


RESEARCH ARTICLE  

Incentivizing Responses in International Organization Elite Surveys: Evidence from the World Bank

Mirko Heinzl¹ , Catherine Weaver² and Ryan Briggs³

¹London School of Economics and Political Science, London, UK, ²The University of Texas at Austin, Austin, TX, USA and ³University of Guelph, Guelph, ON, Canada

Corresponding author: Mirko Heinzl; Email: m.n.heinzl@lse.ac.uk



Abstract

Scholars of International Organizations (IOs) increasingly use elite surveys to study the preferences and decisions of policymakers. When designing these surveys, one central concern is low statistical power, because respondents are typically recruited from a small and inaccessible population. However, much of what we know about how to incentivize elites to participate in surveys is based on anecdotal reflections, rather than systematic evidence on which incentives work best. In this article, we study the efficacy of three incentives in a preregistered experiment with World Bank staff. These incentives were the chance to win an Amazon voucher, a donation made to a relevant charity, and a promise to provide a detailed report on the findings. We find that no incentive outperformed the control group, and the monetary incentive decreased the number of respondents on average by one-third compared to the control group (from around 8% to around 5%).

Keywords: incentives; elite survey; experiment; international organizations; World Bank

Introduction

Elite survey experiments have become increasingly popular in political science research (Dietrich, Hardt, and Swedlund 2021; Kertzer and Renshon 2022). To recruit respondents, scholars draw typically on small populations of interest and need to achieve relatively high response rates to ensure sufficient statistical power. They have used various techniques to recruit respondents for elite surveys, including cultivating personal contacts and snowballing respondents

  This article has earned badges for transparent research practices: Open Data and Open Materials. For details see the [Data Availability Statement](#).

© The Author(s), 2024. Published by Cambridge University Press on behalf of American Political Science Association. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

(Dietrich 2021); phone calls (Dellmuth *et al.* 2021), recruiting respondents at summits (Jongen and Scholte 2022), and cold emailing (Briggs 2021). Our study focuses on the hardest case: increasing responses to cold emails where interaction with respondents is limited and direct persuasion is most difficult. Recent reviews of elite experiments in political science (Kertzer and Renshon 2022) and international relations (Dietrich, Hardt, and Swedlund 2021) have discussed monetary incentives, like Amazon gift card lotteries, and informational incentives, like offering to share reports on research results, to increase response rates to cold emails. However, we often lack systematic evidence on whether these incentives work – especially for elites involved in international politics. To this end, Kertzer and Renshon (2022, 15) call for “more guidance based on systematic empirical work, and less on anecdotes and hunches about why one approach worked and another did not.” In this paper, we take up this challenge for surveys focusing on one particular respondent population that is of crucial interest for research on international relations: elites working in International Organizations (IOs) – specifically the World Bank.

Research on IOs has increasingly used elite surveys to study questions of international cooperation. We systematically searched through all articles published in 20 journals typically targeted by scholars of IOs¹ and identified 23 articles surveying IO staff that were published since 2005.² While the method is still relatively rare, its popularity is increasing with 50% of articles published in the last five years. Two findings stand out from our review of the existing studies. First, the transparency of authors regarding their response rates and whether they used incentives is rather low. 34% of studies do not list or give sufficient detail to calculate a response rate. Only 2 out of the 23 studies we reviewed discussed which survey incentives they used (Clark 2021; Hooghe 2005). It is impossible to tell from the other articles whether they did not use any incentives or simply did not report them. Second, the response rate varies vastly between different IOs. 17 out of 23 studies focus on the EU or specific EU agencies. These studies attained on average a response rate of 37.7%. The other IO surveys that discuss response rates focus on the AU (10%), the World Bank (4.6%), and staff from 28 IOs contacted through LinkedIn ads. Given the lack of empirical research on survey incentives, the lack of transparency on responses in existing studies and the widely differing response rates, it is difficult to ascertain best practices for incentivizing responses in IO elite surveys. The vast differences between organizations also imply that existing evidence on INGOs (Safarpour, Bush, and Hadden 2022) might not generalize well to IOs.

¹American Journal of Political Science, American Political Science Review, European Journal of International Relations, European Union Politics, Global Governance, Global Policy, Governance, International Organization, International Security, International Studies Quarterly, Journal of Politics, Journal of Common Market Studies, Journal of Conflict Resolution, Journal of European Public Policy, Journal of Public Administration Research and Theory, Public Administration Review, Review of International Organizations, Review of International Political Economy, Security Studies, World Development.

