

Understanding Outcome Bias: Asymmetric Sophistication and Biased Beliefs

Andy Brownback*
University of Arkansas

Michael A. Kuhn†
University of Oregon

April 18, 2018

Abstract

Disentangling effort and luck is critical when judging performance. In a principal-agent experiment, we study how principals' judgments of agent effort are biased by luck, even when *effort is perfectly observed*. We find asymmetric sophistication: agents strategically manipulate this bias by controlling the principal's information about luck, but principals fail to recognize their own bias. Independent third-party punishers are similarly biased, implying the bias cannot be driven solely by emotional response nor by distributional preferences. Luck also affects beliefs: punishers believe lucky agents exert more effort on average than identical, unlucky agents. We propose a model of biased belief updating explaining these results.

JEL Classification: C92, D63, D83

Keywords: Experiment, Punishment, Reciprocity, Outcome Bias, Attribution Bias, Fairness, Luck

*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

†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

Turning observed output into estimates of effort and luck is a crucial, but often nebulous task in many environments. The ambiguity associated with multiple inputs mapping to a single output allows room for stereotypes and biases to impede judgment. For example, Sarsons (2017) finds that referring doctors make different inferences about the skill of specialists depending on the specialist’s gender, *holding outcomes fixed*. A number of studies identify 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.’ In this paper, we explore *why* people exhibit outcome bias, and whether sophisticated individuals can predict and exploit it.

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 paper abstracts away from these practical ways to condition behavior on outcomes. We explore the influence of outcomes on punishment when outcomes possess no signal value. Moreover, our setting involves costly punishment, eliminating the possibility of risk-sharing.

Our baseline treatment replicates prior work in a transparent, principal-agent style interaction in the laboratory. Agents make costly investments to improve the principal’s chance of winning a prize. We call the agent’s investment ‘effort.’ Effort is always observed by the principal. Depending on treatment assignment, principals may or may not observe the outcome of the agent’s effort. Because the outcome is random, conditional on effort, information about the outcome carries no signal value. We refer to this conditional randomness as ‘luck.’ After observing this information, principals can pay to punish agents.

Principals’ decisions in our experiment are influenced by luck *even though the agent’s intentions are perfectly observable*. Conditional on agent effort, principals punish bad luck the same as they would a 1.8 standard deviation (SD) decrease in effort. With this as our starting point, we use a variety of experimental treatments to uncover three novel findings:

Asymmetric Sophistication: Agents predict the outcome-biased behavior of principals and exploit those biases. Agents strategically reveal information about luck prior to the principal’s punishment choice to manipulate the principal’s bias.² They are nearly twice as likely to reveal good luck than bad luck to the principal and 64% are willing to pay to reveal this information. Agents with control over the information significantly reduce their exposure to punishment from bad luck. On the other hand,

¹See Section 2.1 for a review of this literature.

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

principals do not predict their own biases and do not take advantage of commitment opportunities to avoid information that may bias their judgment. 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, and nearly 90% will view their luck if it is costless.

Biased Third Parties: This bias cannot be driven entirely by the principal’s emotion from gaining or losing money, because uninvolved third-party punishers, *who must sacrifice money from their endowment to punish agents*, are also biased by luck. Conditional on the agent’s effort, third-party punishers without a monetary stake in the interaction punish bad luck the same as they would a 1.0 SD decrease in effort.

Biased Beliefs: One potential source of this bias is through the influence of luck on beliefs about the agent’s character. We show that luck affects the beliefs of third-party punishers about the average effort an agent exerts *in interactions they don’t observe*. We used an incentive-compatible elicitation mechanism, asking third-parties to predict the mean effort level of the agent across all of their interactions, past and future. We refer to this mean effort as the agent’s “type.” These predictions are significantly influenced by luck. Bad luck causes third parties to decrease their beliefs about an agent’s type by the same amount as if they observed a 0.26 SD decrease in effort.

While outcome bias itself can potentially be explained by a variety of pathways—emotional responses, distributional preferences, mistaken belief updating, to name a few—our set of findings taken together points in the direction of beliefs. Lucky people are considered to be better intentioned on average than their unlucky counterparts. We suggest a simple model of biased belief updating that can explain our findings. Our model can be considered a contextual application of various theories that manifest similarly in our environment: delegated expertise, attribution bias and correlation neglect, as illustrated by Gurdal et al. (2013), Haggag and Pope (2016), and Enke and Zimmermann (2017) respectively.

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 principals believe luck reveals fixed aspects of the agent’s character that will be exposed in

subsequent interactions.

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)).³ For example, Weber et al. (2001) find 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 character and withhold punishment.

Finally, principals may be *attempting* to punish agents based on their efforts, but misinterpreting signals of that effort because of correlation neglect. 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. If principals in our study exhibit an extreme version of this, they will fail to disregard the outcome, which is highly correlated with the perfect signal of effort.

A potential explanation for both the outcome bias and asymmetric sophistication in our study is that principals hold distributional preferences and reduce agent’s payments 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 could explain some of the behavior of principals, it is inconsistent with behavior of third-parties.⁴ Third parties earn less than half as much money as agents regardless of effort levels or luck. Nonetheless, we observe that 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 the variance between the richest two players. This assertion would also contradict Lefgren et al. (2016), who find that voters prefer redistributive policies that favor unlucky, hardworking subjects over equally poor, lazy ones. Moreover, distributional preferences offer no explanation for why beliefs about the agent’s type would consistently vary with luck. We discuss this further in Section 4.2.

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

³This 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)).

⁴Charness and Rabin (2002) find that costly, Pareto-damaging difference aversion is very 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.

contracting based on outcomes leads to sub-optimal risk sharing and weakens incentives for the agent to exert high effort. Thus, when principals in our study allow luck to influence punishment, they are foregoing the first-best solution. When effort is partially or fully observable, over-reliance on outcomes may decrease welfare from repeated interactions or in societies that benefit from collectively enforced norms of behavior.

Asymmetric sophistication with respect to these biases in punishment 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 politicians’ effort.⁵

Our third-party results indicate that independent arbiters may be just as susceptible to these issues 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.⁶ 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

While our study of outcome bias is unique in economics for its exploration of outcome-bias among third parties, many studies have explored outcome bias among interested parties. It has been documented

⁵This 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. This pattern of behavior is particularly dangerous when biases can be manipulated for a strategic advantage. For example, Gneezy and Imas (2014) find that people are sophisticated when strategically manipulating others’ emotions for their personal gain, and Bartling and Fischbacher (2011) find sophisticated delegation of decision-making in order to shift the attribution of responsibility for the decision.

⁶For 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.

in a variety of real-stakes settings among interested parties by Charness and Levine (2007), Cushman et al. (2009), Gurdal et al. (2013), Rubin and Sheremeta (2015) and de Oliveira et al. (2017). 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 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.⁷

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, under the right conditions, information can also mitigate bias in beliefs about fairness.

The origins of outcome bias lie in the psychology literature. In particular, Baron and Hershey (1988)

⁷A related literature explores how context and stereotypes impact how outcomes are mapped to luck and skill (Eil and Rao, 2011; Sarsons, 2017; Erkal et al., 2017).

coin the phrase ‘outcome bias’ and identify it in a variety of hypothetical circumstances. In two of their experiments (Experiments 3 and 4), they assure subjects that outcomes are not potentially valuable signals, similar to the design of the recent economic experiments. In both Experiments, subjects are uninterested parties offering their stated preference on the decision-making quality of doctors and risk-takers. Related to our study of third-party beliefs, they find that stated beliefs about future competence of doctors are colored by outcome information.

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 overweight outcomes relative to other performance information available. While these observational studies of outcome bias are crucial to ensure the external validity of work on outcome bias, we use the lab to understand a variety of 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):** Undergraduates at the University of Arkansas (UArk) and the University of California, San Diego (UCSD) participated in lab sessions with principals making punishment decisions.
- **Wave 2 (May 2017):** Incentive-compatible elicitation of third-party punishment decisions from subjects recruited through Mechanical Turk (MTurk). Using a modified strategy method, we presented a variety of potential principal-agent interactions and implemented punishment decisions on the following wave.
- **Wave 3 (July and August 2017):** lab sessions with principals and agents interacting at UArk. Punishment of agents was determined using data from Wave 2.
- **Wave 4 (August 2017):** MTurk survey of third-party observers with incentive-compatible belief elicitation about agent type. Third-parties were shown data from agents who had previously completed lab sessions.

3.1 Principal-Agent Environments

For Wave 1, we collected responses from 266 subjects in 16 different sessions. All treatment variations were run using zTree software (Fischbacher (2007)) at both UArk and UCSD. 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.⁸

In every treatment, the principal is endowed with \$7. She can win additional money depending on the 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). Thus, the principal always has a chance to win, but cannot be guaranteed to win. The principal always learns of the agent’s investment choice. Depending on the treatment, the principal may or may not learn about the outcome of the random process 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.⁹

“Full” treatment: In our baseline treatment, the principal is given complete information about the effort investment of the agent and the outcome of the die roll. After receiving this information, the principal decides whether to reduce the agent’s payoffs, and if so, by how much. The agent is told that the principal “may or may not observe the outcome of the die roll.”

“Intent” treatment: In the Intent treatment, the principal is not made aware of the outcome of the

⁸Instructions, comprehension quizzes, and survey questions can be found in Appendix Section A.3.

⁹At 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.

die roll until after her punishment decision is made. To match the Full treatment, the agent is told that the principal “may or may not observe the outcome of the die roll.”

