Outcome bias and risk taking in a principal agent setting - experimental design and pre-analysis plan

Moritz Loewenfeld

April 10, 2023

1 Research Question

I consider a setting of delegated risk taking. Agents choose between a first-order stochastically dominant and a dominated lottery. Principals observe choices and outcomes of both lotteries and then decide whether to award a bonus payment to the agent. The goal of this experiment is to study whether outcome bias (OB), that is a tendency to condition bonus payments on outcomes, can shape the incentives faced by agents and thereby their choices. In particular, I seek to address the following research questions. 1) Can outcome bias in bonus decisions eliminate incentives to choose optimal actions? 2) Do agents anticipate the OB of principals correctly and 3) do they adjust their choices accordingly, i.e. can outcome bias induce more choices of sub-optimal actions and thus decrease welfare?

2 Experimental design

Agents make a number of decisions between two lotteries on behalf of the principals. Principals decide on bonus payments. Participants are randomly and permanently assigned to the role of either principal or agent.

There are two treatments. Treatments are assigned at the session level. In the reward-after treatment, principals make bonus decisions for all possible choice-outcome combinations (strategy method). In the reward-before treatment, principals condition only on the agents' choices, but not their outcomes. Principals make a total of 54 bonus decisions in the reward-after treatment and a total of 18 bonus decisions in the reward-before treatment. See table 1 and 2 for the lottery tasks employed. In both treatments, the bonus amounts to 10 Euros. Whenever the bonus is not allocated to the agent matched to the principal, it is allocated to another, randomly chosen, agent.

The order in which principals make bonus decisions for the different choice or choice-outcome combination is randomized between subjects. Moreover, the display of the choice tasks (order of states, position of lottery G and B) is randomized between scenarios (choice of the agent in reward-before treatment or choice and outcome in reward-after treatment) and subjects.

	Correlation 1					Correlation 2		
	1(1/3)	2(1/3)	3(1/3)			1(1/3)	2(1/3)	3(1/3)
G	H + ϵ	M + ϵ	L + ϵ	$\rightarrow$	G	H + ϵ	M + ϵ	L + ϵ
B	M	L	H		B	L	H	M

Table 1 The first row presents different possible states and their probability of occurring. Rows 2 and 3 display the payoffs of option G (FOSD) and B in the different states of the world.

lottery pair	H	M	L	ϵ
1	1953	1031	109	45
2	1953	1031	109	110
3	1403	688	103	359
4	1403	688	103	523
5 (corr 2 only)	1480	750	50	699

Table 2 Parameter values for the different lotteries. All payoffs are in cents.

Each agent makes 2 choices for each of the 9 lottery pairs. In addition, beliefs are elicited. Agents make a first choice for each of the 9 lottery pairs. Thereafter, their beliefs are elicited and they make a second choice for each lottery pair. In the reward-before treatment, agents are asked, for each choice task, how likely they are to receive the bonus when choosing either lottery in their choice set. In the reward-after treatment, agents are asked to state their beliefs conditional on their choice and the outcome of the lotteries. Beliefs are incentivized using the binarized scoring rule.

Each Agent is randomly paired with one principal. For each pair, one of the actions taken by the agent is randomly chosen. The action of agent and the bonus decision of the principal is implemented.

In a addition to the above, principals also make choices between the lotteries used in part I of the experiment for themselves. In addition, 3 multiple choice lists are used to elicit risk preferences.

With 80% probability, participants are paid based on the principal-agent interaction. With 20% probability, principals are paid based on their choices in the risk tasks, and agents are paid based on their beliefs (random-incentive mechanism).

All subjects will further answer a questionnaire (see section 2.1).

The experiment will be conducted at the lab at Toulouse School of Economics, in April 2023. The target number for participants is 300, with 140 participants in the reward-before and 160 participants in the reward-after treatment, and an equal number of principals and agents. The target is to have at least 8 participants in each session.

Participants will receive a show-up fee of 5 Euros and expected additional earnings of around 10 Euros for an expected duration of 60-75 minutes.

2.1 Questionnaire items

The questionnaire (non-incentivized) will contain the following items

- An extended version of the cognitive reflection test (Frederick 2005, Toplak et al. 2014)
- Willingness to take risk (Dohmen et al. 2011)

- Standard demographics: These include age, gender, field of study, nationality, level of education, and household income.

