Good decision vs. good results: Outcome bias in the evaluation of financial agents

Christian König-Kersting¹ · Monique Pollmann² · Jan Potters² · Stefan T. Trautmann³

Published online: 17 September 2020
© The Author(s) 2020

Abstract

We document outcome bias in situations where an agent makes risky financial decisions for a principal. In three experiments, we show that the principal's evaluations and financial rewards for the agent are strongly affected by the random outcome of the risky investment. This happens despite her exact knowledge of the investment strategy, which can, therefore, be assessed independently of the outcome. The principal thus judges the same decision by the agent differently, depending on factors that the agent has no influence on. The effect of outcomes persists in a setting where principals communicate a preferred investment level. Principals are more satisfied with the agent after a random success when the agent did not follow the requested investment level, than after a failed investment that followed their explicit request.

Keywords Decision under risk · Decisions of agents · Accountability · Outcome bias · Financial advice

1 Introduction

Whenever the quality of a decision is evaluated after its consequences have played out and have become public knowledge, there is a chance of falling prey to outcome bias. *Outcome bias* describes the phenomenon by which evaluators tend to take information about the outcome into account when evaluating the quality of a decision itself (Baron and Hershey 1988). This tendency is problematic for two

✉ Stefan T. Trautmann
s.t.trautmann@uvt.nl

¹ University of Innsbruck, Innsbruck, Austria

² Tilburg University, Tilburg, Netherlands

³ Alfred-Weber-Institute for Economics, University of Heidelberg, Bergheimer Strasse 58, 69115 Heidelberg, Germany

reasons. First, the evaluator has available a different information set than the decision maker, who typically faces uncertainty at the time of her decision. Second, a good outcome might derive from a bad decision, and a bad outcome might derive from a good decision.¹ Evaluation of outcomes may, therefore, be questionable and may lead to suboptimal future decisions if decision makers follow strategies that were successful only by chance (e.g., Bertrand and Mullainathan 2001, for managerial performance; or Sirri and Tufano 1998, for investors' mutual fund choices).²

The consideration of potentially irrelevant outcome information in the evaluation of decision quality has been documented in a wide variety of settings including medical advice, military combat decisions and salesperson performance evaluation (Baron and Hershey 1988; Lipshitz 1989; Marshall and Mowen 1993). In these early studies, participants were asked to evaluate the quality of a decision described in hypothetical scenarios differing in featuring a favorable, an unfavorable, or no outcome at all. Later studies on peer review of scientific publications and strategies in professional football move away from scenarios and towards actual decisions as to the basis for evaluation (Emerson et al. 2010; Lefgren, Platt and Price 2015). Relatedly, there is a strand of literature on allocator-responder games with a 'trembling hand' condition, in which responders can infer allocators' intentions, but actual allocation outcomes may deviate from intentions by chance. Cushman et al. (2009) find that responders hold allocators accountable for unintentional negative outcomes, but knowledge of their agents' intentions moderates the effects. These findings are supported and augmented by further studies, e.g. by Murata et al. (2015) and Sezer et al. (2016).

Investors choosing investment funds that have been successful in the past (mostly by chance) is an important policy issue in finance. However, in the field of financial economics, there is little controlled experimental evidence yet on whether such behavior is potentially related to outcome bias: field data typically cannot separate outcome bias from other effects. Our paper is the first to study outcome bias in such a financial investment context, using controlled lab experiments. We focus on client-advisor relationships typical to investment settings (framed in terms of a principal-agent situation) and show that outcome bias is prevalent in such settings and leads to a biased assessment of the quality of the advisor's investments by the client.

Why do we need to study outcome bias in financial investment settings if there is already broader evidence on the bias? For behavioral experiments, seemingly small changes to decision situations can have pronounced behavioral consequences. Especially in regard to cognitive biases, transferability from one situation to

¹ Consider for example a decision between a safe payment and a prospect with positive expected value larger than the safe option, but of substantial variance. A decision maker instructed to make risk-neutral decisions should choose the risky prospect over the safe option. Yet the outcome might turn out unfavorable and lower than the safe option. A negative evaluation on the basis of the bad outcome seems unwarranted.

² It is important to recognize that outcome effects do not always constitute biases. The literature originating from Baron and Hershey (1988) typically speaks of an outcome bias only if responsibility for the outcome is inappropriately assigned to decision makers. We follow this interpretation.

another, even if they appear to be highly similar, cannot be taken for granted (Crusius et al. 2012). For example, Charness et al. (2010) show that the introduction of mild incentives significantly reduces violations of the conjunction principle compared to an otherwise identical, but unincentivized decision situation. In addition, Lefebvre et al. (2011) highlight that the ratio bias phenomenon is sensitive to changes in the decision making environment as well as the incentive structure. The current paper concerns the robustness of the outcome bias phenomenon. We assess the prevalence and implications of the outcome bias in financial decisions with agency, employing a variety of different incentive conditions and assessment methods by the evaluator.

Specifically, we go beyond previous studies by comprehensively studying outcome bias in financial settings from multiple angles. We show that outcome bias (i) occurs in both monetarily incentivized as well as non-incentivized assessments; (ii) emerges in the evaluation of hypothetical vignettes just as in natural behavior in reaction to decisions by human agents; and (iii) appears even if expressing content is costly for the evaluator. By making sure that the evaluators always know the random process, we demonstrate that outcome bias is not just the result of evaluators having to diagnose outcome information for clues of decision-process quality, but that it also robustly materializes in isolation.³ To the best of our knowledge, we are also the first who identify positivity bias in the context of outcome bias. Such positivity bias is consistent with the above described field evidence on investors following mutual funds successful in the past.

In Experiment 1, the principals' assessments of the agents' decisions have direct monetary consequences for principals and agents, and potentially affect agents' future decisions. We compare a situation where principals can observe both the decision itself and the resulting outcome, to a situation where only the investment decision is known but no outcome information is available yet.⁴ We observe that a tendency toward ex-post outcome-based evaluations exists even in situations where (1) the principal has a clear financial incentive to reward good decisions, not lucky good outcomes; and (2) where there is perfect information about the decision and the situation in which it was made.

To control for potential design-specific social-preference effects that reduce the generalizability of our results, we probe the effect of outcome-based evaluations of known processes in Experiment 2. We find that even in the absence of potential social-preference effects, principals' judgments of agents' observable investment decisions are strongly affected by the random outcome on which the agent has no influence. In particular, principals become satisfied with investment *decisions* after positive *outcomes* even if they initially strongly disliked the decision (in the absence

³ In our experiments, the evaluators do not need to diagnose outcome information for signals of decision-process quality, because both pieces of information are always readily available. This differentiates our study on outcome bias from informational cascade models, in which the underlying process is unknown and players rely heavily on diagnostic information. We thank a reviewer for pointing this out.

⁴ The former condition is similar to experiments in Gurdal et al. (2013) where players were rewarded for choosing a risky or a safe lottery for another player. Counterfactual outcomes were available to judges and had an influence on rewards. Below we discuss Gurdal et al.'s interpretation in terms of blame in the light of our results.

of the outcome information, i.e., before the uncertainty is resolved). This positivity bias is consistent with findings by Casal et al. (2019) but was unexpected given the previous demonstrations of the predominance of negative outcome effects (Gurdal et al. 2013; Ratner and Herbst 2005). The current findings suggest that financial agents seem to benefit from the rule that the *result justifies the deeds*. Inspired by our results in Experiment 1, De Oliveira et al. (2017) also study the role of wealth differences for outcome effects (see our analysis in Sect. 2.4). Consistent with our Experiment 2 results, the authors find no support for a wealth-based explanation of outcome effects.

In Experiment 3, we replace the principals' implicit assessment of the agents' decision strategies by the principal's explicit demand for a certain investment level. After observing the invested amounts (which are indicative of the agents following or not following their clients' requests) and the investment outcomes, principals send messages conveying their satisfaction with the investment decisions to their agents. This setup enhances the situation modeled in Experiment 1 by providing (1) the agents with information about what their clients consider a 'good decision' and thus (2) the principals with an obvious benchmark for evaluating the agents' decisions. We observe that satisfaction with the decision is increased by the agent conforming to the principal's wish, but find an even larger effect of the random outcome on satisfaction with the investment.