²We list the full table of studies we identified in the Appendix.

To test the efficacy of incentives to boost survey response rates, we embedded a preregistered³ experiment⁴ in the contact requests for a conjoint survey experiment with World Bank staff members on decision-making factors in the management of aid projects. We randomized three types of incentives – entering into a lottery for an Amazon gift card, a charity donation, and provision of study results – and their combination in contact emails. Contrary to our expectation, we find that none of the incentives outperforms the control group. The Amazon gift card lottery performs worst, decreasing the likelihood of response by 2.9 percentage points compared to a response rate in the control group of 8.4%. This is a sizeable drop that should caution against gift card incentives for surveys with international organization staff. Furthermore, the finding contradicts the results presented by Safarpour, Bush, and Hadden (2022) for INGOs. Therefore, the results imply that incentives work differently for different (types of) organizations. Finally, we find some exploratory evidence that implies that information treatment may have increased the time respondents spent on the survey. Our findings imply that researchers need to understand the organizational context in which potential respondents are embedded. This will help avoid designing incentives that are misaligned with the preferences or organizational culture of their potential respondents.

Types of incentives in elite experiments

Our first incentive was monetary: the chance to win via lottery⁵ one 100 USD Amazon gift card. The idea of a gift card incentive is to compensate respondents for their time without having to transfer money directly to respondents. Typically, Amazon gift cards are used due to the broad range of products respondents could buy using these gift cards. For example, Clark (2021) uses a 250 USD Amazon gift card lottery to incentivize participation in his survey of staff working at development organizations (Clark 2021). We selected 100 USD, because it is the maximum allowed amount in the World Bank's gift policy in one calendar year.

The second incentive was the promise of a 10 USD donation to a charity (up to 3000 USD) on behalf of each respondent who completes the survey. To our knowledge, we were the first to use a charity incentive. The motivation was that elites may be insensitive to being offered an equivalent (small) amount of money, in part because it can be administratively or technically cumbersome to receive such payments or because there are unknown odds of winning in a lottery treatment. At the same time, elites might be motivated by a donation to a charity that works on a cause that they care about. Our offer to them was thus to make this donation on

³The survey was preregistered at <https://osf.io/mv4be>. We initially planned to send only two reminders but were made aware that our first two emails were spam-filtered. We rewrote the email request in subsequent emails and deleted most links from our email. Once we got around the spam filter, we received over 50 direct email responses from TTLs apologizing for missing the original email and agreeing to complete the survey. To ensure that respondents received at least one reminder in their inbox, we decided to depart from our preregistration by sending an additional reminder. This change was the only departure from the preregistration.

⁴The project was approved by the University of Texas at Austin Institutional Review Board.

⁵We chose a lottery and not an assured monetary payment due to budget constraints as we could not offer all respondents the maximum gift amount the World Bank allows.

their behalf and with no hassle. We selected GiveDirectly, an organization that does cash transfers to the global poor, because it is well-known and respected in the international development community.

Third, we provided informational incentives. Dietrich, Hardt, and Swedlund (2021) suggest promising a short report on the survey results to respondents. The idea behind such informational incentives is that respondents may be motivated by learning something about their areas of work. Informational incentives have been used by scholars conducting elite surveys and government officials in national bureaucracies, for example, by Dietrich (2021) in her research on the views of officials from foreign aid ministries, and Herold *et al.* (2021) in their survey of high-level civil servants working on financial and agriculture policy in more than 100 countries.

We hypothesized that each incentive would have a positive effect on the probability of the respondent starting the survey. In addition, we assumed that different incentives might be combined to increase response rates (Dietrich, Hardt, and Swedlund 2021). To this end, we sent some respondents combinations of either monetary treatment and the information treatment.

Research design

Our survey population was World Bank Task Team Leaders (TTLs), the individuals in charge of designing and managing projects at the organization. TTLs are senior staff members managing large project budgets – the average project cost of a World Bank project approved in 2020 was 324 million USD (World Bank 2020). They work either in the World Bank’s Washington, DC headquarters and go to recipient countries on missions or in the recipient country’s field office. TTLs are responsible for almost all major decisions taken during the implementation period of a World Bank project (Heinzl and Liese 2021).

