

Biased Policy Professionals

by Sheheryar Banuri *
Stefan Dercon**
Varun Gauri***

* School of Economics and CBESS, University of East Anglia

** Blavatnik School of Government and Centre for the Study of African Economies

*** Global Insights Initiative, The World Bank

Abstract

A large literature focuses on the biases of individuals and consumers, as well as “nudges” and other policies that can address those biases. Although policy decisions are often more consequential than those of individual consumers, there is a dearth of studies on the biases of policy professionals: those who prepare and implement policy on behalf of elected politicians. Experiments conducted on a novel subject pool of development policy professionals (public servants of the World Bank and the Department for International Development in the UK) show that policy professionals are indeed subject to decision making traps, including sunk cost bias, the framing of losses and gains, frame-dependent risk-aversion, and, most strikingly, confirmation bias correlated with ideological priors, despite having an explicit mission to promote evidence-informed and impartial decision making. These findings should worry policy professionals and their principals in governments and large organizations, as well as citizens themselves. A further experiment, in which policy professionals engage in discussion, shows that deliberation may be able to mitigate the effects of some of these biases.

JEL classification codes

Keywords

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

Biased Policy Professionals
Sheheryar Banuri; Stefan Dercon; Varun Gauri¹

DRAFT PAPER – PLEASE DO NOT CITE OR CIRCULATE WITHOUT THE AUTHORS’ PERMISSION

October 2017

Abstract: A large literature focuses on the biases of individuals and consumers, as well as “nudges” and other policies that can address those biases. Although policy decisions are often more consequential than those of individual consumers, there is a dearth of studies on the biases of policy professionals: those who prepare and implement policy on behalf of elected politicians. Experiments conducted on a novel subject pool of development policy professionals (public servants of the World Bank and the Department for International Development in the UK) show that policy professionals are indeed subject to decision making traps, including sunk cost bias, the framing of losses and gains, frame-dependent risk-aversion, and, most strikingly, confirmation bias correlated with ideological priors, despite having an explicit mission to promote evidence-informed and impartial decision making. These findings should worry policy professionals and their principals in governments and large organizations, as well as citizens themselves. A further experiment, in which policy professionals engage in discussion, shows that deliberation may be able to mitigate the effects of some of these biases.

¹ Banuri: School of Economics and Centre for Behavioral and Experimental Social Science, University of East Anglia, Norwich, UK, NR4 7TJ (e-mail: s.banuri@uea.ac.uk); Dercon: Blavatnik School of Government and Centre for the Study of African Economies, 120 Walton St, Oxford OX2 6GG (email: stefan.dercon@bsg.ox.ac.uk); Gauri: Global Insights Initiative, The World Bank, 1818 H St NW, Washington, DC, 20433 (email: vgauri@worldbank.org). The authors have no relevant or material financial interests that relate to the research described in this paper. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors and do not necessarily represent the views of the World Bank, its Executive Directors, or the countries they represent.

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

Policy professionals play an essential role in the design and implementation of policies, programs, and projects across the world. While key decisions are typically taken by elected officials or other political appointees, decision makers depend on policy professionals, often civil servants, for policy preparation and advice. These professionals play a central role in translating data and research into policy options and in guiding decision-making. In this paper, we use experiments to study cognitive biases in interpreting data for the purpose of providing advice to decision-makers. The subject pool is novel. We study policy professionals working in the area of international development – in particular, UK civil servants and international civil servants working for the World Bank. The experiments are adapted to the development context from a number of classic studies on biases in decision-making.

The objective and accurate use of data on the part of policy professionals is important for at least two reasons. First, it is almost axiomatic that policies will be more effective when those who design and implement them are able to form accurate beliefs about how the world works and conduct accurate assessments of the costs and effectiveness of policy initiatives. Even if policy professionals sometimes utilize effective heuristics or employ an “ecological rationality,” rather than cost-benefit analysis or expected utility theory or a related set of decision rules, biases in assessing information and calculating value remain important. After all, it is the policy professionals themselves who must adopt and use decision rules, as well as evaluate their effectiveness. Bureaucrats also confront novel cases or situations in which prevailing decision rules do not apply. For these reasons, it is important for policy professionals *to be able* to evaluate data and assess value accurately, even if not every decision requires that capacity.

Second, objectivity, impartiality, and accuracy are often legal or official requirements. Those values are elements of modern bureaucracy, in Weber’s (1946) sense, and often help bureaucracies discharge their business objectively, “according to calculable rules and ‘without regard for persons.’” Many governments and organizations maintain the Weberian idea that a bureaucracy is perfected “the more completely it succeeds in eliminating from official business love, hatred, and all purely personal, irrational, and emotional elements which escape calculation. This is the specific nature of bureaucracy and it is appraised as its special virtue.” (Weber 1946: 216)

For instance, the UK Constitutional Reform and Governance Act of 2010 sets out a civil service code that provides guidance on its four “core values” for UK public sector employees and that should guide part of our subject pool: “integrity, honesty, objectivity, and impartiality.”² The objectivity criterion further elaborates that the principle involves “basing your advice and decisions on rigorous analysis of the evidence.” Similarly, the United Nations sets out standards of conduct for the international civil service and lists “honesty, truthfulness, impartiality and incorruptibility” under the concept of integrity. The code further elaborates that “civil servants do not have the freedom of private persons to take sides or to express their convictions publicly on controversial matters, either individually or as members of a group.”³ World Bank staff, who form the other part of our sample pool, are required to act with “integrity, independence and impartiality” in line with

² Accessed January 2016: <https://www.gov.uk/government/publications/civil-service-code/the-civil-service-code>

³ Accessed May 2016: <http://icsc.un.org/resources/pdfs/general/standards.pdf>

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

their status as employees of an international organization.⁴ The World Bank established its Independent Evaluation Group (IEG) in order to “provid[e] impartial evidence based assessments and lessons on drivers of success and failure.”⁵

To promote the impartial and proper use of evidence, large public organizations recruit and rely on the judgments of well-qualified economists, epidemiologists, environmental scientists, engineers, and other professionals. Having been trained in the natural or social sciences, these professionals are, in theory, well-equipped to conduct and use evidence-based assessments of policies. But there are strong reasons to suspect that professionals in public sector organizations, like most individuals, are subject to and exhibit substantial biases in information processing, assessments of value, and decision making. The psychological literature suggests that biases in decision-making, both on the part of the general public and experts, are widespread (DellaVigna and Pope, 2016; Malmendier and Tate, 2008; Cooke, 1991; Shanteau, 1992; Shanteau, 1988; Englich, Mussweiler, and Strack, 2006; Tversky and Kahneman, 1981; Langfeldt, 2004; Stewart and Stasser, 1995; Bero and Jadad, 1997; Herek, Janis, and Huth, 1987; Calvert, 1985; among others). If bureaucrats and policy professionals are as biased as the general population, meritocracy and good recruiting processes will not protect against partiality and subjectivity.

To further promote objective and impartial decision making, large public organizations implement a number of procedural safeguards. Often, they encourage and sometimes require peer review and deliberation, cost-benefit analysis and other kinds of ex-ante policy scrutiny, and the ex-post evaluation of projects and programs on the basis of randomized controlled trials and other methods. It is unclear how effective these procedural safeguards are. In practice, policy professionals are often adept at identifying sympathetic peer reviewers and decision makers, influencing the direction of a deliberative meeting by setting the agenda in particular fashion, burying contestable assumptions of cost-benefit calculations in a thicket of footnotes and appendices, and carefully curating evaluations most relevant to a particular project or decision. More generally, because most principals, including elected politicians in charge of final decision-making, have little idea how biases affect decision making in their own organizations, they are not able to match procedural safeguards to the most significant or prevalent cognitive biases, let alone measure the effectiveness of those procedures. This means that procedural safeguards may not be triggered, especially in cases where the biases are particularly strong (i.e. where policymakers have a larger incentive to implement their beliefs into policy proposals).

This paper presents the results of a survey designed to identify decision making biases within a sample of development professionals in two prominent organizations: The World Bank and Department for International Development in the UK (the UK government department responsible for development policy and spending the foreign aid budget, “DFID” hereafter). The survey used a series of experiments adapted to the development context and consistent with the kinds of decisions staff are in these organizations are asked to make. The survey focused on confirmation bias, sunk cost bias, and the effects of framing on risk aversion. These decision making areas and biases were

⁴ Staff Manual, accessed August 2016:

http://siteresources.worldbank.org/INTSTAFFMANUAL/Resources/StaffManual_WB_web.pdf, p. 3.

⁵ Taken from the 2015 Independent Evaluation Group annual report (accessed May 2016):

<http://ieg.worldbankgroup.org/Data/ar2015-full.pdf>

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

selected for study because they loom large in development policymaking. This paper is, to our knowledge, the first examination of these cognitive biases in large public sector organizations. The results show that staff engage in biased decision-making, including apparent bias correlated with ideological priors.

The next section introduces the survey and procedures used. The succeeding sections assess, in turn, confirmation bias, sunk cost bias, and the effect of framing on risk aversion. The final section presents the results of a further experiment designed to assess the extent to which deliberation can mitigate some of these biases, and also presents a brief summary and concluding remarks.

Survey design and data collection

The surveys were conducted entirely online, through email invitations via SurveyMonkey.com. The underlying population was all World Bank and DFID full-time staff of professional grade levels⁶, both in headquarters (DC and London, respectively) and in country offices for each organization. A random sample for each organization was selected in the following manner: the employee roster for World Bank and DFID staff (including names, email addresses, staff identification numbers, and location) were obtained and, after the population was split into two groups (headquarters and country offices), representative samples were drawn from each group⁷.

