

Foresighted outcome effect: A micro-foundation of agents' risk aversion in principal-agent relations

Michal Livnat-Lerer*, Raanan Sulitzeanu-Kenan†, Tehila Kogut‡

Abstract: Agents' risk aversion is a long-standing source of concern in principal-agent theory and in the practice of organizations. While standard principal-agent theory assumes that principals adequately infer conclusions from noisy outcomes, behavioral research suggests that their inferences are affected by outcome bias. We take a further theoretical step, and propose that when an agent knows that the principal's evaluation of the agent's decision will be based on outcome knowledge, the agent expects the principal to be overly affected by the outcome, rather than by the merit of the choice. As a result, the agent seeks to minimize the likelihood of an adverse outcome, leading to risk aversion. The results of three laboratory experiments support this hypothesis, suggesting that under outcome-knowledge-based principal-agent relationships, agents anticipate the effect of outcome bias on principals, and adjust their ex-ante behavior by opting for less risky alternatives, a phenomenon we call *foresighted outcome effect*.

Keywords: Principal-agent, Agency theory, Outcome bias, Risk, Foresight, Regret

Many studies of public administration point to the ubiquity of risk aversion of public sector employees. Risk averse individuals are more likely to join the public sector than the private sector (Bonin, Dohmen, Falk, Huffman, & Sunde, 2007; Buurman, Delfgaauw, Dur, & Van den Bossche, 2012). It has been argued that public sector employees may also be relatively more risk averse as a result of the formal procedures, greater exposure to scrutiny, and limited flexibility that define the public sector (Chen & Bozeman, 2012). Others noted the public sector's key mission of maintaining public welfare rather than maximizing yield on investments (Bozeman & Kingsley, 1998). Lastly, Mulgan and Albury (2003) identified a culture of risk aversion and poor skills of risk management among seven key barriers to innovation prevalent in the public sector.

Risk aversion of public sector agents is also a key feature of the agency problem. Principal-agent theory posits that organizational and political relations are essentially founded on arrangements of delegation, whereby one party (the principal) delegates authority to another (the agent), and the welfare of the former is affected by the choices made by the latter (Arrow, 1985; Eisenhardt, 1988). Standard agency theory assumes that both parties are respectively self-interested maximizers; and that while principals are risk-neutral, agents are risk-averse (Jensen, 2001). A long-lasting challenge of principal-agency theory in the last four decades has been to devise effective incentives, to substitute monitoring in order to address the agency problem (Miller & Whitford, 2007, p. 214), and a key hurdle in obtaining such an incentive system is risk aversion on the part of agents. Indeed, it has been claimed that agents' risk aversion (combined with information asymmetry) renders it impossible to replace agent monitoring with an equally efficient system of incentives (Hölmstrom, 1979).

This article identifies and empirically tests a behavioral micro-foundation of agents' risk aversion. In the following section we draw on the literature on outcome bias - a robust psychological phenomenon which leads one to take an outcome into account in a way that is irrelevant to evaluating the

* Federmann School of Public Policy and Government, The Hebrew University;

† Political Science Department and Federmann School of Public Policy and Government, The Hebrew University;

‡ Department of Education, Ben-Gurion University
Address correspondence to Raanan Sulitzeanu-Kenan at (raanan.s-k@mail.huji.ac.il)

Copyright: © 2018. The authors license this article under the terms of the Creative Commons Attribution 4.0 International License.

true quality of the decision (Baron & Hershey, 1998; Clarkson, Emby & Watt, 2002), and points to a highly ubiquitous characteristic of principal-agent relations – the anticipated availability of outcome information – that increases ex-ante risk aversion of agents. Next, we report the findings of three experimental studies that demonstrate this effect (N=286). The nearly universal availability of outcome information in real-life principal-agent relations renders laboratory experimentation uniquely useful, as it enables control of this feature of the interaction (Falk & Heckman, 2009). The first study provides an initial test for the existence of such an effect; the second rules out the role of anticipated-regret as an alternative explanation and replicates the results of the first study; the third study explores the role of the subjective probability of losing as a possible underlying mechanism of this effect.

