

A Service of



Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Gillenkirch, Robert M.; Velthuis, Louis

Article — Published Version Delegated risk-taking, accountability, and outcome bias

Journal of Risk and Uncertainty

Provided in Cooperation with: Springer Nature

Suggested Citation: Gillenkirch, Robert M.; Velthuis, Louis (2023) : Delegated risk-taking, accountability, and outcome bias, Journal of Risk and Uncertainty, ISSN 1573-0476, Springer US, New York, NY, Vol. 67, Iss. 2, pp. 137-161, https://doi.org/10.1007/s11166-023-09414-2

This Version is available at: https://hdl.handle.net/10419/307091

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.



WWW.ECONSTOR.EU

https://creativecommons.org/licenses/by/4.0/

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.





Delegated risk-taking, accountability, and outcome bias

Robert M. Gillenkirch¹ · Louis Velthuis²

Accepted: 24 May 2023 / Published online: 12 September 2023 © The Author(s) 2023

Abstract

In a sequence of experiments, this study investigates how people evaluate others who make risky decisions on their behalf, and how such evaluations affect delegated risktaking. A decision maker acts on behalf of a client who holds the decision maker accountable by way of a subjective evaluation after observing a risky decision's outcome. If evaluation is biased towards the outcome, it may have dysfunctional effects with respect to delegated risk-taking in that decision makers' risk choices are increasingly misaligned with their clients' risk preferences. We find evidence giving support to this conjecture. Across and within three experiments, we test for the effects of different types and degrees of accountability in that we manipulate the information available to clients as well as the consequences which evaluations have for decision makers. Evaluations are biased towards outcomes in all experiments. When evaluations affect decision maker's compensations, a stronger outcome bias in evaluations translates into risk-taking decisions being less frequently aligned with clients' risk preferences. In the same situation, giving clients the opportunity to make peer comparisons increases outcome bias. We further find that clients do not hold decision makers accountable for their risk choices when they cannot observe the risk-taking decision, but have to infer it from observing the outcome. Theoretical and practical implications of the results are discussed.

Keywords Risk-taking \cdot Accountability \cdot Outcome bias \cdot Subjective evaluation \cdot Relative outcome information

JEL classification $D81 \cdot D91 \cdot G41$

Louis Velthuis velthuis@uni-mainz.de

¹ Faculty of Business and Economics, Department of Business, University of Osnabrück, D-49069 Osnabrück, Germany

² Faculty of Law, Management and Economics, Gutenberg School of Management & Economics, University of Mainz, D-55099 Mainz, Germany

Robert M. Gillenkirch robert.gillenkirch@uni-osnabrueck.de

1 Introduction

Many decisions involving risk are not made by individuals themselves, but are delegated to others (e.g., business managers, financial or legal advisors). Then the question arises to what extent the decision maker considers his¹ client's risk preferences. Both his responsibility and whether and how he is held accountable for his decisions and their outcomes by his client will affect his actions (Batteux et al., 2019; Bolton et al., 2015; Pahlke et al., 2015). This study focuses on different types and degrees of accountability (Lerner & Tetlock, 1999; Tetlock, 1985) to reach a better understanding of interpersonal risk-taking. Accountability types refer to what a decision maker is held accountable for (Patil et al., 2014): His actions (process accountability) or his actions' outcomes (outcome accountability). With degrees of accountability, we refer to the consequences for the decision maker, which may range from mere criticism to serious material or immaterial consequences (Lerner & Tetlock, 1999). We consider clients' evaluations, a way of holding decision makers accountable that is present in many situations and that is subjective in nature. Even though an evaluation made by a client has the potential of giving a decision maker an incentive to act in her best interest, it is prone to outcome bias, i.e., to the client holding the decision maker accountable for events he is not responsible for (Baron & Hershey, 1988; Brownback & Kuhn, 2019; Gurdal et al., 2013). In order to investigate both the determinants of outcome bias in evaluations of delegated risk choices and the effects that these evaluations have on the respective choices, we conducted three laboratory experiments.

The experiments model a setting in which a client authorizes a decision maker to choose between risky alternatives on her behalf. Before choosing, the decision maker receives information about which alternative his client would prefer. After observing the outcome, the client subjectively evaluates the decision maker. While having this basic setting in common, the three experiments test for the effects of different types and degrees of accountability in that they manipulate the information available to clients when evaluating decision makers as well as the consequences which evaluations have for decision makers. In Experiment 1, we benchmark outcome bias in evaluations and how evaluations affect interpersonal risk-taking. The setting is such that clients are informed about both the decision maker's risk choice and the outcome of the decision before they make an evaluation. Evaluations are mere feedback: They do not have any material consequences for decision makers. In *Experiment 2* as well as in *Experiment 3*, we alter the degree of accountability by introducing monetary consequences of evaluations for decision makers. In Experiment 3, we additionally change the information available to the client in that the client does not observe both the decision and the outcome, but the outcome only.

A common characteristic of delegated risky decision making in the real world is an information environment where individuals can observe outcomes of other individuals in their domain, which allows them to make peer comparisons. Peer

¹ For ease of exposition, we use female pronouns for the client and male pronouns for the decision maker.

comparisons may affect delegated risk-taking directly, but also indirectly through their effect on clients' evaluations, if these become more biased towards outcomes. To test for such an indirect effect of peer comparisons, we contrast, in all three experiments, a baseline situation where clients can observe their own outcomes only with a situation where they additionally receive relative outcome information.

This study integrates and complements both research on interpersonal risk-taking (e.g., Pahlke et al., 2012, 2015; Batteux et al., 2019, 2020; Polman & Wu, 2020) and on outcome bias in subjective evaluations of risk-taking decisions (e.g., Baron & Hershey, 1988; Gurdal et al., 2013; König-Kersting et al., 2021). With respect to clients' evaluations, results of Experiments 1 and 2 show that while evaluations are sensitive to both decisions and outcomes in a situation where they represent mere feedback, they are more strongly biased towards outcomes and reflect almost no decision accountability when they have monetary consequences for decisions makers. Results of Experiment 2 further indicate that, when evaluations have monetary consequences for decision makers, outcome bias increases when an evaluator receives information about others' outcomes that allow her to draw peer comparisons. With respect to delegated risk-taking, results from Experiment 1 show that evaluations represent a form of accountability that has the potential to increase the likelihood of delegated risk choices being aligned with clients' risk preferences. However, results of Experiment 2 provide evidence for dysfunctional consequences of outcome bias in evaluations for interpersonal risk-taking: When clients' evaluations become increasingly biased towards outcomes, decision makers more likely make choices that are not in their clients' interest. Finally, results of Experiment 3 show that clients do not hold decision makers accountable for their decisions, but focus on outcomes only, when decisions are unobservable, even though they can make a probabilistic inference from the observed outcome to the decision.

Our study models a variety of real-world relationships. Examples are delegated investment management, where professional fund managers seek for positive evaluations to attract more funds, or financial risk management, where executives expect to be held accountable by their shareholders for their risk-management decisions. Other examples refer to relationships between financial, legal, or medical advisors and their clients. Furthermore, our experimental manipulations of the information environment model important characteristics of practical situations. While in some situations of delegated risk-taking a client can observe the risk choices made on her behalf, in many situations the actions taken remain hidden to her. Also, outcomes of risk-taking decisions can be frequently compared to the outcomes of comparable decisions made by others: Investment portfolios' returns can be compared against the market, and professionals receive awards or star ratings based on such comparisons. Corporate managers are not only held accountable for absolute performance, but are also evaluated against industry averages.

By integrating the two research streams on interpersonal risk-taking and outcome bias, our study makes contributions to theory and has implications for practice. We complement preceding studies of outcome bias in that we systematically manipulate both the information environment and the consequences which evaluations have for the decision maker. Furthermore, we provide evidence about how decision makers react when they are confronted with outcome bias in the evaluations they receive. Such evidence is crucial for deepening our understanding of the role of subjective evaluations in agency relationships. We conclude from our finding that increases in outcome bias are associated with increases in misaligned risk choices that a client's own behavior may contribute to a conflict of interest between her and the decision maker. Our study has practical implications in that it informs practitioners about determinants of outcome bias that relate to their information environment and in that it sheds light on potential dysfunctional effects of subjective evaluations. Even though such effects have been investigated in other contexts (e.g., Bol, 2011; Prendergast & Topel, 1993), our study is among the first to address delegated risk-taking. Our findings suggest that evaluators, when they hold decision makers accountable, should keep in mind the causal relationship between risk-taking decisions and outcomes. Otherwise, they may give misleading feedback that contributes to decision makers taking unwanted risks, or being overly conservative. In the following we first review both streams of literature to which we contribute, then we point out research gaps and position our study.