We invited respondents to participate in a survey that focuses on “the opinions of TTLs at the World Bank on project design decisions.” To recruit respondents, we drew on a dataset of the names of TTLs reported on the World Bank website. From this, we generated the email addresses of 4,949 TTLs in charge of at least one World Bank project between 2000 and 2020 based on the World Bank email template.

To understand the effectiveness of the different incentives, we randomly allocated respondents into one of six groups. The TTLs received identical emails that only differed in the promised incentive, as described in Table 1. The final group was a control group.

Each of the six groups contained around 825 email addresses. Approximately 53% of emails failed to reach the recipient because people had left the World Bank, the email address was incorrect, their inbox was full, or the server rejected our email. Therefore, our population of interest included 2328 World Bank TTLs who received our email request.⁶ We sent three reminders to ask for participation in our survey, and the survey was online for 38 days (08th March 2022 to 15th April 2022).

⁶We did not preregister how we would handle bounced emails, but excluded them since World Bank staff who were not reached never had the choice to answer the survey. We also conduct a robustness check in the Appendix that re-estimates the models using the full sample of respondents.

Table 1. Wording of the five incentives used in this study

Group	Treatment
Information	To express our gratitude, we will send a report on our findings to you if you are interested in receiving them.
Amazon	To express our gratitude, we will give out a 100 USD Amazon gift card to one randomly selected respondent.
Charity	To express our gratitude, we will donate 10 USD (up to a maximum of 3000 USD) to people living in poverty for each respondent that answers our survey. The donation will be made through GiveDirectly (https://www.givedirectly.org/).
Information + Amazon	To express our gratitude, we will give out a 100 USD Amazon gift card to one randomly selected respondent. We will also send a report on our findings to you if you are interested in receiving them.
Information + Charity	To express our gratitude, we will donate 10 USD (up to a maximum of 3000 USD) to people living in poverty for each respondent that answers our survey. The donation will be made through GiveDirectly (https://www.givedirectly.org/). We will also send a report on our findings to you if you are interested in receiving them.

In line with official World Bank studies and studies by academic researchers, we expected a very low response rate when cold emailing TTLs. Comparable surveys with World Bank staff typically have response rates of around 5–6% (Briggs 2021). 16.9% of emailed respondents clicked on the survey link, and our overall response rate is 7.6%.⁷

The efficacy of different incentives in elite experiments

We estimated linear probability models with robust standard errors.⁸ The unit of analysis was the contacted World Bank staff member. The dependent variable was a binary variable indicating whether an individual completed the survey. We created binary variables for each treatment, coded as one if the treatment was included in the contact email and zero if not. We present descriptive statistics on the sample of respondents that answered the survey (by educational background, main sector of work, main region of work, nationality, and gender) in Figures A1–A5.

Figure 1 illustrates the overall response rate for each group estimated using a linear probability model including all five treatment conditions. We visualize this model to ease the interpretation of the results but note that we did not preregister this specification. The preregistered results are substantively very similar and we present them in Table 2. The response rate in the control group was around 8%. The Amazon incentive decreased the likelihood of a response by 3.5 percentage points on average. We find the information and charity treatments and the interaction terms have near-zero estimated effects. Using an equivalence test (two one-sided

⁷The overall 7.6% response rate is calculated based on 178 completed surveys out of 2328 received emails. If we calculate our response rate based on emails sent (but bounced and so never received), then our response rate is 3.6%.

⁸Replication data for this article can be found in Heinzel, Weaver and Briggs (2023).

Table 2. Regressing survey completion on different types of incentives

	(1)	(2)	(3)	(4)
Information	-0.0040 (0.0107)	-0.0072 (0.0137)	-0.0050 (0.0129)	-0.0123 (0.0196)
Amazon	-0.0286* (0.0128)	-0.0333* (0.0164)	-0.0285* (0.0128)	-0.0356* (0.0178)
Charity	-0.0029 (0.0137)	-0.0027 (0.0137)	-0.0042 (0.0179)	-0.0076 (0.0193)
Information * Amazon		0.0099 (0.0215)		0.0151 (0.0256)
Information * Charity			0.0029 (0.0231)	0.0102 (0.0274)
Constant	0.0842*** (0.0111)	0.0856*** (0.0118)	0.0846*** (0.0117)	0.0880*** (0.0136)
<i>N</i>	2328	2328	2328	2328
<i>R</i> ²	0.002	0.003	0.002	0.003

Robust standard errors in parentheses; **p* < 0.10, ***p* < 0.05, ****p* < 0.01, *****p* < 0.001.