“Force” treatment: The Force treatment introduces a new decision stage for the agent. After the agent makes his investment decision but before the principal makes her punishment decision, the agent observes the outcome of the die roll and whether the principal has won the \$6 prize. The agent then chooses whether to pay p_1 to reveal the outcome to the principal. The price, p_1 , is drawn from the set $\{-\$0.25, \$0, \$0.25\}$, with probabilities $\{\frac{1}{4}, \frac{1}{2}, \frac{1}{4}\}$, respectively. That is, the agent may have to pay to show the outcome, may be paid to show the outcome, or may show the outcome for free. The principal is informed that she “may or may not observe the the dice roll and whether or not [she] won” prior to making her punishment decision. There is no reference to the agent’s control over the principal’s information. Importantly, the agent knows that the principal will not be informed of the agent’s control over the information.¹⁰

“Commit” treatment: The Commit treatment transfers control over the principal’s information to the principal herself. Similar to the baseline treatment, the principal observes the agent’s investment level but is now given the opportunity to pay p_2 to view the outcome of the die roll. The price, p_2 , is drawn from $\{-\$0.25, \$0, \$0.25\}$ with probabilities $\{\frac{1}{4}, \frac{1}{2}, \frac{1}{4}\}$, respectively. Agents are told that the principal “may or may not observe the number that the dice rolled,” to match both Full and Intent. The agent is not made aware of the principal’s control over this information. The principal is told that the agent has not been made aware of who controls the principal’s information.

Payoffs: The agent’s payoff in the Full, Intent, 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:

$$\Pi_A = 13 - \frac{x}{2} - y - p_1 \times \text{SHOW} ,$$

¹⁰While 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 A.3, we analyze the first and last 6-periods separately to show that the effect sizes are not changed with experience.

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

Conditional on the outcome, the principal’s payoff in the Full, Intent and Force treatments is given by:¹¹

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

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 becomes:

$$\Pi_P = 7 + 6 \times \text{WIN} - \frac{y}{4} - p_2 \times \text{VIEW},$$

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

3.2 Third Party Punishment Environments

For 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 these third parties as “Orange” players. Third parties viewed 16 different possible combinations of agent effort and principal outcomes and determined the amount of punishment they wished to impose. These combinations appeared alongside randomly generated rolls of the die.¹² 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.¹³

In Wave 3, we recruited new principals and agents to interact at UArk. 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.¹⁴ Instead, we probabilistically applied the punishment decisions from the third parties to the agents based on their effort level and luck.

¹¹Prior to receiving information about the outcome, expected payoff is a more appropriate measure for the principal. This is characterized by: $E(\Pi_P) = 7 + 6 \times \frac{1+x}{6} - \frac{y}{4}$.

¹²The 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.

¹³Principals were given \$7 in endowment and received approximately \$3 in expected earnings from the roll of the dice (mean investment 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.

¹⁴We asked principals to make guesses just like we will later detail for the third parties. The results are presented in appendix Table A.4.

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

“Commit” treatment: Prior to punishment, third parties have the opportunity to pay p_3 to view the principal’s outcome. We draw prices, p_3 , from $\{-0.10, 0, 0.10\}$ with probabilities $\{\frac{5}{16}, \frac{6}{16}, \frac{5}{16}\}$, respectively. With positive probability, the third party punishment choice is applied to a lab interaction with the same characteristics.

“Force” treatment: Agents in Wave 3—the subsequent, in-person 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 have about the principal’s outcome. Just like the principal-agent environment, the agent can pay p_1 to show the principal’s outcome to the third party. p_1 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 show the outcome to the third party, they are matched to a punishment decision featuring the same effort and outcome from a third party who observed the outcome of the interaction. If the agent chooses not to show the third party the outcome, they are matched to a punishment decision featuring the same effort and outcome from a third party who did not observe the outcome of the interaction.¹⁵

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

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

where y again represents the punishment choice.

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

A third party’s payoff in the Commit 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 if the principal won or not.

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

$$\Pi_A = 13 - \frac{x}{2} - y - p_1 \times \text{SHOW} ,$$

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

3.3 Eliciting Beliefs about Agents

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 outcomes. 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 investment choice of the agent within 0.10. For one randomly selected round, we compare the third party’s guess with the actual mean investment 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 investment.

4 Hypotheses

The subgame perfect solution to our principal-agent interaction involves no investment or punishment.¹⁷

However, we observe substantial punishment and investment, consistent with results on reciprocity such

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 investment across all rounds 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 for payment, they received a \$5 bonus. We analyze these data in the same way as the third-party guesses in appendix Table A.4.

¹⁷See Appendix Section A.1 for more detail.

as the “gift exchange” (Akerlof, 1982). This can be explained by a model of reciprocity, altruistic enforcement of efficient investment as a social norm, or distributional preferences. 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 the desire to enforce norms, 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 parties (third parties) are consistent with this motivation. In this section, we explore a model of belief updating that can explain the role of luck in punishment decisions that derive from norm-enforcement.

4.1 Belief Updating

Principals and third parties must estimate how deserving of punishment each agent is, perhaps due to their lack of altruism or poor adherence with norms. This can be inferred based on an agent’s investment choice, $x \in \{0, 1, 2, 3, 4\}$. Let \bar{x} represent the agent’s mean investment level (his “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 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 investment choice. Information about the outcome of the die—a noisy signal of the agent’s investment choice—is completely supplanted 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 investment choices between the two nearest integer values with probabilities such that their average investment equals \bar{x} .¹⁹ For example, an agent with $\bar{x} = 2.25$ will make investment 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 investment choices throughout the experiment.

Figure 1 plots the specific posterior beliefs generated after observing each possible value of x . Notice that the censoring at the upper and lower bounds cause the taller peaks of the probability distributions for $x = 0$ and $x = 4$.

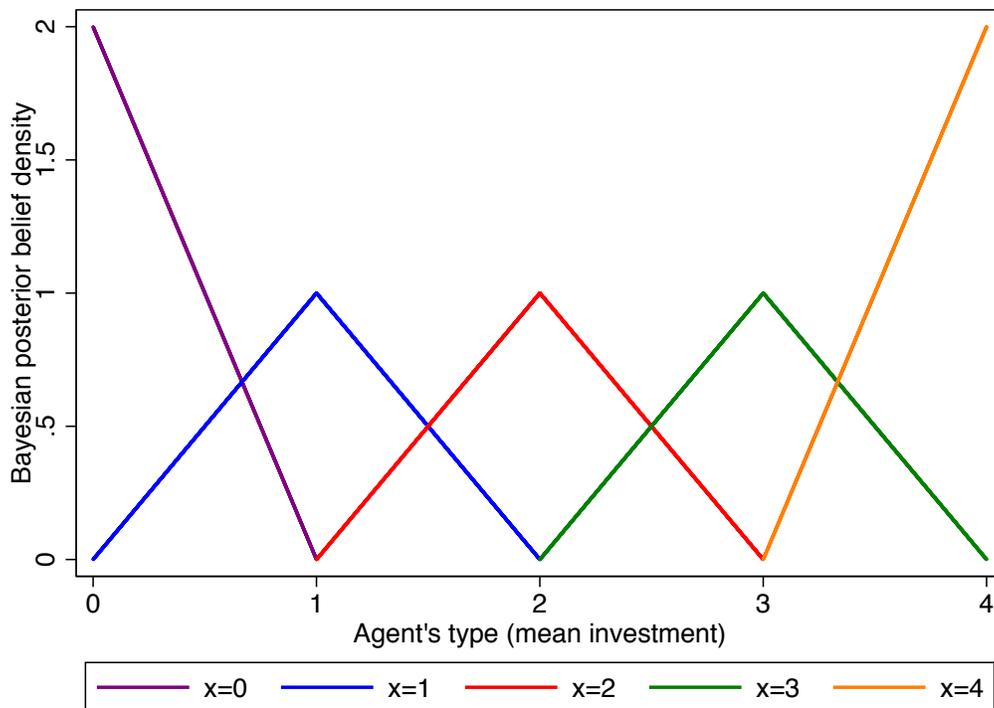


Figure 1: Posterior belief distribution, \hat{x} , conditional on observed investment, x .

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.