1.1 Risk-taking for others and accountability

When individuals make risky decisions on behalf of others instead of their own, depending on the perspectives they take (Tunney & Ziegler, 2015), they may act according to their own risk preference, or may reveal a shift in behavior towards taking more or less risk (Batteux et al., 2019, 2020; Füllbrunn et al., 2022; Polman & Wu, 2020).² Evidence is inconclusive: Studies have observed a significant shift towards more risky choices (e.g., Pollmann et al., 2014; Rigoli et al., 2018; Sun et al., 2017), but also shifts to being more cautious with other people's money (e.g., Chakravarty et al., 2011; Eriksen et al., 2020). Prior research has also investigated determinants of risk-taking for others such as the decision frame (Batteux et al., 2019), responsibility (e.g., Bolton et al., 2015; Füllbrunn & Luhan, 2020), or social distance (e.g., Batteux et al., 2017; Montinari & Rancan, 2018; Zhang et al., 2017).

As the objective of our study is to make a connection between interpersonal risktaking decisions and the evaluations of such decisions, our study focuses on how accountability affects risk-taking.³ Accountability refers to the implicit or explicit expectation of a decision maker to be held accountable by others for his actions (Lerner & Tetlock, 1994, 1999). It can only arise if individuals are able to link an observed outcome to these actions, or to observe the actions themselves. Then, it implies that the decision maker receives feedback from others, is called to justify his actions and/or their outcomes, or faces consequences such as rewards or punishments. Studies addressing the effect of accountability on interpersonal risk-taking find evidence for shifts in risk-taking behavior with increased accountability (Lefebvre & Vieider, 2014; Pahlke et al., 2012; Weigold & Schlenker, 1991). When accountability

 $^{^2}$ We restrict this overview to choices with purely monetary outcomes, thus excluding studies about medical or other non-financial decisions made, e.g., by doctors on behalf of patients or by parents on behalf of their children.

³ There is a vast body of literature on the effects of increased accountability on decision making in various other contexts than risk-taking for others. For a review see Aleksovska et al. (2019).

implies monetary consequences, the decision maker tends to focus on his reward prospects when making his decisions. In case of a discretionary reward, this implies trying to anticipate the judgements underlying the rewarding decision. Even though discretionary rewards have been investigated in the context of interpersonal risk-taking (de Oliveira et al., 2017; Gurdal et al., 2013; König-Kersting et al., 2021; Pollmann et al., 2014), few studies have addressed the effects of such rewarding decisions on decision makers' behaviors.⁴

1.2 Outcome bias

Investigations of outcome bias (building on Fischhoff, 1975) cover a wide range of fields including not only financial, but also ethical, judicial or medical decisions (e.g., Arkes et al., 1981; Kamin & Rachlinski, 1995; Gino et al., 2010). Baron and Hershey (Baron & Hershey, 1988, 1992) were the first to conceptualize the effects of outcomes on evaluations. According to their definition of outcome bias, an evaluation is biased towards the outcome whenever the client has all relevant information to evaluate the decision, observes the decision itself, and still uses the outcome for the evaluation. In contrast, when the decision is unobservable, the client rationally uses the outcome to make an inference about the information not available to her. Now, the evaluation reflects outcome bias when it underweights or even contradicts the inference. Studies analyzing determinants of outcome bias have found that outcome bias is sensitive to social comparisons (Sezer et al., 2016), to comparisons between actual outcomes and potential outcomes of foregone alternatives (Baron & Hershey, 1988; van Dijk & Zeelenberg, 2005; Seta et al., 2015), to prior expectations (Schkade & Kilbourne, 1991), or to perceptions of the controllability of outcomes (Ghosh, 2005).

Our study is related to investigations of outcome bias in the evaluations and/or discretionary rewards clients give to decision makers who make risky choices on their behalf. In their seminal study on outcome bias, Baron and Hershey (1988) find that clients (hypothetically) blame decision makers both for choosing a risky alternative after a bad outcome and for choosing a safe alternative after observing a better outcome of a risky option. In a delegated portfolio management setting, both Asparouhova et al. (2015) and Anufriev et al. (2019) find that fund managers attract more funds from investors when they have been successful in the recent past, an effect that is consistent with outcome bias in investors' evaluations of fund managers. De Oliveira et al. (2017) and König-Kersting et al. (2021) investigate discretionary rewarding decisions. Rewards that increase in outcomes are not only consistent with outcome bias, but also with distributive fairness preferences, as clients may simply share their wealth with their decision makers. The results of both de Oliveira et al. (2017) and König-Kersting et al. (2021) are supportive for outcome biases to actually affect clients' decisions beyond distributive fairness, though. In Gurdal

⁴ Exceptions are investigations of delegated portfolio management (e.g., Agranov et al., 2014; Asparouhova et al., 2015; Kirchler et al., 2018, 2020), where portfolio managers compete for clients' funds and may take excessive risks to outperform other managers.

et al. (2013), clients both make evaluations and grant discretionary rewards; results show that both are biased towards outcomes.

1.3 The present research

This study asks two research questions: What determines outcome bias in subjective evaluations? And: What are the effects of outcome biases in evaluations on interpersonal risk-taking? With respect to the first research question, our study narrows a research gap in that we conduct a series of experiments in which we vary the information environment of the relationship between client and decision maker as well as the consequences evaluations have for the decision maker. With respect to the second research question, our analysis breaks new ground. A key construct behind our research questions is accountability, because both require an analysis of the relative emphasis that clients place on the decision process versus the outcome in their evaluations (Lerner & Tetlock, 1999; Patil et al., 2014). Process versus outcome accountability refers to outcome bias because the less a client focuses on the decision and the more she focuses on the outcome, the stronger is her potential outcome bias. But types accountability are also crucial for connecting evaluations with risk-taking decisions, as a decision maker will try to anticipate what he will be held accountable for (König-Kersting et al., 2021). In the following, we discuss four aspects of our study's setting which are of particular importance for investigating our research questions. In doing so, we position our study in the literature.

Knowledge about clients' preferences Few studies on interpersonal risk-taking have investigated situations in which a decision maker has information about his client's risk preference (Bolton et al., 2015; Kling et al., 2023; König-Kersting et al., 2021). If the decision maker has no such information, there is only a minimum degree of accountability, as a client cannot legitimately hold the decision maker accountable for his decision when the decision maker had no information about the client's preferences in the first place. We thus consider a situation where the decision maker is informed about his client's risk preference before he makes a risk choice on her behalf.

Observability of decisions In analyzing how decision makers react to financial incentives or competitive pressure (e.g., Andersson et al., 2020; Kirchler et al., 2018; Sheedy et al., 2019), previous studies have considered mechanisms that are exogenous to the client-decision maker relationship. In our study accountability is endogenous, as it comes from the client's subjective evaluation. When the client observes both the outcome and the decision, it is easy for her to hold the decision maker accountability can only arise from a probabilistic inference about the decision. The latter situation has high external validity in that delegated risk-taking in reality is usually subject to some information asymmetry between decision maker and client. Contrasting the two situations thus can provide important insights.

Information about other outcomes When decision makers can observe others' outcomes, their behavior is affected in a way that is consistent with an activation of competitive preferences (e.g., Dijk et al., 2014; Kirchler et al., 2018; Lindskog et al., 2022; Wang, 2017). We are interested in the effect of such information on the evaluations clients give to decision makers and thus contrast a situation in which clients receive information about other clients' outcomes with a situation where there is no such relative outcome information. We have three reasons for doing so. First, observing others' outcomes is typical for real-life situations of delegated risk-taking. Second, we expect the manipulation to provide insights into our first research question addressing the determinants of outcome bias because we expect relative outcome information being present to increase outcome bias in evaluations. Third, a manipulation that increases outcome bias in evaluations affect interpersonal risk-taking.

Monetary consequences of evaluations With few exceptions, prior studies have modeled the relationship between client and decision maker such that the client rewards, but does not evaluate, the decision maker.⁵ Our series of experiments allow us to investigate outcome bias in evaluations and how (potentially biased) evaluations affect delegated risk-taking decisions with monetary consequences of evaluations being both present and absent.

