

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

Christian König-Kersting¹, Monique Pollmann², Jan Potters², and Stefan T. Trautmann^{1*}

¹*University of Heidelberg*; ²*Tilburg University*

21 December 2017

Abstract: We document outcome bias in situations where an agent makes risky financial decisions for a principal. In two 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. The principal thus judges the same decision by the agent differently, depending on factors that the agent has no influence on.

Highlights:

- Random outcomes affect principals' judgment of investments by agents despite knowledge of the exact investment strategy
- Principals are satisfied with an investment after seeing a good payoff despite otherwise negative views of the investment
- Outcome bias affects the decision for monetary awards by the principals, not just their judgments

JEL: D81, G29, C91

Keywords: Decision under risk, decisions of agents, accountability, outcome bias, financial advice

* Corresponding author: Alfred-Weber-Institute for Economics, University of Heidelberg, Bergheimer Strasse 58, 69115 Heidelberg, Germany; Phone: +49 6221 54 2952; Fax: +49 6221 54 3592; email: s.t.trautmann@uvt.nl; We are grateful to 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.

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 (Baron and Hershey, 1988). This tendency is problematic for two 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 (e.g., Bertrand and Mullainathan, 2001, for managerial performance; or Sirri and Tufano, 1998, for investors' mutual fund choices).

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 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).

Seemingly small changes to decision situations can have pronounced behavioral consequences. Especially in regard to cognitive biases, transferability from one situation to 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

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.

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 different incentive conditions and assessment methods by the evaluator.

In two experiments, we document evidence on outcome bias in the evaluation of financial agents who take investment decisions for another person. In the first experiment, 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.² 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 a second study. 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 of the outcome information, i.e., before the uncertainty is resolved). This positivity bias is surprising, given the previous demonstrations of the predominance of negative outcome effects (Gurdal et al., 2013; Ratner and Herbst, 2005). In contrast, the current findings suggest that financial agents seem to benefit from the rule that the *result justifies the deeds*. Recognizing, that past experience can bias future evaluations (cf. *rater bias* in Müller and Weinschenk, 2015), systematically biased assessment of the quality of agents' decisions may be undesirable.

² 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.

The remainder of the paper is laid out as follows. The next two sections describe the design and the results of the first experiment. Section 4 discusses these results and the role of social preferences in outcome bias. Subsequently, section 5 presents the second experiment, aimed at testing the robustness of the results. Section 6 concludes with a general discussion of the role of outcome bias in financial agency.

2. Experiment 1: Design

We use data from Pollmann et al.'s (2014) experiment on risk taking by agents under accountability. 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 receive an additional fixed payment of 100 points each. 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. Points not transferred remain with the principal. The agent receives this reward in addition to her fix 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.³ 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. Then the uncertainty about the risky asset is resolved. 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.

³ 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 treatments are discussed in Pollmann et al. (2014).

The described tasks (investment – reward) are statically repeated five times with fixed principal-agent pairs. 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. 34 principal-agent pairs participated in treatment REWARD BEFORE and 33 principal-agent pairs participated in treatment REWARD AFTER. In the computerized experiment (z-Tree, Fischbacher, 2007), roles were assigned randomly, partner identities were kept secret, and decisions were made anonymously with no communication between principals and agents.⁴

3. Experiment 1: Results

Comparing the rewarding behavior of principals in treatment REWARD AFTER in situations in which the risky asset yielded a positive random outcome to situations in which it yielded a negative one, we observe substantial outcome effects. Pooling observations from all rounds, we find average rewards of 28.78 when favorable outcomes are observed, versus 10.54 when unfavorable rewards are observed.⁵ As a placebo test, we make the same comparison for treatment REWARD BEFORE. Here we find average rewards of 18.72 for favorable random outcomes, versus 18.94 for unfavorable ones.⁶