Rational belief updating not only eliminates any punishment based on luck, it also eliminates any benefit to control of information about the outcome of the die. 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) on outcomes to choice variables (like effort) (Haggag and Pope, 2016). We incorporate these possibilities by deriving \hat{x} for principals who believe the outcome and x represent independent signals of \bar{x} . This causes \hat{x} to vary systematically with luck: conditional on good luck, \hat{x} stochastically dominates \hat{x} conditional on bad luck, and the likelihood of punishment adjusts accordingly. Figure 2 shows how uniform priors will be updated into the biased posterior beliefs after

observing good and bad outcomes if outcomes and x are believed to be independent.²⁰

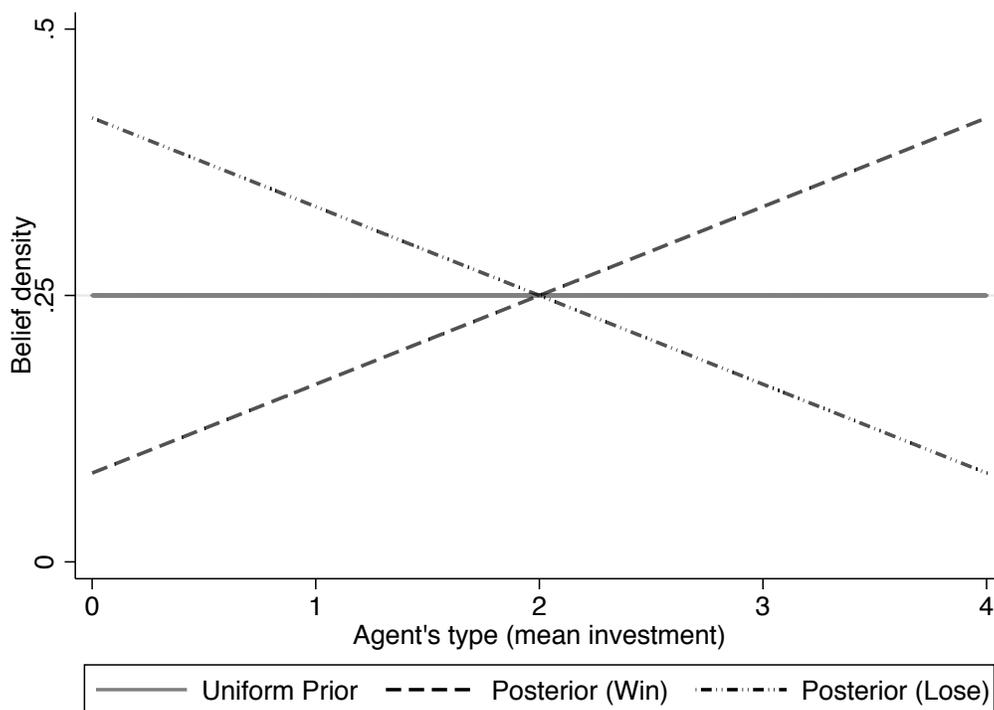


Figure 2: Uniform prior along with posterior belief distributions. Posteriors are conditional good and bad luck, respectively.

After updating \hat{x} based on the outcome, the principal observes x and updates \hat{x} again. The now-non-uniform distribution of beliefs skews the posterior belief distributions associated with each value of x . This skewness can be seen in Figure 3, which plots the posterior belief distributions for each possible value of x for a principal that had observed good luck. Each distribution is now skewed to the right reflecting the higher beliefs after observing good luck.²¹

Conditional on x , any attribution of the outcome to the agent's type produces a bias in \hat{x} . Principals who punish based on the agent's type will, accordingly, be biased by news about outcomes. The behavior predicted by this model is consistent with any form outcome bias, but we will be able to distinguish this bias from other possible mechanisms like emotional responses or distributional preferences by eliciting interim beliefs, \hat{x} .

Under biased belief updating, principals wanting to punish accurately can place a positive value on

²⁰The 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.

²¹The ordering of the signals is unimportant for generating these 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 3.

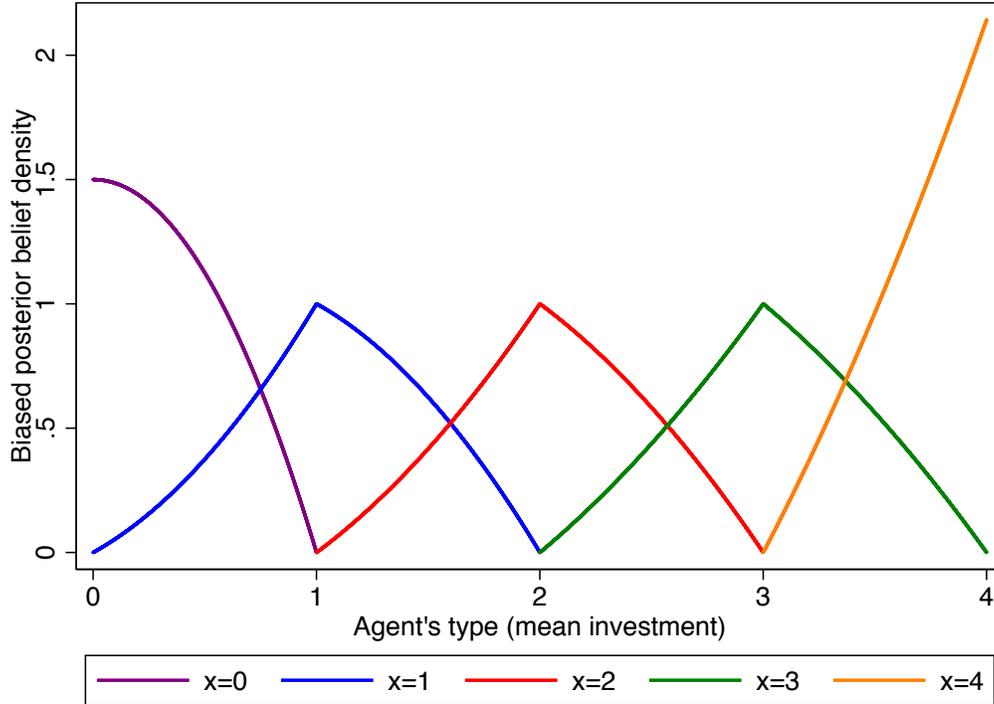


Figure 3: Posterior belief distribution conditional on x after observing good luck.

information about outcomes for a variety of reasons. Agents unambiguously benefit from control over what information is revealed to principals or third parties. Agents should be willing to pay to reveal positive outcomes and should be willing to forgo payments to keep negative outcomes hidden. Hence, such a model can predict a pattern of asymmetric sophistication. Appendix Section A.2 offers an example and more details on Bayesian and biased belief updating in our study.

4.2 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 investment and punishment in our study because knowledge of relative wealth is 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 motives respond to the revelation of the conditionally random outcome.

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 if effort is very low. 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 investment and luck, neither of which influence the third-parties' relative wealth ranking. Punishment by third parties would therefore have to be concerned with the variance only between the two richer players rather than the variance 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 believe updating from the previous section.

5 Principal-Agent Environment Results

We begin our results with the principal punishment environment. First, we explore the Full and Intent treatments to replicate prior results on outcome bias and blame. We then use the Force and Commit treatments to test for sophistication about and exploitation of this bias among principals and agents.

5.1 Full Treatment

The Full treatment presents principals with all available information prior to their punishment decisions. The influence of luck on the punishment amount is clearly visible, exposing potential outcome bias from the principals. We regress punishment on indicator variables for each level of agent investment and predict residual punishment—that is, punishment unexplained by investment. We plot residual punishment over the course of the session separately for good and bad luck in Figure 4. The impact of luck is strong (the standard deviation of punishment is about \$1.43 in this sample) and consistent throughout the session, only tapering off in the final period.

Table 1 presents linear random effects and Tobit models to estimate the impact of good luck on the principal's punishment decision.²² Standard errors are clustered at the subject (principal) level. We include a linear investment variable to control for the agent's choices.²³ Columns (2) and (4) allow the effect of increased investment to depend on the principal winning the prize. Holding fixed the outcome, all specifications indicate that higher investment is associated with a significant decrease in punishment. Thus, punishment does reflect some measure of reciprocity. However, the coefficient on winning the prize

²²The Tobit model adjusts for censoring at \$0 and \$4 of punishment.

²³Table A.1 in the Appendix presents a similar analysis showing that these results are robust to using indicator variables for each investment level rather than a linear control variable.

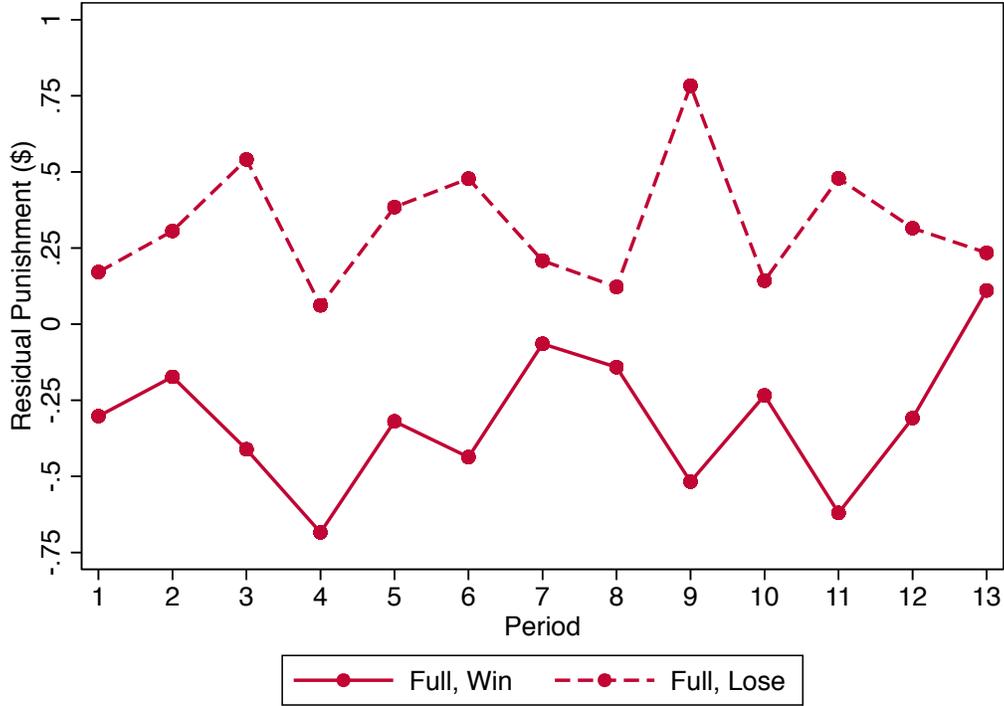


Figure 4: Average punishment as a function of the principal’s outcome in the Full treatment.

