**Title**

Essays in Development Economics and Political Economy



**ISBN**

9798189270727

**Author**

Cheung, Chiman

**Publication Date**

2026-05-01

Peer reviewed|Thesis/dissertation

# Essays in Development Economics and Political Economy

By

Chiman Cheung

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Agricultural and Resource Economics

and the Designated Emphasis

in

Political Economy

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Francis Annan, Co-Chair

Professor Ernesto Dal Bó, Co-Chair

Professor Frederico S. Finan

Professor Aprajit Mahajan

Spring 2026

Abstract

Essays in Development Economics and Political Economy

By

Chiman Cheung

Doctor of Philosophy in Agricultural and Resource Economics

Designated Emphasis in Political Economy

University of California, Berkeley

Professor Francis Annan, Co-Chair

Professor Ernesto Dal Bó, Co-Chair

This dissertation studies how information, institutions, and social preferences shape collective responses to public problems. Across three empirical settings, the chapters examine why individuals and communities do or do not act on information that is socially valuable: information about environmental harm, public infrastructure quality, and public-health risks. The chapters share a common concern with the political and social conditions under which information changes behavior. Information may fail to travel when local leaders have incentives to suppress it, may be strategically distorted by political actors before reaching voters, or may be filtered through durable cultural norms and individual preferences. Together, the dissertation shows that information is not merely a technical input into decision-making; its effects depend on who controls it, how it is delivered, and the social environment in which people interpret and act on it.

Chapter 1 studies whether information can catalyze community collective action against environmental harms from artisanal and small-scale gold mining in Ghana. In many mining-affected communities, residents are exposed to severe long-term health risks but remain poorly informed about the connection between mining pollution and chronic conditions. I evaluate a cluster randomized intervention that delivered a documentary about these harms either privately to traditional chiefs or publicly to both chiefs and community members. The results show that private leader-mediated information raises chiefs'

knowledge but does not reliably generate community learning or collective action, especially when chiefs have conflicts of interest. Public dissemination, by contrast, increases community learning and engagement even when leaders are conflicted. However, translating information into regulatory consensus still depends on local leadership incentives. The chapter shows that information cannot reliably mobilize collective action even when delivery channels build broad-based awareness and create public accountability; collective action emerges only when leaders' incentives are also aligned.

Chapter 2 examines political manipulation of information disclosure in India's rural road quality audit system. Public audits are intended to inform voters about the quality of government-provided infrastructure, but the chapter shows that state-level audits can themselves become tools of political economy. Using administrative data from Bihar's rural road program around the 2015 state election, I compare quality reporting in constituencies aligned and not aligned with the state ruling party. The evidence is consistent with strategic manipulation: state audit reports become differentially favorable to aligned constituencies before elections, while national audits, which are more insulated from state political influence, do not display the same pattern. Additional evidence points to manipulation through auditor assignment and inspection timing. The chapter highlights a core tension in information-based accountability: disclosure systems can improve electoral accountability only when the institutions producing information are themselves protected from political capture.

Chapter 3 asks whether culture and individual preferences travel with people and predict public-health behavior in a shared environment. Using an original survey of international students at UC Berkeley during the early COVID-19 pandemic, I study precautionary behaviors among students from different origin countries who faced a common university and local policy context. The chapter links individual-level measures of preferences, including individualism, risk-taking, reciprocity, and patience, to precautionary behavior, while also matching respondents to origin-country cultural norms. The results show that both respondent-level preferences and origin-country norms predict precautionary behavior. More individualistic and risk-tolerant respondents report less precaution, while more positively reciprocal and patient respondents report more. Origin-country norms also retain predictive power after accounting for matched individual preferences, suggesting that cultural background may matter through broader channels than measured individual traits alone. The chapter contributes to understanding how social preferences and cultural environments shape behavior during collective-risk crises.

# Acknowledgments

I owe my deepest professional gratitude to my Berkeley mentors and colleagues. Francis Annan and Ernesto Dal Bó, my dissertation co-chairs, shaped this dissertation and my development as an economist with extraordinary generosity, patience, and intellectual care. I had the unexpected good fortune of becoming Francis's first Ph.D. student after he joined Berkeley ARE in my fourth year. What began as a research assistantship in Ghana, where I saw textbook development economics unfold firsthand, became one of the central mentorships of my graduate training. Francis has been extraordinarily generous with both his time and his care: in hours-long discussions, in early-morning and late-night conversations, in Berkeley and in Ghana, and in his steady responsiveness whenever I needed guidance. It has been a blessing in disguise, and a pleasure, to be his student. My first encounter with Ernesto Dal Bó was during Berkeley's Ph.D. admissions visit day in March 2019. As an aspiring political and development economist, I already knew his work before arriving. Since then, I have somehow become even more amazed by his wisdom every time he speaks, which, fortunately for his students, is quite often. More seriously, Ernesto never looked down on me during periods when I felt uncertain or directionless, especially after a difficult COVID year. He consistently encouraged me to persevere, remain ambitious, and keep searching for meaningful questions. I am also grateful to Frederico Finan and Aprajit Mahajan, my dissertation committee members: to Fred for his direct and clarifying guidance, and to Aprajit for pushing me toward careful econometric work and sharper empirical thinking.

I am especially indebted to Elisabeth "Betty" Sadoulet, my oral examination chair, who passed away in 2025. This dissertation may fall short of the standard of care, curiosity, and seriousness that she brought to research, but I hope it carries forward something of what she taught me. Betty read everything I sent her, from early proposals and survey instruments to drafts and write-ups, and always returned them with detailed, generous, and exacting feedback. She modeled a way of doing development economics that was intellectually demanding without losing sight of people, context, and humility. Her example continues to shape how I think about research.

Across the individual chapters, I benefited from comments and conversations with Shifrah Aron-Dine, Sam Asher, Matilde Bombardini, Lucas Davis, Alain de Janvry, Benjamin Faber,

Marco Gonzalez-Navarro, Kelsey Jack, Ethan Ligon, Matt Lowe, Jeremy Magruder, Edward Miguel, Dilip Mookherjee, Ajay Shenoy, Francesco Trebbi, Pedro Vicente, Brian Wright, and Guo Xu, as well as participants in the Berkeley Development Lunch, ARE Development Workshop, and BPER Lab. I am especially thankful to my cohortmates and friends Becky Cardinali, Abdoulaye Cisse, Yuen Ho, Jacob Lefler, Leïla Njee Bugha, Elif Tasar, Arthur Wardle, and Shuo Yu. In particular, Abdou and Yuen continued to offer generous feedback, encouragement, and intellectual support even after they had graduated a year earlier and were thousands of miles away. I am also grateful for job-market support and encouragement from Maximilian Auffhammer, Sofia Berto Villas-Boas, Tamma Carleton, Jeffrey Perloff, Joseph Shapiro, Aaron Smith, and David Zilberman. Carmen Karahalios and Diana Lazo provided essential administrative support and steady kindness throughout the program.

The fieldwork and institutional support behind this dissertation involved many people beyond Berkeley. For "Information and Collective Will Against Environmental Harms: Experimental Evidence from Ghana's *Galamsey*," I am indebted to Michael Kodom and Peter Quartey at the Institute of Statistical, Social and Economic Research at the University of Ghana for implementation support; to Dela Kwasi Dzudzor for excellent research assistance; to Erastus Asare Donkor at JoyNews for permitting the use of his documentary, *Poisoned for Gold*; to Bright Kwaku Bobby for local field support; and to Father Joseph Kwame Blay and Elizabeth Allua Vaah at the Ghana Environmental Advocacy Group for guidance on culturally sensitive data collection. I owe special gratitude to Zixu Chen, whose companionship carried me through the final and most exhausting leg of this project. This research was supported by the J-PAL Governance Initiative, the Weiss Fund, the International Growth Centre, the Center for Effective Global Action, the Network for a New Political Economy, the Andrew and Mary Thompson Rocca Fellowship, and the Lau Graduate Fellowship in Climate Equity. For "Political Economy of Information Disclosure: Evidence from Indian Rural Road Quality Audits," I thank Dilip Mookherjee, whose early guidance helped shape the idea and pushed me to think rigorously, for the first time, about the connection between theory and empirics. For "Do Culture and Preferences Travel? Precautionary Behavior During COVID-19," I thank the Berkeley International Office for support with survey recruitment and administrative enrollment benchmarks.

Finally, I thank the teachers, mentors, friends, and family members who sustained me through the long arc of this work. I am grateful to Betty Poon, my high school economics teacher, and to Kaivan Munshi for introducing me to development economics. I also thank Xiaobo Zhang and Albert Park for offering me research assistant opportunities despite my limited research background at the time, and Dilip Mookherjee for taking me under his wing at Boston University and helping prepare me to become an independent researcher during the year before Berkeley. Adrian Law deserves special mention for remaining close during the isolation and uncertainty of COVID. Most of all, I thank my family for giving me the foundation from which all of this became possible, and Lucia Choi, my partner and soon-to-be wife. Long

distance is often measured in miles and time zones, but Lucia consistently found ways to make the distance feel smaller. Her unwavering support, patience, and relentless effort to build a shared life across continents carried me through the hardest stretches of graduate school. I feel more grounded when she is near, and no less supported when she is far away. This dissertation is better because of the people who taught me, argued with me, protected my confidence when it was thin, and stayed close.

# Contents

**1 Information and Collective Will Against Environmental Harms: Experimental Evidence from Ghana's *Galamsey* 1**
1.1 Introduction . . . 1
1.2 Setting and Intervention . . . 7
1.2.1 Artisanal and Small-Scale Gold Mining . . . 7
1.2.2 Study Community Selection and Household Sampling . . . 8
1.2.3 Key Feature 1: Imperfect Information — Limited Awareness of Long-Term Health Risks . . . 9
1.2.4 Key Feature 2: Traditional Local Chiefs in *Galamsey* . . . 10
1.2.5 Intervention: The *Poisoned for Gold* Documentary . . . 13
1.3 Research Design . . . 14
1.4 Data . . . 15
1.4.1 Town Hall – Structured Community Activities . . . 15
1.4.2 Household Surveys . . . 17
1.4.3 Chief Surveys . . . 18
1.5 Results . . . 18
1.5.1 Empirical Specification . . . 18
1.5.2 Leaders as Information Gatekeepers . . . 19
1.5.3 Community Learning . . . 20
1.5.4 Collective Action I: Community Consensus On Mining Regulation . . . 21
1.5.5 Collective Action II: Sustained Community Engagement . . . 24
1.5.6 Accountability and Perceptions of Leadership . . . 25
1.5.7 Pessimism About Collective Action . . . 26
1.5.8 Additional Results and Robustness Checks . . . 27
1.6 Conclusion . . . 27
Appendix . . . 42
A Community Sampling . . . 42
B Validation of the Randomized Response Measure . . . 42
C Town Hall Implementation Details . . . 42

**2 Political Economy of Information Disclosure: Evidence from Indian Rural Road Quality Audits** **65**
2.1 Introduction . . . 65
2.2 Institutional Background . . . 68
2.2.1 PMGSY . . . 68
2.2.2 Roles of MLAs in PMGSY . . . 69
2.3 Data . . . 70
2.3.1 PMGSY Audit Data . . . 70
2.3.2 Bihar State Assembly Election Data . . . 71
2.3.3 Probabilistic Matching . . . 71
2.4 Empirical Design . . . 72
2.4.1 Potential Outcome Framework in the Absence of Pure Controls . . . 72
2.4.2 Empirical Specification . . . 73
2.5 Secondary Findings . . . 73
2.5.1 No Differential Extensive Margin . . . 73
2.5.2 No Differential Targeting on Observables . . . 74
2.6 Main Results: Differential Quality Reporting . . . 74
2.6.1 Mechanism #1: Auditor Assignment . . . 75
2.6.2 Mechanism #2: Manipulation of Audit Timing . . . 75
2.6.3 Robustness: Parallel Trend After Election . . . 76
2.6.4 Heterogeneity by Political Competitiveness . . . 76
2.7 Voter Responses . . . 77
2.8 Conclusions . . . 77
**3 Do Culture and Preferences Travel? Precautionary Behavior During COVID-19** **90**
3.1 Introduction . . . 90
3.2 Data and Summary Statistics . . . 94
3.2.1 Precautionary Behaviors . . . 95
3.2.2 Individual Preferences . . . 95
3.2.3 Country Norms . . . 96
3.2.4 Individual Controls and Auxiliary Variables . . . 97
3.2.5 Country-Level Controls . . . 97
3.3 Estimation Strategy . . . 97
3.4 Results . . . 99
3.4.1 Main Preference and Norm Results . . . 99
3.4.2 Auxiliary Belief and Information Evidence . . . 102
3.5 Conclusion . . . 102
Appendix . . . 109
D Exact Wording for Individual Preference Measures . . . 109
**References** **118**

# List of Figures


| | | |
|---|---|---|
| 1.1 | Baseline Household Awareness of Health Conditions Related to *Galamsey* | 29 |
| 1.2 | Baseline Chief Awareness of Health Conditions Related to *Galamsey* | 30 |
| 1.3 | Distribution of Chiefs' Conflicts of Interest Across Communities | 31 |
| 1.4 | Relationship Between Chiefs' Conflicts of Interest and Mining Prevalence | 32 |
| 1.5 | Study Flow and Experimental Design | 33 |
| 1.6 | Overview of Data Collection and Measurement Points | 34 |
| 1.7 | Average Treatment Effects on Chiefs' Learning | 35 |
| 1.8 | Average Treatment Effects on Community Members' Learning | 36 |
| 1.9 | Heterogeneous Treatment Effects on Sustained Community Engagement | 37 |
| 1.10 | Heterogeneous Treatment Effects on Accountability Norms Toward Local Chiefs | 38 |
| 1.A1 | Distribution of the Share of Mining Households Across Communities | 46 |
| 1.A2 | Jema Template — Example of a Customary Bylaw on Small-Scale Mining Regulation | 47 |
| 1.A3 | Sample First Ballot on Community Mining Regulation | 48 |
| 1.A4 | Sample Second Ballot on Commitment to Sustain Collective Action | 49 |
| 1.A5 | Treatment effects on Health Learning by Town Hall Attendance | 50 |
| 1.A6 | Average Treatment Effects on Chiefs' Reasons to Take Part in Collective Action | 51 |
| 2.1 | State Audit Results Before and After the 2015 Bihar Election | 79 |
| 2.2 | Share of Roads Inspected Before the Election by Inspection Outcome | 80 |
| 2.3 | Event Study of Differential State Audit Reporting | 81 |
| 3.1 | Precautionary Behavior and Matched Traits | 104 |
| 3.A1 | Precautionary Behavior and Country-Only Cultural Dimensions | 114 |
| 3.A2 | Precautionary Behavior and Matched Traits in Pooled Specifications | 116 |

# List of Tables


1.1 Heterogeneous Treatment Effects on Chiefs' Attendance and Dissemination Efforts . . . 39
1.2 Heterogeneous Treatment Effects on Community Mining Regulation Preferences . . . 40
1.3 Average Treatment Effects on Pessimism About Collective Action . . . 41
1.A1 Heterogeneous Treatment Effects on Community Members' Learning . . . 52
1.A2 Voter Turnout at Town Halls by Treatment and Leader Type . . . 53
1.A3 Heterogeneous Treatment Effects on Community Mining Bylaw Measures 54
1.A4 Heterogeneous Treatment Effects on Chiefs' Persuasive Efforts . . . 55
1.A5 Heterogeneous Treatment Effects on Chiefs' Mobilization Efforts . . . 56
1.A6 Effects on Community Mining Regulation Preferences (Split-Sample by Occupation) . . . 57
1.A7 Effects on Community Mining Regulation Preferences (Weighted) . . . 58
1.A8 Effects on Community Mining Regulation Preferences (Robust to Misclassification) . . . 59
1.A9 Heterogeneous Treatment Effects on Commitment to Sustain Collective Action . . . 60
1.A10 Heterogeneous Treatment Effects on Post-Town Hall Discussion . . . 61
1.A11 Heterogeneous Treatment Effects on Pessimism About Collective Action . 62
1.A12 Heterogeneous Treatment Effects on Attitudes Toward Other Community Members . . . 63
1.A13 Heterogeneous Treatment Effects on National Civic Action . . . 64

2.1 Summary Statistics Around Election: State Quality Audits . . . 82
2.2 Audit Frequency Around Election . . . 83
2.3 State Audit Targeting Around Election . . . 84
2.4 Audit Outcomes Around Election . . . 85
2.5 State Audit Manipulation Mechanism . . . 86
2.6 State Audit Outcome Heterogeneity . . . 87
2.7 Voter Response to Inspection Outcome . . . 87

2.A1 Summary Statistics Around Election: National Quality Audits . . . . . 88
2.A2 National Audit Targeting Around Election . . . . . 89

3.1 Summary Statistics for the Stayed-in-Berkeley Sample . . . . . 106
3.2 Representativeness of the Berkeley International Students Survey Sample 107
3.3 Selection into Staying in Berkeley . . . . . 108
3.A1 Survey Wording: Social Orientation Preferences . . . . . 109
3.A2 Survey Wording: Preferences and Uncertainty . . . . . 111
3.A3 Auxiliary Belief and Information Evidence: Individual Preferences . . . . 112
3.A4 Auxiliary Belief and Information Evidence: Country Norms . . . . . 113

# Chapter 1

# Information and Collective Will Against Environmental Harms: Experimental Evidence from Ghana's *Galamsey*

## 1.1 Introduction

Resource extraction is among the largest human interventions in natural systems. While the benefits are immediate and concentrated among resource system users, the degradation accumulates silently across atmospheres, watersheds, food chains, and generations. These negative externalities represent a canonical violation of the First Welfare Theorem, a case where decentralized market transactions among atomistic agents fail to yield an efficient equilibrium. Addressing such market failures requires environmental regulation. While this function is conventionally assigned to the state, state capacity is often limited in the rural areas of developing regions where resource extraction is concentrated. This institutional vacuum highlights the importance of understanding community collective action, in which local leaders and community members jointly define and enforce rules for managing natural resources.

This paper studies community collective action in the artisanal and small-scale gold mining sector (ASGM), which employs 14-19 million miners across more than 80 countries. ASGM is the largest source of mercury pollution worldwide, responsible for 38 percent of annual human-caused emissions. The problem is especially severe in developing regions: in South America and Sub-Saharan Africa, ASGM accounts for over 80 percent of mercury emissions

(UNEP, 2019). In Ghana, where this study takes place, ASGM is known as *galamsey*.[1] In response to the environmental and health harms *galamsey* causes, the Government of Ghana has launched repeated crackdowns, including multiple waves of military operations and a two-year national ban on small-scale mining, yet these efforts have not produced sustained reductions in environmental damage.

Collective action to mitigate environmental harms from *galamsey* has emerged in some communities. In the farming community of Jema, located in the Western North Region, traditional leaders and community members established a community mining regulation, in the form of a customary bylaw, that banned mining and introduced regular assembly meetings to monitor compliance and coordinate responses. This local initiative has proven effective: Jema successfully repelled multiple attempts by *galamseyers* to begin mining within its territory and remains free of *galamsey*. As the only unaffected community in its district, the case of Jema is widely cited as a success story.