2 Method

2.1 Basic experimental design and procedures

All experiments share the same basic design. There was a pre-stage and a main stage. The function of the pre-stage was to elicit risk preferences such that in the main stage each decision maker could be informed about his client's preference as revealed from her decision in the pre-stage. In the pre-stage, each participant chose between two alternatives with equal expected outcomes but different levels of risk. Potential outcomes were 10, 20, or 30 experimental currency units (*ecu*, 4 *ecu* = 1 EUR) for both alternatives, but probabilities differed: For the low-risk alternative, the probabilities were 5%, 90%, and 5% for the low, medium, and high outcome, respectively, whereas the respective probabilities were 30%, 40%, and 30% for the high-risk alternative. Participants were told that their choices would play a role in the second part of the study, without giving any specific information. At the beginning of the main stage, participants were assigned the role of either client or decision maker. Assignment of roles was random but subject to the group of clients having an equal number of participants preferring the high risk and low risk

⁵ In Gurdal et al. (2013) clients reward and additionally evaluate decision makers. In Baron and Hershey (1988) and in the second experiment of König-Kersting et al. (2021), clients' evaluations have no monetary consequences.



Fig. 1 Outcome distributions of the low risk and high risk alternatives. The Figure shows outcome probabilities for the two alternatives available to decision makers in the main stage of the experiments

alternative in the pre-stage, respectively. Dyads were matched for a single round in a stranger design such that each decision maker faced each client exactly once. In each round, the decision maker chose between two alternatives with equal expected values but different risk levels. The alternatives were similar to those of part one, but outcomes ranged from 0 to 40 *ecu*; Fig. 1 shows probability distributions.

Before making his decision, the decision maker was informed whether his client had preferred low risk or high risk in the pre-stage. A client could not revise this implicit goal communicated to the decision maker. After the risk choice, the outcome was generated by the computer system and presented to the client, potentially with additional information depending on the specific design of the experiment. Then, the client evaluated the decision maker on a scale from 1 to 7; no specifications of evaluations were given except information about 1 being the worst and 7 the best evaluation.⁶ The round ended with the evaluation being communicated to the decision maker, and with client and decision maker receiving a summary report of the round. Over all experiments and treatments, a client received a payment that was equal to the outcome of her own risk choice in the pre-stage plus the outcome from a randomly selected round of the main stage. A decision maker also received the outcome of his risk choice in the pre-stage; the second part of his compensation varied across experiments.

The presentation of experiments in this paper is not equal to the sequence in which they were conducted: We started with Experiments 2 and 3, Experiment 1 was conducted last. Participants were recruited from the authors' universities' student

⁶ We choose this form of an "omnibus evaluation" for two reasons: First, asking clients to simply evaluate the decision maker is the only form of an evaluation that can be implemented in all experiments we conduct. Second, in practice, subjective evaluations of clients are unlikely to be such differentiated that they explicitly address both the process and the outcome of delegated decision-making.

subject pools. Experiments 2 and 3 were run in a university computer lab, whereas Experiment 1 was run online.⁷ All sessions were run with 12 clients and 12 decision makers interacting over 12 rounds in a stranger design. Each experiment had three sessions. Sample sizes were derived from prior research investigating similar settings (Bolton et al., 2015; de Oliveira et al., 2017; König-Kersting et al., 2021).

2.2 Experimental manipulations

The settings modeled in the three experiments address different types and degrees of accountability. First, we manipulate the information environment to investigate whether and how clients' evaluations reflect decision versus outcome accountability. We contrast a situation in which a client can observe both the outcome of the risky decision and the decision itself (Experiments 1 and 2) with a situation where the decision is unobservable (Experiment 3). When the decision and its outcome are observable, a client can hold the decision maker directly accountable for both. Then, any reliance on the outcome in the evaluation indicates outcome bias (Baron & Hershey, 1988). If instead the decision is unobservable to the client, she can hold the decision maker accountable for the decision. Then, feedback becomes ambiguous because a decision maker cannot clearly identify whether the decision or the outcome drives the evaluation.

Second, we vary the degree of accountability in that we contrast a situation in which evaluations have no material consequences (Experiment 1) with a situation where evaluations determine decision makers' financial compensations (Experiment 2 and 3). With such monetary consequences, we expect evaluations to have stronger effects on delegated risk-taking than without. As part of Experiment 1, we also test a *Baseline* condition where clients do not evaluate decision makers. Decision makers' compensations vary across experiments: They receive a fixed payment for the second stage of Experiment 1, whereas their pay for the second stage of both Experiments 2 and 3 depends on the evaluations they receive from their clients. More specifically, the rank which a decision maker's average evaluation has among the average evaluations of all decision makers in his session is determined, and pay is rank-dependent with 80 *ecu* for rank one, 60 for rank two, 40 for rank three, 20 for rank four, 10 for ranks five through eight, and zero for all other ranks.

In all three experiments, two treatment conditions are tested between subjects which differ with respect to the information clients receive about outcomes: With relative outcome information being present (*ROI present*), each client learns her own outcome and is additionally provided with a ranking list of the outcomes of all clients in her peer group, including her own; she does not receive information about

⁷ We ensured that online conditions were close to regular lab conditions: Recruited subjects logged into a video conference (with both micro and camera switched off), where they received all instructions. Procedures ensured anonymity. Manipulation checks indicate that participants were as strongly motivated to participate in the online sessions of Experiment 1 as they were in the lab sessions of Experiments 2 and 3.



Fig. 2 Illustration of the experimental Design

other decision makers' decisions, though. With relative outcome information being absent (*ROI absent*), no such additional information is presented. Figure 2 illustrates the experimental design.

3 Results

3.1 Experiment 1: Benchmarking outcome bias and interpersonal risk-taking

Experiment 1 investigates subjective evaluations and their potential effects on risk-taking decisions in a setting where clients can observe both the outcome and the decision and where evaluations have no monetary consequences for the decision maker, whose compensation for the main stage of the experiment is fixed. The experiment not only tests two conditions with relative outcome information (*ROI*) being either *absent* or *present*, but also a *Baseline* condition were clients do not evaluate decision makers. The experiment provides a benchmark: Clients observe both outcomes and risk-taking decisions, and thus any reliance on the outcome in the evaluation indicates outcome bias (Baron & Hershey, 1988). Furthermore, the degree of accountability is low in that evaluations have no material consequences for decision makers so that they cannot imply any financial incentives for risky or cautious shifts in decisions; if they still affect delegated risk-taking, it is the feedback conveyed by the evaluation itself that affects behavior.

.

Panel A: Evaluations					
EVAL: mean	ROI absent	ROI present	Total		
Overall:	4.82	4.97	4.89		
By risk choice: ALIGNED=0	4.02	4.16	4.09		
ALIGNED = 1	5.17	5.25	5.21		
By outcome ^a : OUTCOME < 20	3.91	4.14	4.03		
OUTCOME > 20	5.63	5.85	5.74		
Groups	36 ^b	36	72		
Panel B: Risk-taking decisions					
ALIGNED = 1: frequencies	ROI absent	ROI present	Total	Baseline	
Overall:	75.7%	66.0%	70.8%	66.7%	
By clients' risk preferences: Low risk	72.7%	63.0%	67.8%	63.2%	
High risk	78.7%	69.0%	73.8%	70.1%	
By preference match: MATCH=0	61.5%	65.2%	63.3%	52.8%	
MATCH=1	87.5%	68.6%	78.1%	76.1%	
Groups	36°	36°	72	36 ^c	

Table 1 Experiment 1: Evaluations and risk-taking decisions

The Table presents descriptive data on clients' evaluations in Panel A and on decision maker's risk choices in Panel B. The unit of observation in Panel A is the respective client's average evaluation, the unit of observation in Panel B is the frequency with which the respective decision maker makes aligned choices on behalf of his clients

- EVAL is an integer between 1 and 7

- ALIGNED is an indicator variable equal to one if the decision maker's risk choice is aligned with the client's risk preference, and zero else

- MATCH is an indicator variable equal to one if the decision maker's risk preference matches the client's risk preference, and zero else

– OUTCOME is the outcome of the decision maker's decision, an integer variable ranging from 0 to 40

- ROI denotes relative outcome information

^aCases where OUTCOME = 20 are excluded from calculations of means

^bN=35 in the *ROI absent* condition for ALIGNED=0, as one client never saw a misaligned risk choice

^cThe number of clients is 36 per condition, with 18 clients preferring low risk and 18 clients preferring high risk

Data In total, 214 subjects participated in the experiment, 72 participants each in the *ROI absent* and the *ROI present* conditions, and 70 participants in the *Baseline* condition.⁸

Evaluations Panel A of Table 1 presents data on evaluations after an aligned (ALIGNED=1) versus misaligned decision, after observing an outcome below (OUTCOME<20) versus above the mean, and overall. EVAL is the average