Note: residual punishment obtained from OLS regression of punishment on indicator variables for each level of investment.

is large and significant in all specifications, revealing strong outcome bias. For example, principals who win after the agent purchases two sides of the die punish 55% less than those who lose at the same investment level. Table A.1 in the appendix shows that the impact of winning on punishment is large and negative across all investment levels, though the effect of the random outcome is largest when the agent invests in just one side.²⁴

5.2 Intent Treatment

The Intent treatment represents a counterfactual world to the Full treatment where principals do not have access to information about the bias-generating random outcomes until after their punishment decisions are made. Without information about luck, we expect that punishment will be much more responsive to the investment of the agent. We again estimate punishment as a linear function of investment and the random outcome. Table 2 compares punishment choices in this treatment with those from the Full treatment using both random effects and Tobit specifications. Each coefficient is interacted with an

²⁴This is roughly consistent with the delegated-expertise model of Gurdal et al. (2013), which predicts that the impact of luck is decreasing in the the agent’s investment level.

Table 1: Impact of Luck and Investment on Punishment in Full Treatment

Model:	Random Effects		Tobit	
	(1)	(2)	(3)	(4)
Investment	-0.36*** (0.08)	-0.44*** (0.11)	-1.06*** (0.31)	-1.06** (0.41)
Win	-0.67*** (0.17)	-0.96*** (0.30)	-3.20*** (0.94)	-3.18*** (1.19)
Investment \times Win		0.14 (0.32)		-0.01 (0.50)
Constant	1.99 (0.29)	2.10 (0.32)	1.99 (0.94)	1.98 (1.06)
Principals	33	33	33	33
Observations	429	429	429	429

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$. Standard errors clustered by individual. The Tobit model adjusts for censoring at \$0 and \$4 punishment.

indicator variable for the Intent treatment, leaving the Full treatment as the omitted category. Columns (1) and (3) constrain the impact of winning to be constant across investment levels, and columns (2) and (4) allow it to vary with investment. In the Intent treatment, punishment is between 73% and 106% more responsive to investment than in Full.

With the increase in sensitivity to effort over luck, agents best-respond by decreasing their rate of zero investment over time. Table 3 captures this behavioral dynamic between the Full and Intent treatments: the rates of zero investment are the same in period one, but the gap grows significantly with time. The rate of zero investment in the 13th period (which always has full information regardless of treatment), is 14% lower in Intent than Full ($p = 0.05$, discrete effect from Probit model). The mean investment in period 13 is \$0.18 higher in Intent, but this is not statistically significant ($p = 0.25$, OLS model).

Table 2: Sensitivity of Punishment to Investment and Outcomes by Treatment

Model:	Random Effects		Tobit	
	(1)	(2)	(3)	(4)
Intent	0.16 (0.39)	0.28 (0.45)	0.88 (1.04)	0.98 (1.25)
Investment	-0.33*** (0.08)	-0.44*** (0.12)	-0.84*** (0.26)	-0.92** (0.37)
Intent \times Investment	-0.24** (0.12)	-0.32* (0.17)	-0.89** (0.44)	-0.96 (0.59)
Win	-0.71*** (0.18)	-1.15*** (0.32)	-2.87*** (0.81)	-3.25*** (1.06)
Intent \times Win	0.77*** (0.20)	0.65 (0.40)	2.78*** (0.86)	2.69** (1.34)
Investment \times Win		0.22 (0.13)		0.19 (0.47)
Intent \times Investment \times Win		0.09 (0.17)		0.07 (0.63)
Constant	1.94 (0.29)	2.11 (0.33)	1.90 (0.77)	2.02 (0.89)
Principals	66	66	66	66
Observations	792	792	792	792

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$, * $\Rightarrow p < 0.10$. Standard errors clustered by individual. The Tobit model adjusts for censoring at \$0 and \$4 of punishment.

Table 3: Likelihood of Zero Investment over Time by Treatment

Model:	Random Effects	Probit
	(1)	(2)
Intent	-0.0003 (0.0751)	0.0039 (0.0642)
Period	0.0003 (0.0042)	0.0003 (0.0035)
Intent \times Period	-0.0097* (0.0055)	-0.0110** (0.0052)
Constant	0.1658 (0.0593)	0.1659 (0.0590)
Agents	66	66
Observations	858	858

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$. Standard errors clustered by individual. Marginal effects reported for the Probit model.

5.3 Force Treatment

Having established that principals in the Full treatment exhibit meaningful outcome bias and a lower sensitivity to the agent’s investment levels, we now address the agent’s ability to forecast and manipulate this bias. We examine whether agents in the Force treatment can strategically withhold or reveal information to lessen their exposure to bad luck and maximize their gains from good luck.

Figure 5, Panel A shows the rates of information revelation by agent’s investment level. For each level of investment, agents are more likely to reveal good luck than bad luck with good luck being revealed 30 percentage points more often, on average. Table 4 estimates the impact of good luck on the likelihood of revealing information, conditional on investment. Columns (2) and (4) demonstrate that this effect size is not statistically different across investment levels in the linear specification, although Figure 5 suggests it may be largest for high investment.

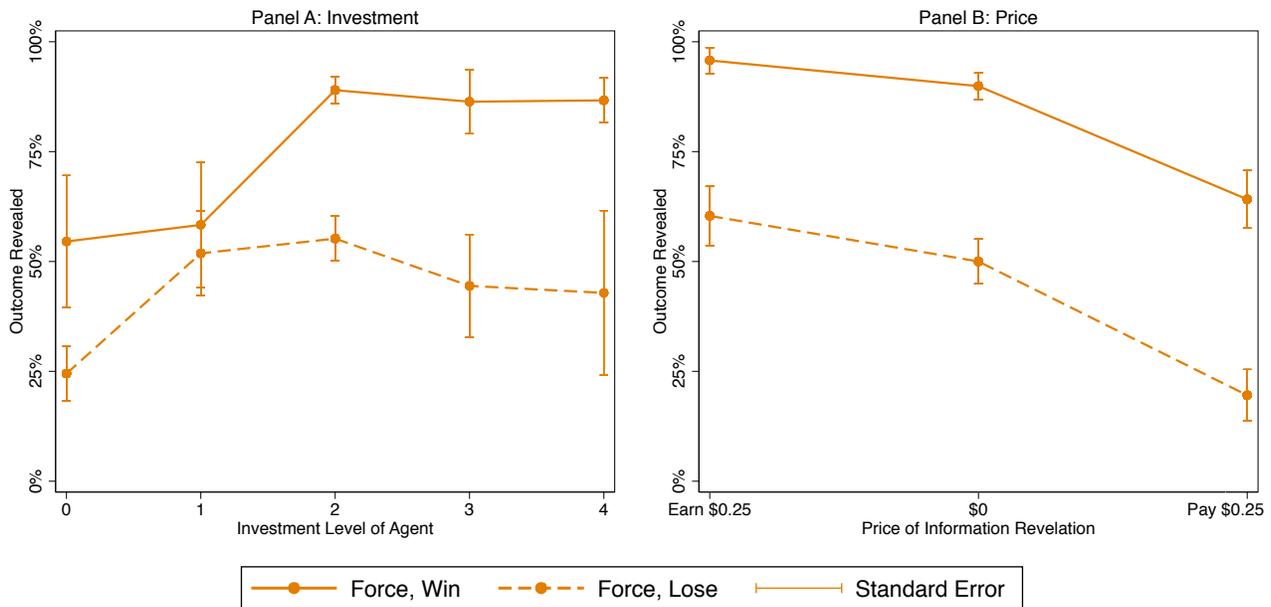


Figure 5: Agent revelation behavior by investment level and price in the Force treatment.

These results suggest that agents are strategic about information revelation, though it does not reveal the value they place on control over that information. To determine this, we plot the demand curve for information revelation among agents separately for winning and losing outcomes in Figure 5, Panel B. We find that 40% of the time, agents are willing to forgo \$0.25 to withhold outcome information from a principal who has lost. Additionally, 64% of the time, agents are willing to pay \$0.25 to reveal that the

Table 4: Likelihood of Revealing the Outcome to the Principal

Model:	Random Effects		Probit	
	(1)	(2)	(3)	(4)
Investment	0.061** (0.024)	0.052 (0.034)	0.076** (0.033)	0.076* (0.045)
Win	0.329*** (0.067)	0.292** (0.124)	0.336*** (0.067)	0.335*** (0.129)
Investment \times Win		0.019 (0.054)		< 0.001 (0.061)
Constant	0.366 (0.077)	0.379 (0.082)	0.331 (0.086)	0.331 (0.098)
Agents	33	33	33	33
Observations	396	396	396	396

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$, * $\Rightarrow p < 0.10$. Standard errors clustered by individual. Marginal effects reported for the Probit model.

principal has won.²⁵ While agents are willing to pay for control over this random information, they do not perfectly leverage this control. Agents only withhold losing outcomes about half of the time when the the revelation decision is costless.²⁶

In a world of rational agents, positive expenditure on information control is strictly welfare decreasing. Information about outcomes is useless, so money spent on its control is wasted. Thus, positive expenditures on information control demonstrate the perception amongst agents that there is value in manipulating the principal’s information. This perception is supported by our comparison between the Intent and Full treatments—there should be value to be gained from hiding bad outcomes and revealing good outcomes. One difference between the impact of exogenous outcome revelation across the Full and Intent treatments and endogenous revelation *within* the Force treatment is that there is within-principal variation in whether the outcome is seen in the Force treatment. While the principals do not know why they sometimes see outcome and sometimes do not, they may suspect the agent’s role, or may treat the information differently. Table A.2 in the Appendix confirms that this change does not eliminate the value of manipulating information in the Force treatment.²⁷ We find a large gain associated with revealing

²⁵The distribution of individual-level mean behavior after winning stochastically dominates the distribution after losing (Kolmogorov-Smirnov test $p < 0.001$), though both distributions span from 0 – 1. 48% of agents always show winning outcomes and 21% never show losing outcomes.