In order to increase the motivation of World Bank staff to participate in the survey, all respondents were offered a free coffee mug. No such incentive was provided to DFID staff. Recruitment was conducted in three waves, with each wave lasting one week. The survey was fairly extensive and took 30-40 minutes to complete. Invitations were sent to 4,724 World Bank and 1,148 DFID staff. At the end of the study period, 2,053 responses were received from the World Bank (response rate 43%) and 825 responses from DFID (response rate 72%)⁸. This yielded an overall sample of 2,878 respondents across the two organizations. The survey adapts classic behavioral experiments, as described in the corresponding sections below. Each treatment in a given experiment was randomly and independently assigned from the other experiments. Table 1 provides summary statistics for our sample, distinguishing between DFID and World Bank respondents.

⁶The specific professional grade levels are GE through GL for the World Bank, and B2 through SCS for DFID.

⁷The samples were drawn based on a 95% level of confidence, and a ± 4 confidence interval.

⁸We define a valid response as the respondent completed at least one question in the survey. Since different vignettes came at different points of the survey, the total number of observations fluctuate across experiments.

DRAFT: PLEASE DO NOT CITE OR CIRCULATE**Table 1: Summary statistics**

	Department for International Development	World Bank
Observations	825	2,053
Age (mean)	42.87	43.71
Female (%)	51.62	41.34
Cognitive reflection test (mean, on scale 0 to 3)	1.76	1.75
Inequality preference (mean, on scale 1 to 10)	4.21	4.91
Risk preference (mean)	3.00	--
Salary Grade (median)	7 (A2)	4 (GG)
Posted at country office (%)	45.69	51.35
Degree		
Economists (%)	15.67	24.05
Development (%)	17.02	4.24
Focus of work		
Poverty and social development (%)	12.01	7.72
Environment and infrastructure (%)	3.91	19.27
Health (%)	3.77	4.16

Note: Inequality preference was based on a 1-10 scale, where 1= "Incomes should be made more equal" and 10= "We need larger difference as incentives for individual effort."

Confirmation bias

A longstanding finding from social psychology holds that the general public uses evidence partially, interpreting findings in light of the symbols or metaphors they invoke, or in accord with the views of respected opinion leaders (Lord, Ross, and Lepper 1979). This is a specific, and socially contextualized, instance of confirmation bias, in which individuals selectively seek, remember, and prefer information in a manner that confirms prior views. The social problems arising from confirmation bias creates of particular concern in public discussions of scientific evidence. Jelveh, Kogut, and Naidu (2015) find substantial sorting of economists into sub-fields and, strikingly, a correlation between the political ideology of economists and the policy-relevant parameter estimates that the economists report. In other words, political preferences appear to be affecting the scientific findings of published economists. Relatedly, Sunstein et al (2006) find that the political composition of judicial panels significantly influences judges' opinions. The same judge, whether appointed by a Democrat or a Republican, exhibits more ideological voting patterns when sitting on a panel with politically homogenous judges than with ideologically mixed panels. Together, these studies show that technical expertise does not by itself resolve the problem of ideologically motivated or socially influenced confirmation bias.

In order to test the presence of confirmation bias among our sample of policy professionals, we designed an experiment in which we asked them to assess data from the evaluation of an intervention. The experiment we conducted was an extension of Kahan et al. (2013). In the original

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

experiment, respondents were asked to evaluate the outcome reported by a study that generated the following frequency table:

Table 2: Data presentation template for confirmation bias vignette

	Good Outcome	Bad Outcome
Individuals taking action X	223	75
Individuals taking action Y	107	21

In Kahan et al. (2013) there were two treatments, both of which presented these exact data but changed the objectives and framing of the two hypothetical studies. One study addressed the effectiveness of a skin rash cream (low prior biases); the other addressed the impact of gun control laws on crime (high prior biases). Respondents were then asked to interpret the conclusion of the study they saw, based on the data presented in the table.⁹ The authors found that respondent accuracy was higher with the skin cream framing than the gun control framing, and that there was evidence that ideology was driving this difference.

Our study used a similar approach but substituted the impact of minimum wage laws for gun control laws because the attitudes toward the minimum wage are likely to exhibit more ideological variation among development practitioners than attitudes toward gun control. The comparison frame in our study focused, like Kahan et al (2013), on skin cream.¹⁰ Again, in each treatment, the numbers presented to respondents were identical, but the labels were changed to reflect the framing. Following Kahan et al. (2013), within each frame we randomly switched the labels for “Good Outcome” and “Bad Outcome” so that the data supported either one or the other policy conclusion (or statement regarding clinical outcomes).

Respondents in the skin cream frame saw the following prompt:

Medical researchers have developed a new cream for treating skin rashes. New treatments often work but sometimes make rashes worse. Even when treatments don't work, skin rashes sometimes get better and sometimes get worse on their own. As a result, it is necessary to test any new treatment in an experiment to see whether it makes the skin condition of those who use it better or worse than if they had not used it. Researchers have conducted an experiment on patients with skin rashes. In the experiment, one group of patients used the new cream for two weeks, and a second group did not use the cream. In each group, the number of people whose skin condition got better and the number whose condition got worse is recorded in the table below. Because patients do not always complete the studies, the total number of patients in each two

⁹ As Kahan et al. (2013) note, the numbers used in the table are such that relying on any heuristic (other than conducting the actual calculation) leads to an incorrect answer. In effect, what this means is that when experts rely on their heuristics and biases, they are likely to get the question wrong. Since our sample are more likely to have biases in favor or against minimum wages, we would expect accuracy to be lower in the minimum wage frame relative to a skin cream frame. Adding pecuniary incentives would be expected to increase accuracy rates overall, and potentially even reduce differences between frames, it is not clear whether this difference would be eliminated as a result of incentives, particularly if one frame is more likely to trigger the heuristic.

¹⁰ We replicate their skin cream treatment, but recognize that gun control laws would work well for a random sample of US citizens, but less well with development professionals, as they are more liberal. We use a “minimum wage” treatment, where we utilize a minimum wage frame.

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

groups is not exactly the same, but this does not prevent the assessment of the results. Please consider two statements about this study:

- (a) People who used the skin cream were more likely to GET BETTER than those who didn't*
- (b) People who used the skin cream were more likely to GET WORSE than those who didn't*

Which statement (above) is the study most consistent with?

The labels on the rows (in table 2) were replaced with “Patients who *did* [did not] use the new skin cream” and the labels on the columns are replaced with “Rash got *better* [worse]”.

Respondents in the minimum wage frame saw the following prompt:

A decentralization reform gave local jurisdictions in an upper-middle income country authority over the minimum wage. Some raised the minimum wage, and others left it unchanged. Because of natural barriers dividing the localities, there was little population mobility in response to the changes. Capital, however, was mobile. Some believe that increasing the minimum wage tends to raise the income of the poorest 40%. Others think that raising the minimum wage slows business growth so much that the incomes of the poorest 40% tend to fall. To examine this question, researchers at a major university measured the number of jurisdictions in which the incomes of the poorest 40% rose, and the number in which the incomes fell, four years after the reform. Please consider two statements about this study:

- (a) The income of the poorest 40% of the population FALLS when the minimum wage is increased*
- (b) The income of the poorest 40% of the population RISES when the minimum wage is increased*

Which statement (above) is the study most consistent with?

As in the original study, respondents were asked to choose which statement best reflected the study findings. The labels on the rows (in table 2) were replaced with “Localities that *did* [did not] increase the minimum wage” and the labels on the columns were replaced with “Income of poorest 40% *rose* [fell]”. Respondents were randomly assigned to either the skin cream or the minimum wage frame, and whether the data supported statement A or statement B.¹¹

Respondents were also asked to assess the quality of the study (on a 10-point scale). In addition to this, the study collected data on the respondents’ ideological orientation with regard to redistribution (adapted from the World Values Survey): *On a scale of 1 to 10, where do your views fall: 1 = “Incomes should be made more equal”; 10 = “We need larger income differences as incentives for individual effort”*. This question measures respondents’ ideological orientation with respect to redistribution, a key predictor of support for/against minimum wage laws.¹²

¹¹ Since we were primarily interested in the responses to the minimum wage frame, we assigned 20% of respondents to the skin cream frame, while 80% were assigned to the minimum wage frame.

¹² Ideological orientation was measured towards the end of the extensive survey. In the survey, respondents were exposed to the vignettes first, then socio-demographic questions were recorded, followed by political orientation. As mentioned, while the treatments within experiments were randomized, the order of the questions was not.

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

If respondents evaluate data objectively and independent of prior beliefs regarding redistribution, they should have offered equally accurate assessments of the minimum wage study and skin cream study frames. Alternatively, if respondents are influenced by their ideologies or values, accuracy in the minimum wage frame should have been lower than in the skin cream frame. In fact, respondents were significantly less accurate in the minimum wage treatments (45% responded with the correct answer) relative to the skin cream treatments (65% responded with the correct answer: two sample proportions test: $p < 0.01$), and relative to random guesses ($p < 0.01$). Respondents exposed to the ideologically charged frame were significantly less likely to interpret the data correctly, compared to the neutral frame, which suggests that development professionals exhibit a bias when interpreting data on ideologically charged interventions. Figure 1 presents these results¹³.

