12*** (0.03)	0.12*** (0.03)	-0.20** (0.09)
Lagged Effort	0.71*** (0.06)		
Effort-level FEs	N	Y	N/A
Clusters (Principals)	98	98	51
Observations	962	962	171

*** $\Rightarrow p < 0.01$, [†] \Rightarrow In columns (1) and (2) the constant is specific to lagged effort of zero, and in column (3) it is specific to lagged effort of four sides purchased (the maximum). Estimates are from linear random-effects models.

Table A2: Likelihood of Revealing the Outcome to the Principal in the *Force* Treatment

Model:	Linear Random Effects			Probit Random Effects		
	(1)	(2)	(3)	(4)	(5)	(6)
Constant [†]	0.37 (0.08)	0.46 (0.06)	0.37 (0.07)	0.35 (0.10)	0.46 (0.09)	0.33 (0.10)
Win	0.33*** (0.07)	0.38*** (0.07)	0.34*** (0.06)	0.34*** (0.05)	0.39*** (0.05)	0.36*** (0.05)
Effort	0.06** (0.02)		0.06** (0.02)	0.05** (0.02)		0.05** (0.02)
Price		-0.74*** (0.16)	-0.49* (0.28)		-0.82*** (0.15)	-0.47* (0.28)
Effort \times Price			-0.12 (0.12)			-0.19 (0.14)
Clusters (Agents)	33	33	33	33	33	33
Observations	396	396	396	396	396	396

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$, * $\Rightarrow p < 0.10$, [†] \Rightarrow in columns (1) and (4), the constant term represents the likelihood when the principal loses, and effort is zero, in columns (2) and (5), the constant term represents the likelihood when the principal loses, and the price is zero, and in columns (3) and (6), the constant term represents the likelihood when the principal loses, effort is zero, and the price is zero. Standard errors clustered by individual. Marginal effects reported for the Probit model.

Table A3: Effect of Information on
Punishment in the *Force* Treatment over Time

	(1)	(2)
Constant (Outcome Hidden, Lose, Effort = 0, Period = 1)	1.68 (0.29)	1.86 (0.33)
Outcome Shown×Win	-0.51** (0.24)	-0.41* (0.23)
Outcome Shown×Win×Period	0.01 (0.04)	0.01 (0.03)
Outcome Shown×Lose	-0.18 (0.34)	-0.18 (0.33)
Outcome Shown×Lose×Period	0.05 (0.05)	0.07 (0.05)
Win	0.27 (0.18)	0.22 (0.17)
Period	-0.03 (0.03)	-0.03 (0.03)
Effort	-0.35*** (0.08)	
Effort-level FEs	N	Y
Clusters (Principals)	33	33
Observations	396	396

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$, * $\Rightarrow p < 0.10$. Standard errors clustered by individual. Estimates are from linear random-effects models.

Table A4: Impact of Luck and Effort on Principal
Guesses about Agent's Mean Effort in the *Force3* Treatment

	(1)	(2)
Constant (Win = 0, Effort = 0)	1.38 (0.12)	1.40 (0.12)
Win	0.11 (0.08)	0.10 (0.09)
Effort	0.29*** (0.06)	
Effort-level FEs	N	Y
Clusters (Principals)	31	31
Observations	372	372

*** $\Rightarrow p < 0.01$. Estimate are from linear random-effects models.

A.4 Laboratory Instructions

Laboratory subjects saw the following instructions before beginning the study:

Page 1:

Welcome to the Behavioral Business Research Lab study. Thank you for participating.

In today's study, you will make a series of decisions that will determine how much you earn. For this reason, it is very important that you read each set of instructions carefully before proceeding. The study will take roughly 30 minutes and you are guaranteed to receive at least \$5 for your participation today. All payments will be made to you in cash at the end of the study.

Page 2:

The following pages explain the decision process you will engage in every round. You will participate in 13 rounds today. **After each round, you will be randomly and anonymously rematched with a different player. You will NOT be matched to the same player for the whole study.** You will learn about your earnings after each round has ended. Your earnings from a given round can never be negative, so you will always have at least your \$5 Participation Payment.