Anticipated Outcome Knowledge: A Source of Agent Risk Aversion

In theoretical analyses of principal-agent relations, scholars assume that principals correctly infer conclusions from noisy outcomes regarding the agent's effort or choice quality. However, outcome bias – a well-established psychological phenomenon – suggests that exposure to outcomes unduly alters evaluations. We take a further step, and propose that people anticipate outcome bias, and ex-ante adjust their behavior. In the context of agency relations, an agent's anticipation to be judged by a principal *after the latter is exposed to the outcome* of the agent's choice (taken under uncertainty), acts to increase the agent's risk aversion. The following section reviews outcome bias.

Outcome Bias

When ex-post evaluations of decisions are carried out after the consequences of the latter are known, they tend to take outcome knowledge into account in a way that is irrelevant to the true quality of the decision (Baron & Hershey, 1988; Clarkson, Emby, & Watt, 2002). Research on outcome bias typically describes a decision made under risk, which is followed by either a negative or a positive outcome. Respondents are then asked to evaluate the quality of the decision, or the decision maker's abilities and characteristics. In five experiments on evaluations of medical and monetary decisions, Baron and Hershey found that outcome information consistently

influenced evaluations of decision quality, the competence of the decision maker, and the willingness to let the decision makers make decisions on their behalf. Respondents' evaluations were more positive when the outcome was favorable than when it wasn't, despite their expressed opinion that one should not consider the outcome when making such evaluations (Baron & Hershey, 1988). Experimental attempts to de-bias the effect of outcome knowledge has shown that merely instructing or warning the evaluator about the potential biasing effect of outcome information was not effective. Only instructions that stressed either the cognitive non-normativeness of the outcome effect or the seriousness and gravity of the evaluation ameliorated the bias significantly (Clarkson et al., 2002).¹

Foresighted Outcome Knowledge within Principal-Agent Relations

Let us return to the central question of this research: Does the expected availability of outcome knowledge in principal-agent relations affect ex-ante behavior, and if so – in what way? In this article we explore principal-agent relations while focusing not on retrospective judgment, but on the *effect of knowing that one's decision will be judged while its outcome is known, on an agent's ex-ante behavior*. We will refer to this situation by the term "*outcome knowledge-based principal-agent relations*" (henceforth OK-based PA relations).

We posit that in OK-based PA relations agents "mirror" their respective principals' anticipated outcome bias. More specifically, when the agent knows that the principal's evaluation of the agent's performance will not be based on outcome knowledge, the agent's decision will be guided by her own risk preference and the risk preference she attributes to the principal. Conversely, when the agent knows that the evaluation will be conducted in the presence of outcome information, she may expect the principal's judgment to be primarily outcome-based, as outcome bias suggests. In this latter case, the agent is expected to "mirror" her principal's anticipated outcome bias, and become more concerned with minimizing the likelihood of the worst outcome – as suggested by the maximin principle (Kameda et al., 2016). Hence our hypotheses are: (H1) OK-based PA relationship increases agents' risk aversion; and (H2) when considering risky choices under uncertainty within OK-based PA relations, agents become more affected by the subjected likelihood of a negative outcome.

Research Design

Study 1

"The paradigmatic problem addressed by principal-agent theory is one in which an agent and a random variable interact to produce an outcome of value to the principal" (Miller & Whitford, 2007, p. 216). Drawing on these necessary conditions, Study 1 was based on an investment game in which subjects were required to make financial decisions involving a choice between a sure and a risky option. The financial outcome of the game directly affected the subjects' and their respective principals' real monetary reward. Three experimental conditions were used: two involved the two types of principal-agent relations – with and without outcome knowledge – while the third was control.