In sum, we demonstrate that outcome bias is present in financial decisions by agents, in which the evaluator is directly monetarily affected by both the decision and its evaluation in terms of monetary rewards paid (Experiment 1). We show that the phenomenon persists in situations where neither income and wealth effects nor social preference considerations can play a role, and separation of outcome and decision-process evaluation is strongly emphasized (Experiment 2). Finally, we highlight that even if principals communicate explicit investment-level demands, they still fall prey to outcome effects if the agents do not follow their demands (Experiment 3). Recognizing, that past experience can bias future evaluations (cf. *rater bias* in Müller and Weinschenk 2015), such systematically biased assessments of the quality of agents' decisions are clearly undesirable.

The remainder of the paper is laid out as follows. Sections 2 to 4 describe the methods and the results of the three experimental studies. Each section also includes a short discussion of the respective results. Section five concludes the paper with a general discussion of the role of outcome bias in financial agency.

2 Experiment 1

2.1 Methods

We use data from Pollmann et al.'s (2014) experiment on risk taking by agents under accountability.⁵ That paper investigates how the behavior of financial agents

⁵ There is a large literature on the effects of accountability in decision making for others (e.g. Charness and Jackson, 2009; Brandts and Garofalo, 2012; Lefevre and Vieider, 2013). Yet, these are concerned with the *agents'* behavior. In this paper, we focus on the other side, namely *principals'* behavior.

differs between situations in which the principals either reward their agents solely on the basis of invested amounts or on invested amounts and outcome information. Pollmann et al. (2014) is not concerned with the behavior of principals, but, by necessity, implements treatments which are suitable to study their behavior as well. That is, the data and analyses presented here have not been reported on before.

The Gneezy and Potters (1997) investment task is used in the experiment, in which decision makers are asked to divide an initial endowment of 100 points between a safe and a risky asset. The safe asset has a return of 0%. In contrast, the risky asset has a return of + 250% with a probability of 1/3 and a return of – 100% with a probability of 2/3, creating a prospect with a positive expected return of + 16.67%.

There are two types of players matched in pairs of two: a principal who is the owner of a 100-point endowment; and an agent, whose task it is to invest the principal's endowment using the above-described technology. The investment portfolio set up by the agent is fully owned by the principal. Both players are aware that they both receive an additional fixed payment of 100 points each, which is independent of the investment decision. After the agent made her investment decision, the principal is given the opportunity to reward the agent by transferring between zero and 100 points from this additional payment to the agent. This ensures that principals can give any reward, independently of how their payoffs from the agents' investment decisions turn out. Points not transferred remain with the principal. The agent receives this reward in addition to her fixed payment of 100 points. Employing a between-subject design, we compare two treatments that differ in terms of the information the principal has available when she is given the opportunity to reward the agent for her decision.⁶ When making her decision of how many points to transfer as a reward in treatment REWARD BEFORE, the principal knows the agent's investment decision (number of points invested in risky and safe), but not the realized return of the risky asset. In treatment REWARD AFTER, both the agent's investment decision and the outcome of the risky prospect are communicated to the principal before she has the opportunity to reward the agent.

The described tasks (investment – reward) are statically repeated five times with fixed principal-agent pairs. This setup increases the importance for the principal to reward investment decisions that are in line with her preferences (not outcomes that are positive) because the same agent will make another investment decision after the reward is given. The idea is that principals can steer their agents towards investment levels they like through their choice of reward.⁷ At the end of each round, payoffs for each player are transferred to her experiment account and cannot be used in the experiment anymore. New endowment and investment funds are provided for each round, ensuring that although wealth is accumulated over time, the decision set remains identical.

⁶ The experiment consisted of two more treatments where participants made investment decisions for their own account, and where they made decisions for others without the possibility of reward. These are discussed in Pollmann et al. (2014).

⁷ In a one-shot version of the game, a purely selfish principal would pay a zero reward. With multiple repetitions, rewards have the potential to change the agent's behavior, such that non-zero rewards can be utility maximizing for the principal even in the absence of social preferences before the final round. Reward payments should never depend on the random outcome.

2.2 Setting and summary statistics

The experiment was programmed in z-Tree (Fischbacher 2007) and conducted at CentErlab, Tilburg (NLD). Roles were assigned randomly, partner identities were kept secret, and decisions were made anonymously with no communication between principals and agents. Participants received instructions in writing as well as on screen and had to complete a set of mandatory comprehension questions. The sessions began only after every participant had correctly answered these questions. The research question was not revealed to participants at any time. Points were exchanged for 0.01€ each at the end of the experiment.⁸

A total of 134 students participated in the part of the experiment relevant for this paper (34 principal-agent pairs in treatment REWARD BEFORE and 33 principal-agent pairs in treatment REWARD AFTER). At the time of the experiment, participants were on average 22.5 years old, 37% of them were female and 36% of Dutch nationality. We asked participants about their major field of studies, which revealed 55% economics, 37% business, and 2% psychology students in our sample. Table 6 in Appendix 2 presents the summary statistics in greater detail.

2.3 Results

Recall that principals are aware that the outcomes are determined randomly and independently of the investment decision by their agent. The outcome information, if available, does not provide meaningful additional information about the agents' decision process. Compared to a situation where only the invested amount is known, additionally learning about the random outcome should not have an effect on reward payments.

Comparing the rewarding behavior of principals in treatment REWARD AFTER in situations in which the risky asset yielded a favorable random outcome to situations in which it yielded an unfavorable one, we observe substantial outcome effects. Pooling observations from all rounds,⁹ we find average rewards of 28.78 (SD = 4.36) when favorable outcomes are observed, versus 10.54 (SD = 1.82) when unfavorable outcomes are observed.¹⁰ As a placebo test, we make the same comparison for treatment REWARD BEFORE. Here we find average rewards of 18.72 (SD = 3.12) for favorable random outcomes, versus 18.94 (SD = 2.47) for unfavorable ones.¹¹

For principals who received information about the investment decision and outcomes, we furthermore see a significantly positive correlation between their own

⁸ Participants could earn an additional 100 points in a belief elicitation task, which is not discussed in this paper. On average, participants earned 7.93€. We provide the complete instructions to this experiment as part of the supplementary material at <https://osf.io/dezqk/>.

⁹ For the development of average rewards over the course of the experiment, refer to Appendix 3.

¹⁰ Two-sample, two-sided *t*-test, $t(163) = -4.56$, $p < 0.001$, $d = -0.76$. When rewards are compared separately for each period, the difference is significant in three out of five periods. We account for the repeated structure in the multivariate analysis below.

¹¹ Two-sample, two-sided *t*-test, $t(168) = 0.06$, $p = 0.956$, $d = 0.01$. The difference in average rewards is neither significant when observations are pooled nor when periods are treated separately.

payoff and the reward they pay to their agent (Fig. 1, left panel, $\delta = 0.45$, $p < 0.001$). We do not find a positive correlation if the principal had to reward the agent before knowing the outcomes of the risky investment (Fig. 1, right panel, $\delta = 0.02$, $p = 0.83$).¹²

To estimate the size of the effect as well as to control for repeated observations and personal characteristics of the participants, we probe these findings in a multivariate analysis. For each treatment, we employ a separate tobit panel regression to regress the size of the reward on the principal's payoff and a constant. In a second step, we test the robustness of the results by including controls for wealth accumulated over the course of the experiment, age, gender, Dutch nationality,¹³ and the field of study. As coefficients are hard to interpret in non-linear models, we report the more convenient average marginal effects in Table 1. The regression analyses confirm that absent information on realized outcomes, there is no effect of the principals' payoffs on rewards. However, once outcomes are available, there is a significantly positive effect of payoffs on rewards: on average, a unit increase in payoff leads to an increase of 0.09 points in reward.¹⁴

Because high payoffs obtain from favorable random draws for the risky investment, we next test whether it is the observation of a success or failure per se that drives the above effect, or whether the effect runs mainly through the size of the outcome. We thus repeat the above analyses, now including as covariates the amount invested in the risky asset, an indicator for a favorable outcome (investment success) and the interaction of these variables. Results are shown in Table 2, models 3 and 4. If both the investment decision and the outcome are observable (REWARD AFTER, Table 2, upper panel), we can report two results: First, the reward in the case of observing a favorable outcome is on average 16.60 points higher than in the case of observing an unfavorable outcome. Second, if the outcome is favorable, the effect of the amount invested on the reward is positive and highly significant. A unit increase in risky investment leads to an average increase in reward of 0.46 points. If the outcome is unfavorable, the effect of the amount invested in the risky asset on the reward is not significantly different from zero. That is, rewards are driven by success in REWARD AFTER, and only in the case of success does the amount invested, and, therefore, the actual payoff to the principal, affect the size of the reward. In the case of a failure, the correlation between the principal's payoff (which then depends inversely on the agent's investment) and the reward is close to zero and non-significant.