Only 1 of the 13 rounds you play today will be used as the *Round that Counts*. Only *The Round that Counts* will determine your payments today. After you have completed all the rounds, the computer will randomly select one of the 13 rounds to be *The Round that Counts* with equal probability. Therefore, you should treat each round as *The Round that Counts*.

A.4.1 Principal-specific Instructions

Page 3 (*Full, Hidden, Force, and Commit* treatments):

There are two types of players in today's study: GREEN players and BLUE players. **You are a BLUE player.**

Including the guaranteed \$5 payment, GREEN players have a budget of \$13 and BLUE players have a budget of \$7. A 6-sided die will be rolled (this will be simulated randomly and fairly by the computer) to determine whether YOU win a \$6 prize. YOU will win the prize if one of YOUR winning numbers comes up. The winning numbers are determined by the GREEN players choice.

The GREEN player can use some of their money to buy winning numbers for YOU. The only winning number YOU begin with is 1. For example, if the GREEN player bought two winning numbers, 2 and 3 would be added to the set of winning numbers, increasing YOUR chance of winning from $1/6$ to $3/6 = 1/2$.

The GREEN player can choose to do nothing (buy 0 numbers) or buy up to 4 numbers. Each number costs the GREEN player \$0.50 to buy.

Page 3 (*Force3* treatment):

There are three types of players in today's study: GREEN players, BLUE players, and ORANGE players. **You are a BLUE player.**

Including the guaranteed payment, GREEN players have a budget of \$13, BLUE players have a budget of \$7, and ORANGE players have a budget of \$3.50. A 6-sided die will be rolled (this will be simulated randomly and fairly by the computer) to determine whether YOU win a \$6 prize. YOU will win the prize if one of YOUR winning numbers comes up. The winning numbers are determined by the GREEN player's choice.

The GREEN player can use some of their money to buy winning numbers for YOU. The only winning number YOU begin with is 1. For example, if the GREEN player bought two winning numbers, 2 and 3 would be added to the set of winning numbers, increasing YOUR chance of winning from $1/6$ to $3/6 = 1/2$.

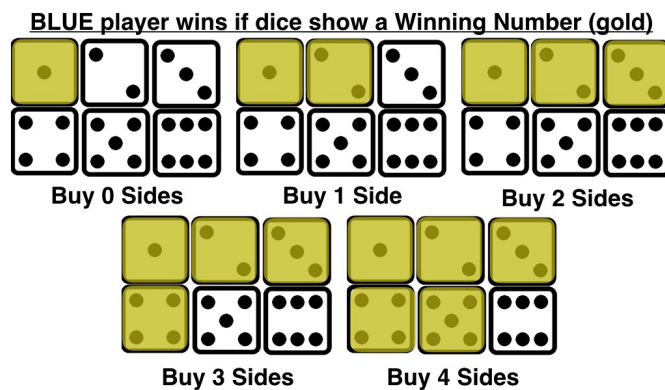
The GREEN player can choose to do nothing (buy 0 numbers) or buy up to 4 numbers. Each number costs the GREEN player \$0.50 to buy.

Page 4:

You are a BLUE Player.

The GREEN player can buy between 0 and 4 winning numbers for YOU. Each number costs the GREEN player \$0.50 to buy. The table and graphic below summarize the GREEN player's options:

Numbers Bought by GREEN Player	0	1	2	3	4
Cost to GREEN Player	\$0	\$0.50	\$1.00	\$1.50	\$2.00
Winning Numbers for BLUE Player	1	1,2	1,2,3	1,2,3,4	1,2,3,4,5
BLUE Player's Chance of Winning	17%	33%	50%	67%	83%



Page 5 (Full treatment):

You are a BLUE player.

YOU will observe how many winning numbers were purchased for YOU by the GREEN player. After YOU observe this, YOU will have an opportunity to punish the GREEN player if YOU choose. **YOU will also observe the outcome of the dice roll before YOU**

make this choice.