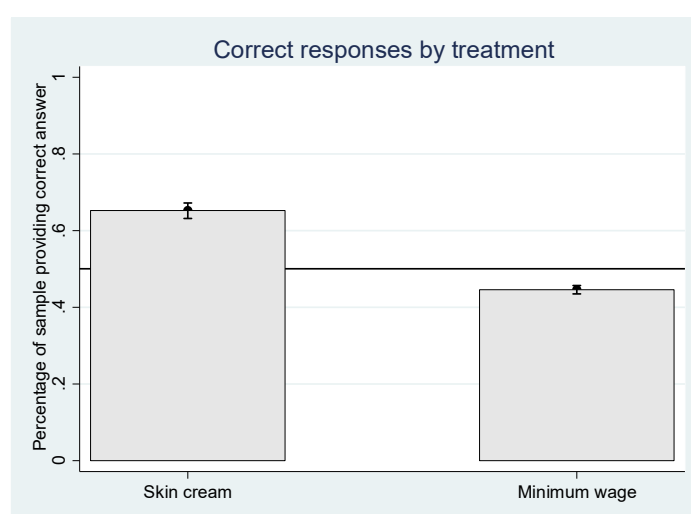


Figure 1: Percentage of respondents reporting correct answer

Table 3 displays the results of probit regressions estimating the likelihood of providing a correct response. Model 1 includes a dummy variable equal to 1 is the respondent was exposed to the minimum wage frame, while model 2 includes controls for mathematical ability using the Cognitive Reflection Test (Frederick, 2005) and respondent assessment of the study (to control for the methods used in the two study frames, one being experimental and the other observational). Model 3 includes controls for socio-demographic variables (age and gender). Model 4 includes organizational controls (dummy variable for whether the respondent works for the World Bank or for DFID, whether the respondent was posted at HQ or a country office, and the salary grade of the respondent). Finally, model 5 includes respondent expertise relevant for the minimum wage frame

¹³ Recall that within each frame (skin cream versus minimum wage), the data support either income (rash) improving, or income (rash) getting worse. 41% of the respondents report the correct answer when the data support income improving (58% accuracy for the rash); while 48% of the respondents report the correct answer when the data support income declining (72% accuracy for the rash getting worse).

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

(whether the subject of the respondents' highest degree was economics or development¹⁴, and whether the respondent currently works on poverty¹⁵).

Table 3: Framing effects on the interpretation of data ^{a,b,c}

Dependent variable: correct response to vignette (= 1)					
	I	II	III	IV	V
Minimum wage frame	-0.526*** (0.06)	-0.532*** (0.07)	-0.531*** (0.07)	-0.548*** (0.07)	-0.524*** (0.08)
Cognitive reflection test score		0.126*** (0.02)	0.117*** (0.02)	0.122*** (0.03)	0.128*** (0.03)
Study rating by respondent (10 = Extremely strong)		0.015 (0.01)	0.015 (0.01)	0.023 (0.01)	0.037** (0.02)
Age (in years)			0.000 (0.00)	-0.003 (0.00)	-0.003 (0.00)
Female			-0.100* (0.05)	-0.071 (0.06)	-0.020 (0.07)
Organization (0 = DFID; 1 = World Bank)				0.034 (0.08)	-0.048 (0.11)
Posting (0 = HQ; 1 = Country office)				0.005 (0.06)	0.005 (0.07)
Respondent grade in organization (1 = Junior; 9 = Senior)				0.050** (0.02)	0.034 (0.03)
Economics major (Subject of highest degree)					0.188** (0.08)
Development major (Subject of highest degree)					0.047 (0.12)
Poverty expert (Respondent works on poverty)					-0.051 (0.11)
Constant	0.391*** (0.06)	0.156* (0.09)	0.202 (0.16)	0.058 (0.19)	0.043 (0.23)
Log likelihood	-1825.8	-1605.5	-1572.5	-1461.4	-1103.8
Pseudo R2	0.020	0.029	0.030	0.035	0.035
P	0.000	0.000	0.000	0.000	0.000
Observations ^d	2689	2386	2339	2185	1651

Notes:

¹⁴ In this and the subsequent experiments we report, the substantive results remain unchanged when all degree options are included, and not only degrees in economics or development.

¹⁵ The poverty expert variable was constructed based on responses to the question: "What best describes the focus of your work?". The variable takes on a value of 1 if the subject responded with poverty-relevant responses: "Poverty"; "Social development" or "Social protection" (World Bank), or "Economics Advisor"; "Humanitarian Advisor"; "Livelihoods Advisor"; or "Social Development Advisor" (DFID).

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

- ^a Probit regressions. The dependent variable takes on a value of 1 if the respondent selected the response supported by the data, and 0 otherwise.
- ^b Table reports coefficients with standard errors in parentheses.
- ^c * 10%, ** 5%, *** 1% significance level.
- ^d The number of observations drops across specifications due to nonresponses on certain questions by respondents.

Table 3 shows support for the finding that ideologically charged questions reduce accuracy in the interpretation of data among development professionals. Respondents were 20% less likely to provide the correct response when exposed to the minimum wage frame ($p < 0.01$), across all specifications. As expected, respondents with higher scores on the Cognitive Reflection Test (related to higher mathematical ability) were significantly more likely to provide the correct response ($p < 0.01$). In addition, economics majors were significantly more likely to provide the correct answer in both study frames, suggesting that, at least to some extent, social science training helps professionals interpret data accurately ($p < 0.05$). Furthermore, respondents that assessed this study to be of a higher quality were also more likely to provide the correct response, though this variable was not robust to changes in model specification. Finally, there were no differences between DFID and World Bank respondents ($p = 0.66$) and between duty stations (HQ vs. country office: $p = 0.94$)¹⁶.

Recall that the minimum wage frame question had two treatments with two different correct answers (depending on the labelling of the columns). Half of the respondents were provided with data that support the conclusion that the income of the poorest *falls* when minimum wages are increased (“Income Falls”), while the other half were provided with data supporting the conclusion that the income of the poorest *rises* when minimum wages are increased (“Income Rises”). This structure was identical for the skin cream frame, as well. Combining these treatments with the question about equality preference permits a direct test of the influence of ideology on data interpretation.

Table 4 presents the results of probit regressions where the dependent variable takes a value of 1 if the respondent provided the correct interpretation of the data. Models 1 and 2 correspond to the ideological frame (minimum wage), with model 1 reflecting the treatment where the data support the bad outcome (income falling) and model 2 reflecting the treatment where the data support the good outcome (income rising). There are corresponding good and bad outcomes in the skin cream treatments (the skin cream eliminates rash and does not eliminate rash). The estimations include controls for inequality preferences, mathematical ability, the respondent’s assessment of the study, age, gender, and organization. The main relationship of interest is the effect of inequality preferences on getting the correct answer. If respondents rely on their priors, rather than evaluating the frequency table carefully, one should expect to see that the higher the respondents’ preferences for inequality (and hence, their anticipated opposition to minimum wage laws), the *more likely* they are to respond correctly when the data support the *income falling*, and the *less likely* to respond correctly when the data support the *income rising*. Since this preference is irrelevant for evaluating the effectiveness of skin creams, one would expect to see no relationship under in both treatments using the skin cream frame.

¹⁶In addition to the above, we also test for whether the treatment effects vary by training (economics), or by area of expertise (poverty experts). No significant differences are found, indicating that economists are significantly more accurate overall, and that poverty experts show no significant differences across treatments.

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

This is precisely what we find. Model 1 corresponds to the treatment where the data support the finding that minimum wage laws lead to the income of the poorest falling. In that model, there is a positive and significant relationship between inequality preferences and the correct interpretation of the data ($p < 0.05$). Similarly, in model 2, where the data support the finding that minimum wage laws increase the income of the poorest, there is a negative and significant relationship between inequality preferences and correct interpretation of the data ($p < 0.05$). This means that when the priors and the results match, respondents were more likely to interpret the data correctly. By contrast, there is no evidence of this relationship in the skin cream frame. In models 3 and 4, both coefficients are not significantly different from 0 ($p = 0.85$ and $p = 0.77$ respectively). The results are similar across the two institutions across all four treatments ($p > 0.20$).

Table 4: The effect of ideology on the interpretation of data ^{a,b,c,d}

	Dependent variable: correct response to vignette (= 1)			
	Ideological frame		Neutral frame	
	Income falls I	Income rises II	Rash worsens III	Rash improves IV
Preference for inequality (10 = Prefer wealth inequality)	0.036** (0.02)	-0.041** (0.02)	-0.007 (0.04)	-0.011 (0.04)
Cognitive reflection test score	0.047 (0.04)	0.172*** (0.04)	0.216*** (0.08)	0.117 (0.08)
Study rating by respondent (10 = Extremely strong)	-0.011 (0.02)	0.062*** (0.02)	-0.065 (0.04)	0.072 (0.05)
Age (in years)	-0.004 (0.00)	0.002 (0.00)	0.011 (0.01)	0.002 (0.01)
Female	-0.167* (0.09)	-0.098 (0.09)	0.230 (0.18)	-0.170 (0.18)
Organization (0 = DFID; 1 = World Bank)	-0.092 (0.09)	-0.174* (0.09)	0.094 (0.21)	0.090 (0.19)
Constant	0.090 (0.26)	-0.476* (0.26)	-0.472 (0.55)	0.111 (0.59)
Log likelihood	-644.6	-607.5	-150.3	-140.2
Pseudo R2	0.010	0.032	0.042	0.022
P	0.053	0.000	0.041	0.379
Observations	939	921	235	241

Notes:

^a Probit regressions. The dependent variable takes on a value of 1 if the respondent selected the response corresponding to the data, and 0 otherwise.

^b Models 1 and 2 correspond to the ideological (minimum wage) frame, while models 3 and 4 correspond to the neutral frame.