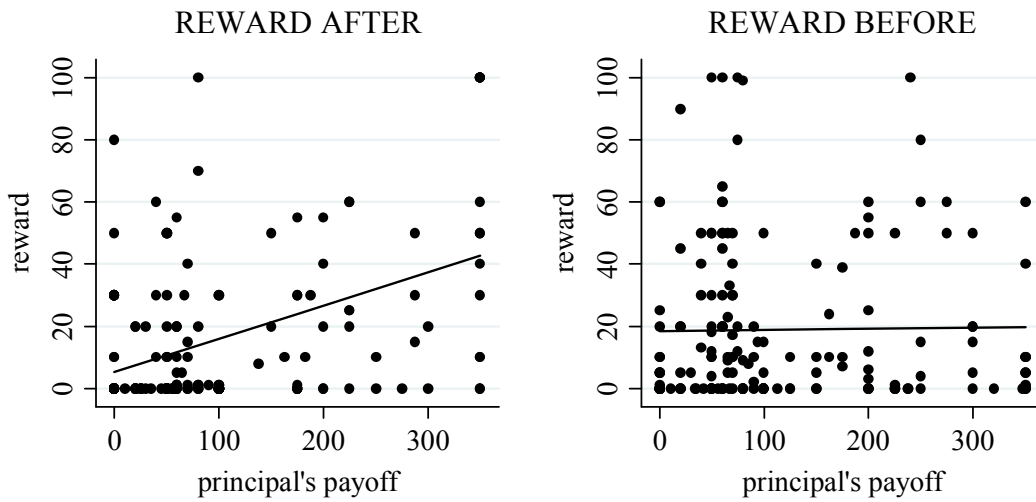
For principals who have received information about the investment decision and outcomes, we furthermore see a significantly positive correlation between their own payoff and the reward they pay to their agent (Figure 1, left panel, $\delta = 0.45$, $p < 0.01$). We do not find a positive correlation if the principal had to reward the agent before knowing the outcomes of the risky investment (Figure 1, right panel, $\delta = 0.02$, $p = 0.83$).

⁴ Points were exchanged for €0.01 each at the end of the experiment. Participants received written instructions and had to complete a set of mandatory comprehension questions. The experiment sessions began only after every participant had correctly answered these questions. 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.

⁵ The difference in average rewards is significant at the 5% level (t-test, $n = 165$). 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.

⁶ The difference in average rewards is neither significant when observations are pooled nor when periods are treated separately.

Figure 1: Relation between principals' payoffs and rewards for the agents



Notes: All periods included; scattered observations with linearly fitted line.

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

	AFTER	AFTER	BEFORE	BEFORE
	Size of Reward	Size of Reward	Size of Reward	Size of Reward
Principal's Payoff	0.0906*** (0.0134)	0.0903*** (0.0129)	-0.0011 (0.0117)	-0.0015 (0.0117)
Controls	No	Yes	No	Yes
# principals	33	33	34	34
# observations	165	165	170	170

Notes: Random effects tobit regression; average marginal effects reported; standard errors in parentheses; *** denotes 1% significance level; controls are: age, gender, field of study and Dutch nationality.

In order 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 age, gender, Dutch nationality,⁷ 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

⁷ 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.

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.⁸

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. 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, conditional on investment success, 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. Conditional on observing an unfavorable outcome, 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 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 investment is observable by the principal (REWARD BEFORE, Table 2 lower panel), we find an modest average increase of 0.09 points per unit increase in the investment. The effect is significant at the 10% level. The placebo test of the effect of the favorable outcome is insignificant. As shown in Table 2, all results are robust to the inclusion of demographic controls.

Figure 2 provides further insights by plotting the marginal effects on rewards for each investment level separately. 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.

⁸ 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.