If only the amount of investment is observable by the principal (REWARD BEFORE, Table 2, lower panel), we do not find a statistically significant effect of

¹² Restricting the correlation analysis to non-zero rewards yields qualitatively similar results. For REWARD AFTER we find $= 0.42$ ($p < 0.001$) and for REWARD BEFORE it is $= -0.03$ ($p = 0.72$).

¹³ The experiment was run in the Netherlands, but with a significant group of foreign students. Since rewarding behavior may vary across different cultural backgrounds we control for Dutch versus foreign students here.

¹⁴ The marginal effect of an increase in payoff on reward is significantly different from zero at all levels of payoff and monotonically increasing from 0.06 to 0.15. Graphs of the marginal effects are available from the authors upon request. All results are robust to using a linear panel OLS regression with standard errors clustered on the individual level instead of the non-linear tobit model.

Fig. 1 Relation between principals' payoffs and rewards for the agents. All periods included; scattered observations with linearly fitted line

Table 1 Relation between principals' payoffs and rewards for the agents

	AFTER	AFTER	BEFORE	BEFORE
Principal's Payoff	0.0906 *** (0.0134)	0.0909 *** (0.0130)	– 0.0011 (0.0117)	– 0.0022 (0.0122)
Controls	No	Yes	No	Yes
# Principals	33	33	34	34
# Observations	165	165	170	170

Random effects tobit regression; dependent variable is the size of reward; average marginal effects reported; standard errors in parentheses

*** Denotes significance at the 0.1% level; controls are: accumulated wealth, age, gender, field of study and Dutch nationality

the invested amount on rewards.¹⁵ The placebo test of the effect of the favorable outcome is also insignificant. As shown in Table 2, all results are robust to the inclusion of our standard set of controls. Restricting the analyses to non-zero rewards only does not qualitatively change the results either (see Table 7 in Appendix 2).

Figure 2 provides further insights by plotting the marginal effects on rewards for each investment level separately (based on specifications 3 and 4). For REWARD AFTER, it shows that in the case of an unfavorable outcome there is indeed an insignificant effect at all investment levels; in contrast, for a favorable outcome the marginal effect is increasing in the investment level. Moreover, the figure confirms

¹⁵ For comparison, we also present results from a regression that only includes the agents' investment and consciously disregards the random outcome information that is available (yet not informative) to the principal in the REWARD AFTER treatment (Table 2, upper panel, models 1 and 2). We do not find evidence for a relationship between invested amounts and rewards unless we tightly control for the principals' characteristics.

Table 2 Relation between agents' risky investment and rewards

REWARD AFTER	(1)	(2)	(3)	(4)
Investment	0.1041 (0.0600)	0.1236 * (0.0591)		
...if favorable outcome			0.4602 *** (0.0785)	0.4990 *** (0.0816)
...if unfavorable outcome			– 0.0531 (0.0506)	– 0.0459 (0.0493)
Favorable Outcome			16.5962 *** (2.6101)	16.1870 *** (2.5861)
Controls	No	Yes	No	Yes
# Principals	33	33	33	33
# Observations	165	165	165	165
REWARD BEFORE			Placebo test	Placebo test
Investment	0.0891 (0.0513)	0.0932 (0.0516)	0.0943 (0.0522)	0.0966 (0.0525)
Favorable outcome			– 0.5752 (2.4931)	– 0.1282 (2.5760)
Controls	No	Yes	No	Yes
# Principals	34	34	34	34
# Observations	170	170	170	170

Random effects tobit regressions; dependent variable is size of reward; average marginal effects reported; standard errors in parentheses

*, *** Denotes significance at the 5%, 0.1% level; controls are: accumulated wealth, age, gender, field of study and Dutch nationality

Fig. 2 Marginal effects of agents' risky investment on reward

the signs of the tobit interaction terms in Table 2. For the case of REWARD BEFORE, the figure shows that investment has a significantly positive effect only at very low investment levels.

Notably, the observed pattern of rewards is not consistent with a general income effect. If higher income, i.e. a higher payoff to the principal, generally translated into higher rewards for the agents, we would observe a negative effect of the size of the investment on rewards for unfavorable investment outcomes: income is decreasing in the investment in this case. The observed pattern is also robust if we restrict the analysis to situations in which principals clearly take the agents' payoffs into consideration by paying non-zero rewards. In particular, the marginal effects of investment level on rewards are close to zero for unfavorable investments, and positive and increasing after an investment success. Taken together, only an income effect conditional on observing a positive outcome is consistent with our results. This conditioning is exactly what outcome bias implies. A similar argument applies to the possibility that instead of income, it is an implicit experimenter demand that drives the outcome bias: available information should be used in the determination of the reward. If this were the case, as we have shown, only positive information about a successful outcome would create experimenter demand; principals are not responsive to experimenter demand caused by increasingly negative outcomes after an unsuccessful investment. That is, this model would imply an outcome bias in participants' assessment of experimenter demand.

2.4 Outcome bias and social comparison

We observe that principals strongly base their rewards on observed outcomes when these are available. In particular, principals reward favorable chance outcomes and additionally reward higher investments conditional on hindsight that larger investments were a good decision. Given that (i) the outcome is not under the control of the agent and (ii) the principal has full information about the agent's decision process (i.e., amount invested in the presence of uncertainty), it seems difficult to justify this focus on outcomes.

Despite our finding that income and wealth effects cannot account for the observed pattern of rewards after favorable and unfavorable investment outcomes, social comparison may still loom large in the current experiment, and may add to the observed outcome effect. To gain some insight into this potential channel behind the observed outcome bias, we analyze the data of Experiment 1 within the context of social preference models. Both ex-ante and ex-post perspectives on fairness may be important in social comparison situations involving risky decisions (Krawczyk and Lelec 2010; Cappelen et al. 2013; Brock et al. 2013). We, therefore, consider the model proposed by Fehr and Schmidt (1999) for the case of observable outcomes (REWARD AFTER, i.e. ex-post), and the model proposed by Trautmann (2009) for the case of unobservable uncertain outcomes (REWARD BEFORE, ex-ante). We assume that after a high payoff to the principal, she might be more inclined to give a higher reward to the agent to make payoffs more equal. That is, we assume that the principals are averse to advantageous inequality. In Appendix 1 we show that the outcome-based model cannot explain the observed patterns of

reward in REWARD AFTER for a fixed distribution of inequality aversion parameters. This reflects our above observation of an absent link between investment and reward after an unsuccessful investment. In contrast, for REWARD BEFORE, distributions of inequality aversion parameters can be constructed to fit the observed pattern of rewards.

Despite the failure of inequality aversion models to account for the pattern of rewards when outcomes are observed, feelings of fairness will obviously be important in many situations outside the lab. It can feel inappropriate not to reward a successful manager despite him profiting from random events occurring in the market. Similarly, a blackjack player may tip the dealer more generously after a good hand. Social comparison motives may thus also loom large in the evaluation of agents outside the current experimental setup and may contribute to outcome bias: it may simply feel inappropriate not to reward an agent after a good result, even if the way the result was obtained would otherwise be judged negatively. Conditional on some reward being appropriate, the size of the reward may in fact depend on social comparison considerations (e.g., equality considerations). However, to probe the generalizability of the outcome bias in financial agency settings where social preferences may be less directly relevant, we conducted a second experiment that excludes social preferences and gives further insights into the interaction of outcome and decision-process evaluations.

3 Experiment 2

3.1 Methods

The second experiment elicits judgments of an agent's investment decision (process), and the resulting investment outcome, by a principal. We employ an unincentivized vignette format in this experiment for two reasons. First, the design allows us to exogenously manipulate different investment levels (risky vs. safe) and different uncertain outcomes (success vs. failure). Second, by directly eliciting measures of satisfaction we prevent social preference issues that become relevant in the allocation of principals' and agents' payoffs with decision-based monetary payoffs.