^c Table reports coefficients with standard errors in parentheses.

^d * 10%, ** 5%, *** 1% significance level.

DRAFT: PLEASE DO NOT CITE OR CIRCULATE**Sunk cost bias**

A major challenge in government agencies involves inertia and path dependency. In particular, bureaucracies exhibit a tendency to continue initiatives even when they have been shown not to work. They have difficulty cutting failed policies and programs, preferring to continue what has previously been authorized and financed. Examples from the policymaking world abound, including the continuing procurement of products and services even after they have been shown to be defective, extending information campaigns even when they do not work, and prolonging wars because the cost of admitting failure is too high (see for example, Levy, 2003; Brockner and Rubin, 1985; Schaubroeck and Davis, 1994; Staw, 1976; McDermott, 2004). The psychology of those involved in designing or implementing policy is likely implicated in some of these decisions.

The sunk cost fallacy is the tendency of individuals to continue a project once an initial investment of resources has been made (Arkes and Blumer 1985). This effect is predicated on the psychology of individuals to not appear wasteful, even though continuing a project that is questionable may be completely dominated by other options. While the initial observation regarding the presence of sunk costs is grounded in prospect theory (Kahneman and Tversky 1979), Arkes and Blumer (1985) find that individuals that incur a sunk cost have a higher expectation regarding the probability of success of the project, relative to those that did not incur a cost.

The second experiment tested for the presence of this cognitive bias in policy professionals. It adapted the sunk cost vignette (scenario 2) used by Garland and Newport (1991)¹⁷ to generate a scenario that would be contextually plausible for those involved in policy:

You are managing a five-year, \$500 million land management, conservation, and biodiversity project focusing on the forests of a small country. The project has been active for four years, \$[450] [350] [250] [150] million of project funds have been disbursed, and the project is [90] [70] [50] [30]% complete.

A new provincial government comes into office and announces that it will develop hydropower in the main river of the forest. This will require major resettlement. At the same time, the new government still wants you to complete the project. How likely is it that, if faced with this situation, you personally would decide to commit the last [50] [150] [250] [350] million dollars to complete the project? Please indicate on a scale of 0-100%.

Respondents were asked to indicate how likely they were to support a project that was unlikely to achieve its development objectives due to a new provincial government coming to power and implementing a new set of policies. (In the development context, resettlement problems, such as the

¹⁷ The original vignette is as follows: "You are the owner and manager of Security Tower, an older downtown office building that overlooks several square blocks in an area that has been slated for urban renewal over the next three years. The City Council has indicated that it would like to create a 'greenway' with grass, trees, and a small lake networked with bicycle and jogging paths. You have begun remodeling your building, anticipating renewed interest in downtown offices, with convenient parking, good access to the cross-town freeway, and a nice view. You have spent _____ of the approximately _____ you had budgeted for remodeling and the project is _____% complete. You have just learned that the 'greenway' plan has been voted down in favor of a sports stadium that will give all 15 floors of your building a view of cement walls and/or parking lots. Additionally, the increased traffic in the area will clog the freeway access for years, even with the plans to widen adjacent streets." – Garland and Newport (1991).

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

one in the vignette, often trigger major and likely insurmountable problems for a project). The treatments varied the amount of funds already disbursed (i.e. the sunk cost) on a \$500 million project. Treatments indicated that 30%, 50%, 70% or 90% of funds were disbursed. The vignette then asked respondents how likely they were (in percent terms) to disburse the remaining funds of the project. On a separate page, they were then asked how likely others in their organization were to commit the remaining funds.

The level of the sunk cost should have no bearing on the decision to commit the remaining funds to the project, yet the experiment finds clear evidence of an effect.¹⁸ Respondents reported, on average, a 40% likelihood of disbursing the remaining funds when 30% of costs were sunk. This likelihood increased to 43% with 50% of costs sunk (3% increase, two sample t-test $p < 0.10$). The likelihood of disbursing the remaining funds increased to 49% when 70% of costs were sunk (6% increase, $p < 0.01$). The 90% sunk cost treatment was somewhat lower than when 70% of costs were sunk, but the difference was not significant (<1% decrease; $p = 0.93$). The likelihood that others would disburse the remaining funds exhibited a very similar pattern. Notably, respondents reported a significantly higher likelihood that others in their organization would disburse the remaining funds in all four treatments (paired t-tests: $p < 0.01$), a finding consistent with the idea that individuals in these organizations may be influenced by sunk costs because they believe that is the norm in their organization. Figure 2 displays these results.




Figure 2: Sunk costs and support for dying projects: Likelihood of self and others disbursing remaining funds

¹⁸ The vignette presented suggests that, in the development context, the expected returns to the biodiversity project fall sharply, after the hydropower plan emerges, and that this holds at all levels of project disbursement. It may be that it is still worthwhile or profitable to implement the project so stopping the project is not necessarily required (i.e. the returns to the project are near zero). However, for reasonable scenarios, whether it is still profitable should not depend on how much has already been spent, so there should not be a relationship between the likelihood of committing the remaining funds and the level of sunk costs. It is possible that, in other scenarios and with other technologies, a new policy could affect the returns to the project in a manner that varies with extant levels of disbursement. A possible case could be projects that are implemented in discrete units. While such exceptions are possible, it is far-fetched to try to explain the observed answers by such scenarios.

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

Tables 5 and 6 present OLS regressions of the likelihood that the respondent would disburse the remaining funds and the likelihood that others in the organization would do the same, respectively. Model 1 includes three dummy variables, one for each treatment (with 30% sunk cost used as the baseline). Model 2 adds controls for socio-demographics (age and gender). Model 3 includes organizational controls (dummy variables for organization and location of posting, and the salary grade of the respondent). Finally, model 4 includes respondent expertise relevant for the framing of the vignette (whether the subject of the respondents' highest degree was economics or development, and whether the respondent was working on environmental issues¹⁹).

Table 5: Likelihood of the respondent to disburse remaining funds ^{a,b,c,d}

Dependent variable: Likelihood (%) of respondent to disburse funds				
	I	II	III	IV
Sunk cost treatment: 50% sunk cost	3.281*	3.296*	3.147*	3.996*
	(1.80)	(1.82)	(1.90)	(2.20)
Sunk cost treatment: 70% sunk cost	8.904***	9.302***	9.578***	9.848***
	(1.84)	(1.86)	(1.93)	(2.22)
Sunk cost treatment: 90% sunk cost	8.744***	8.717***	8.840***	9.535***
	(1.85)	(1.87)	(1.95)	(2.27)
Age (in years)		-0.056	-0.039	-0.029
		(0.07)	(0.08)	(0.10)
Female		-1.829	-1.518	-1.851
		(1.32)	(1.38)	(1.64)
Organization (0 = DFID; 1 = World Bank)			3.994*	0.955
			(2.07)	(2.76)
Posting (0 = HQ; 1 = Country office)			0.477	-0.039
			(1.37)	(1.61)
Respondent grade in organization (1 = Junior; 9 = Senior)			-0.118	-0.585
			(0.53)	(0.70)
Economics major (Subject of highest degree)				1.956
				(1.94)
Development major (Subject of highest degree)				-1.985
				(3.02)
Environment expert (Respondent works on environment)				7.305***
				(2.18)
Constant	39.76***	42.83***	39.38***	41.55***
	(1.31)	(3.51)	(4.46)	(5.36)
R2	0.013	0.015	0.019	0.027

¹⁹ The poverty expert variable was constructed based on responses to the question: "What best describes the focus of your work?" The variable takes on a value of 1 if the subject responded with environment-relevant responses: "Environment" or "Infrastructure" (World Bank); or "Climate and Environment Advisor" or "Infrastructure Advisor" (DFID).

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

P	0.000	0.000	0.000	0.000
Observations	2535	2457	2287	1728

Notes:

^a OLS regressions. The dependent variable is the respondent assessment (on a scale of 0 – 100) of the likelihood that the respondent would commit the remaining funds to complete the project.

^b Table reports coefficients with standard errors in parentheses.

^c * 10%, ** 5%, *** 1% significance level.

Table 6: Likelihood of the others in the organization to disburse remaining funds ^{a,b,c,d}

Dependent variable: Likelihood (%) of others in organization to disburse funds				
	I	II	III	IV
Sunk cost treatment: 50% sunk cost	0.975 (1.55)	0.828 (1.57)	0.887 (1.62)	0.856 (1.87)
Sunk cost treatment: 70% sunk cost	5.630*** (1.58)	6.007*** (1.60)	6.570*** (1.64)	5.973*** (1.88)
Sunk cost treatment: 90% sunk cost	7.566*** (1.59)	7.669*** (1.61)	8.004*** (1.66)	7.029*** (1.92)
Age (in years)		-0.042 (0.06)	-0.083 (0.07)	-0.079 (0.09)
Female		2.455** (1.13)	2.677** (1.18)	2.047 (1.39)
Organization (0 = DFID; 1 = World Bank)			4.572*** (1.76)	3.422 (2.34)
Posting (0 = HQ; 1 = Country office)			-4.086*** (1.16)	-4.977*** (1.36)
Respondent grade in organization (1 = Junior; 9 = Senior)			0.253 (0.45)	-0.032 (0.60)
Economics major (Subject of highest degree)				0.879 (1.65)
Development major (Subject of highest degree)				2.668 (2.56)
Environment expert (Respondent works on environment)				2.057 (1.85)
Constant	47.26*** (1.13)	47.93*** (3.00)	47.52*** (3.79)	49.24*** (4.55)
R2	0.013	0.016	0.027	0.026
P	0.000	0.000	0.000	0.000
Observations	2514	2440	2274	1720

Notes:

^a OLS regressions. The dependent variable is the respondent assessment (on a scale of 0 – 100) of the likelihood that others in the organization would commit the remaining funds to complete the project.