The GREEN player knows that YOU will observe their choice of how many winning numbers to buy, but the GREEN player does not have any information about whether YOU will know the outcome of the dice roll when YOU make YOUR choice about punishing the GREEN player.

YOU can punish the GREEN player by reducing the GREEN players earnings by between \$0 and \$4. It costs YOU \$0.25 to reduce the GREEN players earnings by \$1, and YOU can choose any punishment from \$0 to \$4 in \$1 increments. The table below summarizes YOUR options:

GREEN Player's Punishment	Lose \$0	Lose \$1.00	Lose \$2.00	Lose \$3.00	Lose \$4.00
Cost to BLUE Player	\$0	\$0.25	\$0.50	\$0.75	\$1.00

Page 5 (*Hidden* treatment):

You are a BLUE player.

YOU will observe how many winning numbers were purchased for YOU by the GREEN player. After YOU observe this, YOU will have an opportunity to punish the GREEN player if YOU choose. **YOU will not observe the outcome of the dice roll before YOU make this choice.**

The GREEN player knows that YOU will observe their choice of how many winning numbers to buy, but the GREEN player does not have any information about whether YOU will know the outcome of the dice roll when YOU make YOUR choice about punishing the GREEN player.

YOU can punish the GREEN player by reducing the GREEN player's earnings by between \$0 and \$4. It costs YOU \$0.25 to reduce the GREEN players earnings by \$1, and YOU can choose any punishment from \$0 to \$4 in \$1 increments. The table below summarizes YOUR options: (table from Page 5, *Full*)

Page 5 (*Force* treatment):

You are a BLUE player.

YOU will observe how many winning numbers were purchased for YOU by the GREEN player. After YOU observe this, YOU will have an opportunity to punish the GREEN player if YOU choose. **Before YOU make this choice, YOU may or may not observe the dice roll and whether or not you won.**

The GREEN player knows that YOU will observe their choice of how many winning numbers to buy, but the GREEN player does not have any information about whether YOU will know the outcome of the dice roll when YOU make YOUR choice about punishing the

GREEN player.

YOU can punish the GREEN player by reducing the GREEN player's earnings by between \$0 and \$4. It costs YOU \$0.25 to reduce the GREEN players earnings by \$1, and YOU can choose any punishment from \$0 to \$4 in \$1 increments. The table below summarizes YOUR options: (table from Page 5, *Full*)

Page 5 (*Commit* treatment):

You are a BLUE player.

YOU will observe how many winning numbers were purchased for YOU by the GREEN player. After YOU observe this, YOU will have an opportunity to punish the GREEN player if YOU choose. **Before YOU make this decision, YOU will get to decide whether YOU see the dice roll and whether you win or lose.**

The GREEN player knows that YOU will observe their choice of how many winning numbers to buy, but the GREEN player does not have any information about whether YOU will know if you won or lost and what the dice rolled when YOU make YOUR choice about punishing the GREEN player.

YOU can choose to punish the GREEN player by reducing the GREEN players earnings by between \$0 and \$4. It costs YOU \$0.25 to reduce the GREEN players earnings by \$1, and YOU can choose any reduction from \$0 to \$4 in \$1 increments. The table below summarizes YOUR options: (table from Page 5, *Full*)

Page 5 (*Force3* treatment):

You are a BLUE player.

The ORANGE players will observe how many winning numbers were purchased for YOU by the GREEN player. After the ORANGE player observes this, they will have an opportunity to punish the GREEN player if they choose. **Before the ORANGE player makes this choice, they may or may not observe the dice roll and whether or not you won.**

The GREEN player knows that the ORANGE player will observe their choice of how many winning numbers to buy.