Moreover, subjects in the role of principal are asked to which extend their bonus decisions were impacted by 1) the agent's choice 2) the obtained outcome 3) a comparison between the obtained and the forgone outcome 4) a tendency to award the bonus to the matched agent rather than to a randomly chosen agent. Agents are asked to which extend their lottery choices were driven by 1) a desire to make good choices 2) maximization of the probability to obtain the bonus and 3) whether they sometimes made choices they thought were not in the best interest of the principal because this might not maximize their reward probability.

3 Model

The experimental design and predictions are based on a model in which bonus decisions depend on counterfactual evaluation, that is a comparison between the payoff the agent obtained with the forgone payoff, the payoff they could have obtained, had they chosen a different lottery. In the model, principals are motivated by reciprocity to reward agents for good choices. However, their perception of what the good choice is biased by observing the outcome.

The principal's utility is described by some value function $v()$. If state s materializes, the principal enjoys an (ex-post) utility $v(x_s^\theta)$ from the payoff x_s^θ yielded by lottery $\theta \in \{G, B\}$ in this state. The ex-ante value of lottery $\theta \in \{G, B\}$ is given by $V(\theta) = \sum_{s \in S} p_s v(x_s^\theta)$.[1]

In the model, if state s materialized, the principal rewards a choice of the lottery G iff

$$\Delta(\tilde{V}_s) = \lambda \underbrace{[v(x_s^G) - v(x_s^B)]}_{\text{Ex-post comparison}} + (1-\lambda) \underbrace{[V(G) - V(B)]}_{\text{Ex-ante comparison}} \geq 0$$

,where $\lambda \in [0, 1]$ denotes the principal's degree of outcome bias. If the agent chose lottery B, the principal's perceived goodness of the agent's choice is $-\Delta(\tilde{V}_s)$. The key features of the model are that the bonus probability is increasing in the quality of choice (choosing the FOSD lottery), increasing in the obtained outcome and decreasing in the forgone payoff.

Denote $P_{lp,j}(\theta^{corr}) = \sum_s p_s P_{j,lp,s}(\theta^{corr})$ the probability of receiving the reward for lottery pair $lp \in \{1, 2, 3, 4\}$, from a principal $j \in J$, after lottery choice $\theta \in \{G, B\}$ under correlation structure $corr \in \{1, 2\}$.[2] The incentives for the agents to choose option G rather than B for a given lottery pair lp under correlation structure $corr \in \{1, 2\}$ set by principal j are given by the difference in probabilities of receiving the bonus when choosing option G instead of option B, that is $I_{lp,j}(G^{corr}) = P_{lp,j}(G^{corr}) - P_{lp,j}(B^{corr})$.

The key implications of the model are the following. First, $I_{lp,j}(G^{corr})$ is weakly decreasing in λ_j. Second, even if $\lambda_j = 1$, agents have positive incentives to choose the dominant lottery under

[1] $v()$ could be, but need not be, a v.N.M utility function in which case the decision maker's preferences would satisfy expected utility theory.

[2] The lottery for $lp = 5$ is included solely to aid the structural estimation of λ and is therefore not included in the discussion.

correlation 1, i.e. $I_{lp,j}(G^1) > 0$ for all $lp \in \{1,2,3,4\}$. However, under correlation 2, there exists $\overline{\lambda}_{lp}$ such that $I_{lp,j}(G^2) < 0$ for all $\lambda_j \in (\overline{\lambda}_{lp}, 1]$. Moreover, the choice tasks are chosen such that $\overline{\lambda}_1 < \overline{\lambda}_2 < \overline{\lambda}_3 < \overline{\lambda}_4$. It further holds true that $I_{1,j}(G^{corr}) \leq I_{2,j}(G^{corr}) \leq I_{3,j}(G^{corr}) \leq I_{4,j}(G^{corr})$, for $corr \in \{1,2\}$.[3] In the experiment, agents are randomly matched to principals. Their incentives to choose lottery G are therefore given by $I_{lp}(G^{corr}) = \frac{1}{|J|}\sum_{j\in J} I_{lp,j}(G^{corr})$. The different lottery pairs allow to gauge the strength of correlation sensitivity, both at the individual and the aggregate level (see section 5 for further discussion).

Within the model, the reward-before treatment can be thought of as forcing $\lambda = 0$. The reward-before condition thus provides a baseline against which to compare the behavior in the reward-after treatment. The above discussed patterns should not occur in this treatment.