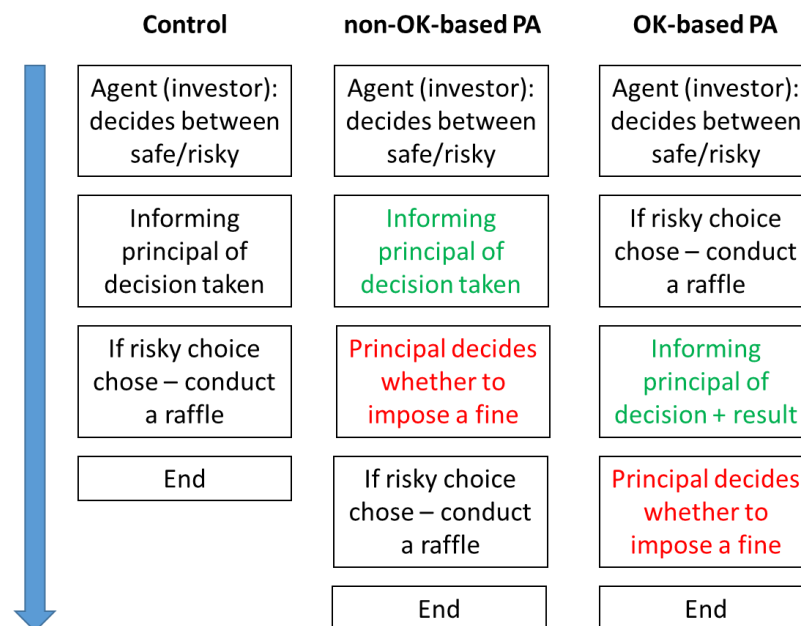
Participants and Design

Eighty-one undergraduate students (51% females; mean age 25.8) were randomly assigned to one of three experimental conditions: (1) Control; (2) Non OK-based PA; and (3) OK-based PA. All subjects were told that they were to participate in an investment game with two players – an investor and a partner. Subjects drew a slip of paper indicating their role in the game, thus laboring under the impression that the roles were assigned randomly. In

actuality, all the notes were inscribed "investor," and the subjects were told that their respective partners are in an adjoining room. The actual setting did not include a real "partner," as we were concerned with the investor's (agent) behavior alone.² Each participant received NIS 50 (about US \$14.5), and was told that this endowment was now jointly owned by herself and her "partner."³ Next, participants were informed of the conditions of the game, which varied depending on the experimental treatment to which they were assigned. As an "investor," each participant was then asked to choose between two alternatives: either to invest the NIS 50 with a 50% chance of receiving NIS 80 (US \$23.2) and a 50% chance of receiving NIS 20 (US \$5.8), or not to invest and keep the original sum. Subjects were told that, irrespective of the outcome, the payoff would ultimately be divided equally between the investor and her partner, thus making the investor's decisions relevant for both. This payoff structure entails that the expected value of both options is equal ($50 = 0.5 \times 80 + 0.5 \times 20$), thus they only differ risk level.⁴ These explanations were identical for all three experimental conditions.

In the *control condition* the investor acted on behalf of herself and her putative partner, and the partner had no power to sanction the investor in any way. In the PA relations without outcome

Figure 1
Experimental design



knowledge condition (non-OK-based PA) the putative partner had the option to sanction the investor (by deducting NIS 10 (\$2.3) from the agent's share), after learning about her decision, but crucially, *before* learning the outcome of her decision.⁵ This condition enabled us to distinguish between a principal-agent relationship that is solely based on the decision, and a principal-agent relationship that is based on the decision and its outcome. The PA relationship with outcome knowledge condition (OK-based PA) was identical to the non-OK-based PA condition, except that the putative partner had the option to sanction the investor *after* learning the outcome of her choice. The structure of each experimental condition was explained orally and graphically to each participant before starting the investment game (see appendix 1 for a detailed description). Our main interest was the difference between the proportions of risk taking under the OK-based PA condition compared with non-OK-Based PA condition. The experimental conditions are described in Figure 1 (see also online appendix 1).