The ORANGE player can choose to reduce the GREEN player's earnings by between \$0 and \$4. It costs the ORANGE player \$0.10 to reduce the GREEN player's earnings by \$1, and they can choose any punishment from \$0 to \$4 in \$1 increments. The table below summarizes the ORANGE player's options:

GREEN Player's Punishment	Lose \$0	Lose \$1.00	Lose \$2.00	Lose \$3.00	Lose \$4.00
Cost to ORANGE Player	\$0	\$0.10	\$0.20	\$0.30	\$0.40

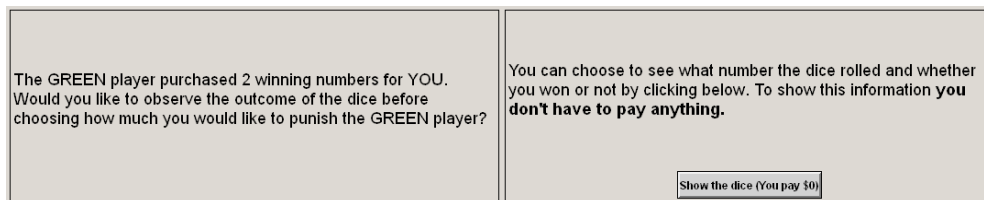
Page 6 (*Commit* treatment):

You are a BLUE player.

After YOU observe the number of winning numbers purchased for you by the GREEN player, YOU will choose whether YOU want to see the outcome of the dice roll before choosing how much you would like to punish the GREEN player.

Specifically, YOU will see a button that says, “Show the dice.” If you click this button, you will observe the number that the dice rolled. If you do not click this button, we will wait to reveal the outcome of the dice roll until after you decide how much YOU want to punish the GREEN player.

Sometimes you will be able to see the outcome of the dice roll for free, sometimes it will cost you \$0.25, and sometimes we will pay you \$0.25 to see the outcome of the dice. Below, is an example of a button you MAY see:



Page 6 (*Force3* treatment):

You are a BLUE player.

Each round, you will be rematched with a randomly chosen GREEN player. For each GREEN player you are matched with, your job is to guess the **average** number of winning sides they will purchase across all 13 rounds of the study.

Because the GREEN player can buy 0 to 4 winning numbers in a round, your guess can be any number from 0.00 to 4.00. The **average** will not necessarily be a whole number, so your guess can have up to two decimal places. **If your guess is within 0.10 of the correct average, you will win a \$1 prize.**

We will not know the average number until the end of all 13 rounds. So, we will wait until the end of the study to tell you if your guess was correct in the Round that Counts. If your guess was correct, we will add \$1 to your earnings.

For example, suppose the GREEN player buys winning numbers, 1, 2, 1, 2, 3, 2, 1, 2, 1, 4, 3, 2, 1. The average is 1.92. If you guessed that the GREEN player's average would be 2.01, you would win \$1. Below, are more *examples* of what would happen depending on your guess and the GREEN player's average number of winning numbers:

GREEN Player's Average Winning Numbers Purchased across All Rounds	0.87	0.87	2.01	2.01	3.23
BLUE Player's Guess	0.50	0.97	1.95	2.12	3.23
BLUE Player's Outcome from Guess	No prize	\$5 Prize	\$5 Prize	No prize	\$5 Prize

A.4.2 Agent-specific Instructions

Page 3 (*Full, Hidden, Force, and Commit* treatments):

There are two types of players in today's study: GREEN players and BLUE players. **You are a GREEN player.**

Including the guaranteed \$5 payment, GREEN players have a budget of \$13 and BLUE players have a budget of \$7. A 6-sided die will be rolled (this will be simulated randomly and fairly by the computer) to determine whether the BLUE player wins a \$6 prize. The BLUE player will win the prize if one of their winning numbers comes up. The winning numbers are determined by YOUR choice.

YOU can use YOUR money to buy winning numbers for the BLUE player. The BLUE player begins with just 1 as a Winning Number. For example, if YOU bought two winning numbers, 2 and 3 would be added to the set of winning numbers, increasing the BLUE player's chance of winning from $1/6$ to $3/6 = 1/2$.

YOU can choose to do nothing (buy 0 numbers) or buy up to 4 numbers. Each number costs YOU \$0.50 to buy.