Yet such successes remain rare. What prevents community collective action from taking shape in more communities? A natural starting point is economic incentives. Mining provides an important source of livelihood in many *galamsey* communities. In my baseline survey, 36 percent of randomly sampled households report having at least one member engaged in mining. In settings with limited alternative employment opportunities and high marginal utility of consumption, communities may therefore be unwilling to restrict an activity that underpins household income, even when it generates substantial local externalities. This explanation alone, which implicitly assumes that communities accurately weigh mining income against its environmental and health consequences when forming collective choices, is likely incomplete, particularly in low-information developing-country settings. Even when an environmental hazard imposes severe and persistent harms that are directly experienced, individuals are often deeply misinformed about their causes and long-term consequences.

This paper therefore turns attention to a precondition for collective action: information about the severity of the problem.[2] The health risks associated with mercury exposure are often long-term, diffuse, and difficult to attribute directly to mining activities. Field interviews indicate that community members often misattribute such chronic health conditions to spiritual or supernatural causes. Baseline survey data quantify this information gap: fewer

[1] *Galamsey* derives from the phrase "gather and sell" in the local Twi language. Because more than 80 percent of ASGM operations are unlicensed, and even licensed ones often violate environmental regulations, the term *galamsey* is also commonly used in media and public discourse to refer to illegal mining.

[2] Classic explanations for weak collective action point to free-riding, whereby individuals can benefit from others' efforts without contributing themselves (Olson, 1965), and coordination failures arising from uncertainty about others' actions (Schelling, 1960). While these forces are undoubtedly relevant, they implicitly assume that individuals hold accurate beliefs about the underlying harm. My focus is on information because it is the margin of experimental variation in this study, though the research design and findings also speak to these classic mechanisms.

than 40 percent of households linked *galamsey* to long-term conditions such as miscarriage, birth defects, kidney or liver disease, or neurological disorders, significantly lower than the 75 percent recognition rate of immediate and visible risks such as malaria or mercury use. This internality — manifested as the underestimation of gradual, cumulative health harms — can dampen motivation for collective action, as individuals may not perceive its urgency or necessity. These facts motivate my focus on information dissemination as a potential catalyst for collective action.

In our setting, a key choice to turn information into learning and action is the delivery channel: whether it is seeded through local leaders via leader-mediated delivery, or broadcast directly to the public through community-wide dissemination. As influential local leaders with deep contextual understanding, traditional chiefs could play a crucial mediating role, legitimizing new information and translating technical concepts into language that resonates with their communities. On the other hand, the very authority that positions chiefs as effective intermediaries also gives them opportunities to extract private gains from the resource system that produces environmental harm. This conflict of interest may leave them reluctant to spread information that goes against their vested interest. These trade-offs give rise to three research questions. First, in the absence of leader conflicts of interest, does seeding information through leaders generate community learning and spur collective action? Second, does broadcasting information publicly overcome the barriers posed by conflicted leadership? Or, instead, does broadcasting induce learning and collective action only when paired with non-conflicted leadership?

I address these questions through a cluster randomized controlled trial in 99 ASGM-affected communities in Ghana. The intervention delivered a professionally produced documentary film, *Poisoned for Gold*, created by a prominent Ghanaian journalist independent of this study. The film directly targets the information gap documented above: scientists, physicians, and affected residents bring to light the cumulative health consequences of long-term heavy-metal exposure, explicitly linking them to the very chronic medical conditions that community members had failed to associate with *galamsey* in the baseline survey. Communities were randomly assigned to one of three groups: a control group with no screening, seeding, in which the documentary was screened privately to traditional chiefs, or broadcasting, in which the same documentary was screened publicly to both chiefs and community members. I pre-registered the stratified treatment randomization by the baseline share of mining households in each community and, importantly, whether the chief has a conflict of interest. Conflicts of interest were measured using household perceptions elicited through a randomized response module and were validated via correlation with chiefs' self-reports to carefully worded questions about their rent-seeking behavior.

To trace the causal chain from information exposure and chiefs' efforts in information dissemination and mobilization to learning and collective action, I collected data at multiple

levels. On the same day the information was delivered, field officers convened structured town halls in all 99 communities. These gatherings, locally known as durbars, are a customary setting for community deliberation and decision-making, making the meeting itself a familiar event for community members. The study measures outcomes both during and after the town hall events. During the town halls, field officers documented several outcomes. First, to assess chiefs' efforts in information dissemination and mobilization, field officers documented their speeches and observable actions.[3] Second, enumerators separately measured learning and actual follow-up discussion by surveying leaders and community members again three weeks after the town hall. Third, as illustrated in the Jema example, collective action in a traditional communal setting typically culminates in community-level decisions that are socially binding for all members. Although such actions are formally enacted by the chief on behalf of the community, they are grounded in a broader participatory process in which community members deliberate and express their collective will in durbars, exert bottom-up pressure on local leadership to uphold that will, and help self-enforce collective decisions through sustained engagement.[4] To capture key dimensions of this collective action process, I embedded several familiar local activities within our town halls. First, I measured collective will through a secret-ballot vote on prospective customary mining bylaws, allowing participants to choose among options of varying regulatory stringency. Second, I measured community engagement, the behavioral dimension of collective action, through volunteer sign-ups for regular follow-up meetings, enrollment in an NGO-moderated WhatsApp group to capture reliance on external coordination channels, and a group bidding exercise to elicit financial contributions toward hosting an NGO workshop. These approaches mirror practices common to successful grassroots initiatives in the region, which rely on durbars to deliberate, build consensus, and mobilize participation. Importantly, the activities had real-world consequences: bylaws voted on could be carried forward for adoption, and volunteer lists served as the foundation for organizing subsequent community meetings.

Turning to the experimental results, I find three key patterns. First, the documentary produced strong learning effects among chiefs in both the seeding and broadcasting groups. Using a standardized index of health knowledge on long-term health conditions *galamsey* can cause, chiefs in both seeding and broadcasting communities scored 0.48–0.57 standard deviations higher than those in control communities.

Second, although chiefs' own learning was similar across arms, the effects on communities differed and depended on whether chiefs had conflicts of interest. Under seeding, 95 percent of

[3]In cases where they excused themselves from the town hall and sent a delegate from the traditional council to represent them, field officers record the speeches and actions of the delegate instead.

[4]This collective process constrains the chief's authority and ensures accountability, as legitimacy depends on community consent and participation. Chiefs who neglect communal welfare or violate their duties risk destoolment — the traditional mechanism for removing a chief from office — which serves as a powerful check aligning leadership with the community's collective interests (Arhin, 1985; Baldwin, 2015).

non-conflicted chiefs made an effort to disseminate the information, whereas conflicted chiefs were 39 percentage points less likely to do so. Despite these dissemination efforts, neither type of leader generated improvements in community members' health-risk learning or detectable shifts in bylaw preferences. I also find no effects on follow-up engagement under conflicted leadership; the impact stops at the chief's palace, where the private screening took place. In communities with non-conflicted leadership, participation in follow-up engagement was significantly lower despite greater mobilization efforts by chiefs during the town hall, suggestive of a crowding-out effect where unilateral leader action dampens individual initiative.

Third, when the information was shared publicly with both chiefs and community members (broadcasting), conflicted chiefs were 41 percentage points less likely to attend the town hall themselves, instead sending a delegate — a 60 percent decrease relative to the 68 percent chief attendance rate in control communities. No such pattern is observed among conflicted chiefs in the seeding treatment, consistent with conflicted leaders seeking to avoid accountability pressures created by public dissemination. Despite this, exposure to the public documentary screening increased community learning by 0.50 standard deviations relative to control communities, regardless of leader type. Learning effects under broadcasting are statistically indistinguishable between town hall attendees and non-attendees, pointing to substantial horizontal information diffusion among community members.[5]

Turning to collective action, under non-conflicted leadership, broadcasting shifted community preferences toward stricter mining bylaws: the vote share for one of the most stringent bylaw options, namely banning mining in forest reserves and river bodies as well as amalgam burning near homes, increased by 21.3 percentage points. On follow-up engagements, sign-ups for regular community meetings increased by 0.40 standard deviations. In line with this pattern, the endline survey three weeks later shows more frequent discussions about *galamsey* among community members.

Under conflicted leadership, broadcasting generated comparable levels of follow-up engagement but polarized mining bylaw preferences without a level shift: the dispersion of bylaw preferences increased by about 19 percent relative to control communities. This increase in dispersion is driven by movements at both extremes of the preference distribution. At one end, the vote share for the least stringent option, maintaining the status quo, increased by 4.5 percentage points (a 79 percent increase relative to control communities with conflicted chiefs). This shift is explained by a lack of buy-in from conflicted chiefs, as evidenced by two behaviors: their absence from town hall meetings and their active attempts to counteract the broadcasting effect. Specifically, conflicted chiefs (or their delegates) were 40.7 percentage points more likely than conflicted chiefs in control communities to advocate for maintaining the status quo. At the other end of the preference distribution, the share of community

[5]This stands in contrast to the lack of community learning under seeding, which suggests weak vertical transmission through leaders.

members supporting the most stringent option, banning all mining activities, increased by 12.6 percentage points (a 58 percent increase relative to control communities with conflicted chiefs). This shift is consistent with strong learning among community members from direct exposure to the documentary.

Although broadcasting did not generate consensus on mining bylaws under conflicted leadership, it strengthened community demand for accountability from chiefs. In the endline survey, community members were 42 percent less likely to express unconditional respect for their chiefs and 56 percent more likely to state that they should both respect and question their role—echoing the lower town hall attendance of conflicted chiefs during public screenings. I also find that when local leadership is compromised, an informed public redirects collective action efforts toward external channels that substitute for local initiatives, as reflected in a 39 percent increase in enrollment in the NGO-moderated WhatsApp group.

Together, these results suggest that the conditions for collective action depend on the margin of interest. Building consensus over local mining regulation requires both an informed public and non-conflicted leadership; neither alone is sufficient. By contrast, sustaining community engagement—such as participation in meetings and discussion—can be achieved through broad public information alone, even when local leadership is conflicted. Despite the overall gains in engagement and accountability, broadcasting also fostered pessimism: respondents reported lower perceived efficacy of collective action and reduced willingness, both personally and in their beliefs about others, to engage in collective community efforts. In the conclusion, I discuss the external validity and policy implications of both the key takeaway and the seemingly paradoxical pessimism result.

This paper contributes to three strands of literature. First, it adds to the research on how information influences responses to environmental problems. A large empirical literature has causally identified the effects of information on private adaptation to environmental hazards (Madajewicz et al., 2007; Bennear et al., 2013; Barwick et al., 2024). In contrast, studies examining information and collective action are mostly conceptual or descriptive (Ostrom, 1990; Baland and Platteau, 2000), and the few well-identified studies rely primarily on lab experiments. This paper moves beyond private adaptation at the individual level and provides field-experimental causal evidence on how information affects collective will and community engagement to address environmental harms at the community level.

Second, I contribute to the political economy of leadership. This paper adds to studies of leadership in environmental contexts (Jack and Recalde, 2015; Kosfeld and Rustagi, 2015). A closely related paper, Armand et al. (2020), shows in the context of Mozambique's natural gas discovery that when information reaches citizens, it can discipline leaders' rent-seeking and reduce violence, whereas information confined to elites can exacerbate capture. This paper extends this evidence by explicitly measuring and stratifying treatments by leaders' conflicts of interest, an underexplored source of heterogeneity. Through the lens of our

information intervention and a detailed data collection design that captures both the words and actions of leaders, I demonstrate how conflicts of interest shape their responses to interventions. These findings have implications for future programs that engage local leaders — not only for information interventions, but for development strategies more broadly. This contribution is particularly timely given the growing interest in involving traditional chiefs as agents of development (Acemoglu et al., 2014; Henn, 2023; Balán et al., 2022; Basurto et al., 2020). I add to this expanding literature by examining chiefs' roles in collective responses to environmental harms, a policy domain where they play a particularly prominent yet understudied role.

Third, I extend research on video-based information interventions. Most existing studies evaluate impacts on individual outcomes: knowledge, attitudes, and private behavior (Hussam et al., 2023; Banerjee et al., 2019; Riley, 2021). Recent work shows that documentaries can also foster cooperation across group lines (Siddique et al., 2024). Natural experiments from the U.S. and China suggest that environmental documentaries can shift demand for environmental quality (Jacobsen, 2011; Tu et al., 2020). In terms of information content, a survey experiment shows that videos focusing solely on climate impacts fail to raise support for climate policy, while those that also clarify policy mechanisms succeed in doing so (Dechezleprêtre et al., 2025). I build on this literature by embedding a documentary within a field experiment that pairs information about long-term environmental and health harms with a concrete, locally relevant template for action. Specifically, the documentary screening is integrated into a structured town hall that introduces a successful community bylaw as a salient example of feasible regulation and invites participants to deliberate and vote on local mining rules. This design demonstrates the potential of video-based information interventions not only to convey complex, long-term risks that are otherwise difficult to communicate, but also to test when such information, combined with a clear mechanism for action, translates into community collective action.

## 1.2 Setting and Intervention

### 1.2.1 Artisanal and Small-Scale Gold Mining

Artisanal and small-scale gold mining (ASGM) provides livelihoods for an estimated 14–19 million miners in more than 60 countries, indirectly supporting about 100 million people worldwide (UNEP, 2019). Unlike large-scale industrial mining — which is capital intensive, mechanized, and often foreign owned, ASGM is labor intensive, technologically rudimentary, and deeply embedded in rural community economies (Hilson, 2002).

Gold extraction methods in ASGM also differ fundamentally from those in industrial operations. Rather than using cyanidation or gravity separation, small-scale miners apply mercury to amalgamate gold from ore. The amalgam is then heated to vaporize the mercury and

recover purified gold, releasing toxic vapors and leaving residual mercury in water and soil (UNEP, 2019). The Minamata Convention on Mercury obliges countries with significant ASGM activity to reduce, and where feasible, eliminate mercury use through National Action Plans (UNEP, 2013). Despite these commitments, mercury use remains widespread because it is inexpensive, quick to apply, and requires little capital or technical training (UNEP, 2012). As a result, it is widely recognized that eliminating mercury use in ASGM in the short term is infeasible. However, the most severe mercury exposure from amalgam burning is highly localized: within a few kilometers of emission sites, mercury concentrations can reach levels two to four orders of magnitude above background (Szponar et al., 2025). For this reason, policy recommendations emphasize restricting amalgam burning in residential areas, enforcing minimum setback distances of at least 500 meters from homes (UNEP, 2012).

Globally, artisanal and small-scale gold mining (ASGM) is the largest source of anthropogenic mercury emissions, responsible for 38 percent of total human-caused releases. In South America and Sub-Saharan Africa, its share exceeds 80 percent (UNEP, 2019). Chronic exposure to mercury causes severe, long-term health consequences, including neurological and kidney damage, cardiovascular disease, and developmental disorders in fetuses and infants (Esdaile and Chalker, 2018).

In Ghana, gold deposits are concentrated in the Ashanti, Central, Eastern, Western, and Western North Regions. Historically, mining in these areas was conducted manually along riverbeds. Beginning around 2006, however, the introduction of heavy machinery by Chinese miners—including excavators, bulldozers, and floating dredges known locally as *changfa*—dramatically expanded gold production while accelerating environmental degradation (Crawford et al., 2015). These mechanized operations have polluted rivers, stripped forests, and contaminated soils through improper tailings disposal. Because mercury is typically applied in fixed proportion to gold output, increases in production have mechanically translated into higher mercury use and emissions.

ASGM in Ghana is deeply local in both organization and impact. Mining operations are embedded in customary land tenure systems and community networks that regulate access, labor, and dispute resolution (Hilson and Maconachie, 2017). The resulting environmental and health harms, such as reduced farm yields and fish stocks, river siltation, and mercury exposure, are borne primarily by those same local communities. This close alignment between cause and consequence highlights why community collective action is essential: where state enforcement capacity is limited, sustainable resource management depends on local coordination and leadership to internalize environmental costs and mitigate damage.

### 1.2.2 Study Community Selection and Household Sampling

The study was conducted in six districts spanning the Ankobra River Basin, which runs through Ghana's Western and Western North Regions. The Ankobra is among the country's

most mercury-contaminated river systems, reflecting the high intensity of artisanal and small-scale gold mining (ASGM) activity in its watershed. To identify *galamsey* communities, research assistants conducted phone interviews with local politicians to identify rural communities that (i) have ongoing ASGM operations or (ii) face imminent pressure from nearby mining expansion. Applying population eligibility criteria between 400 and 3,000 residents and excluding communities too close to others, I selected 99 communities as my study sites. Appendix A provides additional details on the community selection process, including screening criteria and geographic mapping.

I conducted a baseline survey with 16 households per community in February 2025. Enumerators used a standardized random walk method with a pre-specified skip pattern to ensure spatial coverage. In each selected household, enumerators interviewed the member most knowledgeable about income, expenditure, and household decision-making. Eligible respondents were current residents who intended to remain in the community for at least six months following the baseline. Each household received a 20 GHS airtime top-up (approximately 1.3 USD) as compensation for completing the interview. Baseline data indicate that approximately 36 percent of households had at least one member engaged in small-scale mining, with substantial heterogeneity across communities.[6] In parallel, I conducted a baseline survey with each community's traditional chief.

The baseline surveys establish two defining features of the study setting: (i) limited public awareness of the long-term health consequences of mercury exposure, and (ii) the central, often decisive, role of traditional chiefs in shaping community responses to *galamsey*. These contextual features motivate the subsequent information intervention design.

### 1.2.3 Key Feature 1: Imperfect Information — Limited Awareness of Long-Term Health Risks

To measure awareness of health risks, the baseline household survey listed a series of health conditions and asked respondents whether they believed each was related to *galamsey*.[7] Figure 1.1 summarizes the share of households associating each health condition with *galamsey*. The pattern reveals that communities possess imperfect information in a specific form. Most households recognize the immediate and visible health risks of mining. For example, 74 percent link *galamsey* to malaria, reflecting awareness that shallow pits left by mining operations collect stagnant water that breeds mosquitoes, the malaria vector. Respondents are also

[6] Appendix Figure 1.A1 shows that the interquartile range spans from 13 percent at the 25th percentile to 56 percent at the 75th percentile.

[7] Piloting indicated that asking directly whether *galamsey causes* a health condition led to confusion between the extensive margin (whether it can possibly lead to it) and the intensive margin (the magnitude of its effects). I therefore used a less causal framing: whether each condition is "related to" *galamsey*, which sets a lower bar for identifying awareness.

aware that mercury is used in gold extraction and that it is poisonous (75 percent) and accumulates in the food chain (78 percent).

In contrast, far fewer households associate *galamsey* with long-term health conditions. Fewer than half link mining to miscarriage, birth defects, kidney or liver disease, neurological disorders, or respiratory conditions. This pattern suggests that imperfect information takes the form of under-recognition of gradual, cumulative harms rather than general ignorance. Two placebo conditions, diabetes and HIV/AIDS, serve as checks on indiscriminate attribution. Their low association rates reinforce that the observed pattern reflects genuine informational gaps about mercury-related illnesses rather than indiscriminate attribution of diseases to mining.