Results

The results of Study 1 are presented in Figure 2. The average proportion of risk taking for all subjects was 61.0% ($n=82$, $SE=.054$). Marginally significant difference in risk-taking across the three experimental conditions were found ($\chi^2= 5.508$, $p = .064$): Control: 75.0%; non-OK-Based PA:

67.7%; and OK-based PA: 45.2%. Logit regression analysis (reported in online appendix 2) with two dummy variables – Control and OK-based PA (non-OK-based PA as reference) – shows no significant difference between the level of risk-taking under control and non-OK-based PA ($p=.579$), suggesting that merely introducing a sanction option did not result in a significant reduction in risk-taking. However, adding an expectation for outcome knowledge (OK-based PA) reduces the odds of risk-taking. Agents' awareness that their principals would be able to sanction them *after* learning the outcome of the transaction decreased risk-taking by 22.5 percentage points ($p = .076$).

The results of Study 1 provide initial support to our hypothesis. As noted above, participants in the non-OK-based PA condition made their choices under the assumption that they would be judged on merit alone, when the outcome knowledge was not yet available. Their behavior was guided by the premise that their principal would choose whether to sanction them based on the gap between their choice and the principal's own preference. Since participants lacked *a-priori* knowledge regarding their respective principals' preferences, the resulting distribution reflects both the participants' own preferences and their speculation regarding their respective principals' preferences. Participants under the OK-based PA condition, on the other hand, knew in advance that they

Figure 2
The effect of anticipated outcome knowledge on agents' risk-taking



Note: Raw proportions of risk-taking in each experimental condition. Error bars are ± 1 SE

would be judged based on the results of their choices (rather than solely on the choice itself). The risk-averse behavior that was evident in this condition can be explained by the participants' belief that outcome knowledge would lead to a different judgment than the one expected in the absence of such information (as observed under the non-OK-based PA condition). These results provide an initial support for our hypothesis: expected availability of outcome knowledge in principal-agent relations increases agents' ex-ante risk aversion.

Yet the results of Study 1 may also conform to an alternative explanation. It is possible that the increased risk aversion of individuals under the 'OK-based principal-agent relationship' conditions resulted by their motivation to minimize their principals' ability to compare the outcome of their chosen choice with the outcome of the foregone alternative – as the literature on anticipated regret may suggest (see online appendix 3 for a review). Such a motivation might bring them to anticipate the regret of their principals, and thus to act strategically in order to limit the possibility of such a comparison by opting for the safe option.⁶ To disentangle the two explanations, Study 2 examines whether an OK-based principal-agent relationship increases agents' risk aversion, even when outcome knowledge includes foregone payoffs (i.e., when the outcome of the risky option will be exposed to the principal, even if the agent chooses the safe option), thus eliminating any motivation to opt for the safe option for the purpose of limiting foregone outcome information.

Study 2

Study 2 was based on the same experimental paradigm (the investment game), but included a 2X2 design manipulating outcome knowledge availability ('non-OK-based PA' and 'OK-based PA' conditions as in Study 1), and the availability of information regarding the foregone outcome (available in all cases vs. available only if the risky option is chosen). This design allows us to test whether the fore-sighted outcome effect holds also when anticipated regret is unlikely to affect the agents' choices.

Participants and design

One hundred and eleven undergraduate students (51% female, mean age 24.75) were randomly assigned to one of four conditions manipulating anticipated outcome knowledge, and the availability

of information regarding the forgone option. The former was manipulated as in Study 1. The availability of information regarding the foregone option was manipulated by telling the participants that the raffle will be played regardless of which option was chosen, therefore, enabling the agents and principals (the latter only in the OK-based condition) to compare the outcomes of the chosen option to that of the forgone option. Our main interest was whether participants would still demonstrate increased risk-aversion when forgone outcome information is available – i.e. where anticipated regret is not expected to affect the agents' choices.