Page 3 (*Force3* treatment):

There are three types of players in today's study: GREEN players, BLUE players, and ORANGE players. **You are a GREEN player.**

Including the guaranteed payment, GREEN players have a budget of \$13, BLUE players have a budget of \$7, and ORANGE players have a budget of \$3.50. A 6-sided die will be rolled (this will be simulated randomly and fairly by the computer) to determine whether the BLUE player wins a \$6 prize. The BLUE player will win the prize if one of their winning numbers comes up. The winning numbers are determined by YOUR choice.

YOU can use YOUR money to buy winning numbers for the BLUE player. The BLUE player begins with just 1 as a Winning Number. For example, if YOU bought two winning numbers, 2 and 3 would be added to the set of winning numbers, increasing the BLUE player's chance of winning from $1/6$ to $3/6 = 1/2$.

YOU can choose to do nothing (buy 0 numbers) or buy up to 4 numbers. Each number costs YOU \$0.50 to buy.

Page 4:

You are a GREEN Player

YOU can buy between 0 and 4 winning numbers for the BLUE player. Each number costs YOU \$0.50 to buy. The table and graphic below summarize YOUR options: (table and figure from principal instructions, Page 4)

Page 5 (*Full, Hidden, Commit* treatments):

You are a GREEN player.

The BLUE player will observe how many numbers the YOU purchased for them. After the BLUE player observes this, they will have an opportunity to punish YOU if they want. The BLUE player may or may not observe the number that the dice rolled before making this choice.

The BLUE player can choose to punish YOU by reducing YOUR earnings by between \$0 and \$4. It costs the BLUE player \$0.25 to reduce YOUR earnings by \$1, and they can choose any reduction from \$0 to \$4 in \$1 increments. The table below summarizes the BLUE players options: (table from principal instructions, Page 5, *Full*)

Page 5 (*Force* treatment):

You are a GREEN player.

The BLUE player will observe how many numbers the YOU purchased for them. After the BLUE player observes this, they will have an opportunity to punish YOU if they want.

Before the BLUE player makes their choice about YOUR punishment, YOU will get to decide whether the BLUE player sees the dice roll and whether or not they win. We will describe this choice in more detail later.

The BLUE player can choose to punish YOU by reducing YOUR earnings by between \$0 and \$4. It costs the BLUE player \$0.25 to reduce YOUR earnings by \$1, and they can choose any reduction from \$0 to \$4 in \$1 increments. The table below summarizes the BLUE player's options: (table from principal instructions, Page 5, *Full*)

Page 5 (*Force3* treatment):

You are a GREEN player.

The ORANGE players will observe how many winning numbers YOU purchased for the BLUE player. After the ORANGE player observes this, they will have an opportunity to punish the YOU if they want.

Before the ORANGE player makes this choice about YOUR punishment, YOU will get to decide whether the ORANGE player will see the dice roll (and if the BLUE player won \$6). We will describe this choice in more detail later.

The ORANGE player can choose to reduce the YOUR earnings by between \$0 and \$4. It costs the ORANGE player \$0.10 to reduce YOUR earnings by \$1, and they can choose any reduction from \$0 to \$4 in \$1 increments. The table below summarizes the ORANGE player's options: (table from principal instructions, Page 5, *Force3*)

Page 6 (*Force* treatment):

You are a GREEN player.

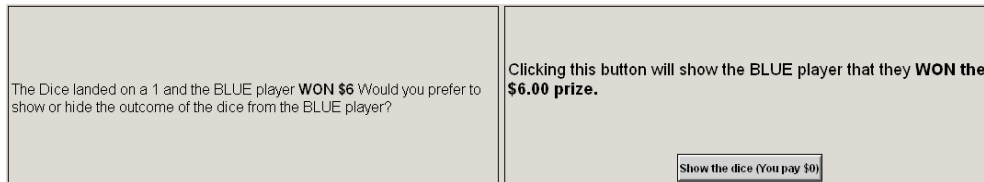
The BLUE player will always know how many winning numbers you purchased for them, but YOU will get to decide whether the BLUE player sees what number the dice rolled before the BLUE player makes their choice about whether or not to punish YOU.