4 Main Hypotheses

All null hypotheses in this section will be tested at the 5% significance level. For within subject tests, I will use the Wilcoxon signed-rank test. For between subject tests, I will use the Wilcoxon rank-sum test.

4.1 Variable construction

Denote $P_{lp,j}(\theta^{corr}, T)$ the probability of receiving the reward for lottery pair $lp \in \{1, ..., 4\}$ when being paired with principal j, after choice $\theta \in \{G, B\}$ and in treatment $T \in \{after, before\}$. In the $before$ treatment, $P_{lp,j}(C^{corr}, before) \in \{0, 1\}$. The probability of receiving the bonus after choosing a given option in the reward-after treatment will be calculated as $P_{lp,j}(C^{corr}, after) = \sum_s p_s P_{lp,j,s}(C^{corr}, after)$, where s are the possible states of the world and p_s the associated probabilities. Hence $P_{lp,j}(C^{corr}, before) \in \{0, 1/3, 2/3, 1\}$.

I define the incentives to choose option G rather than B as the difference in probabilities of receiving the bonus when choosing option G instead of option B, that is $I_{lp}(G^1, T) = P_{lp,j}(G^{corr}, T) - P_{lp,j}(B^{corr}, T)$. In all ensuing tests, I will omit the subscripts j and lp for brevity. Tests will be carried out for the different lottery pairs separately.

4.2 Preliminary: Outcome bias in bonus decisions

Hypothesis 1. ***Outcome bias in bonus decisions:*** *Principals are more likely to award the bonus if*

a) Agents chose their preferred lottery.

b) The obtained outcome is greater than the forgone outcome.

I estimate the following random-effects logit regression model.

$$Bonus_{i,t} = \beta_0 + \beta_1 preferred_{i,t} + \beta_2 obtained\ payoff_{i,t} + \beta_3\{obtained > forgone\}_{i,t} + \epsilon_{i,t} \quad (1)$$

[3] This assumes that v is a v.N.M. utility function with reasonable parameters of risk aversion.

, where $Bonus_{i,t}$ is a dummy variable indicating whether the principal awarded the bonus or not, $preferred_{i,t}$ indicates whether the agent chose the principal's preferred lottery, as measured by her own choices, *obtained payoff*$_{i,t}$ denotes the payoff the principal obtained, in Euro, and $\mathbb{1}\{obtained > forgone\}_{i,t}$ is a dummy variable indicating whether the obtained outcome is higher than the forgone alternative. The regression model will be estimated for the reward-before and the reward-after condition separately.

To test hypothesis 1, I test the following null hypotheses using Wald Chi-Square tests, with standard errors clustered at the subject-level.

- For hypothesis 1a) I test the null hypothesis that $\beta_1 = 0$, against the alternative hypothesis that $\beta_1 > 0$, for both treatments
- For hypothesis 1b) In the reward-after (reward-before) treatment I test the null hypothesis that $\beta_3 = 0$, against the alternative hypothesis that $\beta_3 > 0$ ($\beta_3 \neq 0$). For the reward-before treatment, this hypothesis is not expected to be rejected ("placebo-test").[4]

4.3 OB and incentives to choose lottery G

In hypothesis 2, I collect the hypotheses on how OB affects the agents' incentives to choose the dominant lottery.

Hypothesis 2. ***Outcome bias and incentives-aggregate level:***

a) $I(G^1, before) > 0$ *and* $I(G^2, before) > 0$. *Moreover, no significant difference arises between* $I(G^1, before)$ *and* $I(G^2, before)$.

b) $I(G^{corr}, before) > I(G^{corr}, after)$, *for* $corr \in \{1, 2\}$.

c) $I(G^1, before) - I(G^2, before) < I(G^1, after) - I(G^2, after)$.

I will test the following null hypotheses.

- For hypothesis 2a) I test the null hypotheses that $P(G^{corr}, before) = P(B^{corr}, before)$, for $corr \in \{1, 2\}$, against the alternative hypothesis that $P(G^{corr}, before) > P(B^{corr}, before)$. I also test the null hypothesis that $I(G^1, before) = I(G^2, before)$, against the alternative $I(G^1, before) \neq I(G^2, before)$. Wilcoxon signed-rank tests will be used.
- For hypothesis 2b), I test the null hypotheses that $I(G^{corr}, before) = I(G^{corr}, after)$, for $corr \in \{1, 2\}$, against the alternative that $I(G^{corr}, before) > I(G^{corr}, after)$. Wilcoxon rank-sum tests will be used.
- For hypothesis 2c), I test the null hypotheses that $I(G^1, before) - I(G^2, before) = I(G^1, after) - I(G^2, after)$, for $corr \in \{1, 2\}$, against the alternative that $I(G^1, before) - I(G^2, before) < I(G^1, after) - I(G^2, after)$. A wilcoxon rank-sum test will be used.