Field interviews suggest this unawareness often stems from *misattribution*: individuals often attribute long-term health conditions to spiritual causes or unrelated illnesses rather than prolonged mercury exposure. Figure 1.2 presents chiefs' responses to the same set of health knowledge questions, showing a similar pattern of limited awareness of long-term health risks, albeit with less precise estimates due to the smaller sample size.

While some of these conditions may be low-probability events, the lack of understanding matters because it diminishes perceived urgency and weakens motivation to act. When health risks are invisible, delayed, and poorly understood, community leaders and members have little incentive to coordinate preventive or regulatory action. Whether improving information about these long-term harms can shift beliefs and stimulate collective responses is therefore an empirical question that motivates our information intervention.

### 1.2.4 Key Feature 2: Traditional Local Chiefs in *Galamsey*

The second defining feature of the study context is the pivotal role of traditional local chiefs in community governance and land management.[8] Understanding chiefs' authority and incentives is therefore central to determining how information can be delivered effectively.

Traditional chiefs in Ghana are legally recognized as the custodians of communal lands, known as stool or skin lands, which constitute about 80 percent of the country's land area (Kasanga and Kotey, 2001; Goldstein and Udry, 2008; Ubink and Amanor, 2008). Under Article 267 of the 1992 Constitution, chiefs hold these lands in trust for their communities: they cannot sell them for personal gain but retain substantial discretion over their allocation and use. In contrast, the right to all minerals in their natural state is vested in the President of Ghana

[8]There is a vast literature in political science and a growing literature in political economy examining the roles of chiefs (Acemoglu et al., 2014). Mamdani (1996) and Baldwin (2015) provide influential qualitative and quantitative accounts, respectively, of how chieftaincy institutions interact with citizens and the state across Sub-Saharan Africa. I focus specifically on the role of chiefs in Ghana as it relates to mining land access and the coordination of community responses.

in trust for the people of Ghana under the Minerals and Mining Act, 2006 (Act 703). This means that even when gold or other minerals are found on stool land, ownership of those minerals legally resides with the state rather than the land custodians. To legally commence mining, individuals or entities must obtain a mining license granted by the Minister of Lands and Natural Resources through the Minerals Commission, but this process is often lengthy and administratively cumbersome (Hilson et al., 2022). In practice, many miners bypass formal licensing and seek approval directly from chiefs to operate on stool land, despite chiefs' lack of *de jure* authority to grant mineral rights. Our baseline survey reveals that 93 percent of mining households work in unlicensed sites, underscoring the gap between legal requirements and on-the-ground practice. In effect, chiefs often act as *de facto* gatekeepers for granting or withholding access to land for mining, shaping how local mineral resources are used even in the absence of formal legal authority.

This dual role, as both custodians of collective resources and potential rent seekers, places chiefs at the center of the *galamsey* problem. Their legitimacy and reach make them natural conduits for community information dissemination and coordination. Indeed, in my baseline survey, 48 percent of households identified the traditional chief as the actor best positioned to manage *galamsey*, compared to only 21 percent who named government agencies and 3 percent who mentioned security services.

At the same time, chiefs' control over land and mining access creates strong opportunities for private extraction. Some chiefs receive private payments from miners seeking to operate within their jurisdiction, generating potential conflicts of interest that can undermine collective action. Survey evidence underscores the prevalence of these incentives: 42 percent of chiefs report that they would accept knocking fees — a form of introductory payments made by miners to establish a working relationship with the chief — and 37 percent report receiving ongoing royalties in the form of a share of mining revenues.[9] To summarize, chiefs have the authority to regulate and mobilize collective responses against environmental harms, but the private income creates a disincentive.

To document these leadership incentives empirically, I collected data on chiefs' conflicts of interest three weeks prior to the information intervention. In each of the 99 study communities, I implemented a randomized response (RR) survey with 12 randomly selected community members.[10] The respondents were drawn independently from the baseline household sample. Because the topic is sensitive, the RR method provides a privacy-preserving way to estimate the community-level prevalence of perceptions that the chief would accept money to allow

[9]These questions were designed with NGO partner input for cultural sensitivity, but given the illicit nature of these mining activities, responses are likely subject to downward reporting bias; these figures should therefore be interpreted as lower bounds.

[10]Randomized response is a technique designed to elicit more truthful answers to sensitive questions by providing respondents with plausible deniability. In addition, to protect respondent anonymity, enumerators collected no personally identifying information (e.g., name, age, gender).

*galamsey*.[11] Respondents privately drew one card from a shuffled deck of four playing cards without revealing it to the enumerator. If the card was a heart, they were instructed to answer "Yes" regardless of their true belief; otherwise, they answered truthfully. Adjusting for the 25 percent forced "Yes" probability yields an unbiased estimate of the true affirmative share within each community.[12]

To validate that the RR measure captures meaningful variation in community members' perceptions, I also construct a complementary index based on carefully worded, directly asked questions measuring perceptions of chiefs' conflicts of interest in mining. The RR and directly asked measures are strongly correlated, with correlations of 0.41 at the individual level and 0.65 at the community level. Details on the direct-ask approach and validation results are provided in Appendix B.

Across all communities, the randomized response measure indicates that 52 percent (mean) of households believe their chief would take money to permit *galamsey*. Figure 1.3 plots the distribution of these estimates across communities. In the median community, 53 percent of households hold this view. I classify communities with estimated shares above the median as having *conflicted chiefs*, and those below as *non-conflicted chiefs*. The median conflicted community has an estimated share of 77 percent, compared to 27 percent among non-conflicted ones.

Figure 1.4 plots the relationship between the share of mining households and chiefs' conflicts of interest. Consistent with expectations, chiefs' conflicts of interest are positively correlated with the prevalence of small-scale mining in their communities: chiefs in more mining-intensive areas are more likely to be perceived as accepting payments to permit *galamsey*. However, the figure also reveals considerable variation in perceived chief conflicts with similar levels of mining activity. This heterogeneity underscores that leadership incentives are not mechanically determined by mining prevalence alone. It also mitigates a key identification concern that heterogeneous treatment effects might simply reflect differences in community mining intensity rather than in leaders' incentives. To address this potential confound, I incorporate both measures, chiefs' conflicts of interest and mining prevalence, into the experimental stratification procedure, ensuring balanced assignment across their combinations. All regression specifications control for baseline mining prevalence to isolate variation driven by leadership incentives.

To assess whether my randomized-response (RR) measure captures chiefs' actual vested interests, rather than merely community perceptions, I examine its correlation with chiefs' self-reported rent extraction, including whether they receive "knocking fees" or royalties from miners. Chiefs who reported accepting "knocking fees" were perceived by 13.5 percentage

[11] Enumerators were trained to ensure that respondents understood that the payments in question referred to private transfers to chiefs, rather than formal mining taxes that would be redistributed to the community.

[12] Formally, the true share equals $(p_{\text{obs}} - 0.25)/0.75$, where $p_{\text{obs}}$ is the observed proportion of "Yes" responses.

points more community members as willing to take money to allow *galamsey* ($p = 0.026$). In contrast, self-reported receipt of royalties is associated with a much smaller and statistically insignificant effect of 3.9 percentage points ($p = 0.53$). The RR association with knocking fees remains robust after controlling for the share of mining households, suggesting that community members infer chiefs' vested interests through channels beyond simply observing the extent of mining activity in the community.

Chiefs in Ghana ascend to power through lineage and customary selection rather than popular election. Succession is typically determined within royal families, with kingmakers or councils of elders nominating and enstooling a candidate based on matrilineal or patrilineal descent, moral character, and community standing (Rattray, 1929). Because chiefs are appointed for life and not electorally selected, this system limits formal downward accountability to citizens (Mamdani, 1996; Acemoglu et al., 2014).

Once installed, however, chiefs are constrained by a set of informal and customary accountability mechanisms rooted in communal consensus and traditional councils (Bamfo, 2000; Ayittey, 2010). Their legitimacy depends on maintaining peace, justice, and the welfare of their people. Chiefs who abuse power, mismanage communal resources, or act against community interests face strong social and institutional sanctions, including public shaming, withdrawal of allegiance by elders and sub-chiefs, and, in extreme cases, destoolment (removal from office). Consistent with the view of chiefs as "stationary bandits" embedded within their communities, these diffuse sanctions, such as gossip, reputational loss, and withheld tribute, serve as potential deterrents and can discipline behavior even in the absence of electoral accountability (Arhin, 1985; Baldwin, 2015).

### 1.2.5 Intervention: The *Poisoned for Gold* Documentary

The information intervention is based on the first 29 minutes of an 81-minute educational documentary, *Poisoned for Gold*, produced by investigative journalist Erastus Asare Donkor. For field implementation, this segment was translated and voice-dubbed in Twi, the predominant local language across study sites.[13]

The film adopts an investigative and evidence-based narrative style. It combines on-site footage of polluted rivers, farmland, and fish habitats with laboratory analyses demonstrating contamination from mercury, lead, arsenic, and cyanide. Scientists, physicians, and affected residents discuss the cumulative health consequences of long-term heavy-metal exposure, linking these to concrete medical outcomes such as miscarriage, birth defects, neurological disorders, kidney and liver damage, and respiratory illnesses. The documentary refrains from advocating specific policy actions or political positions; its objective is purely informational,

[13]The corresponding English version is publicly available online: https://www.youtube.com/watch?v=wrhPscbb_Cc&t=636s.

conveying scientific evidence of mining-related environmental and health risks in an accessible way.

The documentary's credibility and production quality are nationally recognized. Donkor, who received Ghana's 2023 Journalist of the Year Award for a separate investigative film, *Destruction for Gold*, produced *Poisoned for Gold* with journalistic rigor. The film has circulated widely through television, radio, and social media, reaching urban and peri-urban audiences. However, its penetration into rural mining communities remains limited: nearly half of the study villages lack reliable mobile phone or television signal. Consistent with this, baseline data show that 26 percent of households had heard of the film or its producer, and only 16 percent had watched any portion of it. This limited prior exposure underscores the low baseline information environment into which the intervention was introduced.

## 1.3 Research Design

The field experiment was conducted in 99 *galamsey* communities. Randomization was conducted at the community level, stratified along two key dimensions: (i) the prevalence of small-scale mining, measured by whether the baseline share of mining households was above or below the sample median of 0.32, and (ii) the presence of conflicted leadership, defined by whether the share of households believing their chief would take money to allow mining exceeded the median of 0.53 (as described in Section 1.2.4). Within each of the four resulting strata, communities were randomly assigned in equal proportions to one of three experimental arms, yielding 33 communities per arm.

The first treatment arm, seeding, involved a private viewing of the documentary *Poisoned for Gold* by the traditional chief and elders in the palace. The second treatment arm, broadcasting, received the same private palace screening for chiefs and elders as in seeding, but added a community-wide showing of the documentary at the beginning of a structured town hall meeting, which invited chiefs and was open to all community members. The control group received no screening prior to or during the town hall. In both treatment arms, screenings occurred on the same day as the town hall.

Town halls, referred to locally as *durbars*, serve as familiar community assemblies where local leaders and community members discuss public issues and coordinate collective tasks. In this study, they provide a natural setting for the information intervention and for measuring short-run collective action responses. All screenings were facilitated by trained field officers who maintained strict neutrality throughout the process. Officers were instructed not to express personal opinions or recommend specific actions, but to refer participants back to the documentary whenever questions arose, ensuring consistency and minimizing experimenter demand effects. Leader-only screenings were held privately in chiefs' palaces, whereas town halls took place in central venues such as churches, schools, or community centers.

The overall study design, timeline, and survey sequence are summarized in Figure 1.5. Baseline surveys were completed in February 2025, followed by randomized response surveys to collect information on chief conflicts of interest in June 2025, and the information intervention and town halls were implemented in July. Endline surveys were conducted three weeks later to measure learning, attitudinal, and behavioral outcomes.

## 1.4 Data

This section details the outcome data collection and measurement points, which are summarized in Figure 1.6.

### 1.4.1 Town Hall – Structured Community Activities

Town halls were planned in all 99 study communities as the main venue for collective action activities. These town halls brought together traditional leaders and community members to deliberate on local mining issues and participate in structured activities. Each town hall followed a standardized sequence: documentary screening (only in communities assigned to the broadcasting arm), a moderated community discussion, chief's address to the crowd, two ballots capturing demand for community mining regulation and commitment to follow-up engagements, and a concluding announcement of results. Implementation details and field protocols are provided in Appendix C.

Importantly, the public screenings in broadcasting communities came as a surprise to participants. All communities received identical, neutrally worded invitations that mentioned a discussion on *galamsey* but made no reference to the documentary. This design aimed to ensure comparable participation across treatment arms and minimized the risk that attendance reflected prior attitudes toward mining.

The first ballot measured participants' demand for community regulation of *galamsey*, in the form of a customary bylaw. Such bylaws are recognized under Ghana's constitution, giving chiefs both legislative and enforcement authority. The Jema case, remarkably the only town unaffected by *galamsey* in its district, demonstrates the potential of this approach. Importantly, broad-based community support is seen as essential: 76.4 percent of baseline households believe a bylaw requires majority community backing to be effective, highlighting the importance of community buy-in.

Before voting, field officers introduced the "Jema Template", a locally recognized case in which a community adopted a customary bylaw to successfully prevent mining.[14] Officers then emphasized that each community has its own priorities and asked participants: "If your

[14]Appendix Figure 1.A2 presents the Jema bylaw template, which outlines the rationale and clauses of the regulation, accompanied by the fingerprints of the chiefs and clan heads who endorsed it.

community were to create a *galamsey* bylaw, what should its focus be?" Participants chose one of five policy options, grouped into three categories: (A) *No mining* — stopping all current and future mining; (B) *Safer and cleaner mining* — B1/ banning mining in forest reserves and water bodies, B2/ banning amalgam burning near homes, or B3/ both; and (C) *No change* — maintaining the status quo. Following the design of ballot used in political elections in Ghana, our ballot includes simple symbols to facilitate participation among low-literacy attendees. This first ballot is anonymous. A sample first ballot is shown in Appendix Figure 1.A3.

The second ballot captured participants' revealed commitment to contribute time and resources toward sustained collective action beyond the meeting. Three real-stakes activities measured different dimensions of engagement. First, participants could volunteer to attend or organize regular community meetings on *galamsey* to track what is working, identify gaps, and organize responses. Second, they could opt to join a WhatsApp group moderated by the Ghana Environmental Advocacy Group (GEAG), an NGO that supports local environmental initiatives. Third, participants took part in a crowd-bidding activity, allocating part of a potential 50 GHS prize (approximately 5 USD) between themselves and a communal fund to host an NGO workshop. The latter measure captures individuals' financial willingness to invest in local collective action. This second ballot is not anonymous. A sample second ballot is shown in Appendix Figure 1.A4.

Before the ballots, chiefs were invited to address the assembly. The chief can designate a delegate from the traditional council to represent them.[15] Two field officers independently coded their speeches using pen and paper structured checklists, recording whether the chief (or delegate) referred to the documentary, mentioned benefits or harms of *galamsey*, or expressed verbal or behavioral support for collective action. For example, a chief could endorse stricter bylaws, encourage meeting attendance by committing to provide logistical support, or contribute to the NGO fund from his personal 50 GHS endowment. Chiefs were informed earlier in the day that they would be invited to speak to avoid surprise. These behavioral and verbal indicators serve as my measures of chiefs' information dissemination and mobilization efforts. Further details on checklist items and coding procedures are provided in Appendix C.

Out of 99 study communities, I encountered noncompliance in five communities. In four communities, two assigned to broadcasting and two in the control group, the town hall could not be organized.[16] In one additional community, the chief neither attended nor sent

[15]Common reasons cited include sickness or not living in the community. I will present results on heterogeneous treatment effect by chief conflict of interest.

[16]In one broadcasting community, the chief passed away a few days before the scheduled town hall, and, under local custom, no public gatherings or visits by outsiders are permitted until after the funeral. Two other communities (one broadcasting and one control) were deemed unsafe for field officers due to strong *galamsey*-related tensions. Finally, one control community was experiencing an unresolved chieftaincy dispute

a delegate, leaving 94 observations available for analyses on chief efforts in information dissemination and mobilization.

In total, among the 95 town halls, I have 4,588 first ballots on the community regulation and 4,505 second ballots on the revealed commitment of time and resources. On average, 41 percent of the participants are from mining households, slightly higher than the 38 percent mining households share estimated from the randomly selected baseline household sample.[17] Turnout rates varied across treatment arms and chief types; I discuss these differences and robustness checks in the result section.

### 1.4.2 Household Surveys

Recall that my hypothesized causal chain begins with information exposure — through either chiefs' dissemination efforts in the seeding arm or direct exposure to the documentary in the broadcasting arm — and proceeds to subsequent learning and collective action responses. The household surveys measure learning through a series of questions testing health-related knowledge and capture collective action behaviors three weeks after the town halls.[18]

The survey also captures a set of secondary outcomes designed to shed light on mechanisms. These include perceptions of *galamsey*, accountability attitudes toward chiefs, willingness (and perceived willingness of others) to participate in community collective action, and attitudes toward mining and non-mining households. For respondents who attended the town hall, a dedicated module asks whether the chief mentioned specific health risks or referenced the documentary, allowing triangulation with the chief speech checklist.

The endline survey was administered in all 99 study communities, including the four where town halls could not be held. Of the 1,586 baseline households, 1,368 were successfully reinterviewed, yielding an attrition rate of 13.7 percent, with no differential attrition across treatment arms. Among endline respondents, 82.1 percent reported awareness of the town hall, and 49.6 percent reported having attended. At the community level, these correspond to an average of 13.8 tracked households out of 16 baseline respondents, 11.3 aware of the event, and 6.9 attending, equivalent to a 43 percent attendance rate for study households.[19] All analyses using household survey data are estimated as intention-to-treat effects.[20]

involving multiple claimants to the stool.

[17]Because some town halls were held outdoors, formal registration at entry was infeasible. As a result, individuals who attended but chose not to cast a ballot, for example, if they perceived the activities as threatening, are not captured in our attendance count.

[18]Due to logistical constraints, no exit survey was conducted immediately after the town halls. The measured health learning effects therefore reflect knowledge that persists for three weeks rather than short-term recall.

[19]Common reasons for officers failing to reach a household included temporary absence from the community or being occupied at farms or mining sites during the day.

[20]Split-sample analyses in Section 1.5.3 suggest that treatment effects on learning among non-attendees are comparable to those among attendees, implying meaningful horizontal spillovers among community members.

### 1.4.3 Chief Surveys

The chief survey parallels the household survey in structure and content, capturing leaders' knowledge, attitudes, and post-intervention actions. It includes measures of health-risk learning, views on mining, as well as questions designed to document chiefs' motivations and engagement in mobilization following the intervention.

In cases where a chief designated a delegate from the traditional council to represent them at the town hall, enumerators interviewed the delegate instead. Among the 96 endline respondents, 79 percent are the substantive chiefs of their communities, and 21 percent are officially designated delegates.[21]

## 1.5 Results

This section begins with the empirical specification. I then present the experimental results by tracing the causal chain from information exposure to chiefs' dissemination and mobilization efforts, learning among chiefs and community members, and subsequent collective action outcomes measured in the town halls and surveys. The final two subsections examine effects on accountability toward chiefs and the emergence of pessimism about collective action.