Results

The results of Study 2 are presented in Figure 3. The average proportion of risk taking for all subjects was 59.5% ($n=111$, $SE=.047$). The two treatment effects were estimated in two logit regression analyses reported in online appendix 2. The first model presents the main effects of anticipated outcome knowledge and forgone outcome information. Anticipated outcome knowledge had a negative effect on the propensity to choose the risky option (no-OK-based PA: $P(risk)=.745$ [.630, .860], OK-based PA: $P(risk)=.447$ [.317, .577], $p=.002$). The main effect of forgone outcome information is positive but statistically insignificant ($p=.362$). To estimate the difference in the effect of outcome knowledge across the two conditions of forgone outcome information, the second model includes an interaction term for the two treatments. The interaction effect is statistically insignificant ($p=.785$), allowing us to reject the hypothesis that regret avoidance accounts for the results. Indeed, the effect of anticipated outcome knowledge is negative and statistically significant both when foregone outcomes are not known (replicating Study 1), as well as when foregone payoffs are known: odds-ratio=.243 [.075, .789] and odds-ratio=.304 [.100, .922], respectively.

Following Cumming (2014) we utilized the *ESCI* software to compute the average effect size of anticipated outcome-knowledge on the likelihood of risk taking, based on Studies 1 and 2. The average effect size based on a random-effect meta-analysis is $-.274$ [-.417, -.132]. This analysis is graphically presented in Figure 4.


In Study 3 we attempt to better understand the underlying mechanism of this effect. It is plausible that, when no outcome knowledge is available to the principal, the agent's decision relies on her

Figure 3
The effect of anticipated outcome knowledge on agents' risk-taking,
across foregone outcome information conditions



Note: Bars represent risk-taking proportions in each experimental condition. Error bars are ± 1 SE.

Figure 4
Average effect size of anticipated outcome knowledge based on studies 1 and 2



Note: The two upper estimates (in black) represent the treatment effect in Studies 1 (top) and 2 (middle), based on the proportions of risk taking, with 95% CIs. The bottom estimate (in red) represents the average effect, based on a random-effect meta-analysis.

personal preference, and her naïve belief about the principal's risk-preference. Research on such beliefs suggest that people are prone to perceive others as more risk seeking than themselves (e.g. Hsee & Weber, 1997; Kogut & Beyth-Marom, 2008). Conversely, we propose that under OK-based PA agents expect that their evaluation would be based on the *outcome* of their choice, rather than on its adherence to the risk preferences of the principal. Agents with such expectations are likely to follow the maximin principle (e.g. Kameda et al., 2016), in the sense that they are primarily motivated to avoid a loss, rather than the prospect of a gain.

Study 3

In this study, our aim was to assess the hypothesis that when considering risky choices under uncertainty within OK-based principal-agent relations, agents are more affected by their subjective likelihood of a negative outcome. Thus, the role of subjective perceptions of losing in predicting participants choices are expected to be greater under OK-based PA, than when non-OK-based PA.

Method

Ninety-four undergraduate students (83% females, mean age 24.44) received a short questionnaire and

were asked to complete it one page at a time. On the first page they were asked to imagine that they were participating in the investment game described in Study 1. Participants were randomly assigned to one of the two conditions: non-OK-based PA or OK-based PA. After reading the description of their respective investment game, participants in both conditions were asked to rate, on a visual analog, seven-inch scale, ranging from "no chance at all" to "full certainty," the likelihood of winning and of losing, if they were to choose the risky option – according to their subjective assessment (each question appeared on a separate page, and their order was counterbalanced). On the last page of the questionnaire, participants were asked to choose between the safe and risky option.