⁸ In two sessions of the *Baseline* condition, two participants acting as decision makers lost internet connection.

evaluation a client makes over all rounds. The data show that evaluation levels exceed the medium category (4), indicating leniency (Bol, 2011; Prendergast & Topel, 1993). Both decisions and their outcomes drive evaluations: Clients not only give significantly higher evaluations after observing an aligned as opposed to a misaligned decision (ROI absent: 5.17 vs. 4.02; ROI present: 5.25 vs. 4.16; p<0.01 in either case),⁹ but evaluations are also clearly higher after above than below average outcomes (*ROI absent*: 5.63 vs. 3.91; *ROI present*: 5.85 vs. 4.14; p<0.001 in either case). The latter result documents outcome bias in evaluations: Even though clients can observe decisions, they hold decision makers accountable for the outcomes of their decisions. Finally, the data show that the treatment manipulation - relative outcome information being absent versus present – has no effect on evaluations, as neither evaluation levels (4.82 vs. 4.97) nor distances (between evaluations after aligned vs. misaligned decisions: 5.17 - 4.02 = 1.15 vs. 5.25 - 4.16 = 1.09; between evaluations after high vs. low outcomes: 5.63 - 3.91 = 1.72 vs. 5.85 - 4.14 = 1.71) differ (substantially) across conditions. We conclude that in Experiment 1, clients hold decision makers accountable for both their decisions and these decisions' outcomes, and that providing clients with relative outcome information has no effect on evaluations.

Risk-taking decisions Panel B of Table 1 shows frequency data on risk choices for the two conditions in which clients' make evaluations, and for the Baseline condition where there are no evaluations. As we expect both the client's and the decision maker's risk preferences to influence the decision, we disaggregate the data by preference matches versus mismatches; a match (MATCH=1) occurs when client and decision maker preferred the same alternative in the pre-stage of the experiment; otherwise, there is a mismatch. Aligned choices are consistently more frequent when risk preferences match, with the differences in frequencies being significant (p < 0.01 in all cases) except for the ROI present condition; here the frequency is equally low for both cases (MATCH=1 and MATCH=0). In comparing the data between the Base*line* and the *ROI absent* condition, we can give a first answer to the question whether subjective evaluations represent an effective means to hold decision makers accountable for their risk-taking. The answer is yes as the data show that the frequency of an aligned decision (ALIGNED=1) is significantly higher in the ROI absent than in the *Baseline* condition (75.7% vs. 66.7%, z_{70} =3.30, p<0.001). However, once *ROI* is present, the frequency of an aligned decision drops to the level of the Baseline condition, and the difference in frequencies between ROI being absent versus present is significant (75.7% vs. 66.0%, z_{70} =3.21, p=0.001). That is, even though peer comparisons do not affect clients' evaluations, they affect decision makers' risk choices.

Additional analyses (untabulated) show that the effect represents a risky shift in behavior, as decision makers considerably more frequently switch to the high risk option when *ROI* is present relative to absent. This observation is consistent with prior findings about peer comparisons affecting risk-taking (e.g., Dijk et al., 2014; Kirchler et al., 2018), but it occurs in a situation where decision makers do not act on behalf of themselves, but on behalf of clients.

⁹ All p-levels reported are two-sided.

3.2 Experiment 2: Effect of outcome bias in evaluations on delegated risk-taking

Experiment 2 increases the degree of accountability compared to Experiment 1 in that we introduce monetary consequences of evaluations. Again, two conditions are tested contrasting an information environment in which clients can make peer comparisons (*ROI present*) or not (*ROI absent*).

Data In total, 144 subjects participated in the experiment, 72 (36 clients and 36 decision makers) in each condition (*ROI absent* and *ROI present*).

Evaluations Table 2 shows descriptive data on evaluations, which convey that (now with monetary consequences) clients no longer hold decision makers accountable for their risk choices, but only for the outcomes of their choices instead: Evaluations are only slightly better after an aligned relative to a misaligned decision (*ROI absent*: 4.72 vs. 4.34; *ROI present*: 4.42 vs. 4.28), but are clearly better after an above average relative to a below average outcome (*ROI absent*: 5.47 vs. 3.69; *ROI present*: 5.69 vs. 3.05; p < 0.001 in either case). This gives clear support to outcome bias driving evaluations. Testing for the effect of *ROI present* relative to *absent*, we see that evaluations become more extreme, as the distance between evaluations after high vs. low outcomes is significantly larger (*ROI absent*: 5.47 - 3.69 = 1.78 vs. *ROI present*: 5.69 - 3.05 = 2.64; $z_{70} = -2.90$, p = 0.003). This indicates that providing clients with relative outcome information increases outcome bias.¹⁰

Risk-taking decisions Table 3 shows frequency data on risk choices in Panel A, and results of fixed-effects probit regressions in Panel B. The data in Panel A show that, as in Experiment 1, aligned choices are consistently more frequent when risk preferences match. While the difference is relatively small and not significant when relative outcome information is absent (*ROI absent*: 78.6% vs. 70.2%), the frequency of an aligned choice drops sharply when preferences do not match instead of match with *ROI present* (73.0% vs. 38.6%, z_{35} =4.71, p<0.001). The overall difference in frequencies of aligned decisions with *ROI* being present versus absent is significant, too (76.1% vs. 56.5%, z_{70} =5.37, p<0.001). In combination with the evaluation data, this result is consistent with outcome bias in evaluations affecting risk-taking decisions: If clients do not hold decision makers accountable for making an aligned risk choice, but increasingly focus on outcomes instead, we can expect decision makers to feel less committed to their clients' risk preferences, and especially so when there is a preference mismatch.

To gain further insights, we analyze how evaluations affect decisions through the feedback that the decision maker receives from evaluations. To do so, we align the outcome scale (0-40 points) with the evaluation scale (1-7) by mapping the outcome into one of seven intervals (0-5, 6-11, 12-17, 18-22, 23-28, 29-34, 29-34, 29-34, 29-34)

¹⁰ As a robustness check, we performed ordered probit regressions estimating clients' evaluations over rounds. Results (untabulated) give support to both evaluations being biased towards outcomes and outcome bias increasing with relative outcome information being present.

72

36

De 2 Experiment 2. Evaluations			
EVAL: mean	ROI absent	ROI present	Total
Overall:	4.62	4.33	4.47
By risk choice: ALIGNED=0	4.34	4.28	4.31
ALIGNED = 1	4.72	4.42	4.57
By outcome ^a : OUTCOME < 20	3.69	3.05	3.37
OUTCOME > 20	5.47	5.69	5.58

 Table 2 Experiment 2: Evaluations

The Table presents descriptive data on clients' evaluations, the unit of observation being the respective client's average evaluation

36^b

- EVAL is an integer between 1 and 7

- ALIGNED is an indicator variable equal to one if the decision maker's risk choice is aligned with the client's risk preference, and zero else

- OUTCOME is the outcome of the decision maker's decision, an integer variable ranging from 0 to 40

- ROI denotes relative outcome information

^aCases where OUTCOME = 20 are excluded from calculations of means

 $^{b}N=35$ in the *ROI absent* condition for ALIGNED=0 (as in Experiment 1)

35-40), and then define a FEEDBACK measure that indicates how strongly the client accounts for the decision in her evaluation. FEEDBACK is equal to the evaluation minus the outcome interval after an aligned decision and equal to the outcome interval minus the evaluation after a misaligned decision. A positive value of FEEDBACK can be interpreted as the client rewarding an aligned and punishing a misaligned choice, respectively, whereas a negative value indicates that an aligned choice is discouraged and a misaligned choice encouraged, respectively.¹¹

Panel B of Table 3 shows results of fixed-effects probit regressions estimating the likelihood of an aligned choice in the two treatment conditions, respectively, and a model pooling the data. All three models show a significant effect of FEEDBACK: When decision makers experience that clients actually hold them accountable for their risk choices, their propensity to make an aligned choice significantly increases. A preference match has a positive effect on the likelihood to make an aligned choice when *ROI* is present, but the effect is not significant when *ROI* is absent. The pooled model further shows that *ROI* being present implies a negative shift in the likelihood of an aligned choice, and that FEEDBACK has a significantly stronger effect on

Groups

¹¹ For example, if the outcome is 20 (outcome interval=4), FEEDBACK=+3 both in the case that the decision maker has received the best evaluation 7 after making an aligned choice, and in the case that he received the worst evaluation 1 after a misaligned choice; if the choice was misaligned and the decision maker still received an evaluation of, say, 5, we code FEEDBACK=1 as the evaluation contradicts the conjecture that the client would punish a misaligned choice. When FEEDBACK=0, the decision maker has no reason to assume that the evaluation does not focus solely on the outcome.