In this experiment, we present hypothetical scenarios involving a financial advisor who is tasked to allocate \$10,000 between a safe and a risky asset for the participant. The scenario is identical to the Gneezy-Potters task used in Experiment 1.¹⁶ We employ two possible allocations, with either low (\$1500; denoted by L) or high (\$8,500; denoted by H) investments in the risky asset and the remainder being invested in the safe asset. In addition to the general scenario and the description of the two assets, we present the agents' investment decision and, depending on the treatment, the outcome of the risky investment. Participants are asked to indicate separately their satisfaction with the investment decision (*decision-process satisfaction*) and, if known, the outcome (*outcome satisfaction*) on a 7-point Likert

¹⁶ We familiarize participants with the investment situation by having them calculate the payoffs for different outcomes of a hypothetical \$5,000 investment in each type of asset.

scale:¹⁷ “How satisfied are you with the **investment decision** the adviser took for you?” and “How satisfied are you with the **outcome** of the investment decision the adviser took for you?” (emphasis in the original). That is, the survey carefully distinguished between the decision to invest a certain level, and the success or failure of the investment. Without explicitly asking for both aspects separately, participants might have construed the term “decision” in a way that comprises the resulting outcome (Blank et al. 2015). By separating the two aspects, participants can signal discontent with a decision that does not fit their risk appetite, while at the same time acknowledging their happiness about the outcome (or vice versa).

The scenario, the advisor’s decision, as well as outcome information are presented on the same screen as the questions regarding participants’ satisfaction. The experiment concludes with a short questionnaire collecting age, gender, education level and current occupation. Table 3 provides an overview of the six between-subjects conditions as well as the respective number of observations. In the current experiment, the ‘unknown’ condition (denoted by ?) corresponds to a situation of REWARD BEFORE. That is, participants know the invested amount (H or L) but not the outcome and only give their decision-process satisfaction rating. The ‘favorable’ (+) and ‘unfavorable’ (–) conditions correspond to the situation of REWARD AFTER, in Experiment 1. In these conditions, participants were asked for both decision-process satisfaction and outcome satisfaction ratings after learning about both the invested amount and the random outcome (investment success or failure).

3.2 Setting and summary statistics

In total, 297 volunteers, recruited on Amazon Mechanical Turk, completed the online experiment and received a compensation of \$0.50 each for their participation, which took a little more than 5 min on average. The actual survey was implemented using SoSci Survey (Leiner 2014). As part of the study description on Amazon Mechanical Turk, we mentioned being “interested in how people judge certain situations”,¹⁸ but neither revealed the research question nor that there were different conditions. We made sure that participants could only take part in the study exactly once and restricted the sample to participants from the US to avoid language barriers and ensure a minimum of homogeneity in the cultural background.

With an average age of close to 39 years, our online sample is older and more heterogeneous than the student samples participating in the laboratory experiments. At the same time, online participants are also much more diverse in their academic background. Only 3% and 14% are trained in economics and business, respectively, while 4% are psychologists. Females comprise 47% of the sample.

¹⁷ Our Likert scales range from „very dissatisfied“ (1) to „very satisfied“ (7). Numbers are not shown.

¹⁸ The instructions to Experiment 2 are available in the supplementary material.

Table 3 Six Treatments in Experiment 2

Outcome investment	Unknown		Favorable		Unfavorable	
	Name	# obs	name	# obs	name	# obs
Low	L ?	51	L +	50	L -	48
High	H ?	51	H +	48	H -	49

3.3 Results and discussion

Outcome satisfaction ratings for the four treatments in which the outcome of the investment decision was available to the participants are shown in Table 4 (upper panel) and Fig. 3 (upper part). Recall that participants in the two known outcome treatments (favorable and unfavorable) were asked to give both, a rating of outcome satisfaction and a rating of decision-process satisfaction. As expected, participants indicate significantly higher satisfaction with favorable compared to unfavorable outcomes for both low and high investment amounts in the risky asset.

Next, we consider participants' satisfaction with the investment *decision-process* itself, rather than with the random *outcome*. Table 4 (lower panel) and Fig. 3 (lower part) summarize the findings. As a first result we find support for the common observation of risk aversion in the current investment setting with potential losses: mean satisfaction with the decision is generally higher for low investment compared to high investment in the absence of outcome information (two-sample, two-sided *t*-test, $M = 4.39$ vs. $M = 2.39$, $t(100) = 6.10$, $p < 0.001$, $d = 1.21$). Rating patterns in the unknown outcome treatments further support this observation: For the low investment in the risky asset, the distribution of ratings is almost uniform, while it is clearly skewed towards a negative evaluation in the high investment case (see Fig. 3).

Now, we consider decision-process satisfaction ratings across the different outcomes for each investment level. While outcome satisfaction is expected to be affected by the randomly determined favorable or unfavorable outcome, decision-process satisfaction is not. Irrespective of the outcome, the hypothetical agent made the same decision based on the same information set. Participants were given information on both the decision decision-process and the outcome and had the possibility to indicate satisfaction separately for the outcome and the decision-process. Absent outcome bias, principals' satisfaction with the decision-process should not be affected by the randomly determined outcome. However, consistent with outcome bias, we observe significantly higher ratings of the same decision after a randomly-obtained favorable investment outcome compared to an unfavorable investment outcome, for both investment levels.

Comparing the evaluation of the investment decision in the presence of outcome information to the situations where participants judged the decision-process in the absence of outcome information, we observe that favorable outcomes have a strongly positive effect, while unfavorable outcomes have a more modest negative effect on decision-process judgments. These results are confirmed in a multivariate analysis (Appendix 2). Observed outcomes have an effect both on outcome

Table 4 Investment outcome and process satisfaction

Treatment	L?	L+	L-	H?	H+	H-
Outcome	n.a.	6.24	3.38	n.a.	6.58	1.35
		***			***	
Process	4.39	5.58	4.54	2.39	4.85	1.73
	***		***		***	

					*	

Means and two-sided t-tests on satisfaction ratings reported.

*, *** Denote significance at the 5%, 0.1% level.

Outcome test statistics:

L+ vs. L-: $t(96) = -10.52$, $p < 0.001$, $d = -2.13$.

H+ vs. H-: $t(95) = -21.65$, $p < 0.001$, $d = -4.40$.

Process test statistics:

L? vs. L+: $t(99) = 3.67$, $p < 0.001$, $d = 0.73$.

L+ vs. L-: $t(96) = -3.46$, $p < 0.001$, $d = -0.70$.

L? vs. L-: $t(97) = 0.45$, $p = 0.657$, $d = 0.09$.

H? vs. H+: $t(97) = 6.87$, $p < 0.001$, $d = 1.38$.

H+ vs. H-: $t(95) = -8.99$, $p < 0.001$, $d = -1.83$.

H? vs. H-: $t(98) = -2.31$, $p < 0.050$, $d = -0.46$.

satisfaction and on decision-process satisfaction. The effect is stronger for outcome satisfaction, but still economically and statistically significant for decision-process judgments. Positive effects for favorable outcomes on decision-process judgments are more pronounced than the negative effects of unfavorable outcomes, for both investment levels. The latter effect is consistent with rewarding behavior in Experiment 1 that also hints at a positive bias. The absolute difference in average rewards between REWARD BEFORE (unknown outcome) and REWARD AFTER (known outcome) is larger for favorable (9.92) than for unfavorable outcomes (8.32); however, the difference is not significant.¹⁹

Experiment 2 results also challenge a possible explanation of the effect in terms of experimenter demand. As we observed, outcomes are very salient, both favorable and unfavorable ones. If participants believe that every piece of information provided by the experimenter is relevant for the situation at hand and should inform their decision, we expect that explicitly mentioning successful or unsuccessful investment outcomes should have a comparable effect on decision-process evaluations. This is not the case. Favorable outcomes clearly have a stronger effect

¹⁹ REWARD BEFORE, $M = 18.86$ vs. REWARD AFTER, favorable outcome, $M = 28.78$: two-sample, two-sided t-test, $t(222) = -2.35$, $p < 0.05$, $d = -0.37$. REWARD BEFORE, $M = 18.86$ vs. REWARD AFTER, unfavorable outcome, $M = 10.54$: two-sample, two-sided t-test, $t(279) = 2.96$, $p < 0.01$, $d = 0.36$. Note that the identification of positive bias is more difficult in Experiment 1 because of the endogenous amount of investment.

Fig. 3 Investment Outcome and Process (Decision) Satisfaction. Satisfaction ratings from 1 “very dissatisfied” to 7 “very satisfied”

and in the low investment condition, the comparison between unknown outcomes and unfavorable outcomes is even insignificant. That is, if there is experimenter demand, it would be highly asymmetric, and driven by outcomes. Information per se does not seem to affect decision-process evaluations.

In addition, recall that we ask participants to evaluate their satisfaction with the decision process and the outcome in two separate questions. We do not only make the difference salient and give participants the opportunity to cleanly distinguish between the two aspects but even specifically demand them to do so. Clearly, participants' outcome satisfaction is driven by the observed outcome. Yet, even if explicitly asked, they are to a large degree unable to prevent outcome information from affecting their decision-process evaluation, as soon as outcome information is available.