^b Table reports coefficients with standard errors in parentheses.

^c * 10%, ** 5%, *** 1% significance level.

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

Tables 5 and 6 confirm the pattern of figure 2. Respondents were 4% more likely to disburse the remaining over the baseline, which is weakly significant ($p < 0.10$). Respondents were significantly more likely (nearly 10%) to disburse the remaining funds in the 70% and 90% sunk cost treatments over the baseline ($p < 0.10$), and were also significantly more likely to disburse given 50% sunk costs ($p < 0.05$ in both cases). There was no significant difference in responses between 70 and 90% sunk costs level, possibly due to probability overweighting or satiation. An alternate specification (not reported) found evidence of a linear trend: there was a 3.5% increase in the likelihood for each 20% increase in sunk costs ($p < 0.01$). Finally, environmental experts – plausibly more committed to complete environmental projects - were 7.3% more likely to disburse the remaining funds across all treatments.²⁰

Table 6, concerning the expected responses of others, shows a similar pattern. And as in the responses regarding own behavior, a linear trend was observable in the alternate specification (not reported): a 2.6% increase in the likelihood for each 20% increase in sunk costs ($p < 0.01$). Environmental experts, unlike the sample as a whole, did not believe that others were more or less likely to disburse than they were themselves. Respondents posted in country offices, possibly more likely to confront this scenario than HQ staff, reported that others in their organization were 5% less likely to disburse the remaining funds.

There were a few modest differences between the samples in the two organizations, with World Bank respondents exhibiting a higher likelihood of disbursing, for themselves and others, but these findings were not robust to specification. There were nearly identical linear trends across the organizations.

In sum, there was strong evidence for the presence of sunk cost bias among development policy professionals. Moreover, respondents' believed that others were more likely to disburse than they were themselves.

Framing and Risk

The framing of information affects the perceptions of risk, as well as decisions to take risky decisions on behalf of others. The latter phenomenon appears linked to social preferences (Eckel and Grossman 2002; 2008; Song, 2008; Chakravarty et al 2011; Anderson et al 2012; Bradler 2009). Several of these studies show that an individual's risk tolerance increases when taking risk on behalf of others, in comparison to the willingness to take risk when making decisions for oneself. In other words, individuals become more risk seeking when they are playing with house money, rather than their own. But these findings seem contrary to the common perception that public sector institutions, and the bureaucrats and professionals in them, are unusually risk averse, and more risk averse than the individuals who compose them.

To explore this issue, we conducted two experiments that examined the effect of framing on risk preferences. The first was a replication, in our novel subject pool, of a classic experiment from

²⁰Environment experts display a similar trend across treatments. Using a specification interacting the environment and treatment dummies (not reported), we find no significant differences within the lowest three sunk cost treatments, and a marginally significant increase in the 90% sunk cost treatment (11.5%: $p < 0.10$).

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

Tversky and Kahneman (1981) that examined whether loss or gain framing affects willingness to take on risk, but importantly, with a health context relevant for policy professionals during the Ebola virus epidemic (2014). The second, conducted with the DFID sample only, directly compared whether respondents were more or less willing to take on risk when making decisions for their organizations than for themselves.

In order to investigate the effect of framing on risky decision-making among policy professionals, we replicated the original Tversky and Kahneman (1981) experiment, and asked respondents the following question:

Suppose your country is preparing for a new disease that is expected to infect 12,000 people. Scientists have come up with two treatments – let's call them Treatment 1 and Treatment 2. Here is the internationally validated scientific evidence on the effectiveness of the treatments.

[Gain frame]

If people take Treatment 1, then 4,000 people will be saved. [Safe choice]

If people take Treatment 2, then there is 1/3 probability that 12,000 people will be saved and 2/3 probability that no one will be saved. [Risky choice]

Which treatment do you think the health authorities should implement?

[Loss frame]

If people take Treatment 1, then 8,000 people will die. [Safe choice]

If people take Treatment 2, there is 1/3 probability that no one will die and 2/3 probability that 12,000 people will die. [Risky choice]

Which treatment do you think the health authorities should implement?

This vignette asked respondents to decide between two alternative medical treatments with the same expected value, but one was safe and the other entailed risk. The experiment randomly assigned respondents to either a frame emphasizing gains (“will be saved”) or losses (“will die”). If framing does not matter, then respondents’ choices are guided by their own risk preferences, and one would expect no difference in the proportion of respondents selecting the risky policy. However, if respondents are prone to biases arising from framing (and prospect theory more generally), one would expect respondents to be more risk-seeking in losses relative to gains (Kahneman and Tversky, 1981)²¹.

Figure 3 reports the results of the experiment across the two samples (pooled). 22% of the respondents assigned to the gains frame (n=1,314) selected the risky policy option while 65% of the

²¹While we conducted this experiment in both the World Bank and DFID, we only collected risk preferences in the DFID sample, as explained in experiment 4 below. Hence we do not control for risk preferences in the subsequent analysis. Controlling for risk restricts our sample to DFID only, but our core result remains robust in the DFID sample alone.

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

respondents assigned to the losses frame (n=1,277) did the same. This difference in proportion is significant (two sample proportions test: $p < 0.01$).²²

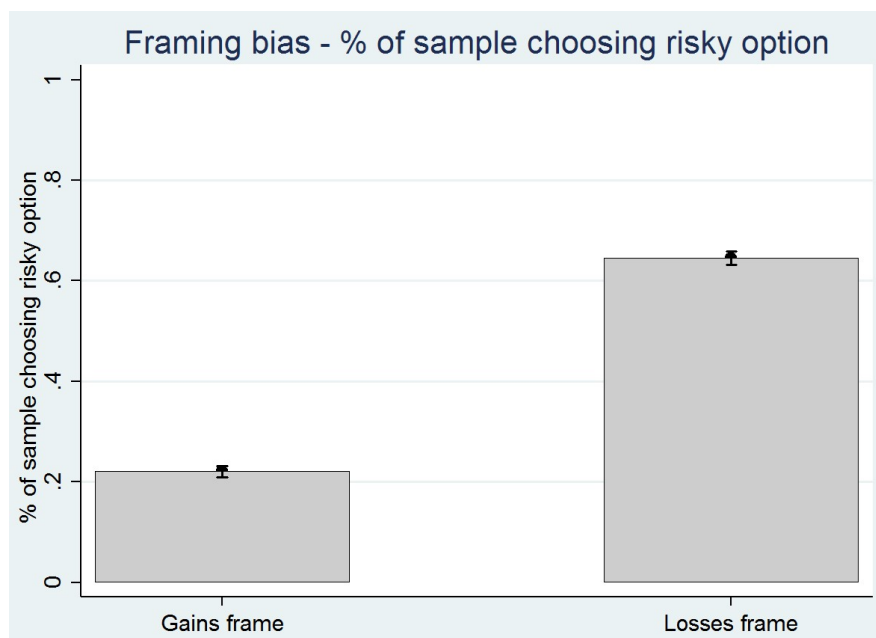


Figure 3: Percentage of respondents choosing risky policy option under different frames

Table 7 displays the results of probit regressions estimating the likelihood that a respondent chose the risky policy option. Model 1 includes a dummy variable equal to 1 if the respondent was exposed to the losses frame, while model 2 includes controls for socio-demographics (age and gender). Model 3 includes organizational controls (dummy variable for whether the respondent works for the World Bank or for DFID, whether the respondent is posted at HQ or a country office, and the salary grade of the respondent). Model 4 includes respondent expertise relevant for this vignette (whether the subject of the respondents' highest degree was economics or development, and whether the respondent was working on health²³).

²²These results are even more striking when compared to the original results reported in Tversky and Kahneman (1981): They are extremely similar. In the original paper, 28% of the sample chose the risky policy in the gains frame (22% of our sample did the same); while 78% chose the risky policy in the losses frame (65% of our sample did the same).

²³The health expert variable was constructed based on responses to the question: "What best describes the focus of your work?" The variable takes on a value of 1 if the subject responded with the health-relevant response: "Health" (World Bank), or "Health Advisor" (DFID).

DRAFT: PLEASE DO NOT CITE OR CIRCULATE**Table 7: Likelihood of respondent choosing risky policy option^{a,b,c,d}**

Dependent variable: Respondent chooses risky policy (=1)				
	I	II	III	IV
Losses frame (=1)	1.142*** (0.05)	1.146*** (0.05)	1.170*** (0.06)	1.212*** (0.06)
Age (in years)		0.002 (0.00)	0.000 (0.00)	0.006 (0.00)
Female		-0.047 (0.05)	-0.069 (0.06)	-0.046 (0.07)
Organization (0 = DFID; 1 = World Bank)			0.140* (0.08)	0.058 (0.11)
Posting (0 = HQ; 1 = Country office)			-0.052 (0.06)	-0.067 (0.07)
Respondent grade in organization (1 = Junior; 9 = Senior)			-0.010 (0.02)	-0.037 (0.03)
Economics major (Subject of highest degree)				-0.010 (0.08)
Development major (Subject of highest degree)				0.048 (0.12)
Health expert (Respondent works on health)				-0.103 (0.17)
Constant	-0.770*** (0.04)	-0.839*** (0.14)	-0.777*** (0.18)	-0.864*** (0.22)
Log likelihood	-1524.0	-1475.0	-1361.5	-1020.1
Pseudo R2	0.139	0.141	0.147	0.158
P	0.000	0.000	0.000	0.000
Observations	2591	2512	2337	1768

Notes:

^a Probit regressions. The dependent variable takes on a value of 1 if the respondent chose the risky option.^b Table reports coefficients with standard errors in parentheses.^c * 10%, ** 5%, *** 1% significance level.