The BLUE player will NOT know that YOU chose whether or not the outcome is revealed. The BLUE player will simply see the outcome without any information about why they see it.

Specifically, YOU will see a button that says, “Show the dice.” If you click this button, the BLUE player will observe the number that the dice rolled and if they won. If you do not click this button, we will wait to reveal the outcome of the dice roll until **after** the BLUE player has decided how much to punish YOU.

Showing what the dice rolled may cost you money, earn you money, or be entirely free. We will tell you how much it will cost you or earn you on the button that says, “Show the Dice.”

Again, sometimes it will be free for you to show the outcome of the dice to the BLUE player. Other times, it will cost you \$0.25 to show the outcome. Sometimes, we will pay YOU and extra \$0.25 to show the outcome of the dice to the BLUE player. Below, you can see one example of a type of question you *MAY* face:



Page 6 (*Force3* treatment):

You are a GREEN player.

The ORANGE player will always know how many winning numbers you purchased for the BLUE player, but YOU will get to decide whether the ORANGE player sees what number the dice rolled before the ORANGE player makes their choice about whether or not to punish YOU.

The ORANGE player will NOT know that YOU chose whether or not the outcome is revealed.

Specifically, YOU will see a button that says, “Show the dice.” If you click this button, the ORANGE player will observe the number that the dice rolled and if the BLUE player won. If you do not click this button, we will wait to reveal the outcome of the dice roll until **after** the ORANGE player has decided how much to punish YOU.

Showing what the dice rolled may cost you money, earn you money, or be entirely free. We will tell you how much it will cost you or earn you on the button that

says, “Show the Dice.”

Again, sometimes it will be free for you to show the outcome of the dice to the ORANGE player. Other times, it will cost you \$0.25 to show the outcome. Sometimes, we will pay YOU and extra \$0.25 to show the outcome of the dice to the ORANGE player. Below, you can see one example of a type of question you *MAY* face:

The Dice landed on a 1 and the BLUE player WON \$6 Would you prefer to show or hide the outcome of the dice from the ORANGE player?	Clicking this button will show the ORANGE player that the BLUE player WON the \$6.00 prize. <input type="button" value="Show the dice (You pay \$0)"/>
--	--

A.5 Laboratory Comprehension Quiz

All laboratory subjects saw the following quiz questions prior to beginning the study. These questions are as read by the green player (the agent). For the blue player (the principal), the questions would be identical, but from the other perspective (i.e. “you” becomes “the GREEN player” and “BLUE” becomes “you”). In the *Force3* treatment, ORANGE is substituted for BLUE in the punishment questions, and the punishment costs are re-scaled.

- If you buy 3 winning numbers for the BLUE player, what is the probability that they win the \$6 prize?
- If you buy 0 winning numbers for the BLUE player, what is the probability that they win the \$6 prize?
- If the BLUE player chooses to spend \$1 of their money to punish you, how much will they reduce your earnings by?
- How much would it cost the BLUE player to reduce your earnings by \$3?
- If the BLUE player does nothing, how much will your earnings be reduced by?
- How often are you randomly rematched with a new partner?
- Which rounds count for your earnings?

In the *Force* and *Force3* treatments, the agent also sees the following questions (with ORANGE substituted for BLUE in *Force3*):

- Who will decide if the BLUE player sees the outcome of the dice before making their choice about punishing YOU?
- Will the BLUE player know that YOU chose to show them the outcome of the dice?

In the *Force3* treatment, the principal also sees the following question:

- How accurately do you need to guess the GREEN player’s average winning numbers if you want to win the extra \$1 prize?

In the *Commit* treatment, the principal also sees the following question:

- Who will decide if you see the outcome of the dice before making YOUR choice about punishing the GREEN player?