### 1.5.1 Empirical Specification

I estimate treatment effects using ordinary least squares (OLS), restricting the sample to observations from the town halls and the endline household and chief surveys. For household or chief $i$ living in community $c$, the outcome variable is denoted by $Y_{ic}$. My main estimating equation is:

$$\begin{aligned} Y_{ic} = \alpha &+ \beta_1 \mathit{Seeding}_c + \beta_2 \mathit{Broadcasting}_c \\ &+ \phi_1 \mathit{Conflicted\ Chief}_c + \phi_2 \mathit{High\ Mining\ Comm}_c + \delta \mathit{Mining\ HH}_{ic} + \lambda_D + \varepsilon_{ic} \end{aligned} \tag{1.1}$$

where $\mathit{Seeding}_c$ and $\mathit{Broadcasting}_c$ are indicator variables for assignment to the seeding and broadcasting treatments, respectively, with the control group omitted. $\mathit{Conflicted\ Chief}_c$ is an indicator for communities where the share of households perceiving the chief would accept

[21] Three chiefs are missing from the endline sample: two from broadcasting communities (one due to the chief's passing before the town hall and one deemed unsafe for fieldwork) and one from a control community experiencing a chieftaincy dispute. Of the endline respondents, 71 percent are the same individuals interviewed at baseline. In cases where a different person was interviewed, this typically reflects that a delegate attended the town hall and therefore was selected to complete the endline survey, while the chief had completed the baseline interview.

money to allow *galamsey* is above the sample median. *High Mining Comm*$_c$ is an indicator for communities with an above-median share of mining households at baseline. *Mining HH*$_{ic}$ is a household-level indicator for whether any household member is engaged in mining. $\lambda_D$ denotes district fixed effects, and the error term $\varepsilon_{ic}$ is clustered at the community level to account for within-community correlation. The coefficients of interest $\beta_1$ and $\beta_2$ capture the average treatment effects of seeding and broadcasting relative to the control group.

To examine heterogeneity by leadership incentives, I extend the specification as follows:

$$\begin{aligned} Y_{ic} = \alpha &+ \beta_1 \textit{Seeding}_c + \beta_2 \textit{Broadcasting}_c \\ &+ \gamma_1 \ \textit{Seeding}_c \times \textit{Conflicted Chief}_c \ + \gamma_2 \ \textit{Broadcasting} \times \textit{Conflicted Chief}_c \\ &+ \kappa_1 \ \textit{Seeding}_c \times \textit{High } \widehat{\textit{Mining}} \textit{ Comm}_c \ + \kappa_2 \ \textit{Broadcasting} \times \textit{High } \widehat{\textit{Mining}} \textit{ Comm}_c \\ &+ \phi_1 \textit{Conflicted Chief}_c + \phi_2 \textit{High } \widehat{\textit{Mining}} \textit{ Comm}_c + \delta \textit{Mining HH}_{ic} + \lambda_D + \varepsilon_{ic} \end{aligned} \tag{1.2}$$

Here, $\widehat{\text{HighMining}}\text{Comm}_c$ is the demeaned version of the high-mining-community indicator, obtained by subtracting its sample mean. This centering ensures that the estimated treatment effects are evaluated at the average level of mining prevalence. The coefficients $\beta_1$ and $\beta_2$ capture the treatment effects of seeding and broadcasting in non-conflicted communities (those with below-median share of households perceiving the chief would accept money to allow *galamsey*). The interaction terms $\gamma_1$ and $\gamma_2$ measure how these effects differ in communities with conflicted leadership versus non-conflicted leadership, such that $\beta_1 + \gamma_1$ and $\beta_2 + \gamma_2$ represent the corresponding treatment effects in conflicted-chief communities. As in the average treatment effect specification, all standard errors are clustered at the community level.

### 1.5.2 Leaders as Information Gatekeepers

I begin the analysis by examining the first step in the causal chain: how chiefs themselves responded to the information treatment through their attendance, communication, and information-sharing efforts during the town halls. Table 1.1 reports the results.

Column (1) shows that when the documentary was screened publicly to both leaders and community members (broadcasting), conflicted chiefs were 41 percentage points less likely to attend the town hall in person ($\beta_2 + \gamma_2$), a 60 percent decline relative to the control mean of 68 percent chief attendance rate. No such pattern appears among conflicted chiefs in the seeding arm. This behavioral avoidance is consistent with conflicted leaders seeking to evade the accountability pressures created by public information dissemination. In Section 1.5.6, I provide survey evidence corroborating this interpretation.

Column (2) examines whether chiefs who attended made an effort to disseminate the information received. In both the seeding and broadcasting treatments, over 95 percent of non-conflicted chiefs referenced the *Poisoned for Gold* documentary during their speech, indicating broad engagement with the intervention content. In contrast, conflicted chiefs were 30–39 percentage points less likely to make any reference to the film ($\gamma_1$ and $\gamma_2$), demonstrating a relative reluctance to share the information they received.

Column (3) considers the intensity of information dissemination, measured by the number of distinct health conditions mentioned by the chief or delegate.[22] Chiefs in the seeding treatment who were non-conflicted mentioned significantly more health conditions than those in the control group, while chiefs in the broadcasting treatment also showed a positive, albeit statistically insignificant, increase. In both treatments, conflicted chiefs mentioned fewer health conditions than non-conflicted chiefs. Although the estimates for differential treatment effects are imprecise, the direction of effects closely aligns with theoretical predictions that vested interests reduce information sharing.

I also examine the broader content of chiefs' speeches. In supplementary analyses (to be included in the online appendix), I report results on whether chiefs mentioned the perceived benefits of *galamsey*, its environmental harms, and mercury exposure pathways.

### 1.5.3 Community Learning

I next examine the second step in the causal chain: whether information exposure translated into greater awareness of health risks from *galamsey*. Data for this analysis come from endline surveys three weeks after the intervention. The endline household survey listed a series of health conditions and asked respondents whether they believed each was related to *galamsey*.[23] For each condition, responses were standardized, and a composite health-knowledge index was constructed using generalized least-squares weighting following Anderson (2008).[24]

Figure 1.7 presents results for chiefs' own learning. The documentary produced strong learning effects among chiefs in both the seeding and broadcasting groups, with increases of 0.48 and 0.57 standard deviations on the health-knowledge index, respectively, relative to control.

Turning to community members, Figure 1.8 shows that despite the information dissemination efforts made by chiefs documented in Table 1.1, there were no detectable improvements in

[22]The enumerator checklist included the same set of health conditions as shown in Figure 1.1, excluding placebo items, with space for additional conditions mentioned.

[23]Responses were recorded on a five-point scale ranging from "strongly believe unrelated" to "strongly believe related." Unlike the binary (yes/no) format used at baseline, this scale increases measurement precision and statistical power.

[24]This method combines several related outcomes into a single index in a way that maximizes statistical power to detect a treatment effect.

health-risk learning among community members in seeding communities, regardless of leader type.[25] In contrast, broadcasting led to substantial learning gains: community members scored 0.50 standard deviations higher on the health-knowledge index relative to control communities.

The learning effect was similar across communities led by conflicted and non-conflicted chiefs, suggesting that direct exposure to the documentary, rather than leader mediation, was crucial for community learning. This pattern lends support to the growing literature on video-based information interventions, which emphasizes the effectiveness of visual and narrative formats for conveying complex technical information. In this context, video delivery substantially enhanced learning beyond what local chiefs could transmit, highlighting the limits of relying on these leaders — who often lack the communicative capacity due to age and lower formal education — to convey complex technical knowledge.[26] While the external validity of this result is context-dependent, it underscores the importance of assessing information complexity and the communicative capacity of local intermediaries when designing delivery channels.

These results are reported as intention-to-treat estimates and include both town hall attendees and non-attendees.[27] The estimated effects therefore capture both information acquired during the town hall through direct exposure to the documentary and indirect post–town hall spillovers from attendees to non-attendees. Appendix Figure 1.A5 splits the sample by attendance status and shows similar learning effects for attendees and non-attendees. The presence of learning among non-attendees points to horizontal information diffusion through post–town hall interactions among community members. By simultaneously informing a critical mass of residents, broadcasting generated shared knowledge that facilitated peer-to-peer discussion and reinforced learning beyond the formal meeting. Consistent with this mechanism, I show in Section 1.5.5 that broadcasting significantly increased the frequency of community discussions about *galamsey*, providing a natural channel through which information spread beyond those directly exposed to the documentary.

### 1.5.4 Collective Action I: Community Consensus On Mining Regulation

Before turning to how information shaped preferences for community mining regulation, I first examine participation in the town hall ballots, which provide the data for the community mining regulation analysis. In the control group, the average attendance was 51.7 participants

[25] Appendix Table 1.A1 presents results on heterogeneous treatment effects on learning.

[26] Because it is considered inappropriate to ask chiefs about their age or education, I lack systematic data to conduct subsample analyses. However, the near-zero effects in seeding suggest that leader-mediated transmission was largely ineffective across the sample.

[27] This approach accounts for potential differential selection into town hall participation, which I examine in the next subsection.

per community. Control communities with conflicted chiefs saw 13.6 fewer attendees on average. Broadcasting communities also experienced lower attendance, about 9.8 fewer participants than control communities, partially due to lower participation among mining households. By contrast, seeding communities with conflicted chiefs saw 10.0 more attendees than control communities with conflicted chiefs, largely due to increased participation among non-mining households.[28]

In subsequent analyses, I address these participation differences by estimating robustness checks that (i) separate results by occupation and (ii) reweigh voters in treatment communities to match the mining versus non-mining household composition observed in control communities.[29]

Table 1.2 examines how the information treatments influenced collective decision-making on community mining regulation, the first dimension of collective action measured during the town halls. Specifically, I analyze voting outcomes from the ballot where participants selected their preferred bylaw governing local mining activities.

The dependent variable in Columns (1) and (2) is a binary indicator equal to one if a town hall attendee voted for one of the two most stringent bylaw options — either banning all mining activities or jointly banning mining in forest reserves and water bodies together with amalgam burning near homes. Column (1) shows that seeding had virtually no effect on preferences for stricter mining regulation ($\beta_1 = 0.026$), while broadcasting increased the share voting for these stringent options by 10 percentage points ($\beta_2$). Column (2) investigates the heterogeneity by chief type. The positive average effect of broadcasting is concentrated in communities led by non-conflicted chiefs: an estimated 16.5 percentage point increase ($\beta_2$). In contrast, the effect is near zero in communities with conflicted leadership ($\beta_2 + \gamma_2 = 0.029$, $p = 0.51$).

Columns (3) and (4) analyze preference dispersion to capture whether information reduced or widened disagreement within communities. The outcome measures the absolute distance between an individual's preferred bylaw stringency and the community mean stringency.[30] Column (3) shows no average effect of either treatment. Column (4) reveals that in communities with conflicted chiefs, broadcasting increased polarization by roughly 19 percent relative to control communities ($\beta_2 + \gamma_2 = 0.125$, $p = 0.038$). Measure-by-measure analyses

[28]Appendix Table 1.A2 reports detailed estimates of turnout differences across treatment arms and leader types.

[29]These weighting estimators allow treatment effects on voting choices (intensive margin), conditional on participation, to vary by occupation, while assuming homogeneous treatment effects within each occupational group. I also plan to conduct sensitivity analyses to assess how strong any differential intensive-margin effects across occupations would need to be to overturn the main result; these analyses will be included in a later version of the paper.

[30]Bylaw intensity is coded as: 3 = ban all mining, 2 = ban both mining in forests/river and amalgam burning near homes, 1 = ban one of these, and 0 = no change.

in Appendix Table 1.A3 indicate that this divergence reflects movements at both extremes of the preference distribution: the share voting for the least stringent option, maintaining the status quo, increased by 4.5 percentage points (79 percent relative to control communities), while the share voting for the most stringent option, banning all mining activities, increased by 12.6 percentage points (58 percent relative to control communities).

Analyses of chiefs' persuasive efforts during the town halls provide additional insight into these patterns. In control communities, 84 percent of chiefs (or their delegates) supported a stringent bylaw.[31] Appendix Table 1.A4 shows that conflicted chiefs in broadcasting communities (or their delegates) attempted to counteract the broadcasting effect: they were 40.7 percentage points more likely than conflicted chiefs in control communities to advocate for maintaining the status quo. ($p = 0.005$). This pattern mirrors the attendance results in Section 1.5.2, where conflicted chiefs were significantly less likely to attend town halls when the documentary was shown publicly.

Taken together, evidence from chiefs' persuasive efforts and community voting suggests the presence of two distinct types of community members: those whose votes reflect their own evaluation of the net benefits of *galamsey*, and those whose choices are influenced by their beliefs about the chief's stance. The latter group may have voted for maintaining the status quo upon observing — or inferring from the chief's absence — that the leader did not endorse the bylaw initiative. The high baseline rate of chiefs supporting stringent bylaws limits my ability to detect whether non-conflicted chiefs changed their own policy preferences or whether their choices directly influenced constituents. Overall, the findings indicate that conflicted chiefs can impede consensus formation by signaling opposition to stricter regulation, while there is no clear evidence that non-conflicted chiefs' persuasive efforts amplified support for regulation.[32]

Overall, these results suggest that information disseminated publicly can shift community preferences toward stricter environmental regulation when delivered under non-conflicted leadership, but may instead deepen divisions when leadership incentives are misaligned. These findings are robust to occupation-specific estimations (Appendix Table 1.A6), weighting to balance voter composition (Appendix Table 1.A7), and a "donut" specification that excludes 33 communities in the middle tercile of the conflict-of-interest distribution (Appendix Table 1.A8).

[31] Chiefs and delegates used the same ballot as community members but were informed that their ballot choices would be announced publicly before the audience, framed as an opportunity to demonstrate leadership and shape community preferences.

[32] In seeding communities, non-conflicted chiefs were 22.9 percentage points more likely to support a stringent bylaw ($p = 0.028$; Appendix Table 1.A4), and their communities were 8.1 percentage points more likely to vote for one ($p = 0.186$; Table 1.2). These results suggest, albeit weakly, that chiefs' persuasion may have modestly influenced community regulation preferences.

### 1.5.5 Collective Action II: Sustained Community Engagement

I now turn to the second dimension of collective action: whether the information treatments fostered continued engagement beyond the initial deliberation on mining regulations. I measure community engagement along two complementary dimensions: (i) stated time commitments made during the town halls, and (ii) actual engagement in follow-up discussions as reported in endline surveys conducted three weeks later.

Figure 1.9 presents the results.[33] The first outcome is participants' decision to join an NGO-moderated WhatsApp group, which provides an external communication channel for coordinating community responses to *galamsey*. The second outcome is an index combining commitments to attend and organize regular follow-up community meetings, as measured from ballots cast at the town halls. The third outcome, constructed from endline survey data, is an index of actual post-event engagement, combining the frequencies with which respondents discussed *galamsey* with mining and non-mining community members. All outcomes are standardized.

Three key patterns emerge. First, in seeding communities with non-conflicted chiefs, follow-up engagement declined across all measures. Yet these same chiefs appeared more proactive during the town hall, with large point estimates for encouraging regular meetings (14.8 percentage points), offering logistical support (22.5 percentage points), and encouraging WhatsApp participation (17.4 percentage points). While only the effect on logistical support is statistically distinguishable from zero, the consistently positive magnitudes across these actions suggest increased leader initiative, albeit measured imprecisely (Appendix Table 1.A5). Thus, a plausible interpretation is that when community members have limited understanding of the problem's urgency, chiefs' unilateral mobilization efforts backfired, crowding out rather than encouraging participation, as community members perceive leaders' actions as a strategic substitute for their own.

Second, broadcasting significantly increased commitments to attend and organize regular follow-up meetings, as well as actual post-event engagement, regardless of the chief's conflict status. This pattern holds despite conflicted chiefs (or their delegates) being 53.5 percentage points less likely to encourage meetings and 31.9 percentage points less likely to provide logistical support during the town hall ($\beta_2 + \gamma_2$ in Appendix Table 1.A5). Thus, for sustaining engagement, informing the broader public alone appears sufficient — unlike building consensus, which requires both an informed public and non-conflicted leadership.

Third, broadcasting also increased enrollment in the NGO-moderated WhatsApp group, but

[33]The heterogeneous treatment specification is reparameterized in all coefficient plots so that effects by leadership type (conflicted vs. non-conflicted) can be directly read from the figure. The underlying regressions remain identical to the main specification. Appendix Tables 1.A9 and 1.A10 reports the corresponding results using non-standardized outcome measures.

only in communities with conflicted leaders, by 11.6 percentage points (a 39 percent increase relative to the 30.1 percent control mean). This pattern suggests that when local leadership credibility is compromised, an informed public may seek alternative coordination channels beyond the realm of traditional authority.

For completeness, I also examine the financial contribution measure elicited during the town halls in Appendix Table 1.A9. I find no significant average or heterogeneous treatment effects on this outcome. Nonetheless, the level of voluntary contribution is notable: on average, participants allocated 26.8 GHS out of a 50 GHS endowment to bring our NGO partner to their community for an additional workshop.

### 1.5.6 Accountability and Perceptions of Leadership

While the previous sections focus on short-run town hall and follow-up engagement outcomes, this section examines community members' attitudes toward holding traditional leaders accountable, an outcome with potentially longer-term implications for local governance dynamics. To measure demand for accountability, respondents were presented with two statements and asked which reflected their views.[34] The two statements are: (1) "When it comes to *galamsey*, community members should be more active in questioning the actions of local chiefs." and (2) "When it comes to *galamsey*, we should have more respect for the authority of local chiefs." Respondents chose one of four mutually exclusive options: "question only," "respect only," "both question and respect," or "neither".

Figure 1.10 presents the results. Seeding had no detectable effect on accountability attitudes. In broadcasting communities with non-conflicted chiefs, respondents were 7.6 percentage points less likely to say that people should question their chiefs' actions (relative to the 49.7 percent control mean), though this difference is statistically insignificant ($p = 0.14$), suggesting no erosion of authority legitimacy. In contrast, under conflicted leadership, broadcasting strengthened accountability norms: respondents were 12.3 percentage points less likely to express unconditional respect for chiefs (relative to the 33.2 percent control mean; $p = 0.006$) and 10.9 percentage points more likely to choose the balanced position that communities should both respect and question their chiefs (relative to the 15.9 percent control mean; $p = 0.077$).

Self-reports from the endline survey with chiefs paint a consistent picture: chiefs in broadcasting communities rated pressure from community members significantly higher as a motivation to participate in collective action (Appendix Figure 1.A6). These patterns echo earlier findings in Section 1.5.2 that conflicted chiefs were less likely to attend the town hall with a public screening — consistent with their attempts to avoid the accountability pressures that public

[34]The design of this question follows Acemoglu et al. (2014), modified to focus on *galamsey* and to reference local chiefs explicitly rather than "leaders" or "authorities."

information dissemination can generate.

### 1.5.7 Pessimism About Collective Action

Despite the overall positive effects of broadcasting, information interventions that convey "bad news" may also have unintended consequences. In the endline survey I measure pessimism in collective action in two ways. Column (1) of Table 1.3 reports responses to the question: "Comparing the two approaches — community collective action and top-down government intervention — which do you think is more promising or effective in addressing the *galamsey* issue?" Columns (2)–(6) measure willingness to take part in a collective community effort, both for the respondent personally and for their perceptions of others' willingness, including their close friends or neighbors, non-mining households, mining households, and the local chief.