In sum, we find clear evidence for the outcome bias in the judgment of agents' investment decisions. Investment decisions were fully observable, and social preference effects were excluded by design.

4 Experiment 3

4.1 Methods

In the first two experiments, principals had to judge how satisfied they were with the investment decision without having to commit to what they consider a good decision *ex-ante*. Although this is a realistic feature in many applied settings, it might have amplified the outcome focus if people construct their preference *ex-post*. In Experiment 3, we reduce this uncertainty about the principals' *ex-ante* preference by letting them state their preferred investment strategy to their agents, who afterwards take a Gneezy-Potters (1997) investment decision for them.²⁰ Note that the investment decision by the agent may be influenced by her own financial interest, which may deviate from the principal's preferences (details are given below). We restrict principals to selecting one of five investment strategies (very conservative, conservative income, balanced, growth, aggressive growth; following the Mutual Fund Dealers Association of Canada (2014)), rather than having them communicate an explicit investment share to their agents. Although strategies are ordered in an unambiguous way, there is variation over the exact interpretation of these verbal categories in terms of the investment share of the risky asset, allowing for "translation errors" in the communication between principal and agent. This allows principals to give agents the benefit of the doubt in case these do not implement the strategy as perceived by the principal. We want to test how principals' satisfaction with the agents' investment decisions is influenced by the

²⁰ Experiment 3 uses data of a larger study analyzing the behavior of financial advisors and clients in an advice relationship. While advisors' behavior is analyzed in a companion paper (Kling et al., 2019), the current paper focusses on the clients' assessment of investment outcomes. The supplementary material reproduces the part of the instructions that is relevant for the decision process evaluation. Further details of the experiment and the instructions for other parts of the study are part of Kling et al. (2019).

outcome of the investment and in particular, whether favorable outcomes make it more acceptable that the agent did not follow the principal's request.

In the first part of the experiment, we ask participants for their perception of the five investment strategy terms. The task is to map the five individual strategies into investment shares (0–100% of total wealth) into a hypothetical risky asset. We keep the risky asset non-specific on purpose, as this is also a feature of real-life risk-classification terms. At the beginning of part two, participants learn about the Gneezy-Potters investment task. Participants consecutively play as both principals and agents (referred to as “client” and “financial advisor” in the experiment). Participants play as principals first and individually choose their preferred verbal investment strategy to be communicated to their agent. Subsequently, roles are switched and participants now take the financial investment decision as agents. On the decision screen, they are reminded of the structure of the risky asset, their compensation, and the principals' investment preferences. While adjusting the amount invested for the principal between 0 and 10€ (in steps of 0.10€, using a slider), they can observe a table of potential payoffs to their principal as well as to themselves which updates in real time. We vary two aspects in the decision by the agents. First, we vary whether an agent serves one or five principals. Second, we vary the incentive structure of the agent (Fixed, Co-Investment, Limited Liability).²¹ We pool these conditions in the current analysis as they regard the agents' rather than the principals' behavior.²²

After all decisions have been made, the roles (principal or agent) are randomly selected to determine financial payoffs for each participant. Principals then see the following information: (i) the (verbal) investment level they demanded; (ii) the actual investment made by their agent; (iii) whether their investment was successful or unsuccessful; and (iv) their payoff. Importantly thus, they are prominently reminded of their preferred strategy, inducing a strong demand for judging on the basis of whether the agent implemented the request of the principal. Principals then pick one of four pre-determined messages to indicate their dissatisfaction or satisfaction with their agents' investment decision: “I am [very dissatisfied / dissatisfied / satisfied / very satisfied] with your investment decision.” Finally, agents learn about the investment results and receive the message sent by their respective principals. The experiment concludes with a short questionnaire on demographics.

²¹ Under Fixed incentives, agents always receive the same fixed amount, independent of their investment for the agent. Under Co-Investment incentives, agents participate to a limited degree in both the gains and the losses of the principal. Finally, under Limited Liability incentives, agents only participate in the gains, but not in the losses. Note that principals were fully aware of the incentive schemes of their agents.

²² Table 10 in Appendix 2 shows that there are no significant interaction effects between the incentive schemes and the random outcome in terms of decision satisfaction. Furthermore, we explicitly control for the agents' incentive schemes as well as principals' risk preferences in Table 11. Our results remain unchanged.

4.2 Setting and summary statistics

Experiment 3 was conducted at AWI-Lab, Heidelberg (GER). The experiment was programmed in oTree (Chen et al. 2016). Sessions and the participant pool were managed with hroot (Bock et al. 2014). The research question was not revealed to participants neither as part of the invitation to partake nor as part of the experiment itself. Each session lasted about 45 min and participants earned an average of €11.85. In total, 324 participants took part in the experiment, yielding 162 observations. The average age of our participants at the time of the experiment was 23 years, 56% were female, and 29% indicated to be studying economics. Approximately 3% were studying psychology.

4.3 Results and discussion

Table 5 shows the satisfaction ratings of principals with their agents' investment decision (decision-process satisfaction). We say an agent *follows* the principal's wish if the invested amount falls into the range of investment shares that the principal associated with the communicated investment strategy.²³ We observe that principals are significantly more satisfied with the investment decision if the result is favorable than if it is unfavorable, replicating the relevance of outcomes. As we would expect, principals are also more satisfied if the agent implements their desired investment ("followed") than if she did not. However, this effect is only significant if the outcome was unfavorable. That is, for favorable investment outcomes, we do not observe a significant effect of the desired investment strategy anymore; the investment outcome moderates the effect of whether the agent implemented the principal's request. Quite strikingly, testing differences along the diagonal reveals that a decision which is in line with the principal's preference but results in an unfavorable random outcome is even seen as significantly less satisfactory than a decision which is at odds with the principal's explicit wish but by mere chance resulted in a favorable outcome (two-sample, two-sided Mann–Whitney-*U* test, average ratings 0.6 vs. 1.1, $z = -2.37$, $p < 0.05$). We also observe that even in the worst case of an unfavorable outcome when the agent did not follow, we only obtain a neutral assessment of -0.2 ($z = -0.97$, $p = 0.33$, Wilcoxon test). This observation is consistent with the positivity bias we also document in Experiments 1 and 2.

A multivariate analysis confirms the initial observations (Appendix 2, Table 11). We regress investment satisfaction on indicators for agents following the principals' wishes, for observing favorable outcomes, and their interaction. Decision-process satisfaction, expressed through the messages sent to agents, is significantly positively affected by observing a favorable outcome of the random draw. There does not seem to be a main effect of observing the agent follow the principal's wish. Testing the linear hypothesis of equality of the respective coefficients, we conclude

²³ While we can cleanly separate our principals into these categories for the analysis, we did not explicitly tell participants whether the investment was compatible with the stated preference in the experiment. Instead, principals had to infer whether the agent followed their wish from the observed investment amount.

Table 5 Decision-process satisfaction

		Investment profile		Difference
		Followed	Not followed	
Outcome	Favorable	1.6	1.1	0.5
	Unfavorable	0.6	− 0.2	0.8**
	Difference	1.0***	1.3***	

The table shows the average satisfaction of principals with the investment decision of their financial agent separated by the agent following or not following the principal's wish and by the random outcome. The scale is − 2 to + 2 for *very dissatisfied* to *very satisfied* with the decision. Statistical significance is based on one-sided, two-sample Mann–Whitney-*U* tests

*, **, *** Denote 5%, 1%, 0.1% significance levels

Favorable: Followed (22) vs. Not Followed (26), $z = -1.68$, $p = 0.092$

Unfavorable: Followed (49) vs. Not Followed (65), $z = -2.72$, $p < 0.01$

Followed: Favorable (22) vs. Unfavorable (49), $z = -4.12$, $p < 0.001$

Not Followed: Favorable (26) vs. Unfavorable (65), $z = -3.79$, $p < 0.001$

that the effect on decision-process satisfaction of observing a favorable outcome is stronger than the effect of recognizing that an agent behaved in the principal's interest.²⁴ The results are unaffected by the inclusion of control variables for principals' risk preferences, agents' incentive schemes, and demographics. Recognizing that principals might appreciate their agents' efforts to implement their desired investment level differently depending on the alignment of interests (through incentive schemes), we also include interactions of the 'follow indicator' and the agents' incentive schemes. The coefficients of the interaction terms remain statistically insignificant.