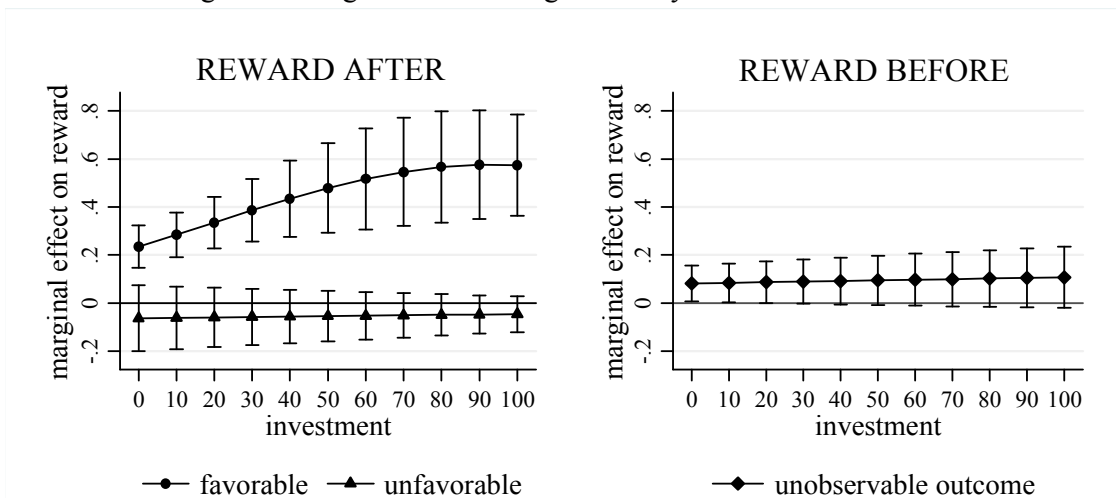
Moreover, the figure confirms the signs of the tobit interaction terms in Table 2. For the case of REWARD BEFORE, the figure shows that investment has a significantly positive effect only at very low investment levels.

Table 2: Relation between agents' risky investment and rewards

REWARD AFTER	Size of Reward	Size of Reward
Investment if observing...		
...a favorable outcome	0.4602*** (0.0785)	0.4845*** (0.0798)
...an unfavorable outcome	-0.0531 (0.0506)	-0.0457 (0.0494)
Favorable Outcome	16.5962*** (2.6101)	16.1321*** (2.5799)
Controls	No	Yes
# principals	33	33
# observations	165	165
REWARD BEFORE	Size of Reward	Size of Reward
Investment	0.0943* (0.0522)	0.0985* (0.0519)
Favorable Outcome	-0.5752 (2.4931)	-0.6896 (2.4094)
Controls	No	Yes
# principals	34	34
# observations	170	170

Notes: Random effects tobit regression; average marginal effects reported; standard errors in parentheses; */*** denotes 10%/1% significance levels; controls are: age, gender, field of study and Dutch nationality.

Figure 2: Marginal effects of agents' risky investment on reward



Notably, the observed pattern of rewards is not consistent with a general wealth effect. If more wealth, 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: wealth 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 agent's payoff 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 a wealth effect conditional on observing a positive outcome is consistent with our results. This conditioning is exactly what outcome bias implies.

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 wealth level 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. We consider the model proposed by Fehr and Schmidt (1999) for the case of observable outcomes (REWARD AFTER), and the model proposed by Trautmann (2009) for the case of unobservable uncertain outcomes (REWARD BEFORE). 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 A 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.

5. Experiment 2: Design and Results

The second experiment elicits judgments of an agent's investment decision, 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.

5.1. Experimental Design

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 1.⁹ We employ two possible allocations, with either low (\$1,500) or high (\$8,500) 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

⁹ 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.

separately their satisfaction with the investment decision and, if known, the outcome on a 7-point Likert scale.¹⁰ 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 corresponds to a situation of REWARD BEFORE, and the ‘positive’ and ‘negative’ conditions correspond to the situation of REWARD AFTER, in Experiment 1. 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 minutes on average. The actual survey was implemented using SoSci Survey (Leiner, 2014).

Table 3: Six Treatments in Experiment 2

Investment \ Outcome	unknown		positive		negative	
	name	# obs.	name	# obs.	name	# obs.
Low investment	L?	51	L+	50	L-	48
High investment	H?	51	H+	48	H-	49

5.2. 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 Figure 3 (upper half). As expected, participants indicate significantly higher satisfaction with positive compared to negative outcomes for both low and high investment amounts in the risky asset.

Table 4: Investment Outcome and Decision Satisfaction

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


Notes: Means and two-sided t-tests on satisfaction ratings reported; **/** denotes 5%/1% significance levels.

¹⁰ Our Likert scales range from „very dissatisfied“ (1) to „very satisfied“ (7). Numbers are not shown.


Figure 3: Investment Outcome and Process (Decision) Satisfaction

Outcome Satisfaction

Low Investment




High Investment




Process Satisfaction

Low Investment



High Investment



Note: Satisfaction ratings from 1 “very dissatisfied” to 7 “very satisfied”.