Results

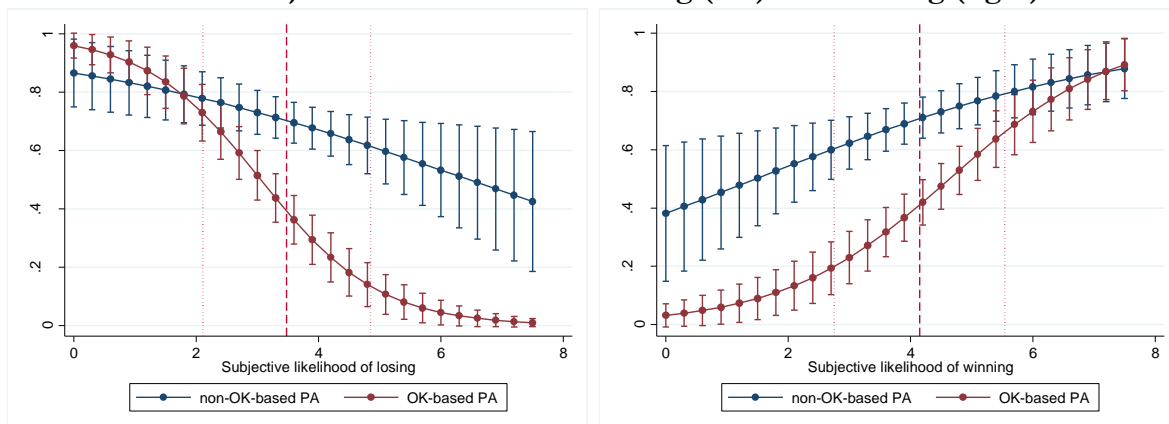
In line with hypothesis 1 and the results of Studies 1 and 2, there was a significant difference between the percentage of risk taking under the non-OK-based PA and the OK-based PA conditions ($\chi^2=5.49, p=.019$): 68.9% and 44.9%, respectively.

There were no significant differences between the two conditions in subjective likelihood estimates of either losing or winning ($t(92)=.92, p=.360$; and $t(92)=1.00, p=.319$, respectively). We next examine the role of participants' subjective winning and losing likelihoods in their risk-taking choices under the two conditions. Two logistic regression analyses were conducted on participants' choices: one with the interaction between OK-

based PA condition and subjective likelihoods of losing; and the other with the interaction between OK-based PA condition and subjective likelihoods of winning, as predictors. These analyses are graphically presented in Figures 5 and 6, respectively. While in both analyses subjective likelihoods of the two outcomes appear more strongly associated with risk taking under the OK-based PA condition, this difference is marginally significant only for losing likelihood ($p=.071$) and statistically insignificant for winning likelihood ($p=.278$). However, as evident from Figures 5 and 6, the associations between both losing and winning subjective likelihoods and risk taking are statistically significant only under OK-based PA ($b=-1.040, p=.002$ and $b=.738, p=.010$, respectively). Participants who evaluated the likelihood of losing as low were more likely to choose the risky option in both conditions. However, those who evaluated the likelihood of losing as high avoided the risky choice mostly in the OK-based PA condition and less so in the non-OK-based PA condition. Similarly, participants who evaluated the likelihood of winning as high were more likely to choose the risky option in both conditions, yet participants with lower expectations of winning tended to avoid the risky option, especially under the OK-based PA condition.

The results of Study 3 replicate the results of Studies 1 and 2 in the context of a hypothetical scenario. In addition, they suggest that evaluations of the likelihood of losing play a greater role in de-

Figure 5 & 6
Probability of risk taking under non-OK-based PA and OK-based PA, across subjective likelihood of the losing (left) and winning (right)



Note: Dashed vertical line represents mean losing/winning-likelihood and dotted lines represent one SD below and above the mean (see: Dawson & Richter, 2006). Error bars are ± 1 SE.

cisions under outcome-based principal-agent relations. Knowing that one's decision will be evaluated based on its outcome appears to increase the weight of the perceived likelihood of losing in the decision, which in turn decreases risk-taking.

Discussion

The three studies reported here examine behaviors motivated by actual monetary incentives (Studies 1 & 2), and scenario based decisions (Study 3). The results of these studies reveal a consistent pattern, according to which outcome-based principal-agent relations increase risk aversion among agents. This effect holds when forgone outcome information is available, allowing us to reject regret-avoidance as an alternative explanation.