From table 7, we find strong evidence for the presence of biases arising as a result of framing. Using marginal effects (not reported), respondents were 45% more likely to select the risky option when the decision problem was framed as a loss. The result was robust to a number of additional controls, including socio-demographics, organization specific variables, and prior expertise. Respondents from the World Bank were marginally more likely to select the risky option ($p < 0.10$), but this was not robust to alternate specifications. In sum, framing dramatically affects how development professionals perceive risk. While our results are very similar to Tversky and Kahneman (1981), we do note that it is striking that policy professionals, who provide advice based on similar types of evidence, are just as susceptible to the framing of the crisis. We also note that

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

the period of our study overlaps with the Ebola virus epidemic of 2014, suggesting that the manner in which the crisis was framed could have had an impact on the decisions made by policy professionals (and hence, policymakers).

Obviously, an important issue is how the organizational context, which involves complex of factors, including reputation, career growth, and social preferences, affects risk taking. Are policy professionals more risk averse when undertaking policy decisions, relative to decisions made for themselves?²⁴ To examine this question, we offered a test for risk taking by including a risk preference measure at two points in the DfID survey. Respondents were provided with the following prompt and then asked to select one of five policy proposals, each increasing in expected value and variance.

You are the Head of Office for a large country. You have a budget of £100m to spend on a vaccination programme and your team have presented you with five proposals on how to implement the programme. The expected number of beneficiaries reached is shown in the table below. The probability of the “things go wrong” scenario is 50% while the probability of the “all goes well” scenario is also 50%. Therefore, both scenarios have an equal chance of occurring.

Expected number of beneficiaries by proposal and scenario (in millions)

Choice	Scenario 1 - things go wrong	Scenario 2 - all goes well
Proposal 1	26	38
Proposal 2	22	46
Proposal 3	18	54
Proposal 4	14	62
Proposal 5	10	70

Figure 4: Risk-taking for the organization – gamble choices

In addition to the above, to measure their own levels of risk aversion, respondents were also asked to select a gamble for themselves:

You are invited to play one of six games of chance. In each game a two-sided coin is tossed. If it lands on heads you receive the low payoff. If it lands on tails, you receive the high payoff. Which game would you play? You can only choose one.

²⁴ There are a number of papers investigating the role of agents undertaking risky decisions on behalf of principals, often with mixed results. Chakravarty et al. (2010); Polman (2012); and Pollmann et al. (2014) all find that agents exhibit *greater risk tolerance* (i.e. lower risk aversion) when undertaking risky investments on behalf of principals. On the other hand, Eriksen and Kvaloy (2010) and Fullbrunn and Luhan (2014) find that agents exhibit higher risk aversion when undertaking decisions on behalf of principals. Importantly, Pollmann et al. (2014) show that accountability reduces the difference between principal and agent risk-taking, and suggest that this mechanism can be used to discipline agent behaviour in these types of settings..

DRAFT: PLEASE DO NOT CITE OR CIRCULATE**Payoffs by game and outcome**

Choice	Low payoff (£)	High payoff (£)
Game 1	30	30
Game 2	26	38
Game 3	22	46
Game 4	18	54
Game 5	14	62
Game 6	10	70

Figure 5: Risk-taking for self – gamble choices

Note that aside from the first choice in figure 5, the risk-free option, the outcomes were identical to the prompt where respondents were asked to select a proposal in a professional capacity²⁵. Both vignettes clearly stated that each outcome had a 50% chance of occurring. While the order of the vignettes was fixed, the vignettes occurred at very different points in the survey (questions 2 and 19 for organizational and personal risk-taking, respectively). This mitigated spillovers from one vignette to the next (since the gambles were nearly identical). This within-respondents' setup allows for a direct comparison of risk-taking in two different domains.

Figure 6 displays the number of respondents that selected each gamble for the organization (Y-axis) and for self (X-axis)²⁶. If respondents evaluated the gambles identically for themselves and their organizations, then all choices would fall on the 45-degree line (dotted line in the figure)²⁷. But the figure shows that the linear trend is less than 1, which indicates that respondents were more likely to be risk-averse in their decisions for the organization than for themselves.

²⁵Our risk preference measure adds a risk-free choice in the manner of previous risk preference measures (for example, see Eckel and Grossman, 2008 and Banuri and Keefer, 2016). This was done to keep risk preference measures comparable with other studies. The risk-free option is dropped in the organizational risk-taking vignette for sake of context.

²⁶Note that the gambles in the risk preference measure have been rescaled from 1-6 in figure 5, to 0-5 in figure 6 to bring them in line with the risk proposals for the organization (figure 4).

²⁷With this exception of those risk-averse respondents that would choose the risk free option (gamble 1 in figure 5), but theoretically, these respondents would also choose the lowest risk choice for their organization (proposal 1 in figure 4).

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

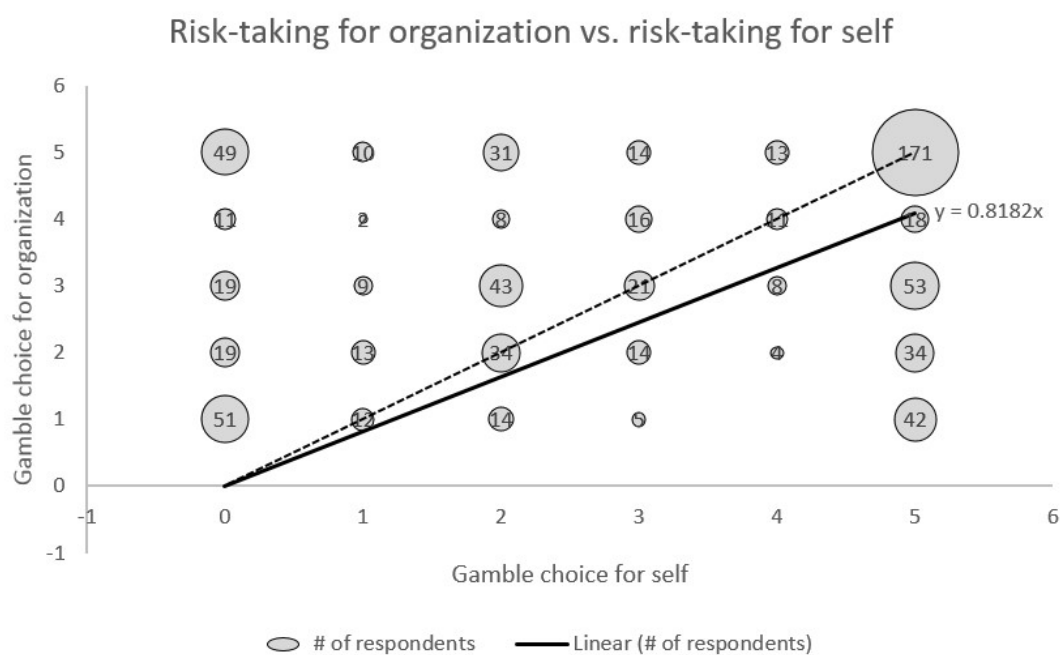


Figure 6: Respondents risk-taking behavior in their professional capacity relative to risk-taking for self

Table 8 (model 1) presents analyses to confirm this finding. Model 1 presents results from a tobit model, accounting for censoring at the lower end of the data, with risk-taking for the organization as the dependent variable. The model suppresses the constant to test whether the coefficient is different from 1 (the null hypothesis of no difference between the two questions). The table shows that the coefficient is lower and significantly different from 1 ($p < 0.01$). This result provides evidence that subjects exhibit greater risk aversion when evaluating policy choices, relative to what they would choose for themselves.²⁸

Table 8 (models 2-5) estimates OLS regressions with additional correlates. Model 2 simply tests for a relationship between risk-taking for self and for the organization. Model 3 controls for respondent evaluation of the importance of their response in the vignette for their career²⁹, in addition to

²⁸ The organizational risk-taking vignette had a control (baseline risk) and three treatments which introduced additional layers of risk: Three-fourths of respondents received an additional line indicating that there was (a) media interest (reputational risk); (b) secretary of state interest (professional risk); and (c) risk of misappropriation (fiduciary risk). The idea was to examine whether respondents chose different proposals in response to oversight and fiduciary risks. To test for differential responses to treatment, we estimated model 2 in table A.1 (appendix A), which interacts the risk preference measure with the control and each treatment separately. In this model, each coefficient is below 1, and all are significantly different from 1 ($p < 0.01$ in all four treatments). The results do not differ across treatments: that is, whether the organizational risk is framed as financial, reputational, or fiduciary does not seem to make a difference in choices. Because the results are not significantly different across these treatments, we pool all treatments together for the subsequent analysis.

²⁹ The actual text of the question is: 'I care about selecting the right proposal, as it will reflect positively on my career.' If you were in this situation, would you agree or disagree with this statement?