Next, we consider participants' satisfaction with the investment *decision* itself, rather than with the random *outcome*. Table 4 (lower panel) and Figure 3 (lower half) 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 (4.39 vs. 2.39, $p < 0.01$). 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 Figure 3).

We now consider decision satisfaction ratings across the different outcomes for each investment level. In the absence of outcome bias we would expect there to be no differences in decision-process satisfaction ratings. Participants were given information on both the decision and the outcome, and had the possibility to indicate satisfaction separately for outcomes and decisions. However, consistent with outcome bias, we observe significantly higher ratings of the same decision after a randomly-obtained good investment outcome compared to a bad investment outcome, for both investment levels. Comparing the evaluation of the decision process in the presence of outcome information to the situations where participants judged the process in the absence of outcome information, we observe that good outcomes have a strongly positive effect, while negative outcomes have a more modest negative effect on decision-process judgments.

These results are confirmed in a multivariate analysis (Appendix B). Observed outcomes have an effect both on outcome 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 good outcomes on process judgments are more pronounced than the negative effects of bad 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.¹¹

On 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.

6. Conclusion

We observe a clear outcome bias in principals' evaluations and rewards for financial agents in risky investment decisions. 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 from the outcome, as would be the case in situations with asymmetric information. 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.

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 good outcomes rather than bad 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

¹¹ Average rewards: REWARD BEFORE 18.86; REWARD AFTER, favorable outcome 28.78; REWARD AFTER, unfavorable outcome 10.54. Note that the identification of positive bias is more difficult in Experiment 1 because of the endogenous amount of investment.

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. In the context of financial decision with symmetric information about the investment and a large random component in outcomes, it is important for principals to understand and prevent outcome bias. Future research may fruitfully focus on the information formats that reduce outcome bias in financial agency.

Appendix A

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

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

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 Figure A1. They crucially depend on the level of investment in the risky asset and the parameter of advantageous inequity aversion.

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

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

$$x_A = 100 + R$$

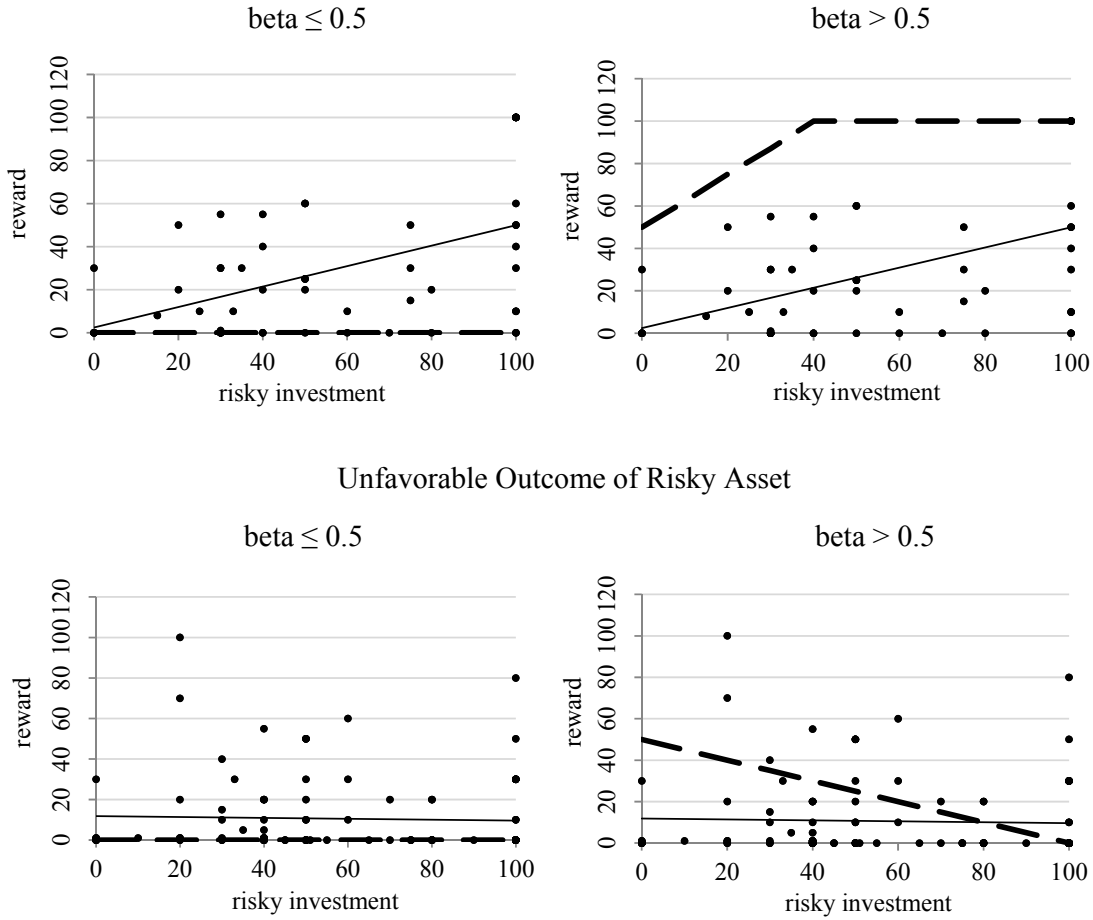
As a result, the utility function of the principal becomes