²⁶In appendix Table A.3, we show that the effect of control over information does not diminish with time. However, depending on the agent’s beliefs about how principals interpret missing information, they may be indifferent between revealing and withholding the outcome. For example, if agents believe that principals think missing information is equivalent to a loss, then revelation is irrelevant.

²⁷Conditional on investment and outcome, whether a principal sees the outcome is random—it is simply a function of the random revelation price and agent characteristics which cannot influence the principal’s punishment choice. Therefore, we

wins and a less robust gain from withholding losses, especially at low investment levels.²⁸

Additionally, we estimate the overall impact of information control in reducing punishment variation between good and bad outcomes. Less random variation in punishment is welfare increasing for any risk-averse agent and thus reflects added power for agents. To measure the reduction in risk exposure, we compare the influence of the outcome on punishment between Full and Force in Table 5. Controlling for the agent's investment decision, luck has less influence on the agent's punishment in the Force treatment. Column (1) shows that the effect of winning on Punishment is 73% smaller in the Force treatment than in the Full treatment. This difference is significant at the 5% level, but it is not significant in other specifications due to larger standard errors.

Though these results demonstrate a sophistication about 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 the outcome would decrease the principal's selected punishment.²⁹ When asked what motivated them to hide the outcome of the dice, 58% indicated that they did so because they thought the outcome would increase the principal's selected punishment.³⁰ These responses suggest that most agents have a keen awareness of the disproportionate impact of outcomes even when investment levels are observed.

can identify the causal impact of showing the outcome to the principal on the punishment decision even though information revelation is endogenous to the agent. Alternatively, we can use the price 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 instrumental variables yields a local average treatment effect at the margin where revelation is influenced by its price, we prefer the estimates in Table A.2.

²⁸Considering costly punishment for low investment 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)).

²⁹These 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.

³⁰These individuals either responded that they hid the outcomes when they had invested little and did not want the principal to know they had lost, or that it was to hide bad luck when they invested heavily (or both). 12% of agents never hid the outcome, and 9% failed to respond.

Table 5: Sensitivity of Punishment to Investment and Outcomes by Treatment

Model:	Random Effects		Tobit	
	(1)	(2)	(3)	(4)
Force	-0.39 (0.39)	-0.42 (0.45)	-0.68 (1.31)	-0.28 (1.51)
Investment	-0.34*** (0.08)	-0.44*** (0.12)	-1.01*** (0.31)	-1.09** (0.44)
Force×Investment	-0.03 (0.11)	-0.01 (0.16)	-0.46 (0.54)	-0.75 (0.64)
Win	-0.71*** (0.18)	-1.14*** (0.32)	-3.48*** (1.00)	-3.87*** (1.29)
Force×Win	0.52** (0.22)	0.64 (0.44)	1.89 (1.17)	0.83 (1.78)
Investment×Win		0.21 (0.13)		0.20 (0.57)
Force×Investment×Win		-0.05 (0.17)		0.60 (0.80)
Constant	1.95 (0.29)	2.11 (0.33)	1.91 (0.93)	2.03 (1.07)
$H_0: \text{Win} + \text{Force} \times \text{Win} = 0$	$\chi^2(1) = 1.70$ $p = 0.19$	$\chi^2(1) = 2.67$ $p = 0.10$	$\chi^2(1) = 4.69$ $p = 0.03$	$\chi^2(1) = 5.27$ $p = 0.02$
Principals	66	66	66	66
Observations	792	792	792	792

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$. Standard errors clustered by individual. The Tobit model adjusts for censoring at \$0 and \$4 of punishment.

5.4 Commit Treatment

In the Commit treatment, principals learn of the agent’s investment level and then are given the opportunity to selectively reveal the outcome of the die—whether they win or lose—prior to making a 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 sophisticated principals to commit against outcome bias. Leveraging the variation in the price of information revelation, we can measure this demand for commitment. Figure 6, Panel B plots the demand curve for information over outcomes among principals. Principals choose to view the information 71% of the time, on average, and even pay to view the information 30% of the time. Surprisingly, demand for information is non-monotonic in prices—principals are less likely to view the information when they are paid to do so

than when it is free.³¹ Both the willingness to pay to see the random outcome and the non-monotonicity cast doubt on the sophistication of principals about the decision to reveal the information.

One motivation for revealing the outcome before punishment could be an overwhelming curiosity preventing the principals from waiting until the next screen to observe the outcome. To test this hypothesis, we consider the relationship between demand for information and the amount of uncertainty over the outcome at a given investment level. Curious principals should reveal information most frequently when uncertainty is highest—when agents choose intermediate investment levels. Figure 6, Panel A shows exactly the opposite: information revelation is highest when investment is $x = 0$ and $x = 4$ and uncertainty over the outcomes is minimized. Regressing an indicator for outcome revealed on an indicator for maximum or minimum investment yields a marginal effect on the probability of 0.11 ($p = 0.064$).³² This behavior is inconsistent with any sophistication from the principals about their own outcome bias. Rather, this behavior is more consistent with principals *wanting* to condition their punishments on luck.

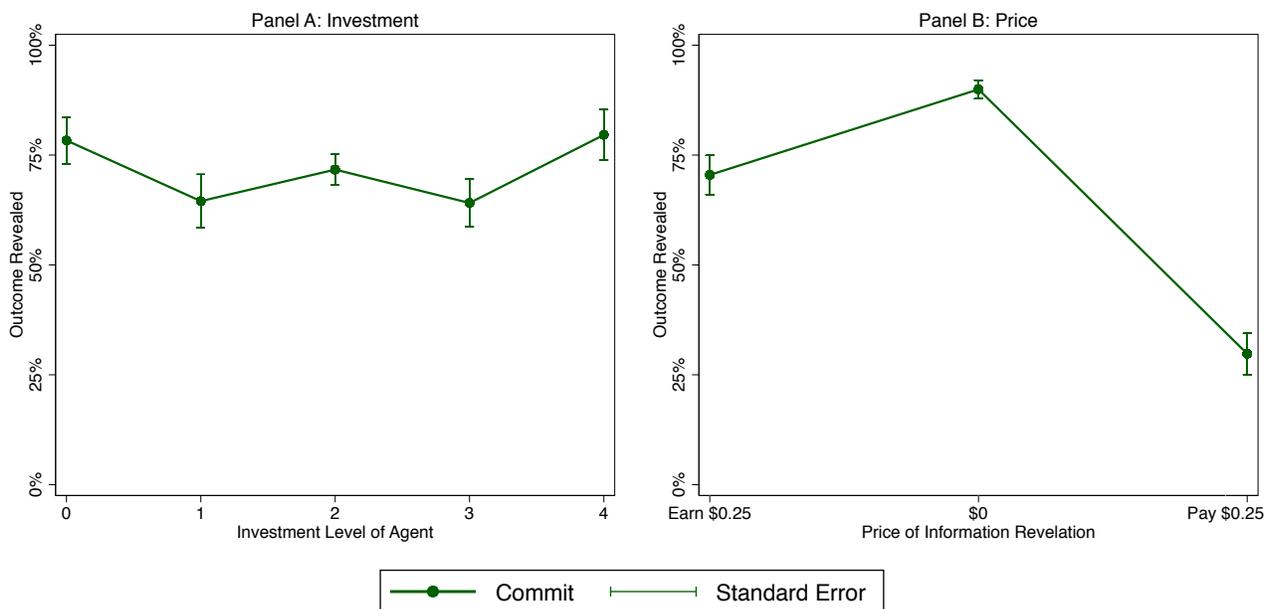


Figure 6: Principal revelation behavior by investment level and price in the Commit treatment.

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 behavior of agents. We compare how sensitive punishment decisions are to winning between the

³¹The non-monotonicity between a price of $-\$0.25$ and $\$0$ is statistically significant at a 5% level. This estimate is from a probit model with standard errors clustered by individual.

³²Estimated using a probit regression with standard errors clustered by individual (S.E. = 0.06). Differences in outcome viewing rates between non-extreme investment levels are not statistically significant.

Commit and Full treatments.³³ Table 6 shows that punishment decisions are weakly less sensitive to luck in Commit compared to Full. Punishment in the Commit treatment—as in the Full treatment—is significantly lower for principals who won than for those who lost, conditional on investment. However, punishment in the Commit treatment is also less sensitive to the agent’s *investment* than in Full. The relationship between agent investment on the principal’s punishment is not statistically significant in any specification in Intent. Giving principals control over their information set appears to have weakened their ability (or desire) to respond to agent effort.