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

respondent gender and age. Further models control organization-specific variables, such as current post and grade (model 4), and educational background (model 5).³⁰

Table 8: Risk-taking for the organization ^{a,b,c,d}

Dependent variable: risk proposal chosen for the organization					
	I	II	III	IV	V
Risk preference measure (0 = Risk averse; 5 = risk seeking)	0.805*** (0.02)	0.194*** (0.03)	0.192*** (0.03)	0.198*** (0.03)	0.211*** (0.04)
Vignette reflects on career (5 = Strongly agree)			0.117** (0.05)	0.144*** (0.05)	0.193*** (0.06)
Age (in years)			0.017*** (0.01)	0.025*** (0.01)	0.018** (0.01)
Female			-0.112 (0.11)	-0.122 (0.11)	-0.134 (0.14)
Posting (0 = HQ; 1 = Country office)				0.006 (0.11)	0.145 (0.14)
Respondent grade in organization (1 = Junior; 9 = Senior)				-0.098*** (0.03)	-0.046 (0.04)
Economics major (Subject of highest degree)					-0.057 (0.20)
Development major (Subject of highest degree)					-0.105 (0.19)
Constant	-- --	2.787*** (0.10)	1.761*** (0.35)	1.921*** (0.40)	1.617*** (0.48)
R2	--	0.064	0.080	0.096	0.099
Left censors	124	0	0	0	0
Observations	749	749	724	671	480

Notes:

^a Model 1: Tobit regressions suppressing the constant. The dependent variable is the proposal selected by the respondent in the risk-taking vignette (organization). Model accounts for censoring at the lower level as the risk preference measure has an additional (risk-free) option.

^b Models 2-5: OLS regressions. The dependent variable is the proposal selected by the respondent in the risk-taking vignette (organization).

^c Table reports coefficients with standard errors in parentheses.

^d * 10%, ** 5%, *** 1% significance level.

We find a strong positive relationship between our risk preference measure and risk-taking for the organization, as expected in model 2 ($p < 0.01$). Risk seeking respondents were significantly more likely to be risk seeking for the organization as well. Respondents who were more likely to agree that their response is important for their career were more likely to take higher risks for the organization ($p < 0.01$). Older respondents were more likely to be risk-seeking ($p < 0.05$). Interestingly, there were

³⁰The results are robust to using a tobit specification with lower level censors.

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

no significant differences by gender ($p=0.34$). There was some evidence for a relationship between organizational risk-aversion and seniority, but this result was not robust to controlling for educational background. In sum, policy professionals appear to take fewer risks for their organizations than for themselves. Further research might help explain whether reputational or career concerns at these organizations drive this phenomenon, whether framing risk in organizational terms elicits widespread risk aversion among policy professionals, or whether other factors are at play.

Summary and discussion

These experiments with World Bank and DFID staff are the first to systematically document the presence of biases among policy professionals. The data suggest that despite the non-political charter of many bureaucracies, such as the World Bank and DFID, and despite the fact that public institutions are designed to address, in Weber's language, instrumental/purposive rationality rather than values/belief rationality, significant biases in decision-making are evident. This finding should worry policy professionals and their principals in governments and large organizations, as well as citizens themselves.

As mentioned at the outset, some bureaucracies implement procedures that attempt to address cognitive biases in decision making. It is unclear how effective these procedures are, in significant part because research on the biases of policy professionals and bureaucrats is scarce. Further exploration of the extent and magnitudes of these biases could also inform the design of bureaucratic procedures to address them. Among the procedures that may go some way in mitigating biases among bureaucrats are "red teaming" major decisions (e.g., implementing mock adversarial arguments, as in playing the "devil's advocate," or war games to identify the strengths and weakness of different courses of action or points of view), "dogfooding" products and services (e.g., sampling the products and services that consumers or citizens use before rollout), prediction tournaments, and group deliberation (Tetlock and Gardner 2015, World Bank 2015).

Hastie and Sunstein (2015) explore the effects of group deliberation and interaction on the quality of decision making. They argue that group interaction is helpful for "eureka" problems (when the right answer, once announced, is clear to all), and can mitigate the effects of availability heuristics and anchoring. On the other hand, group deliberation can increase polarization and reduce variance in the views that individuals hold ("groupthink"), largely as a result of informational and reputational cascades – individuals do not want to appear to be out of step with the views of others. Group deliberation also tends to increase the confidence, or overconfidence, of individuals in their own judgments and beliefs.

To explore the effects of one potential intervention, group deliberation, we conducted a follow-up experiment with a small sample of DFID policy professionals. First, the policy professionals were exposed, individually, to the confirmation bias, sunk cost, and loss versus gain framing experiments

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

described above. Then they met in pairs and were asked, again, to provide answers following a brief period of deliberation.³¹

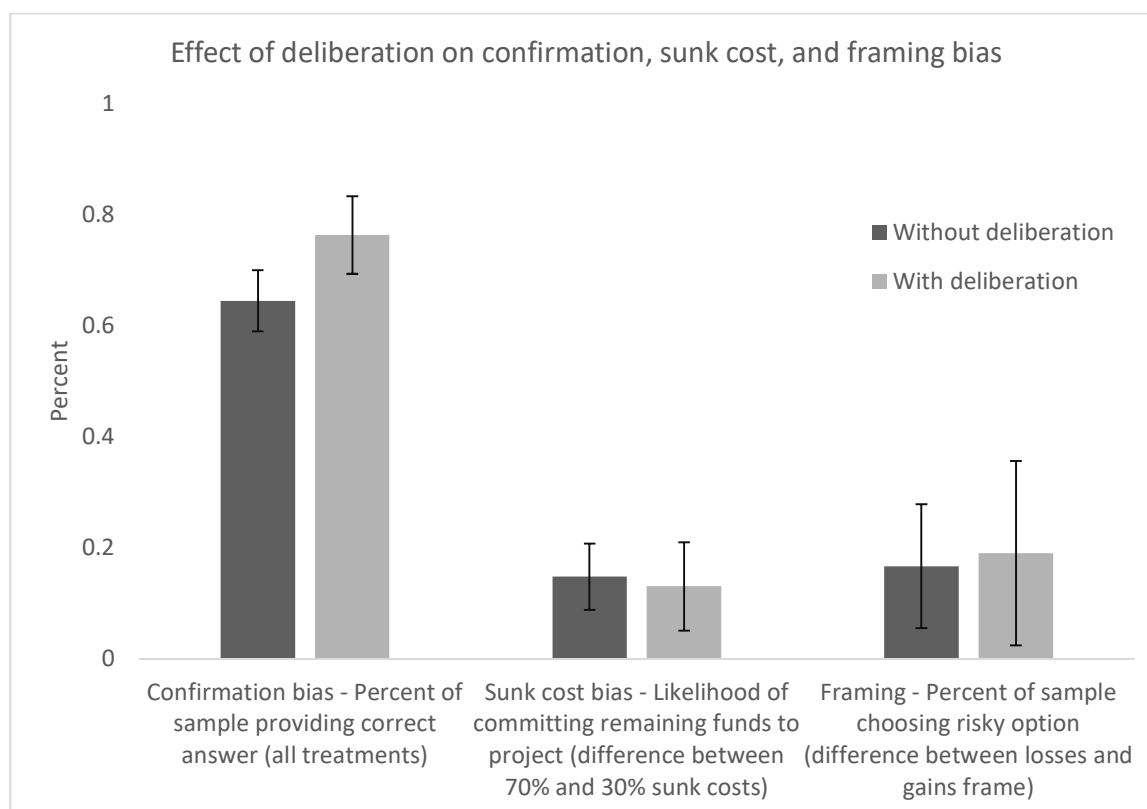


Figure 7: Effects of deliberation on confirmation, sunk cost, and framing bias

Figure 7 presents the main results of the experiment for our three biases of interest. The confirmation bias result (left) pools both treatments together and finds an overall increase in accuracy of 12%.³² For sunk cost bias (center), without deliberation, we find that increasing the level of sunk cost (from 30 to 70%) increases the likelihood of committing remaining funds to the project

³¹The deliberation experiment was conducted during an annual retreat for economists working for DfID. Unlike the World Bank and DfID staff survey reported earlier, this experiment was conducted using pen and paper (rather than online). This (plus the small sample) restricted the number of treatments that could be employed for the experiment. Hence, there were two versions of the survey, each containing a single sequence of treatments for the confirmation, sunk cost, and framing bias experiments: Version 1 contained 70% sunk cost, the gains frame, and confirmation bias vignette data supporting income falling as a result of the intervention. Version 2 contained 30% sunk cost, the losses frame, and confirmation bias vignette data supporting income rising as a result of the intervention. 81 participants were randomly allocated to either version 1 or 2 of the survey at the beginning of the retreat. In the afternoon, the participants that were given version 1 of the survey were randomly paired and given version 1 again, and asked to discuss their answers (and the same for version 2). Hence, each individual completed the same version of the survey twice, one individually, and once with a partner. Three participants left the retreat early to meet other commitments, so there were a total of 39 groups in the partner phase of the survey.

³²This increase in accuracy, while modest, could also be explained by experience: respondents participate in the survey twice, and may have engaged in the calculation once again independent of any deliberation effects. Our design does not allow us to test for the effects of experience independently of deliberation, however.

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

by 15% in the absence of deliberation. We observe no change with deliberation, increasing the level of sunk cost increases the likelihood of committing remaining funds to the project by 13%. Finally, for the framing vignette (right), participants exposed to the losses frame were 17% more likely (relative to gains) to choose the risky policy option in the absence of deliberation. With deliberation, there is no real change in this difference: participants were 19% more likely to choose the risky policy option in the losses frame (relative to gains). Hence, we find that group deliberation shows some evidence of mitigating biases associated with confirmation bias. It did not, however, have any effect on the biases associated with sunk costs, or framing.