Panel A: Descriptive statistics				
ALIGNED = 1: frequencies	ROI absent	ROI present	Total	
Overall:	76.2%	56.5%	66.3%	
By clients' risk preferences: Low risk	76.9%	63.4%	70.1%	
High risk	75.5%	49.5%	62.5%	
By preference match: MATCH=0	70.2%	38.6%	54.4%	
MATCH=1	78.6%	73.0%	75.8%	
Groups	36 ^a	36 ^a	72	
Panel R. Prohit regressions				

Table 3 Experiment 2: Risk-taking decisions

Panel B: Probit regressions

Dependent variable: ALIGNED

Independent Variables: coefficient (p-value)	ROI absent	ROI present	Pooled model
MATCH	0.214 (0.174)	1.021 (<0.001)	0.216 (0.159)
FEEDBACK (lag 1)	0.165 (0.002)	0.317 (<0.001)	0.165 (0.001)
ROI present			-1.017 (<0.001)
MATCH $\times ROI$ present			0.822 (<0.001)
FEEDBACK × ROI present			0.153 (0.048)
ROUND	-0.018 (0.473)	-0.002 (0.933)	-0.009 (0.577)
N (groups)	396 (36)	396 (36)	792 (72)

The Table presents descriptive data on decision makers' risk choices in Panel A and results from random effects Probit regressions in Panel B. The unit of observation in Panel A is the frequency with which decision makers make aligned choices on behalf of the respective client. The unit of observation in Panel B is the respective decision maker's choice in the respective round

- ALIGNED is an indicator variable equal to one if the decision maker's risk choice is aligned with the client's risk preference, and zero else

- MATCH is an indicator variable equal to one if the decision maker's risk preference matches with the client's risk preference, and zero else

- FEEDBACK is equal to the evaluation minus the outcome interval (mapping outcomes from 0 to 40 into seven intervals) when the decision maker made an aligned choice and equal to the outcome interval minus the evaluation when he made a misaligned choice. The variable is lagged for one period to capture how the decision maker's choice is affected by the evaluation he has received in the preceding round

- ROUND is the experimental round

- *ROI* denotes relative outcome information, *ROI present* is an indicator variable equal to one of relative outcome information is present, and equal zero if it is absent

- "X" denotes interactions

^aThe number of clients is 36 per condition, with 18 clients preferring low risk and 18 clients preferring high risk

behavior with *ROI* being present relative to absent; that is, holding a decision maker accountable for his risk choice has a stronger effect when *ROI* is present. The results indicate that outcome bias in evaluations mediates the effect of *ROI* being present relative to absent on delegated risk-taking. We formally test for the mediation following the three-step procedure suggested in Baron and Kenny (1986). In the first step, we test whether the average FEEDBACK significantly differs between the two treatment conditions and find that this is the case (*ROI absent* vs. *ROI present*: 0.560 vs. 0.176, z_{70} =3.79, p<0.001). In the second step, we estimate the pooled probit model of Table 3, Panel B, without including FEEDBACK (results untabulated), while the third step includes FEEDBACK and is thus equivalent to the pooled model displayed in Panel B of Table 3. The regression results show that the mediation is partial, but not total, as the indicator variable for *ROI* being present keeps having a significantly negative coefficient in the pooled regression model: Risk-taking decisions are less frequently aligned when *ROI* is present, both because of a direct effect of such information available, and because of the mediation.

3.3 Experiment 3: Exploring whether evaluations reflect Bayesian inferences

Experiment 3 varies the form of accountability: Now the decision is hidden to the client, who can only make an inference from the observed outcome to the unobserved action. As there is a low risk and a high risk alternative, the Bayesian inference from outcome to the decision is such that, when the client prefers high risk, *both high and low* outcomes are indicative of the decision maker having made an aligned decision, whereas, when the client prefers low risk, the same outcomes indicate a misaligned choice. More precisely, outcomes below 16 and above 24 are less likely for the low risk alternative and more likely for the high risk alternative, whereas outcomes within the two levels are more likely for low than for high risk (see Fig. 1). Applying Bayes' rule, a client who prefers low risk (high risk) infers from an outcome between 16 and 24 that the decision was more likely aligned than misaligned (more likely misaligned than aligned), whereas an outcome above 24 or below 16 indicates a misaligned (an aligned) choice.

Data As in Experiments 1 and 2, the assignment of participants' roles was such that the group of clients should have an equal number of decision makers preferring the high risk and low risk alternative. However, this did not work out in one session of the condition with *ROI present*, where only five individuals preferred the low-risk alternative in the pre-stage. Data come from 143 student participants, 72 in the *ROI absent* and 71 in the *ROI present* condition.¹²

Evaluations Panel A of Table 4 presents evaluation data by clients' risk preferences and by outcome ranges. If evaluations reflected Bayesian inferences, clients preferring

¹² One client failed the manipulation checks in the post-experimental questionnaire by stating that "neither decisions nor evaluations would matter anyway". We eliminated the evaluations of the respective participant.

ROI absent	ROI present	Total
4.44 [36]	4.37 [35]	4.40 [71]
2.47 [32]	2.23 [31]	2.35 [63]
4.58 [36]	4.39 [35]	4.49 [71]
5.72 [32]	6.41 [33]	6.07 [65]
4.50 [18]	4.32 [17]	4.41 [35]
2.49 [15]	2.06 [15]	2.28 [30]
4.68 [18]	4.42 [17]	4.55 [35]
5.69 [14]	6.46 [15]	6.09 [29]
4.38 [18]	4.41 [18]	4.40 [36]
2.45 [17]	2.39 [16]	2.42 [33]
4.49 [18]	4.37 [18]	4.43 [36]
5.75 [18]	6.36 [18]	6.05 [36]
ROI absent	ROI present	Total
62.5% [36]	66.9% [35]	64.7% [71]
70.4% [18]	54.4% [17]	62.6% [35]
54.6% [18]	78.7% [18]	66.7% [36]
45.2% [36]	63.6% [28]	53.3% [64]
81.0% [36]	75.4% [30]	78.4% [66]
	ROI absent 4.44 [36] 2.47 [32] 4.58 [36] 5.72 [32] 4.50 [18] 2.49 [15] 4.68 [18] 5.69 [14] 4.38 [18] 2.45 [17] 4.49 [18] 5.75 [18] ROI absent 62.5% [36] 70.4% [18] 54.6% [18] 45.2% [36] 81.0% [36]	ROI absent ROI present 4.44 [36] 4.37 [35] 2.47 [32] 2.23 [31] 4.58 [36] 4.39 [35] 5.72 [32] 6.41 [33] 4.50 [18] 4.32 [17] 2.49 [15] 2.06 [15] 4.68 [18] 4.42 [17] 5.69 [14] 6.46 [15] 4.38 [18] 4.41 [18] 2.45 [17] 2.39 [16] 4.49 [18] 4.37 [18] 5.75 [18] 6.36 [18] ROI absent ROI present 62.5% [36] 66.9% [35] 70.4% [18] 54.4% [17] 54.6% [18] 78.7% [18] 45.2% [36] 63.6% [28] 81.0% [36] 75.4% [30]

Table 4 Experiment 3: Evaluations and risk-taking decisions

The Table presents descriptive data on clients' evaluations in Panel A and on decision maker's risk choices in Panel B. The unit of observation in Panel A is the respective client's average evaluation, the unit of observation in Panel B is the frequency with which decision makers make aligned choices on behalf of the respective client

- EVAL is an integer between 1 and 7

- ALIGNED is an indicator variable equal to one if the decision maker's risk choice is aligned with the client's risk preference, and zero else

- MATCH is an indicator variable equal to one if the decision maker's risk preference matches with the client's risk preference, and zero else

- OUTCOME is the outcome of the decision maker's decision, an integer variable ranging from 0 to 40

- ROI denotes relative outcome information

low risk would respond to outcomes close to the mean with the highest evaluations, whereas clients preferring high risk would do the opposite. In stark contrast, the data reveal that evaluation patterns are almost identical for the two groups: Irrespective of a client's risk preference, average evaluations are significantly higher for medium than for low as well as for high than for medium outcomes. This gives clear support to outcome accountability, but almost no decision accountability reflected in evaluations. The data further show a stronger outcome bias for *ROI* being present than absent, as evaluations are a little lower for low and considerably higher for high outcomes; a corresponding test of differences in evaluations after high vs. medium outcomes is

significant (data pooled for all clients, *ROI absent*: 5.72 - 4.58 = 1.14 vs. *ROI present*: 6.41 - 4.39 = 2.02, $z_{63} = -3.43$, p < 0.001), whereas a test of differences in evaluations after medium vs. low outcomes is not (*ROI* absent: 4.58 - 2.47 = 2.11 vs. *ROI present*: 4.39 - 2.23 = 2.16).