The results reveal a consistent pattern. Seeding has small and statistically insignificant effects across all but one outcome. By contrast, broadcasting has statistically significant negative effects across all outcomes.[35]

Respondents in broadcasting communities were 10.8 percentage points less likely to view collective community action as more effective than government intervention. This seems paradoxical given the positive collective action results documented in Sections 1.5.4 and 1.5.5. A plausible interpretation is that while the information shock was strong enough to prompt initial engagement, it also reinforced the perception that *galamsey* is an intractable problem beyond local capacity to solve.

Respondents in broadcasting communities were also 11.2 percentage points less willing to participate themselves (albeit from a high control mean of 91.6 percent), and perceived others as substantially less willing to engage. The magnitude of this perceived decline ranges from 11.9 to 15.8 percent relative to control means, depending on the reference group considered. This result helps distinguish between two distinct Bayesian updating mechanisms. In one model, individuals who know others were also exposed to the documentary may positively update their expectations about others' future actions — believing that others' increased awareness of the risks will spur greater collective effort. In the alternative model, individuals interpret the documentary's portrayal of environmental degradation as evidence of others' past inaction and, assuming path dependence, infer that such inaction will persist. My findings align with the latter: exposure to vivid evidence of environmental harm led to more pessimistic beliefs about others' willingness to act. This result sheds light on the rapidly growing literature on misperception, much of which elicits and models the causes and consequences of second-order beliefs (Bursztyn and Yang, 2022). I revisit the policy implications of the pessimism results in the concluding section.

[35]Appendix Table 1.A11 shows that these pessimistic responses do not vary systematically by chief type.

### 1.5.8 Additional Results and Robustness Checks

Before concluding, I document several additional results and robustness checks. First, using a set of dictator games, I find no evidence of backlash or social exclusion between mining and non-mining households, nor any evidence on perceived backlash (Appendix Table 1.A12). This provides reassurance that increasing awareness of the harms of *galamsey* did not generate social divisions or stigmatization of miners, a potential unintended consequences that could undermine community cohesion.

Second, I examine whether the intervention affected individuals' willingness to engage in collective action at the national level (Appendix Table 1.A13). Respondents were asked (i) whether they would like to add their name to a petition opposing a pro-mining legislative instrument, and (ii) how much of a 20 GHS survey incentive they wished to donate to an NGO supporting grassroots anti-*galamsey* activism.[36] I find no significant treatment effects on either outcome, suggesting that while the intervention influenced local preferences and actions, it did not spill over into broader, national-level collective action.

## 1.6 Conclusion

An extensive literature shows that when individuals receive new information about environmental hazards or their health impacts, they often respond through private adaptation and by revising their policy preferences toward state intervention. Yet there remains limited field evidence that information can causally jump-start community collective action.

In this setting, traditional chiefs serve as the *de facto* local governance authority, filling the vacuum left by weak state institutions. This stands in contrast to much of the recent literature, which randomizes chiefs' involvement in policy domains where the state remains the dominant governance actor. Although randomizing chiefs directly is infeasible in my setting, I exploit experimental variation in information dissemination channels to examine how leadership characteristics, particularly conflicts of interest, shape collective responses to new information.

Specifically, I randomize two channels for delivering a documentary on the health impacts of artisanal and small-scale gold mining: seeding the information through chiefs and broadcasting the same information directly to both chiefs and community members. Which channel performs better is theoretically ambiguous. The seeding approach leverages chiefs' local authority and network centrality, which could legitimize information and facilitate diffusion.

[36]The petition, part of an ongoing effort by the Ghana Coalition Against Galamsey, called on the President to revoke Legislative Instrument 2462, which permits mining in forest reserves under certain conditions. In December 2025, four months after the study concluded, the Ghanaian government repealed this legislative instrument.

The broadcasting approach, by contrast, bypasses potentially conflicted leaders who might suppress or distort information. The experiment thus reveals how the interaction between dissemination channels and leader incentives shapes both learning and community collective action.

On information transmission, while non-conflicted chiefs exerted greater effort in passing on information than conflicted chiefs, these differences did not translate into differences in community learning. Instead, learning depended primarily on the dissemination channel: I observe null learning under seeding but strong health-risk learning under broadcasting. In this context, the information was likely too complex to be accurately transmitted by traditional chiefs, many of whom are older and less formally educated, and more effectively conveyed through direct, video-based communication. Although the external validity of this result must be assessed on a case-by-case basis, it brings to light several underappreciated factors critical to designing effective information interventions: message complexity, the capacity of information conduits, and the technology used for transmission.

On collective action, broadcasting communities, regardless of leader conflicts of interest, showed higher commitment to and actual participation in follow-up engagements. However, conflicts of interest critically shaped how learning translated into demand for stricter local environmental regulation. An informed public increased support for stricter community mining bylaws only when leaders were non-conflicted; where leaders were conflicted, it polarized preferences. The findings point to an important distinction between consensus and engagement. Sustaining community engagement does not depend on leader alignment and responds strongly to broad public information. In contrast, reaching agreement on stricter local mining rules requires both an informed public and non-conflicted leadership; neither alone is sufficient. Although an informed public failed to build consensus under conflicted leadership, the emergence of accountability pressures, where informed citizens begin to question their chiefs, suggests that the arc of collective action may eventually bend toward community welfare.

Finally, the pessimism findings offer a nuanced lesson for the design of information campaigns. Information not only raises awareness but also exposes structural barriers to collective action. "Bad news" campaigns that reveal the extent of environmental harm can reduce ignorance but risk inducing fatalism unless paired with concrete pathways for action. in my setting, while the intervention included real avenues for collective action, such as bylaw ballots and volunteer sign-ups for follow-up meetings, these proved insufficient to counteract the rise in pessimism. Future interventions should therefore pair information dissemination with tangible responses, ensuring that information empowers rather than discourages communities.

# Figures

Figure 1.1: Baseline Household Awareness of Health Conditions Related to *Galamsey*

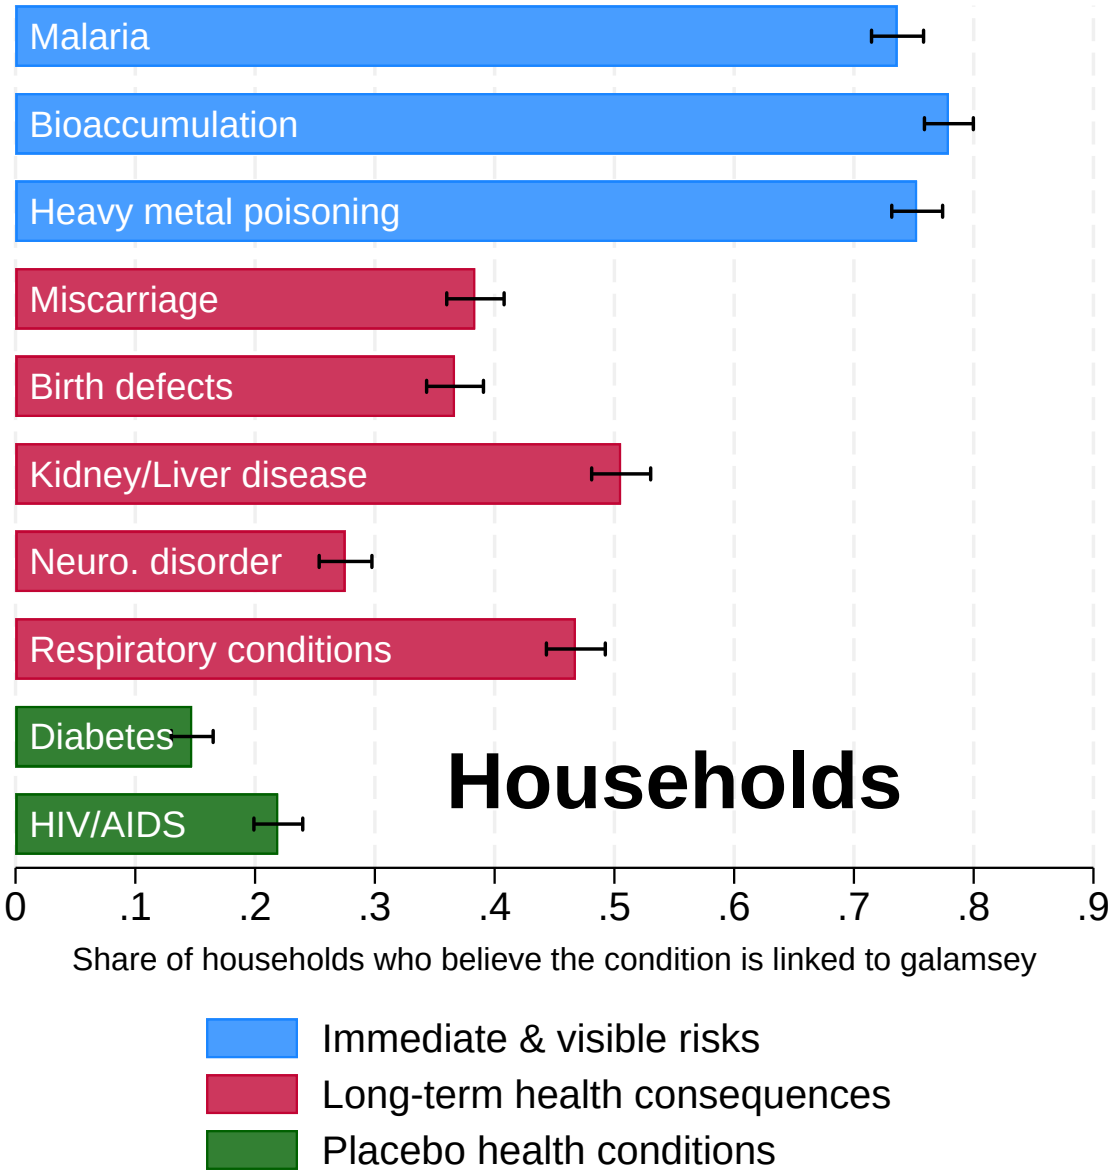


**Notes:** Figure shows the share of households that believe each health condition is related to artisanal mining activities. Bars indicate 95% confidence intervals. Blue bars denote immediate and visible risks, red bars long-term health consequences, and green bars placebo conditions.

Figure 1.2: Baseline Chief Awareness of Health Conditions Related to *Galamsey*

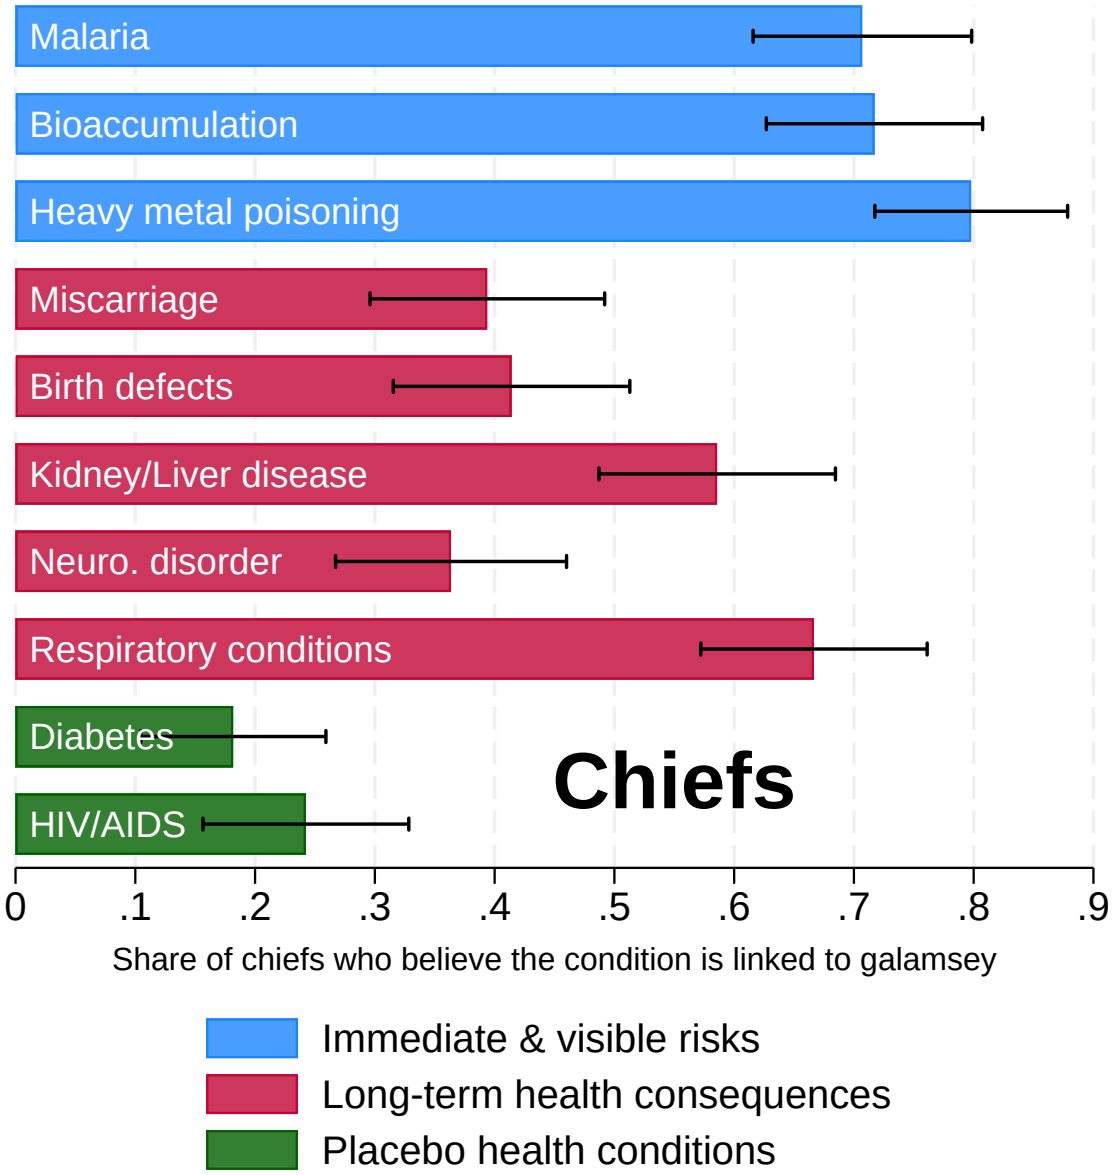


**Notes:** Figure shows the share of chiefs that believe each health condition is related to artisanal mining activities. Bars indicate 95% confidence intervals. Blue bars denote immediate and visible risks, red bars long-term health consequences, and green bars placebo conditions.

Figure 1.3: Distribution of Chiefs' Conflicts of Interest Across Communities

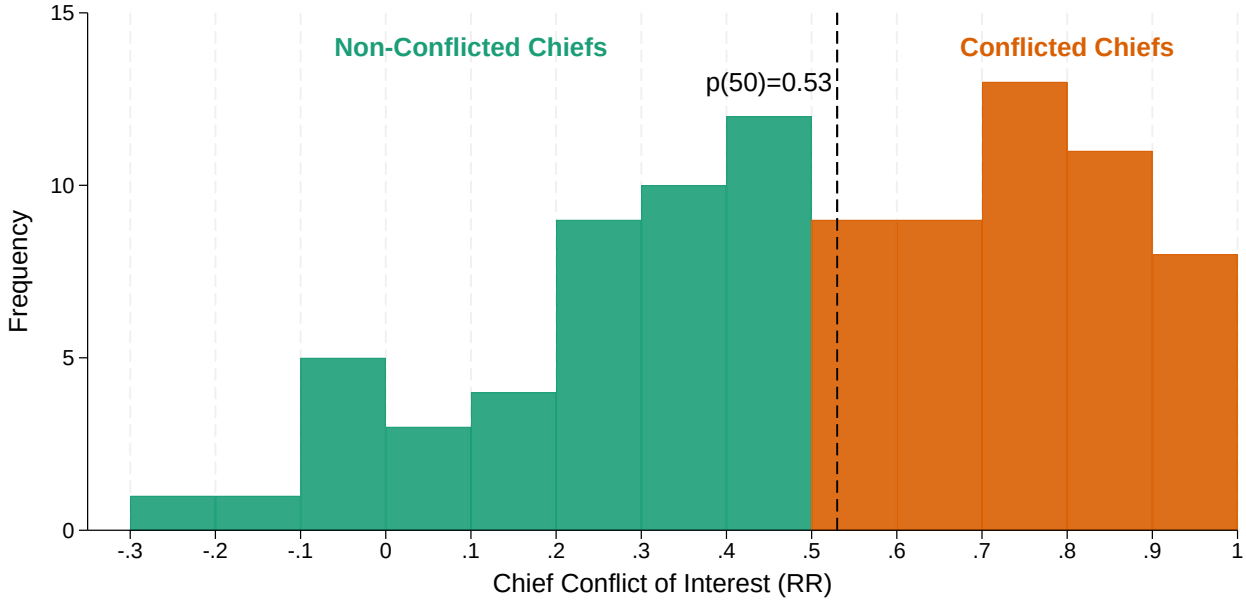


**Notes:** Figure shows the distribution of community-level estimates of chiefs' conflicts of interest, measured using a randomized response protocol. The dashed line indicates the sample median (0.53). Communities with estimated shares above this threshold are classified as *conflicted*, and those below as *non-conflicted*.

Figure 1.4: Relationship Between Chiefs' Conflicts of Interest and Mining Prevalence

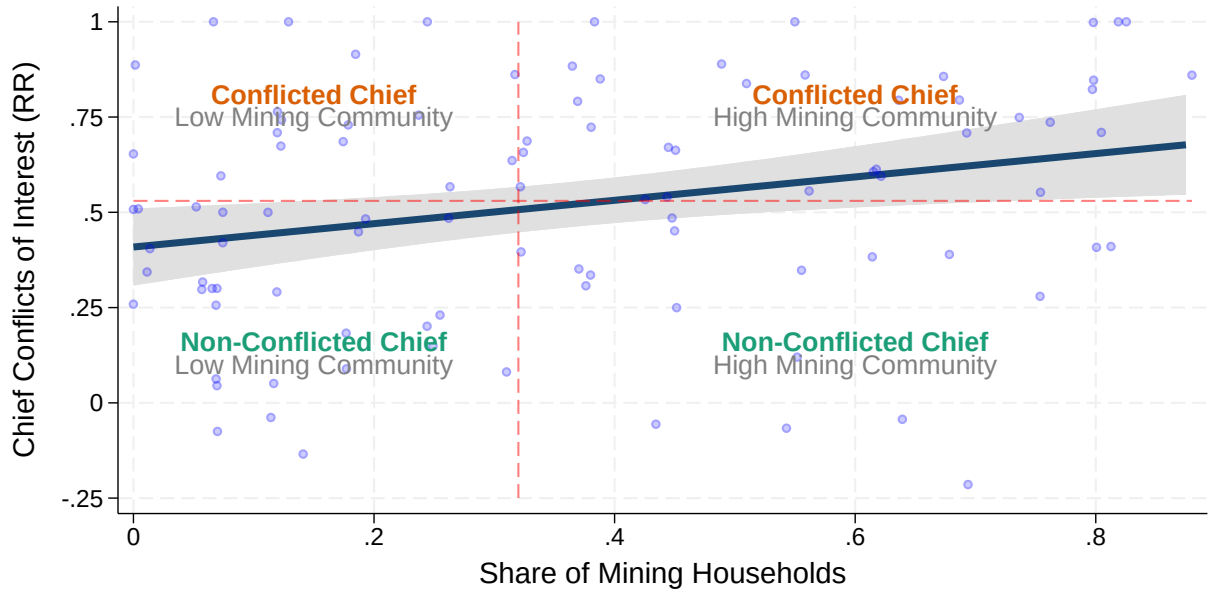


**Notes:** Each point represents a community. The solid line plots the fitted linear relationship between the share of mining households and chiefs' conflicts of interest, with the shaded band indicating the 95 percent confidence interval. The vertical and horizontal dashed lines mark the sample medians of the share of mining households (0.32) and chiefs' conflicts of interest (0.53, as measured via randomized response), respectively.