A.6 Online Instructions

mTurk subjects saw the following instructions before beginning the study:

Page 1:

Thank you for participating.

In today’s study, you will make a series of decisions that will determine how much you earn. For this reason, it is very important that you read each set of instructions carefully before proceeding. The study will take around 10 minutes and you are guaranteed to receive at least \$3 through Amazon’s Mechanical Turk for your participation today. All additional payments will be made to you as “Bonus Payments” upon verification of your submission.

You will make 16 total “Decisions” today. Each Decision relates to a random and anonymous group of players. Your Decisions will NOT relate to the same players for the whole study.

Only 1 of the 16 Decisions you make will be used as the Decision that Counts. Only The Decision that Counts will determine your payments today. After you have made all Decisions, the computer will randomly select one of the 16 Decisions to be The Decision that Counts with equal probability. Therefore, you should treat each Decision like it is The Decision that Counts.

Page 2:

You will observe data from a series of interactions between BLUE and GREEN players. These are in-person interactions conducted in a research laboratory and are NOT simulated or hypothetical interactions.

After you observe an interaction, you will be given an additional budget of \$0.50 which you can use to punish the GREEN player for his or her actions if you want to.

GREEN players start with \$13. BLUE players start with \$7.

BLUE players also start with a 1/6 chance to win \$6 more: a standard 6-sided dice is rolled and if a 1 is shown, they will win. GREEN players can use their own money to increase the chance that the BLUE player wins. For every \$0.50 the GREEN player spends, the BLUE player gets another winning number, and their chance of winning goes up by 1/6.

For example, if the GREEN player spent \$1.00 of their money to help the BLUE player, they would add 2 and 3 to the set of winning numbers and increase the blue player’s chance

of winning the \$6 prize from $1/6$ to $3/6 = 1/2$.

The GREEN player can choose to do nothing (buy 0 numbers) or buy up to 4 numbers for the BLUE player. The table and graphic below summarize the GREEN player's choice and its consequences: (table and figure from principal instructions, Page 4)

Page 3 (*Full3* treatment):

There are three types of players in today's study: GREEN players, BLUE players, and ORANGE players. **You are an ORANGE player.**

As an ORANGE Player, you will observe a series of interactions between BLUE and GREEN players. Specifically, you will see how many winning numbers the GREEN player purchased for the BLUE player. After you observe this, you will be given an additional budget of \$0.50 which you can use to punish the GREEN player for his or her choice if you want to. **YOU will also observe the outcome of the dice roll before you make this choice.**

All interactions between the GREEN and BLUE players are in-person interactions conducted in a research laboratory and are NOT simulated or hypothetical interactions.

The GREEN player knows that you will observe how many winning numbers they purchased for the BLUE player, but the GREEN player does not have any information about whether you will know the outcome of the dice roll when you make your choice about punishing the GREEN player.

If you determine that the GREEN player should be punished, you can reduce the GREEN player's earnings by between \$0 and \$4. It costs you \$0.10 to reduce the GREEN player's earnings by \$1, and you can choose any punishment from \$0 to \$4 in \$1 increments.

After the computer randomly selects one of your Decisions at the Decision That Counts, we will then apply that punishment decision towards a real, in-person interaction between GREEN and BLUE players. The table below summarizes your options:

GREEN Player's Punishment	Lose \$0	Lose \$1	Lose \$2	Lose \$3	Lose \$4
Cost to YOU	\$0	\$0.1	\$0.2	\$0.3	\$0.4

Page 3 (*Commit3* treatment):

There are three types of players in today's study: GREEN players, BLUE players, and ORANGE players. **You are an ORANGE player.**

As an ORANGE Player, you will observe a series of interactions between BLUE and GREEN players. Specifically, you will see how many winning numbers the GREEN player purchased for the BLUE player. After you observe this, you will be given an additional budget of \$0.50 which you can use to punish the GREEN player for his or her choice if you want to. **Before you make this decision, YOU will get to decide if you want to see the dice roll and if the BLUE player won or lost.**

All interactions between the GREEN and BLUE players are in-person interactions conducted in a research laboratory and are NOT simulated or hypothetical interactions.