Standard principal-agent theory assumes that principals adequately infer conclusions from noisy outcomes. However, behavioral research suggests that their inferences are affected by outcome bias. Taking a further theoretical step, we propose that when an agent knows that the principal's evaluation of the agent's decision will be based on outcome knowledge, the agent expects the principal to be overly affected by the outcome, rather than by the merit of the choice. As a result, the agent's decision seeks to minimize the likelihood of an adverse outcome, leading to risk aversion. The differences in risk taking under the non-OK-based and OK-based PA conditions in the three experiments indicate the presence of what we term *foresighted outcome effect*, which refers to the agent's assumption that her decision, although made under uncertainty, would be judged by its outcome, that in all likelihood swayed the agent's ex-ante choice. Further support for the greater weight assigned to the anticipated loss under OK-based principal-agent relations is provided by the results of Study 3. Note that subjective probabilities of losing in the lab may differ from subjective priors of bureaucrats in realistic settings, as the latter often possess vast stores of information on probable outcomes that likely shapes their likelihood estimations. However, in order to mitigate this external validity concern, our setting explicitly provides full information regarding the objective probability of losing/losing. Despite this information, subjective probabilities were found to vary, and were predictive of agents' choices.

These findings carry important implications for many social and political settings and are of particular relevance to individuals who face the

need to take decisions under uncertainty in their professional life, including politicians, medical doctors, managers, engineers, etc. Our results indicate that decision-making processes in such situations may be fraught with a fundamental problem. Not only is a retrospective evaluation of others' behavior open to outcome bias, but individuals who know that their decisions could be subjected to such an evaluation are likely to behave defensively, following a pattern which mirrors the biased evaluation. While defensive behavior in various socio-political contexts is well documented (Hood, 2011), the current work reveals that its behavioral catalyst in principal-agent relations may be outcome knowledge.

Notes

1. A close concept to outcome bias is hindsight bias, which refers to one's tendency to retrospectively overestimate the probability of a particular outcome, after learning that it did indeed happen (Blank, Musch, & Pohl, 2007; Fischhoff, 1975).
2. The experiments were approved by the Ethical Review Board at the Hebrew University.
3. This payoff structure was chosen to simulate the agent's intrinsic motivation for successful work performance (Vroom, 1964). Studies have shown that bureaucrats typically hold distinct value preferences that likely instill a sense of extraordinary ownership in the public product (e.g., Nalbandian & Edwards, 1983), and the PSM literature suggests that they are imbued with an ethic to serve the common good, which characterizes organizational behavior in public agencies (Barbara, 1990; Brewer & Selden, 1998; Crewson, 1997; Perry & Wise, 1990; Rainey, 1982).
4. Our behavioural measure of risk preference follows the logic of other, more elaborate measures (e.g., Holt & Laury, 2002, 2005), and was simplified for the purpose of being integrated in a stylized principal-agent setting.
5. The sanction introduces an element of accountability hierarchy into the relationship. The sanction power is unilaterally at the hands of the principal, and it serves as a signal of the latter's satisfaction with the agent. The size of the sanction was designed: (1) to be a substantive, yet not dominant incentive; (2) to maintain the equality of expected value of the safe

and risky options; (3) to provide that regardless of incurring a sanction or not, winning or losing in the risky option provides a higher or lower payoff than the safe option, respectively. To clarify that the partner has no incentive to sanction the investor other than for the purpose of signaling discontent, the instructions explicitly state that the sanction meant a reduction in the share of the investor, but this reduction was not to be gained by the partner (see appendix 1).

6. Choosing the risky option under the 'OK-based PA' condition in Study 1 provided principals with information on the outcomes of both the risky option and the safe option, while choosing the safe option provided information on the outcome of the safe option only. This may have led agents to choose the safe option in order to restrict the ability of the principal to compare across potential outcomes.