Figure 1.5: Study Flow and Experimental Design

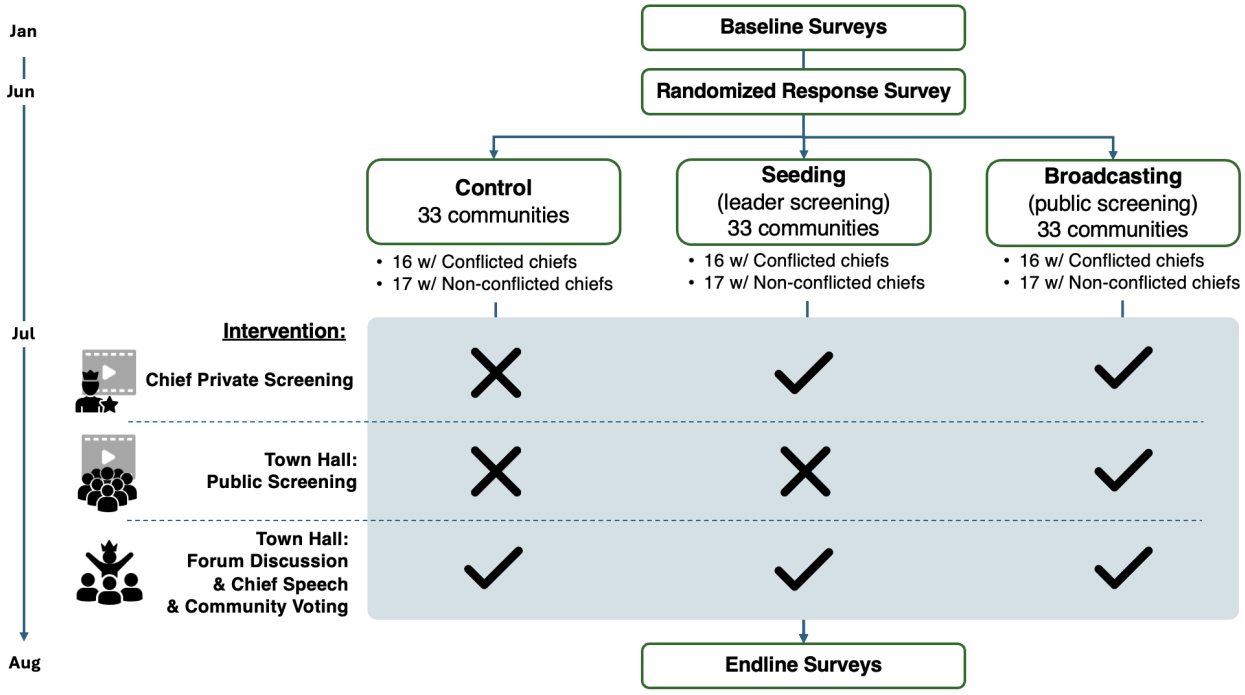

Figure 1.6: Overview of Data Collection and Measurement Points

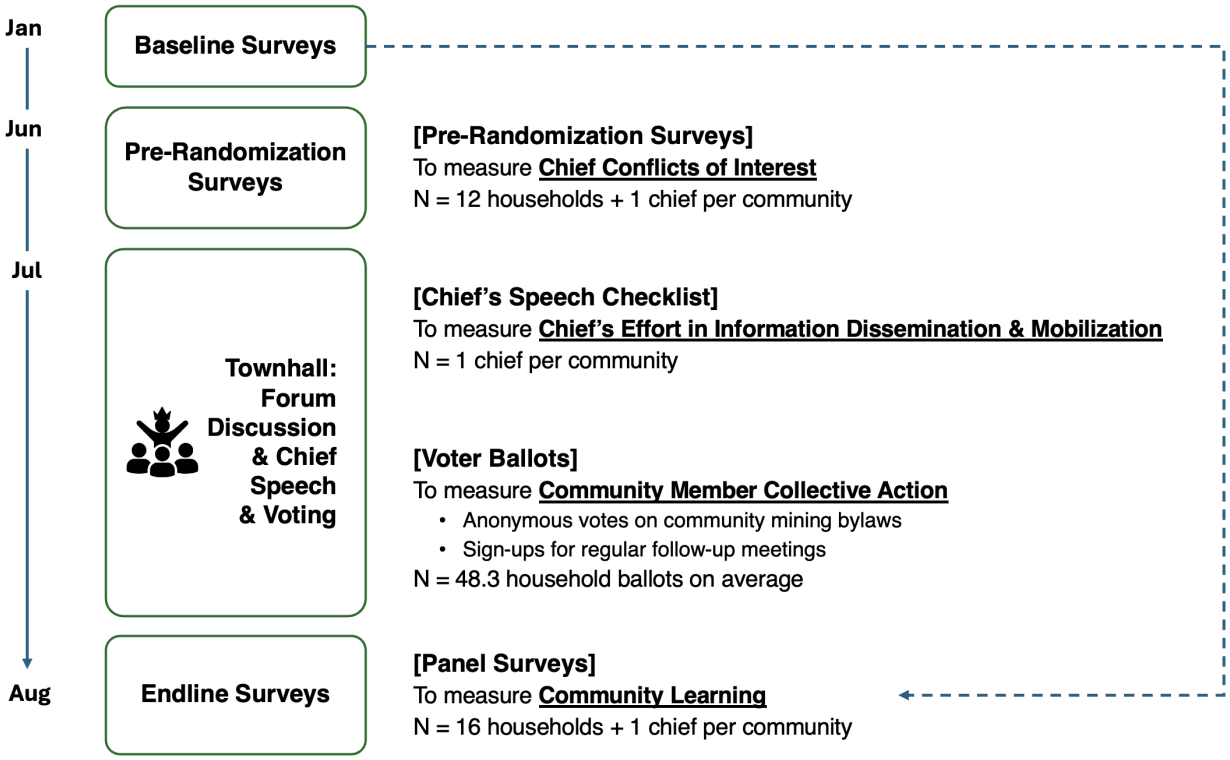

Figure 1.7: Average Treatment Effects on Chiefs' Learning

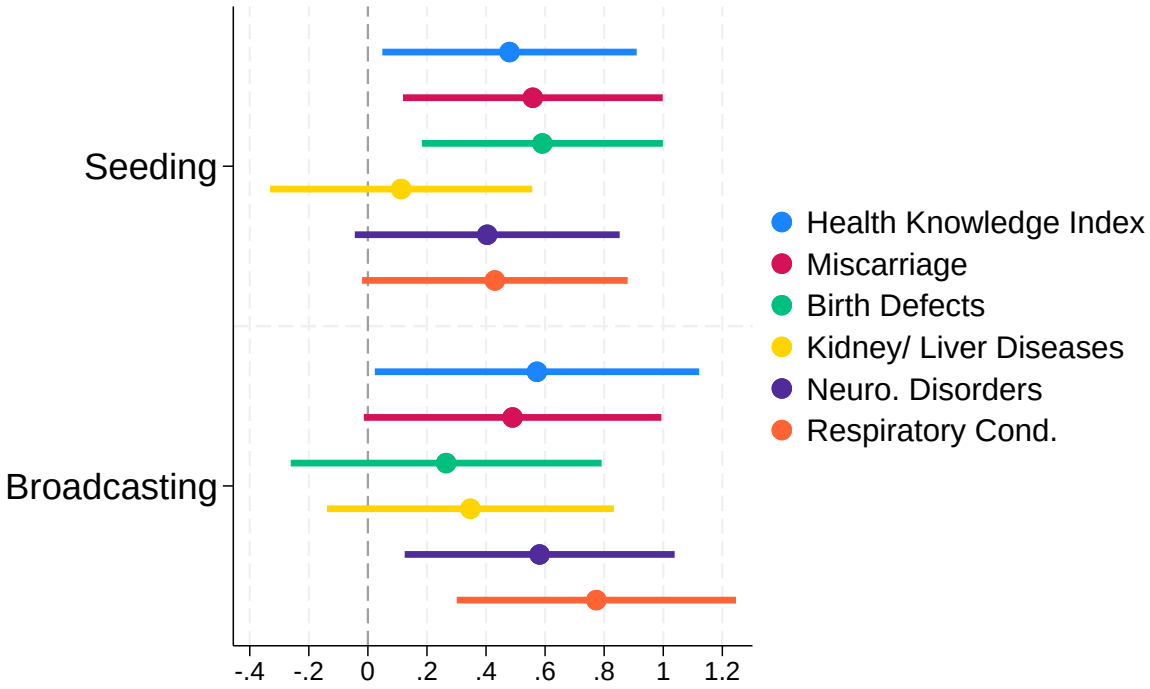


**Notes:** The unit of observation is a chief. The figure reports intention-to-treat effects of seeding and broadcasting on chiefs' knowledge of long-term health risks from *galamsey*. All specifications control for district fixed effects, Conflicted Chief (above-median share of households perceiving the chief would accept money to allow *galamsey*), and High-Mining Community (above-median share of mining households). Points denote coefficient estimates and bars indicate 95% confidence intervals. Standard errors are robust.

Figure 1.8: Average Treatment Effects on Community Members' Learning

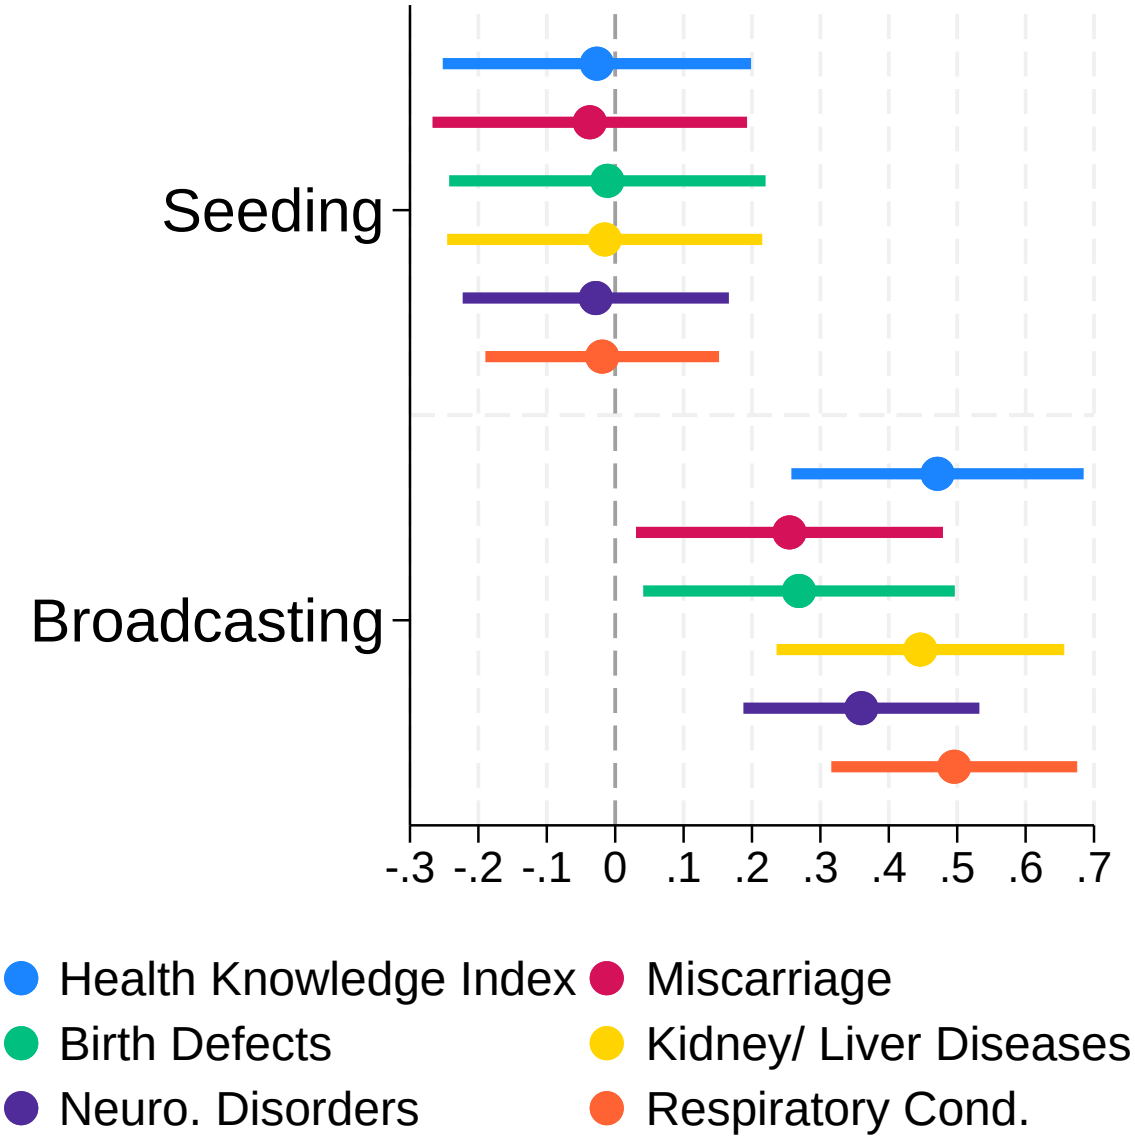


**Notes:** The unit of observation is a household. The figure reports intention-to-treat effects of seeding and broadcasting on household knowledge of long-term health risks from *galamsey*. All specifications control for district fixed effects, Conflicted Chief (above-median share of households perceiving the chief would accept money to allow *galamsey*), and High-Mining Community (above-median share of mining households). Points denote coefficient estimates and bars indicate 95% confidence intervals. Standard errors are clustered at the community level.

Figure 1.9: Heterogeneous Treatment Effects on Sustained Community Engagement

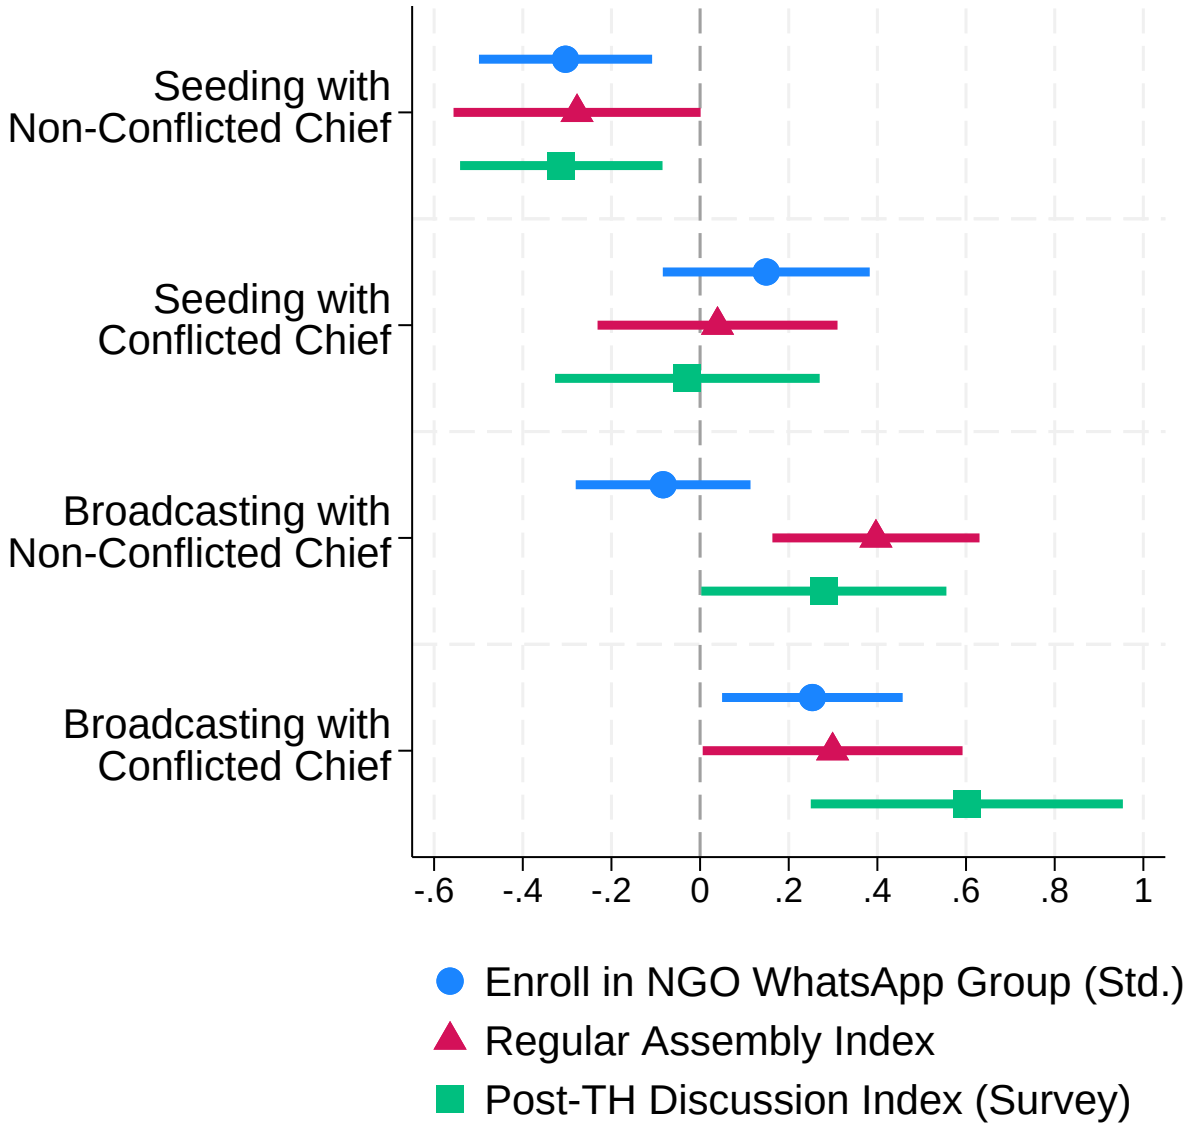


**Notes:** The first two outcomes (joining an NGO WhatsApp group and the regular meeting index) are measured from town hall ballots (n=4,505). The third outcome (discussion index) is measured from the endline household survey (n=1,366). Indexes are standardized using generalized least squares weighting following Anderson (2008). All specifications control for district fixed effects, gender, mining-household status, Conflicted Chief (above-median share of households perceiving the chief would accept money to allow *galamsey*), and interactions between treatment indicators and the demeaned high-mining-community indicator (above-median share of mining households). Points denote coefficient estimates and bars indicate 95% confidence intervals. Standard errors are clustered at the community level.