The reason for this, we suspect, is that the confirmation bias experiment is a “eureka” problem. There is a “right” answer to the minimum wage experiment, in the sense that in each treatment the data are more consistent with one interpretation of the data than with its converse. The experiment regarding the framing of losses and gains is different. The main finding relies on how different respondents (or pairs of respondents) answer differently under different frames, not on whether individuals or pairs provide the correct answer in a given frame. Similarly, the sunk cost vignette does not have a “correct” course of action either. Since the responses then are simply based on preferences, deliberation may well aggregate the preferences within groups, but does not yield any bias mitigation.

To summarize, deliberation seems to be a more effective safeguard for some problems than others. In general, when organizations implement procedural safeguards to reduce biases and improve decisions among policy professionals, they need to be cognizant that the safeguards may be effective for some biases and ineffective (or even counterproductive) for others. Group deliberation might reduce confirmation bias, but have no effect the framing of risk, and could even increase a policy professional’s risk aversion if the group makes risk-aversion in the authorizing environment more salient.

Psychological processes and social contexts can have dramatic effects on the ways in which policy professionals use and interpret data, make project decisions, and take personal and organizational risks. Public sector organizations rarely study and collect data on the biases of their staff. As a result, they are often unaware of the extent to which these influences affect their organizations’ choices and policies, and cannot develop evidence-based approaches to mitigating biases. Efforts to study and understand the decision making processes of policy professionals could help governments and international organizations devise more effective policies and programs, and to live up to legal and professional ideals of objectivity, impartiality, and scientific accuracy.

DRAFT: PLEASE DO NOT CITE OR CIRCULATE**References**

- Arkes, Hal R., and Catherine Blumer. 1985. "The psychology of sunk cost." *Organizational Behavior and Human Decision Processes*, 35(1): 124–140.
- Andersson, Ola, Håkan J. Holm, Jean-Robert Tyran, and Erik Wengström. 2014. "Deciding for others reduces loss aversion." *Management Science*, 62(1): 29–36.
- Banuri, Sheheryar, and Philip Keefer. 2016. "Pro-social motivation, effort and the call to public service." *European Economic Review*, 83(C):139–164.
- Bero, Lisa A., and Alejandro R. Jadad. 1997. "How consumers and policymakers can use systematic reviews for decision making." *Annals of Internal Medicine*, 127(1): 37–42.
- Bradler, Christiane. 2009. "Social preferences under risk-an experimental analysis." *Jena Economic Research Papers*22: 1–35.
- Brockner, Joel and Jeffrey Z. Rubin. 1985. *Entrapment in Escalating Conflicts: A Social Psychological Analysis*. New York: Springer-Verlag.
- Calvert, Randall L. 1985. "The value of biased information: A rational choice model of political advice." *The Journal of Politics*, 47(2): 530–555.
- Chakravarty, Sujoy, Glenn W. Harrison, Ernan E. Haruvy, and E. Elisabet Rutström. 2011. "Are you risk averse over other people's money?" *Southern Economic Journal*, 77(4): 901–913.
- Charness, G. and Sutter, M., 2012. Groups make better self-interested decisions. *The Journal of Economic Perspectives*, 26(3), pp.157-176.
- Cooke, Roger M. 1991. *Experts in Uncertainty: Opinion and Subjective Probability in Science*. New York: Oxford University Press.
- DellaVigna, S. and Pope, D., 2016. Predicting Experimental Results: Who Knows What? (No. w22566). National Bureau of Economic Research.
- Eckel, Catherine C., and Philip J. Grossman. 2002. "Sex differences and statistical stereotyping in attitudes toward financial risk." *Evolution and Human Behavior*, 23(4): 281–295.
- Eckel, Catherine C., and Philip J. Grossman. 2008. "Men, women and risk aversion: Experimental evidence." In *Handbook of Experimental Economics Results*, edited by Charles Plott and Vernon Smith, 1(7): 1061–1073.
- Englich, Birte, Thomas Mussweiler, and Fritz Strack. 2006. "Playing dice with criminal sentences: The influence of irrelevant anchors on experts' judicial decision making." *Personality and Social Psychology Bulletin*, 32(2): 188–200.
- Eriksen, Kristoffer W., Kvaloy, Ola. 2010. "Myopic investment management." *Rev. Finance* 14, 521–542.
- Frederick, Shane. 2005. "Cognitive reflection and decision making." *The Journal of Economic Perspectives*, 19(4): 25–42.
- Füllbrunn, Sascha. & Luhan, Wolfgang J. 2015. "Am I My Peer's Keeper? Social Responsibility in Financial Decision Making." NiCE Working Paper 15-03, Institute for Management Research, Radboud University, The Netherlands.
- Garland, Howard, and Stephanie Newport. 1991. "Effects of absolute and relative sunk costs on the decision to persist with a course of action." *Organizational Behavior and Human Decision Processes*, 48(1): 55–69.
- Hastie, Reid, and Cass Sunstein. 2015. *Wiser: Getting Beyond Groupthink to Make Groups Smarter*. Boston, MA: Harvard Business Review Press.
- Herek, Gregory M., Irving L. Janis, and Paul Huth. 1987. "Decision making during international crises is quality of process related to outcome?" *Journal of Conflict Resolution*, 31(2): 203–226.

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

- Jelveh, Zubin, Bruce Kogut, and Suresh Naidu. 2015. "Political language in economics." Working Paper 14-57, Columbia Business School.
- Kahan, Dan M., Ellen Peters, Erica C. Dawson, and Paul Slovic. 2013. "Motivated numeracy and enlightened self-government." Working Paper 307, Yale Law School.
- Kahneman, Daniel, and Amos Tversky. 1979. "Prospect theory: An analysis of decision under risk." *Econometrica*, 47(2): 263–291.
- Langfeldt, Liv. 2004. "Expert panels evaluating research: decision-making and sources of bias." *Research Evaluation*, 13(1): 51–62.
- Lord, Charles G., Lee Ross, and Mark R. Lepper. 1979. "Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence." *Journal of Personality and Social Psychology*, 37(11): 2098–2109.
- Levy, Jack S. 2003. "Applications of prospect theory to political science." *Synthese*, 135(2): 215–241.
- Malmendier, U. and Tate, G., 2008. Who makes acquisitions? CEO overconfidence and the market's reaction. *Journal of financial Economics*, 89(1), pp.20-43.
- McDermott, Rose. 2004. "Prospect theory in political science: Gains and losses from the first decade." *Political Psychology*, 25(2): 289–312.
- Pollmann, M.M., Potters, J. and Trautmann, S.T. 2014. "Risk taking by agents: The role of ex-ante and ex-post accountability." *Economics Letters*, 123(3), pp.387-390.
- Polman, Evan. 2012. "Self–other decision making and loss aversion." *Org. Behav. Human Decision Process*. 119, 141–150
- Schaubroeck, John, and Elaine Davis. 1994. "Prospect theory predictions when escalation is not the only chance to recover sunk costs." *Organizational Behavior and Human Decision Processes*, 57(1): 59–82.
- Shanteau, James. 1988. "Psychological characteristics and strategies of expert decision makers." *Acta Psychologica*, 68(1): 203–215.
- Shanteau, James. 1992. "Competence in experts: The role of task characteristics." *Organizational Behavior and Human Decision Processes*, 53(2): 252–266.
- Song, Fei. 2008. "Trust and reciprocity behavior and behavioral forecasts: Individuals versus group-representatives." *Games and Economic Behavior*, 62(2): 675–696.
- Staw, Barry M. 1976. "Knee-deep in the big muddy: A study of escalating commitment to a chosen course of action." *Organizational Behavior and Human Decision Processes*, 16(1): 27–44.
- Stewart, Dennis D., and Garold Stasser. 1995. "Expert role assignment and information sampling during collective recall and decision making." *Journal of Personality and Social Psychology*, 69(4): 619–628.
- Sunstein, Cass R. 2006. *Infotopia: How many minds produce knowledge*. New York: Oxford University Press.
- Tetlock, Philip E., and Dan Gardner. 2015. *Superforecasting: The Art and Science of Prediction*. New York: Broadway Books.
- Tversky, Amos, and Daniel Kahneman. 1981. "The framing of decisions and the psychology of choice." *Science*, 211 (4481): 453–58.
- Weber, Max. 1946. *From Max Weber: Essays in Sociology*, edited by Hans H. Gerth and Charles W. Mills. New York: Oxford University Press.
- World Bank. 2015. *World Development Report 2015: Mind, Society, and Behavior*. Washington DC: World Bank.

DRAFT: PLEASE DO NOT CITE OR CIRCULATE

DRAFT: PLEASE DO NOT CITE OR CIRCULATE**Appendix A: Additional tables for review****Table A.1: Risk-taking for the organization vs. risk-taking for self^{a,b,c,d}**

Dependent variable: risk proposal chosen for the organization		
	I	II
Risk preference measure (0 = Risk averse; 5 = risk seeking)	0.805*** (0.02)	
Risk preference X Baseline risk		0.833*** (0.05)
Risk preference X Reputational risk		0.830*** (0.04)
Risk preference X Professional risk		0.762*** (0.05)
Risk preference X Fiduciary risk		0.789*** (0.05)
Log likelihood	-1520.5	-1519.7
Observations	749	749
Left censors	124	124

Notes:

- ^a Tobit regressions suppressing the constant. The dependent variable is the proposal selected by the respondent in the risk-taking vignette (organization).
- ^b Models account for censoring at the lower level as the risk preference measure has an additional (risk-free) option.
- ^c Table reports coefficients with standard errors in parentheses.
- ^d * 10%, ** 5%, *** 1% significance level.