Risk-taking decisions Given that we find hardly any decision accountability in clients' evaluations, we have no reason to expect evaluations to contribute to risk-taking decisions being more aligned with clients' risk preferences. Panel B of Table 4 presents frequency data on risk choices. The pattern differs from what we saw in Experiment 1, as the frequency of an aligned decision (ALIGNED=1) slightly increases from *ROI absent* (62.5%) to *ROI present* (66.9%). The difference is insignificant, though. A closer look at the data (untabulated) shows that clients who preferred low risk in the pre-stage of the experiment receive aligned choices in 70.4% of the cases in the *ROI absent* condition, but only 54.4% of the cases in the *ROI present* condition, a pattern that is very similar to what we observed in Experiment 1. In contrast, the data for clients who preferred high risk clearly contradict our previous findings. Overall, the evidence from Experiment 2 is inconclusive with respect to delegated risk-taking; we comment on this finding in the discussion of our results.

4 Discussion

This study complements prior research on the determinants of outcome bias and extends research on interpersonal risk-taking by investigating how delegated risktaking is affected by subjective evaluations which decision makers receive from their clients. We first discuss the study's contributions along its research questions, then we address practical implications and limitations.

Determinants of outcome bias In all three experiments, clients' evaluations are biased towards outcomes and reflect only limited or even no decision accountability. Results show considerable differences across experiments and conditions, though. The differences we observe support and complement explanations for outcome bias and its determinants given in the literature. Baron and Hershey (1988) argue that a potential driver of outcome bias is overgeneralization: Evaluators believe that outcomes always matter and that higher outcomes are always indicative of better decisions. Overgeneralization leads to attribution errors in that evaluators will erroneously attribute the outcome not to external conditions, but to the decision maker (Heider, 1958; Ross, 1977). Our observation that providing clients with information about other clients' outcomes increases outcome biases is consistent with this explanation, as prior research has successfully used such a manipulation to increase outcome salience (e.g., Kirchler et al., 2018; Schoenberg & Haruvy, 2012), and, ceteris paribus, the relative weight a client gives to the outcome in her evaluation can be expected to increase when the outcome becomes more salient. Our results on the effects of relative outcome information furthermore complement previous findings showing that outcome bias increases when evaluators are confronted with

outcomes of foregone alternatives (e.g., Seta et al., 2015). While relative outcome information in our study is unlikely to activate regret (as clients cannot observe outcomes of foregone alternatives), it still increases outcome salience and represents an additional source of outcome bias instead. Our finding also corresponds to evidence about actual risk-taking behavior being affected by relative outcome information (e.g., Kirchler et al., 2018, 2020). Moreover, our result that outcome bias is more pronounced when evaluations have monetary consequences for decision makers is consistent with the explanation of outcome bias proposed by Gurdal et al. (2013). The authors argue that a client who finds herself in an unfamiliar situation may refer to a similar situation salient to her and will make her rewarding decision as if she were in this similar situation. Likely, in our setting the similar situation is one in which better (worse) outcomes are indicative of better (worse) decisions. Then, the client will be biased towards outcomes, and the effect can be expected to more likely occur when evaluations affect decision makers' compensations than in a situation where there is no such link. The same logic applies to the effect of relative outcome information: In real world situations of risky decision making, the outcomes of peers are often affected by the same external factors, and peer comparisons help to filter such factors out and thus improve the accuracy of the evaluation of the decision. However, the setting of all experiments is such that other clients' outcomes are uncorrelated and thus uninformative even if the decision is unobservable. Comparing the results of Experiments 1 and 2 further suggests that there is an interaction effect of relative outcome information and monetary consequences of evaluations, as providing clients with relative outcome information significantly increases outcome bias in Experiment 2, whereas there is no such effect in Experiment $1.^{13}$

Effects of subjective evaluations on delegated risk-taking Our work also makes contributions to the literature on interpersonal risk-taking, which has not only stressed the importance of different forms of accountability, but also of the decision versus outcome dimensions of accountability. When a decision maker acts as a perfect surrogate, he accepts his client's risk preference to guide his decision. If instead he fails to take his client's perspective, his choice may reflect his own risk preference so that it will be misaligned whenever risk preferences do not match, or he may feel unable to make any reasonable forecast of how his decisions affect evaluations and thus more or less randomizes his choice. Our first finding from Experiment 1 is that evaluations represent a form of accountability that has the potential to increase the likelihood of delegated risk choices being aligned with clients' risk preferences. Our second and main finding from Experiment 2 however is that outcome bias in evaluations, reflecting a lower degree of decision accountability and a higher degree of outcome accountability, has dysfunctional consequences for interpersonal

¹³ Outcome-based evaluations may simply reflect clients' distributive fairness preferences in that they align decision makers' rewards with their own incomes by giving better evaluations for higher outcomes. We took great care in designing the experiments to rule out such a motive and to make sure instead that clients' evaluations were driven by their preferences for procedural fairness (Leventhal, 1980). Post-experimental questionnaire items addressing procedural fairness preferences indicate that we were successful in doing so.

risk-taking, as increases in outcome bias are associated with risk choices becoming less frequently aligned with clients' risk preferences. Our interpretation of this finding is that it is *a client's own behavior* that contributes to a conflict of interest between her and the decision maker. However, even though the mediation analysis from Experiment 2 demonstrates that the feedback a client gives to a decision maker with her evaluation affects the likelihood of an aligned choice in subsequent rounds, the results from Experiments 1 and 2 indicate that relative outcome information also has a direct effect, as the mediation is only partial in Experiment 2 and as we find an increase in the frequency of misaligned decisions with relative outcome information present even though outcome bias does not increase in Experiment 1.

Experiment 3 provides evidence on evaluations and risk-taking decisions when clients cannot observe the risk choice. Here, for evaluations to affect decision makers' behaviors, they have to be able to identify the degree of decision versus outcome accountability in the evaluation, which is difficult under hidden action. For example, with an observable decision, when a decision maker gets a negative evaluation after a low outcome even though his choice was aligned, he can conclude that it was the outcome that led the client's evaluation. In contrast, when the choice is unknown to the client, the decision maker cannot be sure about what the client's evaluation implies. More generally, to infer an outcome bias from evaluations under information asymmetry requires the decision maker to re-enact the client's belief revision, which he cannot easily do without knowing the client's beliefs. Evidence on delegated risk-taking from Experiment 3 indicates that even though decisions are more frequently misaligned when the client's and the decision maker's risk preferences do not match, providing clients with the opportunity to make peer comparisons does not increase the frequency of misaligned decisions.¹⁴

Practical implications, limitations and future directions Our findings have at least two practical implications. First, they may help to design management control systems in hierarchical organizations (e.g., Merchant & van der Stede, 2017). Such systems serve to guide subordinates such that they contribute to successfully carrying out an organization's plans, and holding subordinates accountable for their actions and the outcomes of these actions is a core mechanism of management control. Our findings suggest that superiors, in their evaluations of subordinates, may not sufficiently take into account that performance results from the combination of effort, strategy, and factors outside a subordinate's control. Our setting models risk-taking as a specification of strategy, and we find that participants acting in the role of evaluating superiors fail to consistently account for the subordinate's strategy in their evaluations. Furthermore, we observe that subjective evaluations are sensitive to other outcomes which are in no way informative about the quality of the decision

¹⁴ The result can be explained by the risk preferences of subjects taking part in Experiment 2: While nearly 1/2 of the participants (33 out of 72) in the *ROI absent* condition prefer high risk in the pre-stage of the experiment, the corresponding fraction is 2/3 (48 out of 72) in the *ROI present* condition. This implies that the number of cases in which there is a preference match for high risk between client and decision maker is extremely high in the *ROI present* condition, which has a considerable effect on the results.

that has been made. As such it could be beneficial to deliberately exclude such information from evaluations to make them less prone to outcome bias.