Figure 1.10: Heterogeneous Treatment Effects on Accountability Norms Toward Local Chiefs

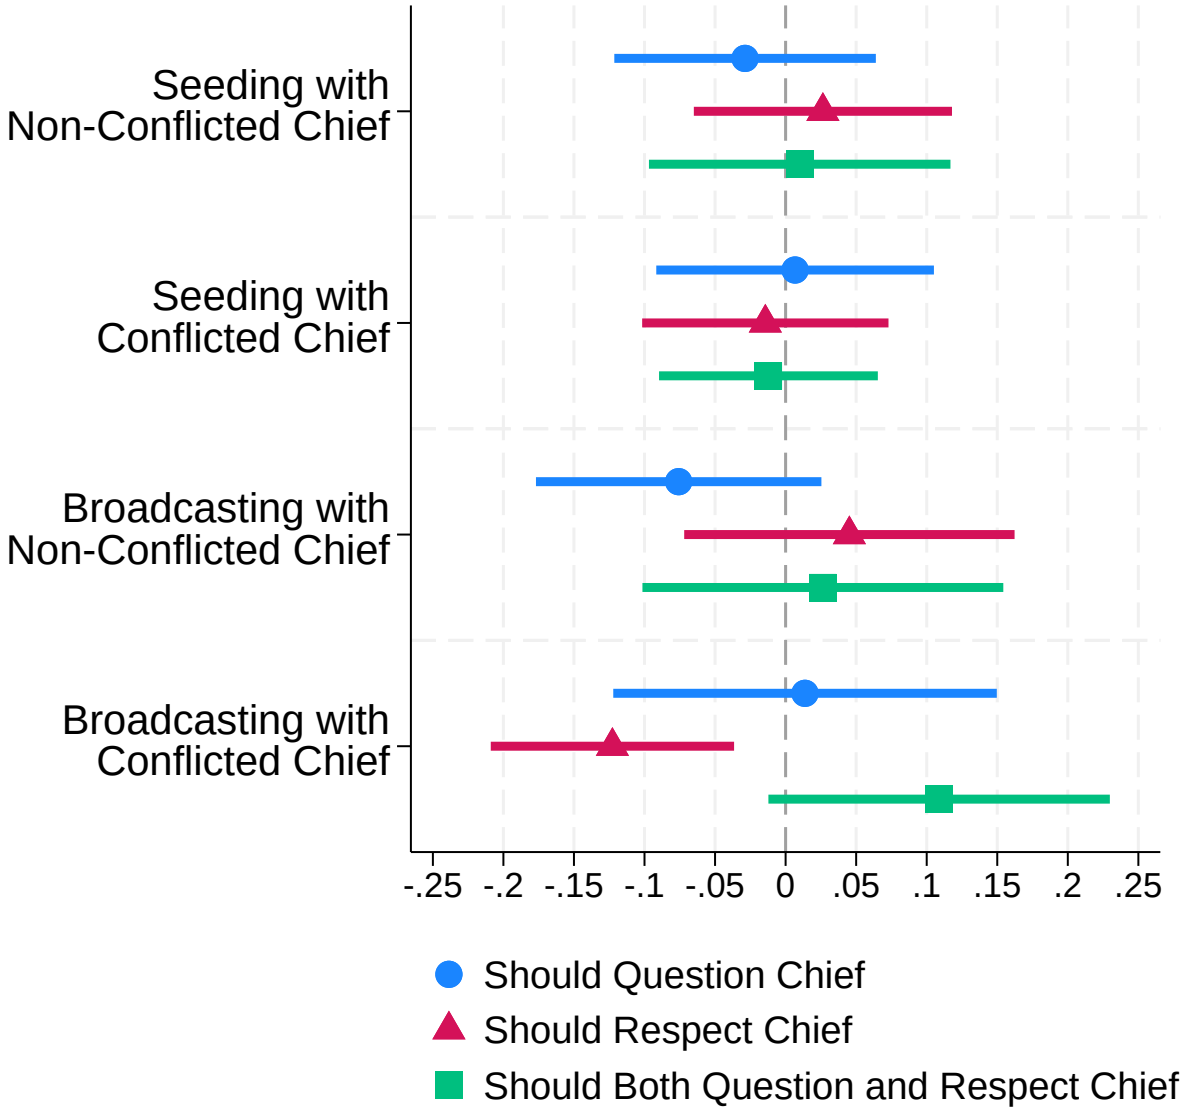


**Notes:** The unit of observation is a household. Outcomes correspond to whether respondents agree that community members should question chiefs, respect chiefs, or both in the context of *galamsey*. All specifications control for district fixed effects, gender, mining-household status, Conflicted Chief (above-median share of households perceiving the chief would accept money to allow *galamsey*), and interactions between treatment indicators and the demeaned high-mining-community indicator (above-median share of mining households). Points denote coefficient estimates and bars indicate 95% confidence intervals. Standard errors are clustered at the community level.

# Tables

Table 1.1: Heterogeneous Treatment Effects on Chiefs' Attendance and Dissemination Efforts

| | Chief Attended Town Hall (1) | Chief Referenced Documentary (2) | No. Health Conditions Mentioned (3) |
|---|---|---|---|
| $\beta_1$ : Seeding | -0.028 | 0.953*** | 0.988*** |
| | (0.190) | (0.086) | (0.359) |
| $\beta_2$ : Broadcasting | -0.040 | 0.987*** | 0.619 |
| | (0.188) | (0.052) | (0.403) |
| $\gamma_1$ : Seeding × Conflicted Chief | -0.049 | -0.386** | -0.364 |
| | (0.258) | (0.159) | (0.549) |
| $\gamma_2$ : Broadcasting × Conflicted Chief | -0.374 | -0.297** | -0.259 |
| | (0.258) | (0.144) | (0.616) |
| Conflicted Chief | 0.154 | -0.062 | -0.180 |
| | (0.189) | (0.053) | (0.389) |
| Ctrl Mean | 0.677 | 0.000 | 0.871 |
| Ctrl SD | 0.475 | 0.000 | 0.922 |
| Obs. | 94 | 94 | 94 |
| R-squared | 0.202 | 0.692 | 0.231 |
| $\beta_1 + \gamma_1\,(est.)$ | -0.076 | 0.567 | 0.624 |
| $\beta_1 + \gamma_1 = 0\,(p\,value)$ | 0.648 | 0.000 | 0.131 |
| $\beta_2 + \gamma_2\,(est.)$ | -0.414 | 0.690 | 0.359 |
| $\beta_2 + \gamma_2 = 0\,(p\,value)$ | 0.021 | 0.000 | 0.424 |

**Notes:** The unit of observation is a town hall. Column 1 is an indicator equal to one if the chief personally attended the town hall rather than sending a delegate. Column 2 equals one if the chief (or delegate) referenced the *Poisoned for Gold* documentary during the speech. Column 3 records the number of distinct health conditions mentioned by the chief (or delegate). Conflicted Chief is an indicator for communities where the share of households perceiving the chief would accept money to allow *galamsey* is above the sample median. All regressions include district fixed effects and interactions between treatment indicators and the demeaned high-mining-community indicator (equal to one for communities with an above-median share of mining households at baseline). Robust standard errors are shown in parentheses. The bottom panel reports estimated linear combinations and corresponding $p$-values for treatment effects in conflicted-chief communities, specifically $\beta_1 + \gamma_1$ for seeding and $\beta_2 + \gamma_2$ for broadcasting.

Table 1.2: Heterogeneous Treatment Effects on Community Mining Regulation Preferences

| | Stringent Measure | | Polarization | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| $\beta_1$ : Seeding | 0.026 | 0.081 | 0.005 | -0.052 |
| | (0.043) | (0.061) | (0.043) | (0.063) |
| $\beta_2$ : Broadcasting | 0.100** | 0.165*** | 0.052 | -0.014 |
| | (0.038) | (0.053) | (0.039) | (0.048) |
| $\gamma_1$ : Seeding × Conflicted Chief | | -0.119* | | 0.124 |
| | | (0.069) | | (0.081) |
| $\gamma_2$ : Broadcasting × Conflicted Chief | | -0.136** | | 0.140* |
| | | (0.065) | | (0.077) |
| Conflicted Chief | 0.013 | 0.098** | 0.008 | -0.080 |
| | (0.034) | (0.040) | (0.043) | (0.057) |
| Ctrl Mean | 0.590 | 0.590 | 0.643 | 0.643 |
| Ctrl SD | 0.492 | 0.492 | 0.471 | 0.471 |
| Obs. | 4588 | 4588 | 4588 | 4588 |
| R-squared | 0.108 | 0.113 | 0.026 | 0.030 |
| $\beta_1 + \gamma_1\,(est.)$ | | -0.038 | | 0.072 |
| $\beta_1 + \gamma_1 = 0\,(p\,value)$ | | 0.389 | | 0.183 |
| $\beta_2 + \gamma_2\,(est.)$ | | 0.029 | | 0.125 |
| $\beta_2 + \gamma_2 = 0\,(p\,value)$ | | 0.508 | | 0.038 |

**Notes:** The unit of observation is a household ballot from the town hall vote on community mining regulation. Columns 1–2 measure whether respondents voted for a stringent bylaw option (banning all mining or both mining in forest/river areas and amalgam burning near homes). Columns 3–4 measure preference dispersion, defined as the absolute distance between an individual's preferred bylaw stringency and the community mean. Conflicted Chief is an indicator for communities where the share of households perceiving the chief would accept money to allow *galamsey* is above the sample median. All regressions include district fixed effects, voter gender, High Mining Community (equal to one for communities with an above-median share of mining households at baseline), and Mining Household. Columns 2 and 4 demean High Mining Community for its main effect and include interactions between treatment indicators and the demeaned high-mining-community indicator. Standard errors are clustered at the community level. The bottom panel reports estimated linear combinations and corresponding $p$-values for treatment effects in conflicted-chief communities, specifically $\beta_1 + \gamma_1$ for seeding and $\beta_2 + \gamma_2$ for broadcasting.

Table 1.3: Average Treatment Effects on Pessimism About Collective Action

| | Perceived or Own Willingness | | | | | |
|---|---|---|---|---|---|---|
| | 3 Closest Friends & Neigh. | 10 Other Non-min. HHs | 10 Other Mining HHs | Chief | Self | Perceived Collective Action Efficacy |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Seeding | 0.011 | -0.066 | -0.476** | 0.013 | -0.024 | 0.003 |
| | (0.062) | (0.154) | (0.196) | (0.031) | (0.026) | (0.042) |
| Broadcasting | -0.365*** | -1.180*** | -0.566* | -0.122*** | -0.112*** | -0.108** |
| | (0.081) | (0.200) | (0.305) | (0.041) | (0.032) | (0.042) |
| Conflicted Chief | -0.063 | -0.275* | -0.064 | -0.038 | -0.005 | -0.038 |
| | (0.069) | (0.146) | (0.262) | (0.037) | (0.026) | (0.041) |
| Ctrl Mean | 2.310 | 7.838 | 4.776 | 0.791 | 0.916 | 0.603 |
| Ctrl SD | 0.782 | 1.886 | 2.668 | 0.407 | 0.278 | 0.490 |
| Obs. | 1366 | 1366 | 1366 | 1368 | 1368 | 1368 |
| R-squared | 0.069 | 0.094 | 0.072 | 0.055 | 0.035 | 0.019 |

**Notes:** The unit of observation is a household. Column (1) reports responses on whether community collective action is perceived as more effective than government intervention. Columns (2)–(6) measure respondents' own willingness and their beliefs about others' willingness to participate in community efforts to address *galamsey*. Conflicted Chief is an indicator for communities where the share of households perceiving the chief would accept money to allow *galamsey* is above the sample median. All specifications include district fixed effects, High Mining Community (equal to one for communities with an above-median share of mining households at baseline), and Mining Household indicators. Standard errors are clustered at the community level.

# Appendix

## A Community Sampling

Prior to the baseline, research assistants conducted phone interviews with 134 out of 174 elected assemblymen across the six study districts. To ensure cost-effectiveness, I excluded small hamlets with populations below 400. To avoid large communities where group size limits the feasibility of collective action, I also excluded urban areas and communities with populations exceeding 3,000. I focused instead on 241 medium-sized rural communities with populations between 400 and 3,000. Among these, I identified 111 communities where assemblymen reported either (1) ongoing small-scale mining activities or (2) a future threat of irresponsible small-scale mining; this constituted the full set of eligible villages for the baseline survey. 12 of these were later dropped because they were too large or geographically proximate, leaving 99 communities in the final randomization sample to minimize cross-community spillovers.

## B Validation of the Randomized Response Measure

After completing the randomized response (RR) protocol, enumerators privately administered a set of direct questions to the same respondents to cross-check perceptions of chiefs' vested interests in mining. These "direct ask" (DA) questions were carefully worded to avoid judgmental framing and to elicit beliefs about whether the chief or elders personally benefit from mining, own or hold stakes in a mining site or concession, or send representatives to collect payments from miners. If respondents reported that no mining was currently occurring in their community, they were instead asked two hypothetical variants: (i) if *galamseyers* offered a *knocking fee* (introductory payment) to begin mining, would the chief and elders accept it? and (ii) if mining commenced and royalties were offered, would they accept them?

I standardized and aggregated responses from these items into a composite DA index of perceived leadership conflicts. The RR and DA measures were strongly correlated, supporting the validity of the RR-based measure used in the main analysis: 0.41 at the individual level ($n = 1,114$) and 0.65 at the community level. These correlations indicate that the RR approach captures meaningful variation in underlying perceptions of leadership incentives while benefiting from the privacy protection necessary for candid responses on sensitive topics. To protect respondent anonymity, no personally identifying information (e.g., name, age, gender) was collected during either the RR or DA modules.

## C Town Hall Implementation Details

Field officers organize a town hall meeting in every community, inviting both traditional leaders and community members to participate. These town halls serve as the primary setting

for measuring collective action. The detailed sequence of activities is as follows: (1) opening remarks, (2) documentary screening (only in broadcasting communities), (3) moderated forum discussion, (4) chief's address to the crowd, (5) two ballots capturing collective action, and (6) ballot result announcement and closing remarks. No prior registration is required for community members. Both survey and non-survey households are welcome to attend, and the event is structured to encourage broad participation.

To maximize participation, the event is promoted in multiple ways. Announcements are made through the local public address system, and formal invitation letters are delivered to key opinion leaders, such as religious leaders, school headmasters, leaders of farmer-based organizations or cooperatives, and youth representatives. The 16 study households in each community also receive invitation letters. On the day of the event, officers conduct targeted outreach to these study households, going door-to-door to personally invite them to attend.

### Opening Remarks

Field officers begin by welcoming participants, introducing themselves and the purpose of the meeting, and outlining the agenda. Officers emphasize their neutrality and the goal of learning from the community. They explain that the session will cover local environmental issues and collective decision-making processes. To minimize experimenter demand effects, I train officers to maintain a neutral tone throughout.

### Documentary Screening (Broadcasting only)

In broadcasting communities, field officers publicly screen the first 29 minutes of *Poisoned for Gold*. Officers briefly introduce the documentary and emphasize that the documentary is aimed at providing unbiased information about potential health consequences of *galamsey* to help the community make informed decisions.

### Moderated Forum Discussion

Following the introduction (and screening, if applicable), participants remain seated as one large group for an open forum. Officers emphasize that all viewpoints are welcome, whether in support of or against *galamsey* or any regulation, and that the aim is to foster mutual understanding rather than judge individuals. The discussion guide prompts participants to share: (1) what they know about *galamsey*, (2) perceived benefits to themselves or the community, (3) problems or negative effects caused by *galamsey*, and (4) their views on the best way for the community to address the issue.

A field officer records every turn of the discussion using a structured tally sheet, documenting the age group (young vs. old), gender, whether the speaker is contributing for the first time,

and whether their overall position is for or against *galamsey*. These tallies form a secondary outcome dataset, pending further analysis.

### Chief's Address

After explaining both ballots but before participants begin filling them, officers invite the chief to address the community. During the chief's remarks, two enumerators independently complete a structured speech checklist, then cross-check their entries after the event. The reconciled version is used for analysis. To reduce the risk of self-censorship and given the sensitivity of the topic, no audio or video recordings are taken.

For the bylaw voting measure, field officers record whether the chief publicly expressed support for any of the bylaw options. For the in-person *galamsey* meeting and WhatsApp group measures, field officers capture verbal influence if the chief explicitly encouraged or discouraged participation. In the case of the in-person meeting measure, field officers also record behavioral support when the chief offers logistical assistance, such as providing a meeting space or appointing a liaison to gather feedback and report to the chief after each meeting.

For the financial contribution measure, field officers document both verbal and behavioral support. Verbal support is recorded if the chief explicitly urges contributions to the NGO invitation bidding fund. Behavioral support is measured by the amount the chief chooses to contribute from a 50 GHS endowment in the same lottery-based game as other participants. The chief's ticket is treated as an automatic winner, and his contribution decision is publicly announced before participants complete their own tickets.

To ensure that chiefs are adequately prepared and not taken by surprise, they are informed during the re-entry process on the day of the town hall that they will be invited to deliver a speech. In seeding and broadcasting communities, this notification is provided following the private viewing of the documentary for the chiefs in their palace.

### The First Ballot

See the main text.

### The Second Ballot

The second ballot records participants' revealed preferences for engaging in sustained collective action beyond the town hall. Participants have to provide their name and phone number for follow-up.

1. **Attending or Organizing Regular *Galamsey* Meetings:** This measure captures the time commitment participants are willing to make to address the *galamsey* challenge.

The first option records willingness to attend regular community meetings focused on monitoring progress, identifying challenges, and discussing solutions. The second option captures willingness to take the more time-consuming role of organizing such meetings, which may involve scheduling, mobilizing participants, and preparing the venue.

2. **Joining an NGO Local WhatsApp Group:** This measure captures willingness to stay engaged via digital communication. Participants are invited to join a WhatsApp group managed by the Ghana Environmental Advocacy Group (GEAG), an NGO that supports communities in protecting land, water, and health from mining-related impacts. This WhatsApp group is exclusive to members of the participant's own community and focuses on the *galamsey* situation in that community, rather than national-level issues.

3. **Crowd-Bidding for an NGO Workshop:** Participants decided how much of a 50 GHS lottery prize to contribute to the community versus keep personally. Participants are informed that GEAG will visit the highest–average-contribution community in each of the six districts to run a planning workshop, introducing case material from Jema. Participants were explicitly told that the relevant "average" would be calculated from only the four randomly drawn winners plus the chief, as opposed to all submitted tickets. Thus, contributing more increases a community's chance of being selected. After choices are made, four winners per community are randomly drawn, and their stated contribution decisions are implemented. Note that the average contribution is computed from the drawn tickets and the chief's ticket. For transparency, the winner's name and contribution amount are announced publicly at the event. Whether this public disclosure creates additional social image costs or benefits depends on local dynamics as perceived by the participants themselves.