It is worth recalling how this experiment differs from the previous two studies. First of all, the invested amounts as well as outcome information are given to all principals. There cannot be any effects of having both outcome and decision-process information versus decision-process information only. Any outcome related effect must stem directly from the realizations of the random outcome draw. Principals evaluate agents who behave in line with their preferences significantly better than those who do not if the outcome was unfavorable. Given their ability to thus identify the relevant benchmark for evaluation, it is even more surprising that their decision-process evaluations are even more strongly affected by the arbitrary (because random) outcome information.

Another aspect in which this experiment differs from our first laboratory experiment is the fact that the evaluation takes place in the form of costless messages, rather than payoff-affecting rewards. Thus, social payoff considerations in the sense that gains from positive random outcomes could be shared between the two participants cannot play a role in the decision process evaluation. As it is a one-shot

²⁴ Wald test on the equality of both coefficients in the preferred ordered logit specification: model 3: $\chi^2(1) = 8.07$, $p < 0.01$; model 4: $\chi^2(1) = 9.43$, $p < 0.01$. The results are qualitatively similar in OLS regressions and testing the linear hypothesis with a two-sided F-test: model 1: $F(1, 154) = 7.14$, $p < 0.01$; model 2: $F(1, 147) = 7.31$, $p < 0.01$.

decision, the costless messages also cannot affect subsequent behavior or instill a “team spirit” between the two participants. This would only be possible Experiment 1, with its dynamic multi-round setting and fixed principal-agent pairs. Ruling out social payoff considerations as well as concerns for the behavior in future rounds narrows down the number of alternative explanations for the effects at play.

5 Conclusion

We observe a clear outcome bias in principals’ evaluations and rewards for financial agents in risky investment decisions. In contrast to previous studies in the context of CEO salaries that have observed financial rewards for luck only if principals are weak (Bertrand and Mullainathan 2001), in the current experiment the effect was fully due to the principals’ decision-making. The outcome focus seems normatively questionable because it rewards lucky behavior on the basis of hindsight, rather than to reward good decisions on the basis of the information available to the agent. Importantly, it exists in settings where the decision process is clear and observable, and, therefore, there is no need to draw inferences about the decision-process from the outcome, as would be the case in situations with asymmetric information. We can only speculate on the degree of outcome-dependence in situations where decision-process information is limited or outright unavailable but would expect it to be even more pronounced. In these situations, outcome-dependence might be separable into a pure outcome bias component and a component reflecting the evaluator’s efforts to infer decision-process information from outcome information. The current experiments do not inform us about the relative strengths of these components.

Studying the potential processes lying behind this outcome focus, we found that social preference effects, which may also loom large in situations outside the lab, might be a relevant aspect. Contingent on an outcome-based trigger to reward (random) successes, social comparison may play a role in defining the size of the rewards. However, outcome bias is relevant also in the absence of social comparison as shown in Experiment 2. Moreover, outcome bias seems more pronounced after favorable outcomes than after unfavorable ones. This suggests that justification is an important aspect and with either the decision or the outcome having a stronger influence depending on which turns out more justifiable. In contrast to Gurdal et al.’s (2013) interpretation, blame might not be the main driver of outcome bias in situations of (financial) agency. Our results also provide an interesting exception to the often observed negativity bias (Baumeister et al. 2001).

Additional channels for the occurrence of outcome biases in the current experiments exist. The observed outcome-biased behavior may derive from the fact that in many situations outcomes are indicative of information available to the decision maker but not to the evaluator (Hershey and Baron 1992), or potentially provide the only available basis for judgments of the decision process (Baron and Hershey 1988). Consequently, a focus on outcomes may be inappropriately transferred to situations in which more or even all relevant information on the decision process is available. Future research may fruitfully focus on the information formats that reduce outcome bias in financial agency.

Acknowledgement We are grateful to Matteo Ploner, our referees, the seminar audiences in Berlin, Frankfurt, Heidelberg, Madrid, Munich, Tilburg, the 8th BNUBS-GATE workshop, and the ESA and the Experimental Finance (EF) world meetings for very helpful comments.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Funding Open Access funding provided by Projekt DEAL.

Appendix 1

This Appendix derives the optimal behavior of an expected-utility maximizing agent with fairness preferences of the form proposed by Fehr and Schmidt (1999) for the case of outcomes (REWARD AFTER), and of the form proposed by Trautmann (2009) for the case of expected outcomes (REWARD BEFORE).

The general two-player variant of the utility function of player i in the presence of comparison to player j in the Fehr and Schmidt model is given by

$$U_i(x_i, x_j) = x_i - \alpha_i \max\{x_j - x_i, 0\} - \beta_i \max\{x_i - x_j, 0\}, i \neq j,$$

where x_i and x_j denote the payoffs for each player and α_i and β_i denote the individual's parameters of inequity aversion. It is assumed that players suffer more from disadvantageous inequality than from advantageous ($\alpha_i \geq \beta_i$) and that players do not like to be better off than others ($0 \leq \beta_i < 1$).

In the case outcomes are observable (REWARD AFTER) and turn out favorable, the payoffs to the principal (x_P) and the agent (x_A) are given by

$$x_P = 100 - R + (100 - I) + 3.5I = 200 + 2.5I - R,$$

$$x_A = 100 + R,$$

where R and I denote the reward paid to the agent and the amount invested in the risky asset by the agent, respectively. The principal maximizes her utility

$$U_P(R) = 200 + 2.5I - R - \alpha_P \max\{-100 + 2R - 2.5I, 0\} \\ - \beta_P \max\{100 - 2R + 2.5I, 0\}$$

by choosing the reward $R \in [0, 100]$ optimally. The resulting expected utility-maximizing rewards are shown in the upper panel of Fig. 4. They crucially depend on the level of investment in the risky asset and the parameter of advantageous inequity aversion.

Fig. 4 Fehr and Schmidt's Outcome Fairness, REWARD AFTER

In the case where outcomes are observable but turn out unfavorable, the payoffs for principals and agents equal:

$$x_P = 100 - R + (100 - I) = 200 - R - I$$

$$x_A = 100 + R$$

As a result, the utility function of the principal becomes

$$U_P(R) = 200 - R - I - \alpha_P \max\{-100 + 2R + I, 0\} - \beta_P \max\{100 - 2R - I, 0\}.$$

The principal maximizes her utility by choosing the reward optimally. The expected utility maximizing rewards again depend on the parameter of advantageous inequity aversion and the risky investment by the agent. They are graphically illustrated in the lower panel of Fig. 4.

Figure 4 shows the qualitative predictions of the model in terms of expected utility-maximizing rewards for modest ($\beta \leq 0.5$) and strong ($\beta > 0.5$) inequality aversion. Clearly, the model cannot explain the observed reward pattern in its strict form assuming the same beta parameter for all participants. If we allow for a heterogeneous distribution of beta parameters, a different distribution of parameters is needed for the case of a successful investment versus an unsuccessful investment. For the favorable outcome, subjects should predominantly have large betas > 0.5 . In contrast, for the unfavorable outcome prediction to match the data, subjects should hold small betas ≤ 0.5 .

Trautmann's (2009) model of expected outcome fairness modifies the Fehr and Schmidt model by replacing the comparisons of realized outcomes with comparisons of expected outcomes. The general utility function for player i in the presence of comparison to player j in the two-player case is given by

$$U_i(x_i, x_j) = x_i - \alpha_i \max\{E[x_j] - E[x_i], 0\} - \beta_i \max\{E[x_i] - E[x_j], 0\}, i \neq j.$$

The assumptions about α_i and β_i remain unchanged.

In treatment REWARD BEFORE, only the amount invested in the risky asset is known to the principal at the time she chooses the reward for the agent. Consequently, she does not know her realized payoff and thus chooses the reward to maximize expected utility based on expected payoffs. The expected payoffs for the principal and the agent are given by:

$$E[x_P] = 100 - R + (100 - I) + \frac{1}{3} * 3.5 * I = 200 + \frac{1}{6}I - R$$

$$E[x_A] = 100 + R$$

Accordingly, the principal maximizes the utility function

$$U_P(R) = 200 + \frac{1}{6}I - R - \alpha_P \max\left\{-100 + 2R + \frac{1}{6}I, 0\right\} - \beta_P \max\left\{100 - 2R + \frac{1}{6}I, 0\right\}$$

by choosing the reward optimally. The resulting expected utility-maximizing rewards depend on the risky investment by the agent and the principal's parameter of advantageous inequity aversion. Predictions are shown in Fig. 5.