Second, it is often argued that ill-designed financial reward systems are a root cause for investment professionals to not act in the interests of their clients (e.g., Chevalier & Ellison, 1997), for managers to take excessive risks (e.g., Bebchuk & Spamann, 2010; Fahlenbrach & Stulz, 2011), or even for managers to act overly conservative (e.g., Gormley & Matsa, 2016; Wiseman & Gomez-Mejia, 1998). However, if clients and shareholders showed outcome bias in their evaluations of these managers in the way we observe in our experiments, their evaluations might contribute to managers' risk-taking decisions being misaligned with the objectives initially communicated to them. Even though this distortion in incentives might be equally problematic as ill-designed financial incentives, it is significantly more difficult to address.

Our research is subject to several limitations which narrow its contributions, but also provide future research opportunities. The most important limitation of our study is likely that delegation of risk-taking is not at the discretion of the client, but compulsory. There is thus no way of knowing whether and why clients would delegate the risk-taking decision, and observing a misaligned choice is not necessarily a problem if the client preferred delegating the decision in the first place.¹⁵ There are two potential reasons why a client would delegate the risk-taking decision even though the decision manager has-as in our setting-no superior information about risky prospects¹⁶: First, the client may be unsure about her own preferences or future decisions. Then, delegating the risk choice and observing outcomes allow her to learn, and an evaluation may reflect the client updating her preferences. Second, risk preferences maybe affected by emotional reactions to risk rather than cognitive evaluations (Loewenstein et al., 2001). Delegating risk-taking then implies that the risk choice is de-emotionalized (Batteux et al., 2020), and the client may evaluate the decision maker being well aware of this effect of delegation. If clients delegate risk-taking decisions for these potential reasons, we have to be careful with our interpretation of outcome bias in clients' evaluations contributing to delegated risk-taking decisions being misaligned with the objectives of clients, as clients may be fully satisfied with the risk-taking decisions made on their behalf.

Another limitation of our findings is that in our setting, not only clients received relative outcome information, but also decision makers. Results from Experiment 1 support the conjecture that being able to make peer comparisons does not only affect clients' evaluations, but also directly decision makers' behaviors. Such a direct effect has been observed before (e.g., Dijk et al., 2014; Kirchler et al., 2018, 2020). While our experimental design does not allow us to clearly disentangle the indirect and direct effects of relative outcome information, modifications of the design would

¹⁵ A related characteristic of our design is that the decision maker was informed about his client's risk preference without the client explicitly directing the decision maker. A corresponding experimental manipulation would have allowed us to analyze how decision versus outcome accountability reflected in clients' subjective evaluations depended on the subjective commitment of clients to initial directions.

¹⁶ If a task includes searching for alternatives and investigating the alternatives' attributes before making a decision, delegating potentially affects both aspects of information search (Liu et al., 2018).

achieve this goal. Finally, we modeled the situation such that the decision maker chooses between two alternatives only. While such a design is standard in investigations of outcome bias, it does not allow to observe gradual shifts towards more or less risky alternatives.

5 Conclusion

This study reports findings from three experiments investigating determinants of outcome bias in clients' subjective evaluations, and the effects these evaluations have on interpersonal risk choices. We provide insights into the determinants of outcome bias and present evidence for outcome bias in evaluations contributing to decision makers making risk choices that are not aligned with their client's risk preferences. Our findings contribute to explaining the behavior of decision makers and their clients in various contexts, and have practical implications in that they shed light on potential dysfunctional effects of subjective evaluations. Fund managers not only have incentive contracts that may induce them to take risks, but, potentially equally importantly, have career incentives that depend on the subjective evaluations they receive from their superiors and private investors. Managers in companies make operating, strategic or financial decisions affecting risk, and the subjective evaluations they receive from their board members, shareholders or other stakeholders contribute to explaining managers' risk-taking behavior. Advisors and other service agents who have significant influence on their clients' decisions are likely to anticipate biases in their clients' evaluations, with the potential results that their recommendations are not in their clients' best interest. There have been numerous examples stressing the role of ill-designed financial incentives in explaining dysfunctional behaviors of, e.g., investment professionals, corporate managers, financial or legal advisors, or sales representatives. In an admittedly broad view, this study provides a complementary explanation for control failures in that it focuses on the bounded rationality of those who make subjective evaluations of decision makers. When evaluators focus on outcomes in a way that they seem to forget about the causal relationship between the actions a decision maker takes and the probabilities of alternative outcomes from such actions, they may implicitly motivate their decision makers to make decisions that are not aligned with the objectives that were communicated to them in the first place. Such decisions may not necessarily represent excessive risk-taking, but also overly conservative behavior; in either case, though, conflicts of interest would likely not be mitigated, but aggravated.

Acknowledgements We thank an anonymous reviewer and the editor for their valuable comments and help in improving the paper.

Funding Open Access funding enabled and organized by Projekt DEAL.

Data availability All experimental data are available from the corresponding author upon request.

Declarations

Conflict of interests None.

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/.

References

- Agranov, M., Bisin, A., & Schotter, A. (2014). An experimental study of the impact of competition for other people's money: The portfolio manager market. *Experimental Economics*, 17, 564–585.
- Aleksovska, M., Schillemans, T., & Grimmelikhuijsen, S. (2019). Lessons from five decades of experimental and behavioral research on accountability: A systematic literature review. *Journal of Behavioral Public Administration*, 2 (2). https://doi.org/10.30636/jbpa.22.66
- Andersson, O., Holm, H. J., Tyran, J. R., & Wengström, E. (2020). Risking other people's money: Experimental evidence on the role of incentives and personality traits. *The Scandinavian Journal of Economics*, 122(2), 648–674.
- Anufriev, M., Bao, T., Sutan, A., & Tuinstra, J. (2019). Fee structure and mutual fund choice: An experiment. *Journal of Economic Behavior & Organization*, 158, 449–474. https://doi.org/10.1016/j.jebo. 2018.12.013
- Arkes, H. R., Wortmann, R. L., Saville, P. D., & Harkness, A. R. (1981). Hindsight bias among physicians weighing the likelihood of diagnoses. *Journal of Applied Psychology*, 66(2), 252–254.
- Asparouhova, E., Bossaerts, P., Čopič, J., Cornell, B., Cvitanić, J., & Meloso, D. (2015). Competition in portfolio management: Theory and experiment. *Management Science*, 61(8), 1868–1888.
- Baron, J., & Hershey, J. C. (1988). Outcome bias in decision evaluation. *Journal of Personality and Social Psychology*, 54(4), 569–579.
- Baron, R. M., & Kenny, D. A. (1986). The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*, 51(6), 1173–1182.
- Batteux, E., Ferguson, E., & Tunney, R. J. (2017). Risk preferences in surrogate decision making. *Experimental Psychology*, 64, 290–297.
- Batteux, E., Ferguson, E., & Tunney, R. J. (2019). Do our risk preferences change when we make decisions for others? A meta-analysis of self-other differences in decisions involving risk. *PLOS One*, 14(5).https://doi.org/10.1371/journal.pone.0216566
- Batteux, E., Ferguson, E., & Tunney, R. J. (2020). Do we make decisions for other people based on our predictions of their preferences? Evidence from financial and medical scenarios involving risk. *Thinking & Reasoning*, 26(2), 188–217.
- Bebchuk, L. A., & Spamann, H. (2010). Regulating banker's pay. Georgetown Law Journal, 98(2), 247–287.
- Bol, J. (2011). The determinants and performance effects of managers' performance evaluation biases. *The Accounting Review*, 86(5), 1549–1575.
- Bolton, G. E., Ockenfels, A., & Stauf, J. (2015). Social responsibility promotes conservative risk behavior. *European Economic Review*, 74, 109–127. https://doi.org/10.1016/j.euroecorev.2014.10.002
- Brownback, A., & Kuhn, M. A. (2019). Understanding outcome bias. Games and Economic Behavior, 117, 342–360. https://doi.org/10.1016/j.geb.2019.07.003
- Chakravarty, S., Harrison, G., Haruvy, E., & Rutström, F. (2011). Are you risk averse over other people's money? *Southern Economic Journal*, 77(4), 901–913.
- Chevalier, J., & Ellison, G. (1997). Risk taking by mutual funds as a response to incentives. Journal of Political Economy, 105(6), 1167–1200.
- 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 Experimental Psychology*, 62, 87–97. https://doi.org/10.1016/j.joep.2017.06.008