To assess for the conscious sophistication of the principals, we asked their motivations behind avoiding the outcome of the dice prior to making their punishment decisions. Only 32% indicated that they did so to avoid punishing an agent who did not deserve to be punished, and only only 12% (a strict subset of the previous group) indicated that did so to make sure they punished deserving agents. This suggests that relatively few principals understood, or cared, that luck would influence their choice.³⁴

³³We can test for sensitivity to information entirely within the Commit treatment by using the randomly assigned price to reveal the outcome as an instrumental variable (IV) to estimate the impact of revealing the outcome on Punishment. We do this separately for winning and losing principals in order to avoid non-linear IV. The coefficient on revealing the outcome is negative for winners (-0.130) and positive for losers (0.761), but both are imprecisely estimated ($p = 0.829$ for winners and $p = 0.304$ for losers, standard errors clustered by individual). This is in part due to the non-monotonicity of the demand curve in the first stage weakening the relevance of the instrument. Thus, we prefer to use the across-treatment technique to analyze sensitivity to luck.

³⁴44% never avoided the outcome, 3% failed to respond.

Table 6: Sensitivity of Punishment to Investment and Outcomes by Treatment

Model:	Random Effects		Tobit	
	(1)	(2)	(3)	(4)
Commit	-1.13*** (0.33)	-1.22*** (0.38)	-3.75*** (1.12)	-3.37*** (1.28)
Investment	-0.33*** (0.08)	-0.44*** (0.12)	-0.97*** (0.30)	-1.04** (0.42)
Commit×Investment	0.27*** (0.10)	0.32** (0.14)	0.85* (0.47)	0.59 (0.59)
Win	-0.71*** (0.18)	-1.15*** (0.32)	-3.33*** (0.95)	-3.72*** (1.24)
Commit×Win	0.42* (0.24)	0.66 (0.40)	1.72 (1.27)	0.56 (1.77)
Investment×Win		0.22 (0.13)		0.20 (0.54)
Commit×Investment×Win		-0.11 (0.17)		0.57 (0.76)
Constant	1.94 (0.29)	2.11 (0.33)	1.91 (0.89)	2.03 (1.02)
H_0 : Investment + Commit×Investment = 0	$\chi^2(1) = 1.38$ $p = 0.24$	$\chi^2(1) = 2.50$ $p = 0.11$	$\chi^2(1) = 0.10$ $p = 0.75$	$\chi^2(1) = 1.32$ $p = 0.25$
H_0 : Win + Commit×Win = 0	$\chi^2(1) = 3.46$ $p = 0.06$	$\chi^2(1) = 3.99$ $p = 0.05$	$\chi^2(1) = 3.06$ $p = 0.08$	$\chi^2(1) = 6.36$ $p = 0.01$
Principals	67	67	67	67
Observations	804	804	804	804

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$, * $\Rightarrow p < 0.10$. Standard errors clustered by individual. The Tobit model adjusts for censoring at \$0 and \$4 of punishment.

6 Third Party Results

We now examine outcome bias among third parties. We use the Full information treatment to estimate the magnitude of outcome bias among third parties. The Force treatment then measures agents' perceptions about outcome bias among third parties. Finally, the Commit treatment tests for sophistication among third parties about their own outcome bias.

Sensitivity to outcomes 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 Third Party Full Treatment

Punishment decisions by third parties with full information are shown in Table 7. They are clearly influenced by luck even though the third parties are not directly affected by either the agents' investment or luck. While punishment levels are lower for third parties than principals, the proportion of their punishments originating from luck versus investment is similar. In all specifications, we find that investment significantly decreases punishment. Thus, similar to principals, punishment reflects some version of reciprocity. Conditional on investment, there is a large and statistically significant impact of the principal winning on punishment in all specifications.

Table 7: Impact of Luck and Investment on Third Party Punishment in Full Treatment

Model:	Random Effects		Tobit	
	(1)	(2)	(3)	(4)
Investment	-0.20*** (0.03)	-0.27*** (0.04)	-0.98*** (0.17)	-1.05*** (0.16)
Win	-0.28*** (0.06)	-0.56*** (0.08)	-1.30*** (0.27)	-1.59*** (0.25)
Investment \times Win		0.14*** (0.03)		0.17 (0.15)
Constant	1.06 (0.13)	1.18 (0.14)	0.37 (0.36)	0.47 (0.34)
Third Parties	99	99	99	99
Observations	1584	1584	1584	1584

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$. Standard errors clustered by individual. The Tobit model adjusts for censoring at \$0 and \$4 of punishment.

6.2 Third Party Force Treatment

We now test for demand among agents to control the information given to the third parties. Figure 7, Panel A shows that, regardless of the agent's investment, they are more likely to reveal good luck to the third parties than bad luck. Figure 7, Panel B then shows the demand curve for information revelation under good and bad luck. The two demand curves have nearly identical responses to changes in price, but good luck increases demand by approximately 24 percentage points at each price.

In Table 8 we estimate the impact of the outcome on agents' information revelation choices. Luck plays a significant role in determining whether agents reveal the outcome to the third parties. Unlike in the principal-agent context, however, the agent's investment no longer affects the information revelation

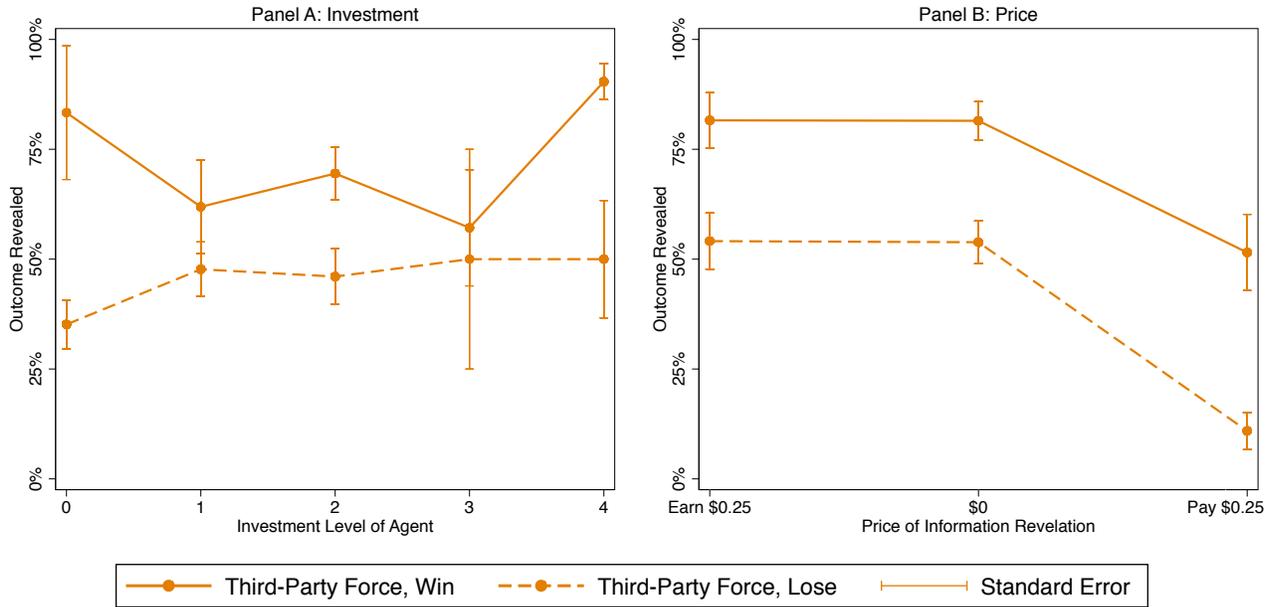


Figure 7: Agent revelation behavior by investment level and price in the Third Party Force treatment.

decision. 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.

Table 8: Likelihood of Revealing the Outcome to the Third Parties

Model:	Random Effects		Probit	
	(1)	(2)	(3)	(4)
Investment	0.026 (0.033)	0.020 (0.049)	0.081 (0.079)	0.042 (0.093)
Win	0.240** (0.095)	0.212 (0.162)	0.704*** (0.170)	0.504* (0.308)
Investment × Win		0.014 (0.062)		0.105 (0.176)
Constant	0.418 (0.062)	0.426 (0.080)	-0.225 (0.167)	-0.180 (0.176)
Agents	31	31	31	31
Observations	372	372	372	372

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$, * $\Rightarrow p < 0.10$. Standard errors clustered by individual. Marginal effects reported for the Probit model.