While the strict form of the model with a unique beta parameter for all principals cannot match the data, assuming a distribution of betas with roughly half of the participants below and above the 0.5 threshold would lead to predictions similar to the actual behavior.

Notes: Scattered observations; linearly fit – solid line; EU-maximizing reward – dashed line.

Fig. 5 Trautmann's Expected Outcome Fairness, REWARD BEFORE

Appendix 2

See below Tables 6, 7, 8, 9, 10 and 11.

Table 6 Summary statistics

Variable	Experiment 1	Experiment 2	Experiment 3
Age (SD)	22.5 (2.8)	38.9 (12.9)	23.3 (4.0)
Female	37.3%	47.5%	55.6%
Dutch	36.2%	–	–
Field of Studies			
Economics	55.2%	2.7%	29.0%
Business	36.6%	13.5%	0.9%
Psychology	1.5%	4.4%	2.8%
Law	3.0%	0.7%	1.8%
Other	3.7%	78.8%	65.4%
Participants	134	297	324
Agents	67	0	324
Principals	67	297	162

Experiment 1 was run in the Netherlands at a university with a large share of foreign students. We did not collect nationality information aside from asking whether participants were Dutch or not. Experiment 2 was run on Amazon Mechanical Turk and was restricted to participants located in the United States of America. Experiment 3 was run in Germany and we did not collect nationality information. In Experiment 1, half of the participants played in the role of agents, the other half in the role of principals. In Experiment 2 everyone took part in the role of the principal. In Experiment 3, participants took on both roles. Everyone made an investment decision as an agent. As payoff relevant roles were determined randomly before principals sent their messages expressing satisfaction with the investment decision, the number of observations is reduced to 162

Table 7 Experiment 1—Relation between agents' risky investment and rewards (non-zero rewards only)

REWARD AFTER	(1)	(2)	(3)	(4)
Investment	0.2733 (0.0847)	0.3214 *** (0.0826)		
...if favorable outcome			0.4988 *** (0.0862)	0.5390 *** (0.0917)
...if unfavorable outcome			– 0.0122 (0.0901)	– 0.0184 (0.0917)
Favorable outcome			17.3052 *** (3.2424)	16.7079 *** (3.3467)
Controls	No	Yes	No	Yes
# Principals	28	28	28	28
# Observations	85	85	85	85
REWARD BEFORE			Placebo test	Placebo test
Investment	0.0066 (0.0783)	– 0.0341 (0.0798)	– 0.0003 (0.0807)	– 0.0222 (0.0816)
Favorable outcome			– 0.8791 (3.7507)	– 2.2769 (3.8649)
Controls	No	Yes	No	Yes
# Principals	30	30	30	30
# Observations	109	109	109	109

Random effects tobit regressions; dependent variable is the size of reward; restricted to non-zero rewards only; average marginal effects reported; standard errors in parentheses

*** Denotes significance at the 0.1% level; controls are: accumulated wealth, age, gender, field of study and Dutch nationality

Table 8 Experiment 2—Multivariate analysis of outcome satisfaction

Low investment	OLS		Ordered Logit	
	Outcome satisfaction		Outcome satisfaction	
Favorable outcome	2.8650 *** (0.2723)	2.9308 *** (0.2659)	3.7128 *** (0.5403)	4.1574 *** (0.5717)
Controls	No	Yes	No	Yes
# Observations	98	98	98	98
High investment	Outcome satisfaction		Outcome satisfaction	
	Outcome satisfaction		Outcome satisfaction	
Favorable outcome	5.2363 *** (0.2419)	5.2752 *** (0.2430)	5.8998 *** (0.8514)	6.1119 *** (0.8907)
Controls	No	Yes	No	Yes
# Observations	97	97	97	97

Base category is unfavorable outcome; standard errors in parentheses

*** Denotes significance at the 0.1% level; controls are age, gender, education level and being an economist

Table 9 Experiment 2—Multivariate analysis of decision satisfaction

Low investment	OLS		Ordered logit	
	Decision-process satisfaction		Decision-process satisfaction	
Unfavorable outcome	0.1495 (0.3211)	0.0831 (0.3179)	0.0759 (0.3559)	0.0106 (0.3607)
Favorable outcome	1.1878 *** (0.3178)	1.0232 ** (0.3158)	1.3982 *** (0.3755)	1.2846 ** (0.3792)
Controls	No	Yes	No	Yes
# Observations	149	149	149	149
High investment	Decision-Process satisfaction		Decision-Process satisfaction	
	Decision-Process satisfaction		Decision-Process satisfaction	
Unfavorable outcome	− 0.6575 * (0.3288)	− 0.7194 * (0.3352)	− 1.1768 ** (0.3951)	− 1.3091 ** (0.4049)
Favorable outcome	2.4620 *** (0.3305)	2.4190 *** (0.3358)	2.2869 *** (0.4006)	2.2573 *** (0.4031)
Controls	No	Yes	No	Yes
# Observations	148	148	148	148

Base category is unknown outcome; standard errors in parentheses

*, **, *** Denote significance at the 5%, 1%, 0.1% level; controls are age, gender, education level and being an economist

Table 10 Experiment 3 – Decision satisfaction, outcome and treatment interactions

	OLS		Ordered logit	
	(1)	(2)	(3)	(4)
Follow	0.7177 ** (0.2255)	0.7072 ** (0.2327)	0.9697 ** (0.3662)	0.9080 * (0.3806)
Favorable Outcome	1.2161 ** (0.3968)	1.1618 ** (0.4014)	2.0845 ** (0.6593)	2.0646 ** (0.6746)
Follow x Favorable	– 0.2316 (0.4097)	– 0.3556 (0.4196)	– 0.0036 (0.6733)	– 0.1375 (0.6977)
Co-Investment	– 0.4242 (0.2708)	– 0.4802 (0.2843)	– 0.6011 (0.4359)	– 0.7556 (0.4682)
Co-Inv. x Favorable	0.1809 (0.5033)	0.3197 (0.5143)	0.2903 (0.8384)	0.5111 (0.8640)
Limited Liability	– 0.4559 (0.2726)	– 0.5414 (0.2825)	– 0.5984 (0.4314)	– 0.7745 (0.4570)
Lim. Lia. x Favorable	– 0.0051 (0.4933)	0.1610 (0.5078)	– 0.3792 (0.7979)	– 0.1403 (0.8329)
Principal's risk appetite indicators	No	Yes	No	Yes
Controls	No	Yes	No	Yes
# Observations	162	162	162	162

Coefficients reported; standard errors in parentheses; dependent variable is decision-process satisfaction from messages sent by principals to agents. Base category for agents' incentive treatments is a fixed payment. Under Co-Investment incentives, the agents' interests are aligned with those of the principals. Under Limited Liability incentives, agents only participate in the upside potential but do not face downward risk from unfavorable outcomes. In models (2) and (4) we include indicators for the riskiness of the investment the principal requested. The base category is the lowest risk appetite and we include individual indicators for the remaining 4 levels

*, ** Denote significance at the 5%, 1% level; controls are: age, gender, and being an economics student

Table 11 Experiment 3—multivariate analysis of decision satisfaction

	OLS		Ordered logit	
	(1)	(2)	(3)	(4)
Follow	0.2651 (0.3397)	0.2659 (0.3488)	0.2866 (0.5482)	0.1888 (0.5625)
Favorable Outcome	1.2768 *** (0.2715)	1.3157 *** (0.2749)	2.0561 *** (0.4771)	2.1784 *** (0.4850)
Follow x Favorable	– 0.2519 (0.4044)	– 0.3490 (0.4155)	– 0.0711 (0.6716)	– 0.1915 (0.6948)
Co-Investment	– 0.5989 * (0.2954)	– 0.6237 * (0.3097)	– 0.8773 (0.4959)	– 1.0635 * (0.5339)
Co-Inv. x Follow	0.4972 (0.4583)	0.5489 (0.4715)	0.7431 (0.7512)	1.0005 (0.7879)
Limited Liability	– 0.8812 ** (0.3089)	– 0.8559 ** (0.3123)	– 1.3837 ** (0.5087)	– 1.4316 ** (0.5203)
Lim. Lia. x Follow	0.8958 * (0.4496)	0.7928 (0.4547)	1.4066 (0.7312)	1.3074 (0.7442)
Principal's risk appetite indicators	No	Yes	No	Yes
Controls	No	Yes	No	Yes
# Observations	162	162	162	162