The GREEN player knows that you will observe how many winning numbers they purchased for the BLUE player, but the GREEN player does not have any information about whether you will know the outcome of the dice roll when you make your choice about punishing the GREEN player.

If you determine that the GREEN player should be punished, you can reduce the GREEN player’s earnings by between \$0 and \$4. It costs you \$0.10 to reduce the GREEN player’s earnings by \$1, and you can choose any punishment from \$0 to \$4 in \$1 increments.

After the computer randomly selects one of your Decisions at the Decision That Counts, we will then apply that punishment decision towards a real, in-person interaction between GREEN and BLUE players. The table below summarizes your options: (table from online instructions, Page 3, *Full3*)

Page 3 (*Guess3* treatment):

There are three types of players in today’s study: GREEN players, BLUE players, and ORANGE players. **You are an ORANGE player.**

Each BLUE and GREEN player will participate in 13 rounds of interactions but will be randomly matched each round. These interactions are in-person interactions conducted in a research laboratory and are NOT simulated or hypothetical interactions.

As an ORANGE Player, we will show you 16 randomly chosen interactions between BLUE and GREEN players. Specifically, you will see how many winning numbers the GREEN player purchased for the BLUE player.

For each GREEN player, we will calculate how many winning numbers they purchase *on average* across their 13 rounds. We will call this their “Average Investment.”

After you observe the interaction between BLUE and GREEN players, you will be given the chance to earn a \$1 bonus if you can guess the Average Investment for that GREEN player within 0.1. **You will also observe the outcome of the dice roll before you make this guess.**

If you guess the GREEN player’s Average Investment on the Decision that Counts within 0.1, you will earn a \$1 bonus. The table below displays a few example scenarios and their bonus amounts:

GREEN Player's Average Investment	2.25	1.35	3.00	0.85	3.50
Your Guess	1.5	1.25	0.75	0.81	3.00
Bonus Amount	\$0	\$1.00	\$0	\$1.00	\$0

Page 4 (*Commit3* treatment):

After YOU observe how many winning numbers the GREEN player purchased for the BLUE player, you will choose whether you want to see the outcome of the dice roll. Then, you will choose how much you would like to punish the GREEN player.

YOU will see a button that says, “Show the Dice and Outcome.” If you click this button, you will observe the number that the dice rolled and if the BLUE player won. If you don’t click this button, you will find out the number that the dice rolled and if the BLUE player won after your punishment decision.

Sometimes you will be able to see the outcome of the dice roll for free, sometimes it will cost you \$0.10, and sometimes we will pay you \$0.10 to see the outcome of the dice. Below, is an example of a button you MAY see:

A rectangular button with a light gray background and a thin border. The text inside the button is "Show the Dice and Outcomes (You pay \$0.10)" in a dark gray, sans-serif font.

A.7 Online Comprehension Quiz

All mTurk subjects saw the following quiz questions prior to beginning the study.

- How many total Decisions will you make in today’s study?
- How many Decisions will determine your parents in today’s study?
- What determines which Decision is *The Decision that Counts*?
- How many winning numbers does the BLUE player start with?
- How much does it cost the GREEN player to buy 2 more winning numbers for the BLUE player?
- If the GREEN player buys 3 winning numbers for the BLUE player, what is the probability that the BLUE player wins the \$6 prize?

In the *Full3* and *Commit3* treatments, subjects also see the following questions:

- If you pay \$0.20 to punish the GREEN player, how much money will the GREEN player lost?
- How much would it cost you to reduce the GREEN player’s earnings by \$3?

In the *Commit3* treatment, subjects also see the following question:

- If you don’t click “*Show the Dice and Outcome*,” when will you see if the BLUE player won?

In the *Guess3* treatment, subjects also see the following question:

- If the GREEN player has an average investment of 1.10 and your guess is 1.01, will you earn a \$1 bonus?