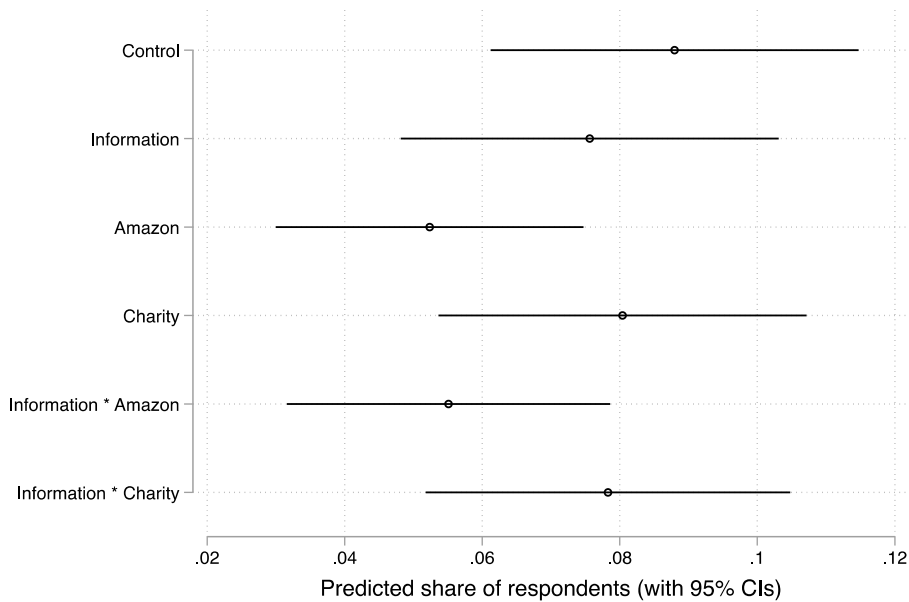


Figure 1. Predicted response rate by treatment groups.

Note: Predicted share of respondents with 95% confidence intervals for each treatment condition based on Model 4, Table 2.

t-test), we can rule out effects larger than 3 percentage points in either direction for the information and charity treatments with $p < 0.05$. The estimates show a $p < 0.05$ decrease for the Amazon and a $p < 0.1$ decrease for the Information * Amazon group. In other words, combining the information and the Amazon treatments does not appear to make a statistically significant difference in the effect size.

Table 2 presents the full results. Model 1 includes only the three treatments, Models 2 and 3 further estimate the interactions, and Model 4 includes the full set of treatments. Models 1, 2, and 3 were preregistered while Model 4 was not. The results are consistent across Model specifications. The Amazon treatment appears to reduce the response rate by between 2.9 and 3.6% and the coefficients are statistically significant ($p < 0.05$). We find the information and charity treatments and the interaction terms have near-zero estimated effects.

We conducted several additional exploratory (not preregistered) analyses reported in the supplementary appendix. First, we re-estimated our models using an alternative dependent variable: link clicks. Around 17% of the World Bank staff members who received the survey clicked on the link. We analyze these link clicks to understand whether incentives impacted motivations to click or finish the survey (Table A1). The results show that the Amazon incentive appears to disincentivize clicking and completion. The effect remains negative, albeit smaller and only marginally significant ($p < 0.1$).

Second, the overall response rate can only tell us so much about whether incentives are useful. Researchers are not just interested in maximizing responses but also in attaining unbiased samples. Therefore, we probed whether the different incentives predict differences in respondents' educational background (Table A2), gender (Table A3) and whether they work differently for people from or working on African countries – where GiveDirectly is primarily active (Table A4). We find a $p < 0.1$ decrease in the likelihood that respondents are women for the information treatment (Model 11). All other coefficients are close to zero and fail to attain statistical significance at conventional thresholds. Furthermore, we re-estimate the models using the full sample of contacted World Bank staff members – including bounced emails. The results remain similar (Table A5).

Moreover, we analyzed the time respondents (in seconds) took to finish the survey (Table A6). We estimated negative binomial models because the data are over-dispersed. The coefficient for the informational treatment is statistically significant ($p < 0.01$), positive, and sizeable. The results imply that people who got the information treatment took considerably more time filling out the survey. Therefore, the exploratory analysis suggests that information treatments may be worthwhile despite the null results on their effects on survey completion as they appear to increase respondent effort.