[4] For clarity: Whenever a clear direction of the alternative is stated (<, or >, one sided tests will be used. If no direction is stated for the alternative ($\neq$), two-sided tests will be used.

4.4 The agent's choices

Denote $F(G^{corr}, treatment)$ the frequency with which agents choose the dominant lottery under correlation $corr \in \{1, 2\}$, in $treatment \in \{before, after\}$. I average the choice frequency for both choices agents make for each choice task.

Hypothesis 3. ***The agents' choices***

a) In the reward-before treatment, agents choose the dominant lottery at a high frequency, under both correlation structures. Moreover, no significant difference arises between $F(G^1, before)$ *and* $F(G^2, before)$.

b) $F(G^{corr}, before) > F(G^{corr}, after)$*, for* $corr \in \{1, 2\}$.

c) $F(G^1, before) - F(G^2, before) < F(G^1, after) - F(G^2, after)$.

I will test the following null hypotheses.

- For hypothesis 3a) I test the null hypotheses that $F(G^1, before) = F(G^2, before)$ against the alternative that $F(G^1, before) \neq F(G^2, before)$, using a Wilcoxon signed-rank test. The hypothesis is not expected to be rejected.

- For hypothesis 3b) I test the null hypotheses that $F(G^{corr}, before) = F(G^{corr}, after)$, against the alternative that $F(G^{corr}, before) > F(G^{corr}, after)$, for $corr \in \{1, 2\}$, using a Wilcoxon rank-sum test.

- For hypothesis 3c) I test the null hypotheses that $F(G^1, before) - F(G^2, before) = F(G^1, after) - F(G^2, after)$, against the alternative that $F(G^1, before) - F(G^2, before) < F(G^1, after) - F(G^2, after)$ using a Wilcoxon rank-sum test.

5 Further analysis

In a preliminary step, I will test for correlation sensitivity in the principals' lottery choices they make for themselves. The hypotheses above are derived under the assumption that the change in the correlation structure does not meaningfully influence the principals' preferences. If the principals' preferences are found to be strongly correlation-sensitive (which is not expected), hypotheses 2 and 3 should not be expected to hold true.

An important question is to what extent agents are capable of anticipating the principals' OB and how the OB impacts their incentives to choose between the different lotteries. I will thus analyze how well agents' beliefs reflect the principals' bonus decisions. Although the working hypothesis is that agents form accurate beliefs, this analysis is descriptive and somewhat exploratory in nature. Therefore, no specific hypotheses are specified here.

Further, the different lottery pairs allow to gauge the usefulness of the model. The model predicts that OB should lead to greater changes in incentives and lottery choices for lottery pairs with lower indices. Therefore, hypotheses 1b), 2 b) and c), and 3 b) and c) (given that agents form accurate beliefs) are most likely to hold true for $lp = 1$, are somewhat less likely for $lp = 2$, and are more unlikely for $lp = 3$ and even less likely for $lp = 4$, since higher and higher levels of outcome bias are required.

I will further estimate the model structurally. For each principal in the reward-after condition, I will estimate an individual level of outcome bias λ. From agents' beliefs in the reward-after condition, I will estimate their perceived level of outcome bias in the population of principals. This exercise will allow quantifying OB and perceived OB within my model. It will also facilitate the study of heterogeneity. In particular, I will test for a correlation between the principals' λ and their performance in the extended CRT. Data from a previous experiment indicated that higher λ are correlated with lower CRT scores, which motivates examining this particular correlation.

I will further explore correlations between demographic variables, questionnaire responses, and subjects' behavior.

References

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J. & Wagner, G. G. (2011), 'Individual risk attitudes: Measurement, determinants, and behavioral consequences', *Journal of the European Economic Association* **9**(3), 522–550.

Frederick, S. (2005), 'Cognitive reflection and decision making', *Journal of Economic perspectives* **19**(4), 25–42.

Toplak, M. E., West, R. F. & Stanovich, K. E. (2014), 'Assessing miserly information processing: An expansion of the cognitive reflection test', *Thinking & reasoning* **20**(2), 147–168.