$$U_p(R) = 200 - R - I - \alpha_p \max\{-100 + 2R + I, 0\} - \beta_p \max\{100 - 2R - I, 0\}.$$

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 Figure A1.

Figure A1: Fehr and Schmidt's Outcome Fairness, REWARD AFTER
Favorable Outcome of Risky Asset



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

Figure A1 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

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

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 Figure A2.

Figure A2: Trautmann's Expected Outcome Fairness, REWARD BEFORE



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

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.

Appendix B

Table B1: 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	OLS		Ordered Logit	
	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

Notes: Base category is unfavorable outcome; standard errors in parentheses; *** denotes 1% significance level; controls are: age, gender, education level and being an economist.

Table B2: Multivariate analysis of decision satisfaction

Low investment	OLS		Ordered Logit	
	Process satisfaction		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	OLS		Ordered Logit	
	Process satisfaction		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

Notes: Base category is unknown outcome; standard errors in parentheses; ***/** denotes 1%/5% significance levels; controls are: age, gender, education level and being an economist.

References

- Baron, J., and 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. and Vohs, K. D. (2001) Bad is Stronger Than Good. *Review of General Psychology*, 5, 323–370.
- Bertrand, M. and Mullainathan, S. (2001). Are CEOs rewarded for luck? The ones without principals are. *Quarterly Journal of Economics*, 116(3), 901–932.
- Charness, G., Karni, E., and Levin, D. (2010). On the conjunction fallacy in probability judgement: New experimental evidence regarding Linda. *Games and Economic Behavior*, 68(2), 551-556.
- Crusius, J., van Horen, F. and 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., and Costa, J. (2009). Accidental Outcomes Guide Punishment in a “Trembling Hand” Game. *PLoS ONE*, 4(8), e669.
- Emerson, G. B., Warme, W. J., Wolf, F. M., Heckman, J. D., Brand, R. A., and 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., and 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., and 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., and Rustichini, A. (2013). Why Blame? *Journal of Political Economy*, 121(6), 1205–1247.
- Hershey, J. C., and Baron, J. (1992). Judgment by outcomes: When is it justified? *Organizational Behavior and Human Decision Processes*, 53(1), 89–93.
- Lefebvre, M., Vieider, F. M., and Villeval, M. C. (2011) The ratio bias phenomenon: fact or artifact? *Theory and Decision*, 71(4), 615-641.
- Lefgren, L., Platt, B., and 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., and 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., and Weinschenk, P. (2015). Rater Bias and Incentive Provision. *Journal of Economics & Management Strategy*, 24(4), 833-862.
- Murata, A., Nakamura, T., Matsushita, Y., and Moriwaka, M. (2015). Outcome bias in decision making on punishment or reward. *Procedia Manufacturing*, 3, 3911-3916.
- Pollmann, M. M. H., Potters, J., and 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., and 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., and Bazerman, M. H. (2016). Overcoming the outcome bias: Making intentions matter. *Organizational Behavior and Human Decision Processes*, 137, 13-26.
- Sirri, E. R. and 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.