Conclusion

Researchers conducting elite experiments often wish to increase their response rates to cold emails through the use of incentives. However, systematic studies of their efficacy are rare (Kertzer and Renshon 2022; Safarpour, Bush, and Hadden 2022) and absent for surveys focusing on IO elites. This study analyzed three types of

incentives used in an elite experiment with World Bank staff members conducted in March and April 2022. We tested whether promising entry in an Amazon gift card lottery, a charitable donation to GiveDirectly, or the promise to receive a detailed report on research findings affect the likelihood that elites respond to email invitations to participate in an academic survey. We showed that the Amazon gift card lottery decreased the response rate for our survey by around one-third compared to the control group (about 5% versus about 8%). The other two tested incentives had near-zero coefficients and were not statistically significantly different from the control group.

We did not test the mechanisms driving this drop in responses. However, the existing literature provides some suggestions. First, gift policies of public sector organizations may render monetary incentives complicated to manage for respondents. While our incentive was designed not to exceed the World Bank's gift policy, elites may be unsure of what the rules are and decide not to participate (Dietrich, Hardt, and Swedlund 2021). The gift policy also severely limits the amount that may be promised. A promise for a chance to attain a 100 USD Amazon gift card might simply not be enough to incentivize well-paid elite bureaucrats. Low amounts may be especially problematic if paperwork is needed to accept gifts. Additionally, email recipients may have judged the likelihood of winning the lottery as very low. The efficacy of monetary incentives could increase with larger or assured incentives. Second, monetary incentives may decrease the reputation of the researchers in the eyes of respondents. Some have suggested that elites find monetary incentives insulting if they do not perceive them as accurate compensation for their time (Renshon 2015). Furthermore, elites are faced with many scam emails, and they need to decide whether they perceive the request as credible. The mention of a gift card might increase their reluctance to click on a link sent to them by someone they do not know. This was confirmed by several TTLs who sent us emails asking whether there would be an alternative way to participate in the survey that did not involve clicking on an email link.

Our experience in administering the survey also reveals the challenges scholars face in conducting elite surveys via email. First, the timing of elite surveys matters. We intentionally sent out our initial request well in advance of the World Bank and IMF Spring meetings and the end of the fiscal year; nonetheless, we received many out-of-office or direct messages that indicated that many TTLs were away on missions or immersed in other deadlines. We also discovered that the World Bank security protocols had effectively redirected nearly all of our email requests to spam folders. Hence, our first two email requests only got a response rate of less than 1%. By deleting additional links from our email request and changing the wording of the email request, we managed to get around the spam filter and received many direct email responses from TTLs apologizing for missing the original email and agreeing to complete the survey.

In total, we received over 100 direct email responses to our survey requests that were not out-of-office messages. From these messages, we were able to glean a few more insights on TTL non-response rates. A very small number ($n = 9$) replied with substantive concerns about the survey design. Others ($n = 33$) noted that they felt that they lacked experience to answer the survey. Most critically, we received 35

emails that inquired as to whether our research was approved by or conducted in collaboration with World Bank management. Despite recent policy reforms intended to enhance World Bank transparency, staff members explained that they had media training that highlighted that they should refrain from sharing their opinions about organizational policies with the outside world.

Our findings imply that differences between organizations may shape the efficacy of incentives scholars can use. While monetary incentives appear to work well in INGOs, they appear to have discouraged responses in our survey (Safarpour, Bush, and Hadden 2022). Researchers should also consult with insiders about what incentives may work in their specific population of interest given different organizational norms and policies about participating in external research (Kertzer and Renshon 2022). Pilot testing may also be optimal if it poses little threat to the ensuing statistical power of the main study or pollutes the targeted population (via unintended sharing or discussion of the survey prior to refinement). This context specificity limits the generalizability of our World Bank study to other IOs. Nevertheless, strict gift policies exist in other IOs and monetary incentives may, thus, not be the best approach.