- Dijk, O., Holmen, M., & Kirchler, M. (2014). Rank matters The impact of social competition on portfolio choice. *European Economic Review*, 66, 97–110. https://doi.org/10.1016/j.euroecorev.2013.11.010
- Eriksen, K. W., Kvaløy, O., & Luzuriaga, M. (2020). Risk-taking on behalf of others. Journal of Behavioral and Experimental Finance, 26, 1–13. https://doi.org/10.1016/j.jbef.2020.100283
- Fahlenbrach, R., & Stulz, R. (2011). Bank CEO incentives and the credit crisis. Journal of Financial Economics, 99, 11–26.
- Fischhoff, B. (1975). Hindsight is not equal to foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, 1(3), 288–299.
- Füllbrunn, S., & Luhan, W. J. (2020). Responsibility and limited liability in decision making for others– An experimental consideration. *Journal of Economic Psychology*, 77. https://doi.org/10.1016/j.joep. 2019.06.009
- Füllbrunn, S., Kvaløy, O., & Luhan, W. (2022). Investing other people's money. In S. Füllbrunn & E. Haruvy (eds.), *Handbook of Experimental Finance*, (pp. 132–144). Edward Elgar Publishing.
- Ghosh, D. (2005). Alternative measures of managers' performance, controllability, and the outcome effect. *Behavioral Research in Accounting*, 17(1), 55–70.
- Gino, F., Shu, L. L., & Bazerman, M. H. (2010). Nameless + harmless = blameless: When seemingly irrelevant factors influence judgment of (un)ethical behavior. Organizational Behavior and Human Decision Processes, 111(2), 93–101.
- Gormley, T. A., & Matsa, D. A. (2016). Playing it safe? Managerial preferences, risk, and agency conflicts. *Journal of Financial Economics*, 122(3), 431–455.
- Gurdal, M. Y., Miller, J. B., & Rustichini, A. (2013). Why blame? Journal of Political Economy, 121(6), 1205–1247.
- Heider, F. (1958). The Psychology of Interpersonal Relations. Wiley.
- Hershey, J. C., & Baron, J. (1992). Judgment by outcomes: When is it justified? Organizational Behavior and Human Decision Processes, 53, 89–93.
- Kamin, K. A., & Rachlinski, J. J. (1995). Ex post \neq ex ante: Determining liability in hindsight. *Law and Human Behavior*, 19(1), 89–104.
- Kirchler, M., Lindner, F., & Weitzel, U. (2018). Rankings and risk-taking in the finance industry. *Journal of Finance*, 73(5), 2271–2302.
- Kirchler, M., Lindner, F., & Weitzel, U. (2020). Delegated Investment Decisions and Rankings. *Journal of Banking & Finance*, 120. https://doi.org/10.1016/j.jbankfin.2020.105952
- Kling, L., König-Kersting, C., & Trautmann, S. T. (2023). Investment preferences and risk perception: Financial agents versus clients. *Journal of Banking & Finance*, 154. https://doi.org/10.1016/j.jbank fin.2022.106489
- König-Kersting, C., Pollmann, M., Potters, J., & Trautmann, S. T. (2021). Good decision vs. good results: Outcome bias in the evaluation of financial agents. *Theory and Decision*, 90(1), 31–61.
- Lefebvre, M., & Vieider, F. M. (2014). Risk taking of executives under different incentive contracts: Experimental evidence. *Journal of Economic Behavior and Organization*, 97, 27–36. https://doi. org/10.1016/j.jebo.2013.10.008
- Lerner, J. S., & Tetlock, P. E. (1994). Accountability and social cognition. Encyclopedia of Human Behavior, 1, 3098–3121.
- Lerner, J. S., & Tetlock, P. E. (1999). Accounting for the effects of accountability. *Psychological Bulletin*, 125(2), 255–275.
- Leventhal, G. S. (1980) What should be done with equity theory? In K. J. Gergen, M. S. Greenberg, & R. H. Willis (eds.), *Social Exchange*. Springer, Boston, MA, 27–55.
- Lindskog, A., Martinsson, P., & Medhin, H. (2022). Risk-taking and others. Does the social reference point matter? *Journal of Risk and Uncertainty*, 64, 287–307.
- Liu, Y., Polman, E., Liu, Y., & Jiao, J. (2018). Choosing for others and its relation to information search. Organizational Behavior and Human Decision Processes, 147, 65–75. https://doi.org/10. 1016/j.obhdp.2018.05.005
- Loewenstein, G. F., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as feelings. Psychological Bulletin, 127(2), 267–286.
- Merchant, K. A., & Van der Stede, W. A. (2017). Management control systems: Performance measurement, evaluation and incentives. Pearson Education, 4th edition.
- Montinari, N., & Rancan, M. (2018). Risk taking on behalf of others: The role of social distance. Journal of Risk and Uncertainty, 57(1), 81–109.

- Pahlke, J., Strasser, S., & Vieider, F. M. (2012). Risk-taking for others under accountability. *Economics Letters*, 114(1), 102–105.
- Pahlke, J., Strasser, S., & Vieider, F. M. (2015). Responsibility effects in decision making under risk. Journal of Risk and Uncertainty, 51(2), 125–146.
- Patil, S. V., Vieider, F. M., & Tetlock, P. E. (2014). Process versus outcome accountability. In M. Bovens, R. E. Goodin, & T. Schillemans (eds.), *The Oxford Handbook of Public Accountability* (pp. 69–89). Oxford University Press.
- Polman, E., & Wu, K. (2020). Decision making for others involving risk: A review and metaanalysis. Journal of Economic Psychology, 77. https://doi.org/10.1016/j.joep.2019.06.007
- Pollmann, M. M. H., Potters, J., & Trautmann, S. T. (2014). Risk taking by decision makers: The role of ex-ante and ex-post accountability. *Economics Letters*, 123(3), 387–390.
- Prendergast, C., & Topel, R. (1993). Discretion and bias in performance evaluation. European Economic Review, 37, 355–365.
- Rigoli, F., Preller, K. H., & Dolan, R. J. (2018). Risk preference and choice stochasticity during decisions for other people. *Cognitive, Affective, & Behavioral Neuroscience, 18*(2), 331–341.
- Ross, L. (1977). The intuitive psychologist and his shortcoming: distortions in the attribution process. In Advances in Experimental Social Psychology, 10, 173–220. https://doi.org/10.1016/S0065-2601(08)60357-3
- Schkade, D. A., & Kilbourne, L. M. (1991). Expectation-outcome consistency and hindsight bias. Organizational Behavior and Human Decision Processes, 49(1), 105–123.
- Schoenberg, E. J., & Haruvy, E. (2012). Relative performance information in asset markets: An experimental approach. *Journal of Economic Psychology*, 33(6), 1143–1155.
- Seta, C. E., Seta, J. J., Petrocelli, J. V., & McCormick, M. (2015). Even better than the real thing: Alternative outcome bias affects decision judgements and decision regret. *Thinking & Reasoning*, 21(4), 446–472.
- 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. https:// doi.org/10.1016/j.obhdp.2016.07.001
- Sheedy, E., Zhang, L., & Tam, K. C. H. (2019). Incentives and culture in risk compliance. Journal of Banking & Finance, 107. https://doi.org/10.1016/j.jbankfin.2019.105611
- Sun, Q., Liu, Y., Zhang, H., & Lu, J. (2017). Increased social distance makes people more risk-neutral. *The Journal of Social Psychology*, 157(4), 502–512.
- Tetlock, P. E. (1985). Accountability: The neglected social context of judgment and choice. Research in Organizational Behavior, 7, 297–332.
- Tunney, R. J., & Ziegler, F. V. (2015). Toward a psychology of surrogate decision making. Perspectives on Psychological Science, 10(6), 880–885.
- van Dijk, E., & Zeelenberg, M. (2005). On the psychology of 'if only': Regret and the comparison between factual and counterfactual outcomes. Organizational Behavior and Human Decision Processes, 97, 152–160.
- Wang, L. W. (2017). Recognizing the best: The productive and counterproductive effects of relative performance recognition. *Contemporary Accounting Research*, 34(2), 966–990.
- Weigold, M. F., & Schlenker, B. R. (1991). Accountability and Risk Taking. Personality and Social Psychology Bulletin, 17(1), 25–29.
- Wiseman, R. M., & Gomez-Mejia, L. R. (1998). A behavioral agency model of managerial risk taking. Academy of Management Review, 23(1), 133–153.
- Zhang, X., Liu, Y., Chen, X., Shang, X., & Liu, Y. (2017). Decisions for others are less riskaverse in the gain frame and less risk-seeking in the loss frame than decisions for the self. *Frontiers in Psychol*ogy, 8. https://doi.org/10.3389/fpsyg.2017.01601

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