6.3 Third Party Commit Treatment

Given the outcome bias we observe in third party punishment choices, we again test for a desire to commit against revealing luck. Similar to the information revelation of the principals, the third parties were presented with the choice to reveal information about the principal's outcomes prior to making their punishment choices. Figure 8, Panel B displays the demand curve for information revelation among third parties. On average, third parties revealed the information 55% of the time and even paid to reveal information 22% of the time. These rates are lower than those of the principals, but still strikingly high for information about luck. The third parties exhibit the same pattern of non-monotonic demand shown by principals in the Commit treatment; third parties are less likely to view the information when they are paid to do so than when it is free.³⁵

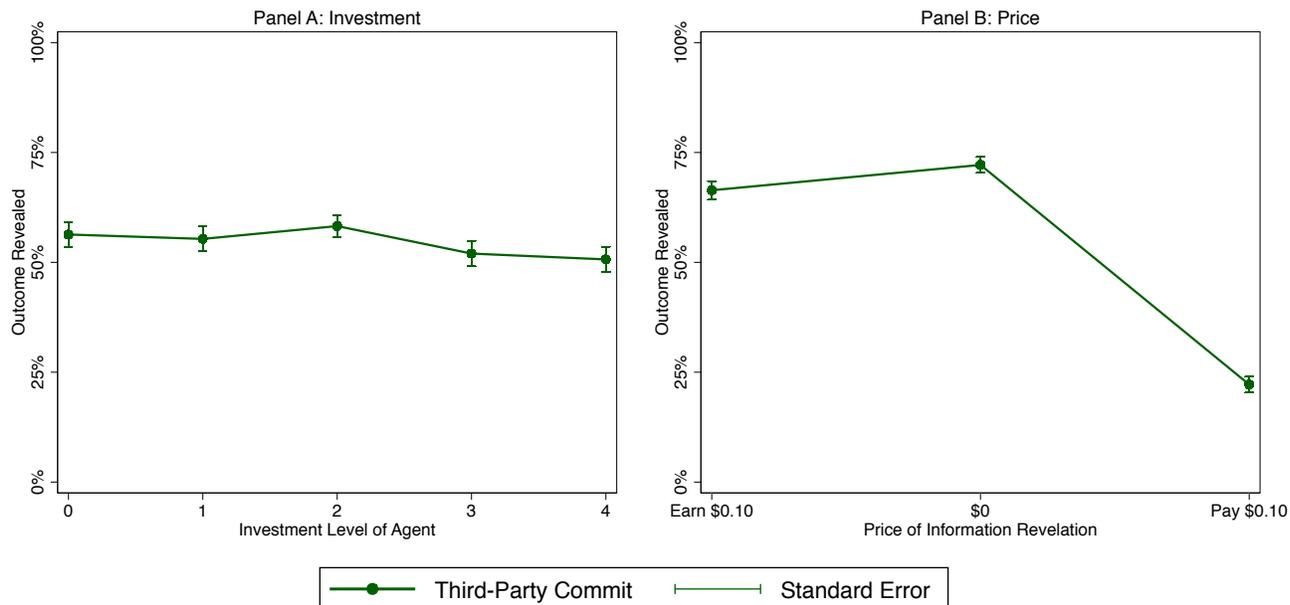


Figure 8: Third party revelation by investment and price in the Third Party Commit treatment.

Information revelation rates at each investment level are in Figure 8, Panel A. This suggests that, like principals, third parties are not motivated by curiosity. Revelation is not related to the amount of uncertainty the third party faces.

Low demand for commitment, on average, does not imply that control over information revelation is a net loss to the third parties. Instead of evidence against sophistication, it could simply be that only the biased agents commit against the information and the rest of the agents are indifferent. If this were the

³⁵The non-monotonicity from a price of -\$0.10 to \$0 is statistically significant at the 10% level but is smaller than among principals. This estimate is from a probit model with standard errors clustered by individual.

case, then we would expect that the ability of third parties to condition their punishment on investment should improve when given control over their information.

To test for changes in the ability of third parties to ignore luck, we compare their punishment decisions in the Commit and Full treatments. Table 9 shows that punishment decisions are significantly less sensitive to the principal's outcome in Commit compared to Full ($p = 0.002$). This decreased sensitivity to luck suggests that control over revelation of luck improves punishment decisions for third parties. Unlike the principals, punishment is still sensitive to investment, making a stronger case for sophistication on the part of the third parties. Third parties subject to bias in their interpretation of random information appear to correct for those biases by either selectively choosing to reveal that information or by recognizing, through the revelation process, the irrelevance of this information.

Table 9: Sensitivity of Punishment to Investment and Outcomes by Treatment

Model:	Random Effects		Tobit	
	(1)	(2)	(3)	(4)
Commit	-0.28* (0.17)	-0.32* (0.18)	-1.00* (0.53)	-0.91* (0.51)
Investment	-0.20*** (0.03)	-0.27*** (0.04)	-1.02*** (0.17)	-1.10*** (0.16)
Commit×Investment	0.04 (0.04)	0.06 (0.05)	0.15 (0.21)	0.08 (0.21)
Win	-0.28*** (0.06)	-0.56*** (0.08)	-1.35*** (0.27)	-1.65*** (0.26)
Commit×Win	0.26*** (0.08)	0.35*** (0.12)	0.95** (0.39)	0.75* (0.38)
Investment×Win		0.14*** (0.03)		0.18 (0.16)
Commit×Investment×Win		-0.05 (0.03)		0.12 (0.20)
Constant	1.06*** (0.13)	1.18*** (0.14)	0.32 (0.37)	0.43 (0.35)
H_0 : Investment + Commit×Investment = 0	$\chi^2(1) = 31.19$ $p = 0.00$	$\chi^2(1) = 39.24$ $p = 0.00$	$\chi^2(1) = 39.86$ $p = 0.00$	$\chi^2(1) = 49.69$ $p = 0.00$
H_0 : Win + Commit×Win = 0	$\chi^2(1) = 0.15$ $p = 0.70$	$\chi^2(1) = 6.54$ $p = 0.01$	$\chi^2(1) = 2.48$ $p = 0.12$	$\chi^2(1) = 11.18$ $p = 0.00$
Third Parties	199	199	199	199
Observations	3184	3184	3184	3184

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$, * $\Rightarrow p < 0.10$. Standard errors clustered by individual. The Tobit model adjusts for censoring at \$0 and \$4 of punishment.

7 Beliefs about the Agent’s Type

In this section, we measure the impact of luck on the beliefs that third parties hold about the agents’ types. We presented third parties with investment and outcomes for 16 principal-agent interactions. We then elicited the third party’s belief about the mean investment choice of that agent across all 13 periods in which the agent made investment decisions. With information about the agent’s investment choice, information about luck should be useless in the belief updating process. Nonetheless, we observe a strong influence of luck on the beliefs elicited about agents’ types. We regress guesses on indicator variables for each investment level and predict the residual guess—that is, the portion of the guess that is not explained by investment. Figure 9 shows the effect of luck on this residual guess about the agent’s mean investment level.

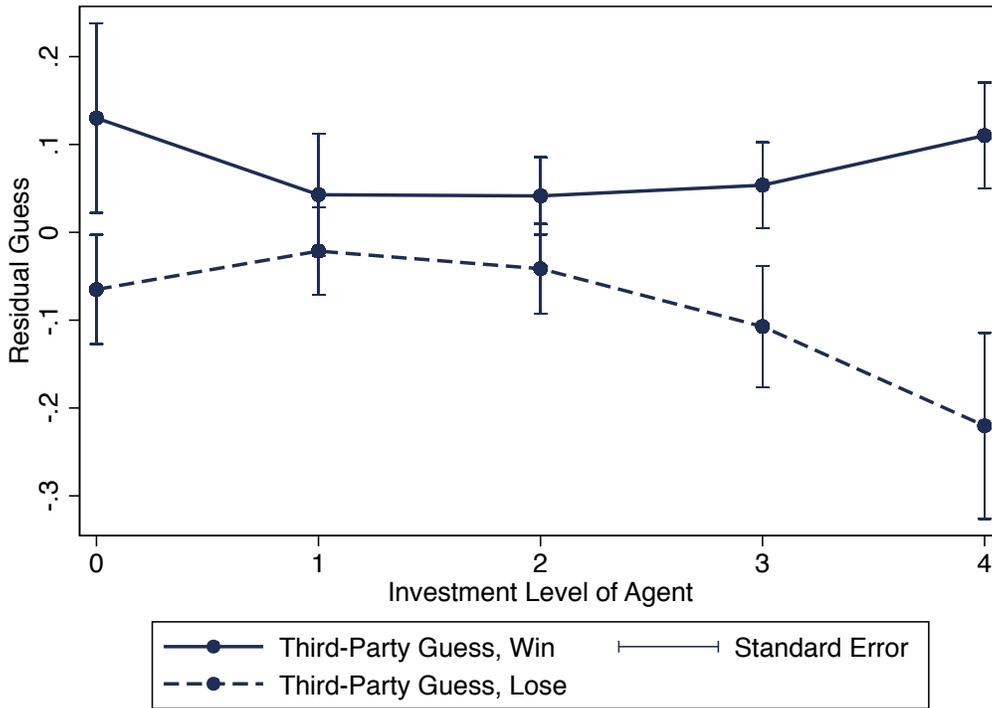


Figure 9: Third Party Beliefs about the Agent’s Mean Investment.

In Table 10, we estimate the impact of the principal’s outcome on beliefs about the agent’s type. Despite complete information about the agent’s investment, good luck has the same effect as observing that the agent increased investment by 0.26 SD that period ($p < 0.01$). That is, good luck is misattributed to investment. While this bias exists for all investment levels, the coefficients on the interaction term indicate that it may be slightly larger at higher investment levels.

Table 10: Third Party Guesses about Agent’s Mean Investment

Model:	Random Effects		Tobit	
	(1)	(2)	(3)	(4)
Investment	0.462*** (0.034)	0.450*** (0.037)	0.511*** (0.018)	0.497*** (0.025)
Win	0.164*** (0.056)	0.113 (0.080)	0.182*** (0.049)	0.129 (0.087)
Investment \times Win		0.025 (0.029)		0.026 (0.036)
Constant:	1.027 (0.082)	1.047 (0.086)	0.922 (0.058)	0.944 (0.066)
Third Parties	100	100	100	100
Observations	1600	1600	1600	1600

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$, * $\Rightarrow p < 0.10$. Standard errors clustered by individual. The Tobit model adjusts for censoring guesses at \$0 and \$4.