Coefficients reported; standard errors in parentheses; dependent variable is decision-process satisfaction from messages sent by principals to agents. Base category for agents' incentive treatments is a fixed payment. Under Co-Investment incentives, the agents' interests are aligned with those of the principals. Under Limited Liability incentives, agents only participate in the upside potential but do not face downward risk from unfavorable outcomes. In models (2) and (4) we include indicators for the riskiness of the investment the principal requested. The base category is the lowest risk appetite and we include individual indicators for the remaining 4 levels

*, **, *** Denote significance at the 5%, 1%, 0.1% level; controls are: age, gender, and being an economics student

Appendix 3

In this appendix, we show the development of average rewards over the course of the five rounds of Experiment 1. We do not find evidence for cooperative behavior between agents and principals in the sense of high initial rewards and a sudden drop in the last round. In both treatments, the average rewards including zeros decrease slightly from round 1 to round 2 (Fig. 6, left panel). At this point, the two treatments start to diverge: In REWARD AFTER, the downward trend continues to round 4 only to pick up again to levels similar to that in round 1 in the last round. In REWARD BEFORE, average rewards increase from round 2 to round 4 only to drop to the same level as the rewards in REWARD AFTER (and the first round) in the last round. If we look at non-zero rewards only, the pattern is not conclusive either (Fig. 6, right panel). While in REWARD BEFORE we first observe a drop

Fig. 6 Experiment 1—Average rewards per round

Table 12 Experiment 1—share of zero rewards

Treatment	Round				
	1	2	3	4	5
BEFORE	29%	24%	35%	47%	44%
AFTER	39%	45%	42%	61%	55%
Pooled	34%	34%	39%	54%	49%

from round 1 to round 2, followed by a recovery until round 4 and a drop in the last round, rewards in REWARD after are basically stable at first and, if anything, show an uptick in the last two rounds of the game Table 12.

Looking at the development of the share of zero rewards does not reveal a clear pattern either. Table 12 presents the shares for the two treatments individually, as well as for both treatments pooled. We see a modest increase in zero rewards from round 1 to 5.

References

- Baron, J., & Hershey, J. C. (1988). Outcome bias in decision evaluation. *Journal of Personality and Social Psychology*, 54(4), 569–579.
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is Stronger Than Good. *Review of General Psychology*, 5, 323–370.
- Bertrand, M., & Mullainathan, S. (2001). Are CEOs rewarded for luck? The ones without principals are. *Quarterly Journal of Economics*, 116(3), 901–932.
- Blank, H., Diedenhofen, B., & Musch, J. (2015). Looking back on the London Olympics: Independent outcome and hindsight effects in decision evaluation. *British Journal of Social Psychology*, 54, 798–807.
- Bock, O., Baetge, I., & Nicklisch, A. (2014). hroot – Hamburg registration and organization online tool. *European Economic Review*, 71, 117–120.
- Brandts, J., & Garofalo, O. (2012). Gender pairings and accountability effects. *Journal of Economic Behavior & Organization*, 83(1), 31–41.

- Brock, M. J., Lange, A., & Ozbay, E. Y. (2013). Dictating the risk: Experimental evidence on giving in risky environments. *American Economic Review*, 103(1), 415–437.
- Cappelen, A. W., Konow, J., Sørensen, E. Ø., & Tungodden, B. (2013). Just luck: An experimental study on risk-taking and fairness. *American Economic Review*, 103(4), 1398–1413.
- Casal, S., Ploner, M., & Sproten, A. (2019). Fostering the best execution regime: An experiment about pecuniary sanctions and accountability in fiduciary money management. *Economic Inquiry*, 57(1), 600–616.
- Charness, G., & Jackson, M. O. (2009). The role of responsibility in strategic risk-taking. *Journal of Economic Behavior & Organization*, 69(3), 241–247.
- Charness, G., Karni, E., & Levin, D. (2010). On the conjunction fallacy in probability judgement: New experimental evidence regarding Linda. *Games and Economic Behavior*, 68(2), 551–556.
- Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree – An open-source platform for laboratory, online and field experiments. *Journal of Behavioral and Experimental Finance*, 9, 88–97.
- Crusius, J., van Horen, F., & Mussweiler, T. (2012). Why process matters: A social cognition perspective on economic behavior. *Journal of Economic Psychology*, 33, 677–685.
- Cushman, F., Dreber, A., Wang, Y., & Costa, J. (2009). Accidental Outcomes Guide Punishment in a “Trembling Hand” Game. *PLoS ONE*, 4(8), e669.
- De Oliveira, A. C. M., Smith, A., & Spraggon, J. (2017). Reward the lucky? An experimental investigation of the impact of agency and luck on bonuses. *Journal of Economic Psychology*, 62, 87–97.
- Emerson, G. B., Warne, W. J., Wolf, F. M., Heckman, J. D., Brand, R. A., & Leopold, S. S. (2010). Testing for the presence of positive-outcome bias in peer review: a randomized controlled trial. *Archives of Internal Medicine*, 170(21), 1934–1939.
- Fehr, E., & Schmidt, K. M. (1999). A Theory of Fairness, Competition, and Cooperation. *The Quarterly Journal of Economics*, 114(3), 817–868.
- Fischbacher, U. (2007). Z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171–178.
- Gneezy, U., & Potters, J. (1997). An Experiment on Risk Taking and Evaluation Periods. *The Quarterly Journal of Economics*, 112(2), 631–645.
- Gurdal, M. Y., Miller, J. B., & Rustichini, A. (2013). Why Blame? *Journal of Political Economy*, 121(6), 1205–1247.
- Hershey, J. C., & Baron, J. (1992). Judgment by outcomes: When is it justified? *Organizational Behavior and Human Decision Processes*, 53(1), 89–93.
- Kling, L., König-Kersting, C., & Trautmann, S. T. (2019). *Investing for Others: Principals' vs (p. 674). AWI Discussion Paper Series, No: Agents' Preferences.*
- Krawczyk, M., & Le Lec, F. (2010). ‘Give me a chance!’ An experiment in social decision under risk. *Experimental Economics*, 13, 500–511.
- Lefebvre, M., & Vieider, F. M. (2013). Reining in excessive risk-taking by executives: the effect of accountability. *Theory and Decision*, 75, 497–517.
- Lefebvre, M., Vieider, F. M., & Villeval, M. C. (2011). The ratio bias phenomenon: fact or artifact? *Theory and Decision*, 71(4), 615–641.
- Lefgren, L., Platt, B., & Price, J. (2015). Sticking with What (Barely) Worked: A Test of Outcome Bias. *Management Science*, 61(5), 1121–1136.
- Leiner, D. J. (2014). Software: SoSci Survey (Version 2.5.00-i), available at <https://soscisurvey.de>
- Lipshitz, R. (1989). “Either a medal or a corporal”: The effects of success and failure on the evaluation of decision making and decision makers. *Organizational Behavior and Human Decision Processes*, 44(3), 380–395.
- Marshall, G. W., & Mowen, J. C. (1993). An experimental investigation of the outcome bias in salesperson performance evaluations. *Journal of Personal Selling & Sales Management*, 13(3), 31–47.
- Müller, D., & Weinschenk, P. (2015). Rater Bias and Incentive Provision. *Journal of Economics & Management Strategy*, 24(4), 833–862.
- Murata, A., Nakamura, T., Matsushita, Y., & Morioka, M. (2015). Outcome bias in decision making on punishment or reward. *Procedia Manufacturing*, 3, 3911–3916.
- Mutual Fund Dealers Association of Canada (2014) MFDA Discussion Paper on the Use of Investor Questionnaires available from: <https://mfda.ca/bulletin/Bulletin0611-C/>
- Pollmann, M. M. H., Potters, J., & Trautmann, S. T. (2014). Risk taking by agents: The role of ex-ante and ex-post accountability. *Economics Letters*, 123(3), 387–390.

- Ratner, R. K., & Herbst, K. C. (2005). When good decisions have bad outcomes: The impact of affect on switching behavior. *Organizational Behavior and Human Decision Processes*, 96(1), 23–37.
- Sezer, O., Zhang, T., Gino, F., & Bazerman, M. H. (2016). Overcoming the outcome bias: Making intentions matter. *Organizational Behavior and Human Decision Processes*, 137, 13–26.
- Sirri, E. R., & Tufano, P. (1998). Costly Search and Mutual-Fund Flows. *Journal of Finance*, 53(5), 1589–1622.
- Trautmann, S. T. (2009). A tractable model of process fairness under risk. *Journal of Economic Psychology*, 30(5), 803–813.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.