We also did not find a benefit in response rate from either charity or information incentives. However, they also did not appear to decrease response rates significantly. Providing information incentives may be appropriate for other reasons. Informational incentives may be advisable to prevent “poisoning the well” for future academic research projects that are seeking participation of the same respondents (Kertzer and Renshon 2022). We also find some (exploratory) evidence that information incentives increase the engagement of respondents and the efficacy of informational incentives may increase for topics of very high interest for respondents. Furthermore, charity incentives may work in other contexts (for example, where elites are not directly involved in development work) and they could also have broader benefits by channeling resources that might have otherwise been used for Amazon gift cards to important social causes. Therefore, we encourage more research into their efficacy, for example, by promising donations to different types of charities and varying the amounts promised for the donations. Future research should also probe the efficacy of different strategies than cold emailing to recruit respondents in elite surveys.

Supplementary material. To view supplementary material for this article, please visit <https://doi.org/10.1017/XPS.2023.39>

Data availability statement. The data, code, and any additional materials required to replicate all analyses in this article are available at the Journal of Experimental Political Science Dataverse within the Harvard Dataverse Network, at: <https://doi.org/10.7910/DVN/IR9CYA>.

Author contribution. Research design and conceptualization: MH (34%), RB (33%), CW (33%); literature review: CW (100%); statistical analysis: MH (70%), RB (30%); writing: MH (40%), RB (30%), CW (30%).

Funding statement. This work was supported by the Robert Strauss Center for International Security and Law, The University of Texas at Austin.

Competing interests. The authors have no conflicts of interest to declare.

Ethics statement. The study received exemption from review from the University of Texas at Austin Institutional Review Board under the study number: STUDY00001310. The research adheres to APSA's Principles and Guidance for Human Subjects Research. More information on the study can be found in the supplementary appendix of this article.

References

- Briggs, Ryan C.** 2021. "Why Does Aid Not Target the Poorest?" *International Studies Quarterly* 65: 739–52.
- Clark, Richard.** 2021. "Pool or Duel? Cooperation and Competition among International Organizations." *International Organization* 75: 1133–53.
- Dellmuth, Lisa, Jan Aart Scholte, Jonas Tallberg, and Soetkin Verhaegen.** 2022. "The Elite-Citizen Gap in International Organization Legitimacy." *American Political Science Review* 116(1): 283–300.
- Dietrich, Simone.** 2021. *States, Markets, and Foreign Aid*. Cambridge: Cambridge University Press.
- Dietrich, Simone, Heidi Hardt, and Haley J. Swedlund.** 2021. "How to Make Elite Experiments Work in International Relations." *European Journal of International Relations First View* 27(2): 596–621.
- Heinzl, Mirko, and Andrea Liese.** 2021. "Managing Performance and Winning Trust: How World Bank Staff Shape Recipient Performance." *Review of International Organizations* 16: 625–53.
- Heinzl, Mirko, Catherine Weaver, and Ryan Briggs.** 2023. "Replication Data for: Incentivizing Responses in International Organization Elite Surveys: Evidence from the World Bank." <https://doi.org/10.7910/DVN/IR9CYA>.
- Herold, Jana, Andrea Liese, Per-Olof Busch, and Hauke Feil.** 2021. "Why National Ministries Consider the Policy Advice of International Bureaucracies: Survey Evidence from 106 Countries." *International Studies Quarterly* 65: 669–82.
- Hooghe, Liesbet.** 2005. "Several Roads Lead to International Norms, but Few via International Socialization: A Case Study of the European Commission." *International Organization* 59: 861–98.
- Jongen, Hortense, and Jan Aart Scholte.** 2022. "Inequality and Legitimacy in Global Governance: An Empirical Study." *European Journal of International Relations* 28: 667–95.
- Kertzer, Joshua D., and Jonathan Renshon.** 2022. "Experiments and Surveys on Political Elites." *Annual Review of Political Science* 25: 1–26.
- Renshon, Jonathan.** 2015. "Losing Face and Sinking Costs: Experimental Evidence on the Judgment of Political and Military Leaders." *International Organization* 69: 659–95.
- Safarpour, Alauna, Sarah Sunn Bush, and Jennifer Hadden.** 2022. "Participation Incentives in a Survey of International Non-Profit Professionals." *Research and Politics* 9(3): 1–6.
- World Bank.** 2020. "World Bank Projects & Operations." <https://projects.worldbank.org/> (March 2, 2020).

Cite this article: Heinzl M, Weaver C, and Briggs R. Incentivizing Responses in International Organization Elite Surveys: Evidence from the World Bank. *Journal of Experimental Political Science*. <https://doi.org/10.1017/XPS.2023.39>