The fact that we use third parties for this analysis—rather than principals—is important for identifying beliefs as the channel through which luck affects punishment. 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 like disappointment or excitement. Observing that luck influences the beliefs that third parties hold about an agent’s type—and, thus, how deserving agents are of punishment—is strong evidence that biased belief updating is a driver of the results in our study.³⁶

8 Conclusion

This paper studies three previously-unexplored aspects of outcome bias in a principal-agent setting: 1) the sophistication of principals and agents about outcome bias, 2) outcome bias among third parties, 3) the association between attribution bias, correlation neglect and outcome bias. We find that agents display strategic sophistication, but principals do not. 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

³⁶We conducted the same test with the 31 passive principals in Wave 3 and found consistent, but underpowered results. Good luck is equivalent to a 0.29 SD increase in investment ($p = 0.22$). Full results can be found in appendix Table A.4.

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. Third parties punish bad luck the same as they would a 1.0 SD decrease in effort. They can also be manipulated by agents with control over information. This finding has broad consequences, since it eliminates the possibility that outcome bias originates simply from emotion associated with good or bad luck or the inequality that outcomes create. 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. However, we 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 third-parties hold about the investment 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 0.26 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 investment—a perfect signal. This model can be thought of as a combination of correlation neglect (Enke and Zimmermann, 2017) and attribution bias (Haggag and Pope, 2016). When individuals fail to recognize the correlation between signals, they end up attributing more of the outcome to effort than they should. Whether that failure is accidental, or explained by a model of delegated expertise Gurdal et al. (2013) is unknown and an avenue for future work. The fact that we observe beliefs trending with luck is clearly inconsistent with a model of distributional preferences, as discussed in Section 4.2.

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

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.

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, learning and experience will be important avenues for future research. In our 13-period study, we do not see significant changes over time in principal behavior. This is consistent with the extended, 25-period sessions in Gurdal et al. (2013). Whether principals learn over a longer horizon, whether role reversal leads to more symmetric behavior, and whether experienced evaluators exhibit such biases all need to be determined through further investigation.

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.
- 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.
- Demski, Joel S. and Sappington, David E.M.** (1987). ‘Delegated Expertise’, *Journal of Accounting Research* 25(1), 68–89.
- de Oliveira, Algela 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.
- 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

Table A.1: Impact of Random Outcome and Agent Investment on Punishment in Full Treatment

	Random Effects	Tobit
	(1)	(2)
Constant: Invest 0, Lose	2.08 (0.36)	4.21 (0.70)
Invest 1	-0.37 (0.32)	-0.39 (1.21)
Invest 2	-0.85** (0.38)	-1.10 (1.32)
Invest 3	-1.42*** (0.39)	-2.83** (1.33)
Invest 4	-1.64*** (0.45)	-5.57** (2.10)
Win Invest 0	-0.68 (0.42)	-1.81 (1.35)
Win Invest 1	-1.08*** (0.29)	-5.36** (2.11)
Win Invest 2	-0.74*** (0.23)	-3.60*** (1.21)
Win Invest 3	-0.27 (0.24)	-1.91 (1.32)
Win Invest 4	-0.48 (0.41)	-2.52 (2.60)
Clusters	33	33
Observations	429	429

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$. Standard errors clustered by individual. The Tobit model adjusts for censoring at \$0 and \$4 of punishment.

Table A.2: Effect of Information Revelation on Punishment in the Force Treatment

Model:	Random Effects		Tobit	
	(1)	(2)	(3)	(4)
Constant: Invest 0, Outcome Hidden	1.540 (0.274)	1.590 (0.318)	0.980 (1.039)	0.832 (1.154)
Investment	-0.345*** (0.077)	-0.379*** (0.107)	-1.417*** (0.439)	-1.272** (0.523)
Outcome Shown X Win	-0.272** (0.134)	-0.702* (0.376)	-2.262* (1.286)	-3.706** (1.860)
Outcome Shown X Investment X Win		0.188 (0.137)		0.589 (0.739)
Outcome Shown X Lose	0.046 (0.160)	0.451 (0.447)	0.488 (0.784)	3.499* (2.050)
Outcome Shown X Investment X Win		-0.230 (0.184)		-1.902* (1.131)
Clusters	33	33	33	33
Observations	396	396	396	396

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$, * $\Rightarrow p < 0.10$. Standard errors clustered by individual. The Tobit model adjusts for censoring at \$0 and \$4 of punishment.

Table A.3: Sensitivity of Punishment to Investment and Outcomes by Treatment and Time

Model:	Random Effects		Tobit	
	(1)	(2)	(3)	(4)
Force	-0.38 (0.39)	-0.65 (1.31)		
Investment	-0.34*** (0.08)	-0.99*** (0.32)		
Force×Investment	-0.03 (0.11)	-0.48 (0.54)		
Win	-0.71*** (0.25)	-3.52*** (1.27)		
Force×Win	0.56** (0.28)	1.72 (1.50)		
Period×Win	0.00 (0.02)	0.01 (0.11)		
Force×Period×Win	-0.00 (0.02)	0.03 (0.13)		
Constant	1.99 (0.30)	2.28 (0.96)		
Principals			66	66
Observations			792	792

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$. Standard errors clustered by individual. The Tobit model adjusts for censoring at \$0 and \$4 of punishment.

Table A.4: Principal Guesses about Agent's Mean Investment

Model:	Random Effects		Tobit	
	(1)	(2)	(3)	(4)
Investment	0.294*** (0.057)	0.274*** (0.069)	0.298*** (0.029)	0.278*** (0.040)
Win	0.113 (0.079)	0.028 (0.168)	0.113 (0.079)	0.028 (0.139)
Investment X Win		0.044 (0.074)		0.043 (0.059)
Constant:	1.378 (0.117)	1.402 (0.127)	1.369 (0.082)	1.393 (0.088)
Principals			31	31
Observations			372	372

*** $\Rightarrow p < 0.01$, ** $\Rightarrow p < 0.05$, * $\Rightarrow p < 0.10$. Standard errors clustered by individual. The Tobit model adjusts for censoring guesses at \$0 and \$4.

A.1 Subgame Perfect Nash Equilibrium

The subgame perfect solution to our principal-agent interaction involves no investment 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 investment level (his "type").

Moving backwards to the investment 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 investment 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 2.

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:

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

$$f(\theta|x) = \frac{f(x|\theta)f(\theta|\text{WIN} = 0)}{g(x)} \quad (7)$$

We incorporate the conditional probability distribution, $f(\theta|\text{WIN} = 1)$, into the updating process to arrive at the biased posterior distribution in Figure 3. To arrive at the opposite bias, we can similarly incorporate $f(\theta|\text{WIN} = 0)$ and find distributions slanted to the left instead of the right.

A.3 Instructions

All 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.3.1 Principal's Instructions

Page 3:

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 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 $3/6 = 1/2$.

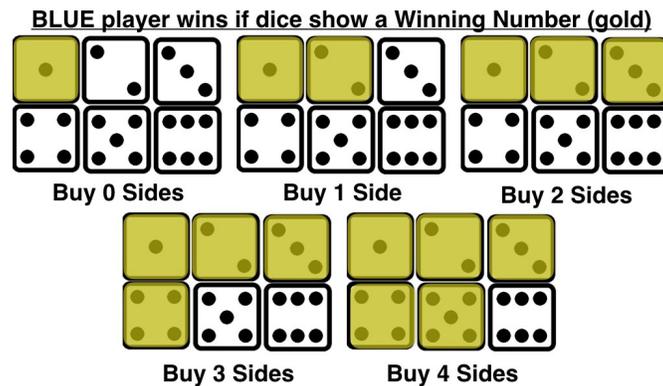
The GREEN player can choose to do nothing (buy 0 numbers) or buy up to 4 numbers. Each number costs the GREEN player \$0.50 to buy.

Page 4:

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

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.

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 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:

<p>The GREEN player purchased 2 winning numbers for YOU. Would you like to observe the outcome of the dice before choosing how much you would like to punish the GREEN player?</p>	<p>You can choose to see what number the dice rolled and whether you won or not by clicking below. To show this information you don't have to pay anything.</p> <p style="text-align: center;">Show the dice (You pay \$0)</p>
--	--

A.3.2 Agent's Instructions

Page 3:

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 $3/6 = 1/2$.

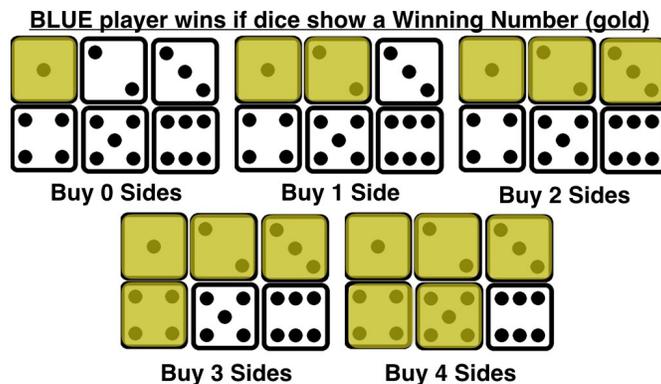
YOU can choose to do nothing (buy 0 numbers) or buy up to 4 numbers. Each number costs YOU \$0.50 to buy.

Page 4:

Recall: 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:

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, Intent, Commit treatments):

Recall: **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.

Page 5 (Force treatment):

Recall: **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 players 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 6 (Force treatment):

Recall: **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:

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 BLUE player?	Clicking this button will show the BLUE player that they WON the \$6.00 prize. <input type="button" value="Show the dice (You pay \$0)"/>
--	---

A.4 Comprehension Quiz

All 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”).

- 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 treatment, the agent also sees the following questions:

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

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?