References

- Arrow, K., J. (1985). The economics of agency. In J. Pratt & R. Zeckhauser (Eds.), *Principals and agents: the structure of business*. Cambridge, MA: Harvard Business School Press.
- Barbara, R. (1990). Employee investment and commitment. *Public Administration Review*, 50, 374–382.
- Baron, J., & Hershey, J. C. (1988). Outcome bias in decision evaluation. *Journal of Personality and Social Psychology*, 54(4), 569–579.
- Blank, H., Musch, J., & Pohl, R. F. (2007). Hindsight Bias: On Being Wise After the Event. *Social Cognition*, 25(1), 1–9.
- Bonin, H., Dohmen, T., Falk, A., Huffman, D., & Sunde, U. (2007). Cross-sectional earnings risk and occupational sorting: The role of risk attitudes. *Labour Economics*, 14(6), 926–937.
- Bozeman, B., & Kingsley, G. (1998). Risk Culture in Public and Private Organizations. *Public Administration Review*, 58(2), 109–118.
- Brewer, G. A., & Selden, S. C. (1998). Whistle Blowers in the Federal Civil Service: New Evidence of the Public Service Ethic. *Journal of Public Administration Research and Theory*, 8(3), 413–440.
- Buurman, M., Delfgaauw, J., Dur, R., & Van den Bossche, S. (2012). Public sector employees: Risk averse and altruistic? *Journal of Economic Behavior & Organization*, 83(3), 279–291.
- Chen, C.-A., & Bozeman, B. (2012). Organizational Risk Aversion: Comparing The Public and Non-Profit Sectors. *Public Management Review*, 14(3), 377–402.
- Clarkson, P. M., Emby, C., & Watt, V. W. -S. (2002). Debiasing the Outcome Effect: The Role of Instructions in an Audit Litigation Setting. *AUDITING: A Journal of Practice & Theory*, 21(2), 7–20.
- Crewson, P. E. (1997). Public-Service Motivation: Building Empirical Evidence of Incidence and Effect. *Journal of Public Administration Research and Theory*, 7(4), 499–518.
- Cumming, G. (2014). The New Statistics Why and How. *Psychological Science*, 25(1), 7–29.
- Dawson, J. F., & Richter, A. W. (2006). Probing three-way interactions in moderated multiple regression: Development and application of a slope difference test. *Journal of Applied Psychology*, 91(4), 917–926.
- Eisenhardt, K. M. (1988). Agency- and Institutional-Theory Explanations: The Case of Retail Sales Compensation. *Academy of Management Journal*, 31(3), 488–511.
- Falk, A., & Heckman, J. J. (2009). Lab Experiments Are a Major Source of Knowledge in the Social Sciences. *Science*, 326(5952), 535–538.
- 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.
- Hölmstrom, B. (1979). Moral Hazard and Observability. *The Bell Journal of Economics*, 10(1), 74–91.
- Holt, C. A., & Laury, S. K. (2002). Risk Aversion and Incentive Effects. *The American Economic Review*, 92(5), 1644–1655.
- Holt, C. A., & Laury, S. K. (2005). Risk Aversion and Incentive Effects: New Data without Order Effects. *The American Economic Review*, 95(3), 902–904.
- Jensen, M. C. (2001). *Foundations of Organizational Strategy*. Cambridge, Mass.: Harvard University Press.
- Kameda, T., Inukai, K., Higuchi, S., Ogawa, A., Kim, H., Matsuda, T., & Sakagami, M. (2016). Rawlsian maximin rule operates as a common cognitive anchor in distributive justice and risky decisions. *Proceedings of the National Academy of Sciences*, 113(42), 11817–11822.
- Miller, G. J., & Whitford, A. B. (2007). The Principal's Moral Hazard: Constraints on the Use of Incentives in Hierarchy. *Journal of Public Administration Research and Theory*, 17(2), 213–233.
- Mulgan, G., & Albury, D. (2003). *Innovations in the Public Sector* (p. 40). London: Cabinet Office.

- Nalbandian, J., & Edwards, J. T. (1983). The Values of Public Administrators: a Comparison With Lawyers, Social Workers, and Business Administrators. *Review of Public Personnel Administration*, 4(1), 114–127.
- Perry, J. L., & Wise, L. R. (1990). The Motivational Bases of Public Service. *Public Administration Review*, 50(3), 367–373.
- Rainey, H. G. (1982). Reward Preferences among Public and Private Managers: In Search of the Service Ethic. *The American Review of Public Administration*, 16(4), 288–302.
- Vroom, V. H. (1964). *Work and motivation*. Oxford, England: Wiley.