**Ballot Result Announcement and Closing Remarks**

After all participants have cast both ballots, officers immediately count the first ballot and announce the results following a standardized script. For the second ballot, they draw four winners at random, implement the contribution amounts indicated on their tickets, and calculate the community's average contribution based on these four amounts and the chief's automatically winning ticket. The average is announced publicly, along with a reminder that the highest-contributing community in each district will receive a follow-up workshop from GEAG.

In closing, officers thank attendees and invite those interested in forming a local group for regular *galamsey* meetings to remain. They record the number of participants who stay and, among them, the number willing to help organize future meetings.

Full field manuals and procedural details are available upon request.

# Appendix Figures

Figure 1.A1: Distribution of the Share of Mining Households Across Communities

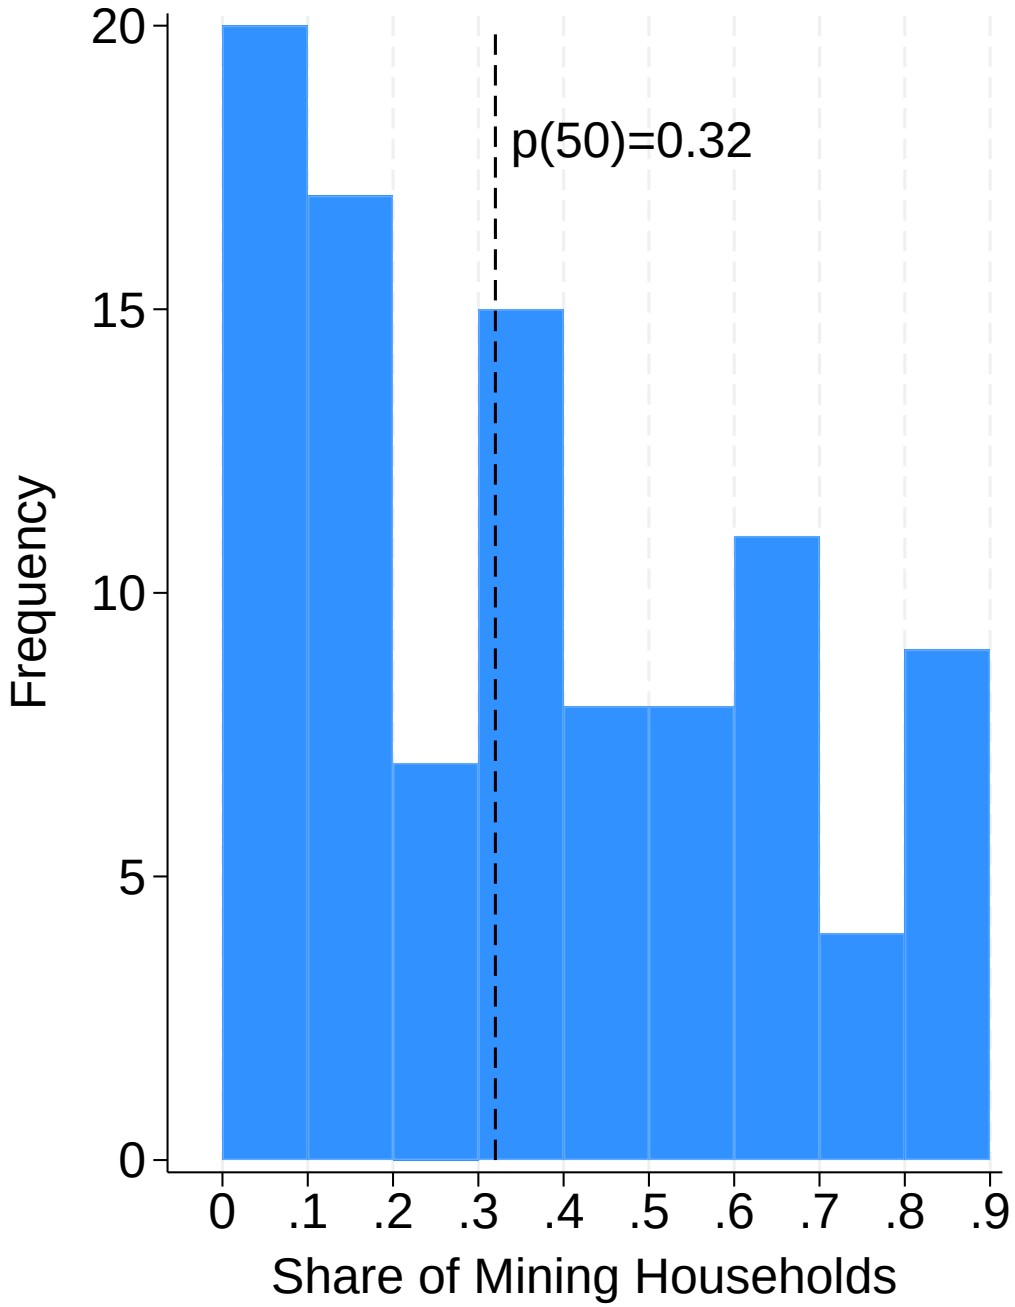


**Notes:** Figure shows the distribution of community-level shares of mining households. For each community, the share is calculated as the proportion of 16 randomly sampled households that report having at least one member engaged in small-scale mining. The dashed vertical line indicates the sample median (0.32).

Figure 1.A2: Jema Template — Example of a Customary Bylaw on Small-Scale Mining Regulation

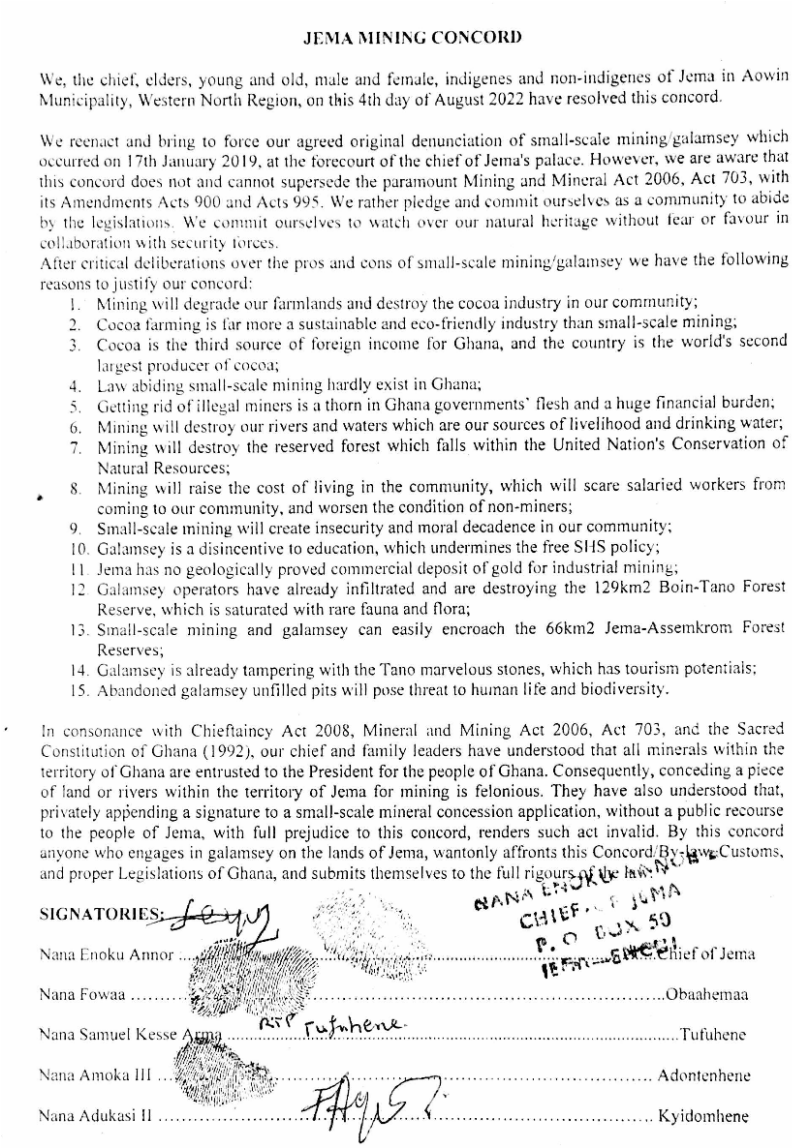

**JEMA MINING CONCORD**

We, the chief, elders, young and old, male and female, indigenes and non-indigenes of Jema in Aowin Municipality, Western North Region, on this 4th day of August 2022 have resolved this concord.

We reenact and bring to force our agreed original denunciation of small-scale mining/galamsey which occurred on 17th January 2019, at the forecourt of the chief of Jema's palace. However, we are aware that this concord does not and cannot supersede the paramount Mining and Mineral Act 2006, Act 703, with its Amendments Acts 900 and Acts 995. We rather pledge and commit ourselves as a community to abide by the legislations. We commit ourselves to watch over our natural heritage without fear or favour in collaboration with security forces.

After critical deliberations over the pros and cons of small-scale mining/galamsey we have the following reasons to justify our concord:

1. Mining will degrade our farmlands and destroy the cocoa industry in our community;
2. Cocoa farming is far more a sustainable and eco-friendly industry than small-scale mining;
3. Cocoa is the third source of foreign income for Ghana, and the country is the world's second largest producer of cocoa;
4. Law abiding small-scale mining hardly exist in Ghana;
5. Getting rid of illegal miners is a thorn in Ghana governments' flesh and a huge financial burden;
6. Mining will destroy our rivers and waters which are our sources of livelihood and drinking water;
7. Mining will destroy the reserved forest which falls within the United Nation's Conservation of Natural Resources;
8. Mining will raise the cost of living in the community, which will scare salaried workers from coming to our community, and worsen the condition of non-miners;
9. Small-scale mining will create insecurity and moral decadence in our community;
10. Galamsey is a disincentive to education, which undermines the free SHS policy;
11. Jema has no geologically proved commercial deposit of gold for industrial mining;
12. Galamsey operators have already infiltrated and are destroying the 129km2 Boin-Tano Forest Reserve, which is saturated with rare fauna and flora;
13. Small-scale mining and galamsey can easily encroach the 66km2 Jema-Assemkrom Forest Reserves;
14. Galamsey is already tampering with the Tano marvelous stones, which has tourism potentials;
15. Abandoned galamsey unfilled pits will pose threat to human life and biodiversity.

In consonance with Chieftaincy Act 2008, Mineral and Mining Act 2006, Act 703, and the Sacred Constitution of Ghana (1992), our chief and family leaders have understood that all minerals within the territory of Ghana are entrusted to the President for the people of Ghana. Consequently, conceding a piece of land or rivers within the territory of Jema for mining is felonious. They have also understood that, privately appending a signature to a small-scale mineral concession application, without a public recourse to the people of Jema, with full prejudice to this concord, renders such act invalid. By this concord anyone who engages in galamsey on the lands of Jema, wantonly affronts this Concord/By-law, Customs, and proper Legislations of Ghana, and submits themselves to the full rigours of the law.

**SIGNATORIES:**

Nana Enoku Annor ..... Chief of Jema

Nana Fowaa ..... Obaahemaa

Nana Samuel Kesse Arma ..... Tufuhene

Nana Amoka III ..... Adontenhene

Nana Adukasi II ..... Kyidomhene

**Notes:** The "Jema Mining Concord" was developed by traditional leaders and community members in Jema to ban all forms of mining. It provides the reference template used in our town hall bylaw voting exercise.

Figure 1.A3: Sample First Ballot on Community Mining Regulation

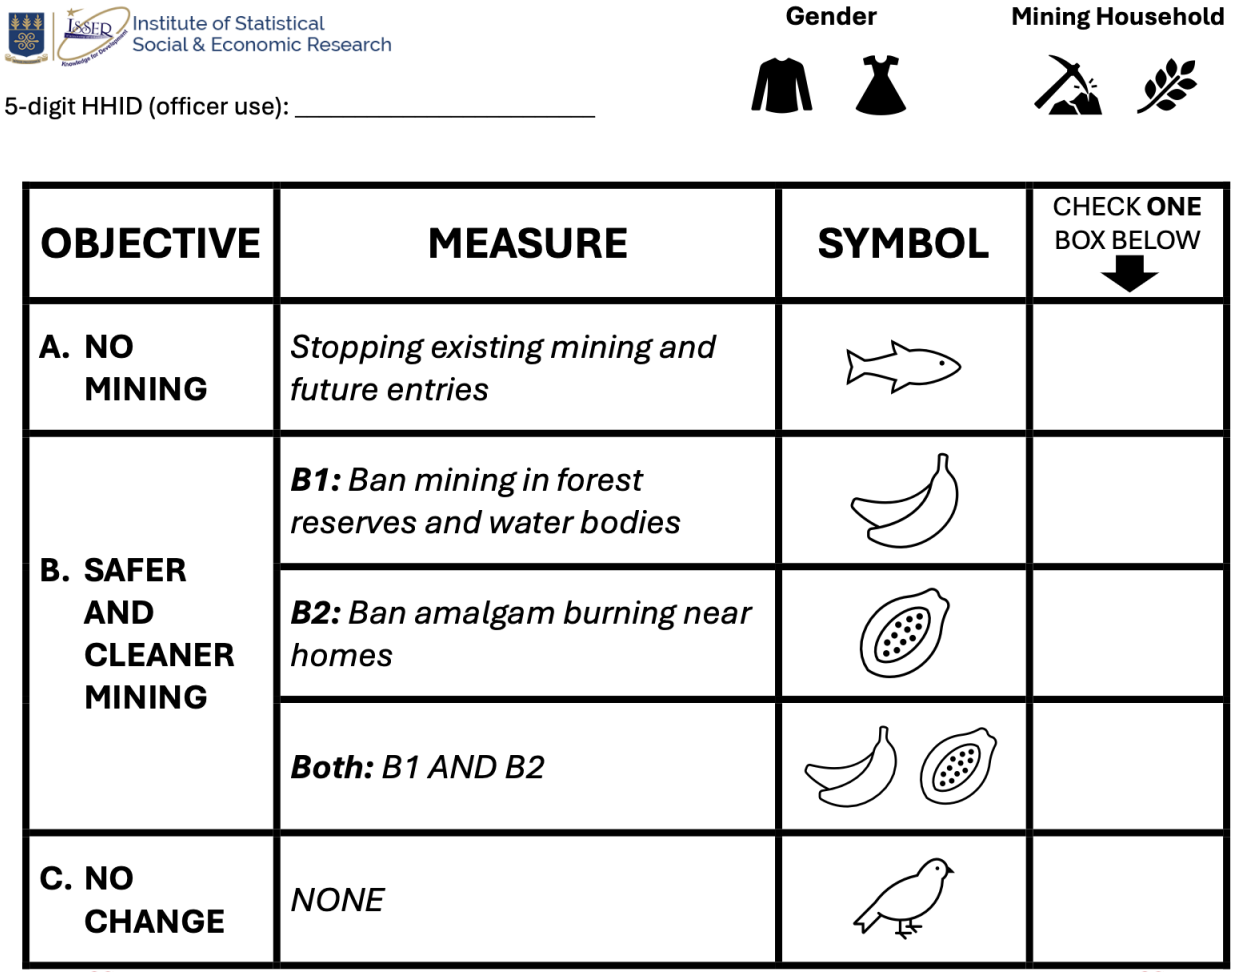

Figure 1.A4: Sample Second Ballot on Commitment to Sustain Collective Action

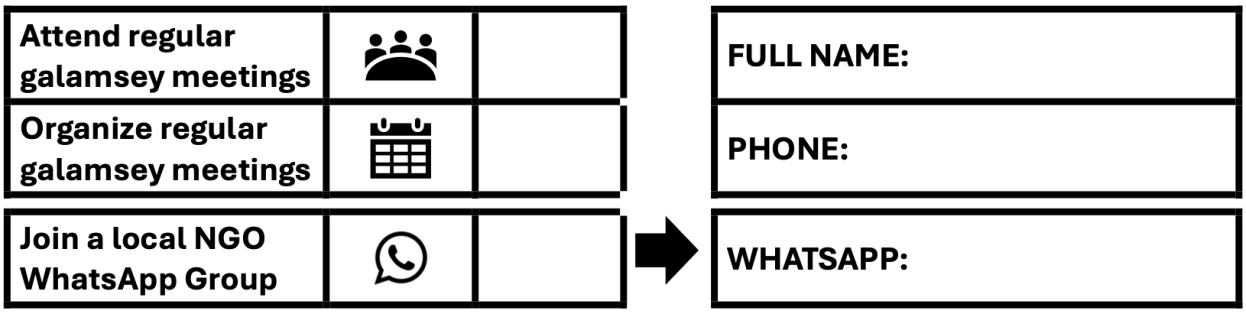


| KEEP (GHS) | 0 | 10 | 20 | 30 | 40 | 50 |
|---|---|---|---|---|---|---|
| GIVE (GHS) | 50 | 40 | 30 | 20 | 10 | 0 |
| CHECK **ONE** BOX ON THE RIGHT | | | | | | |

Institute of Statistical Social & Economic Research

5-digit HHID (officer use): _____

**Gender**

**Mining Household**

Figure 1.A5: Treatment effects on Health Learning by Town Hall Attendance

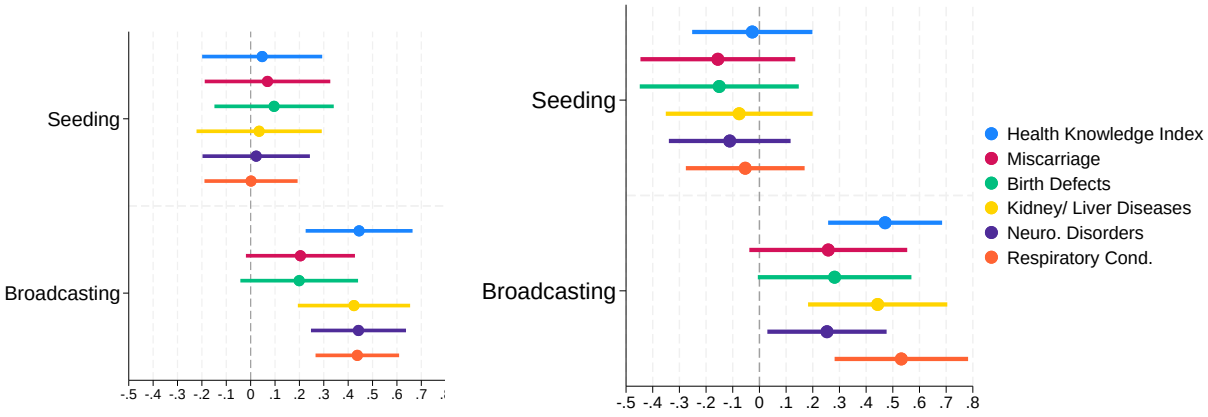


(a) Attendees (b) Non-attendees

**Notes:** Panels report intention-to-treat estimates of treatment effects on the health knowledge index and its components, separately for town hall attendees (left) and non-attendees (right). Points denote coefficient estimates and horizontal bars indicate 95% confidence intervals. Standard errors are clustered at the community level.

Figure 1.A6: Average Treatment Effects on Chiefs' Reasons to Take Part in Collective Action

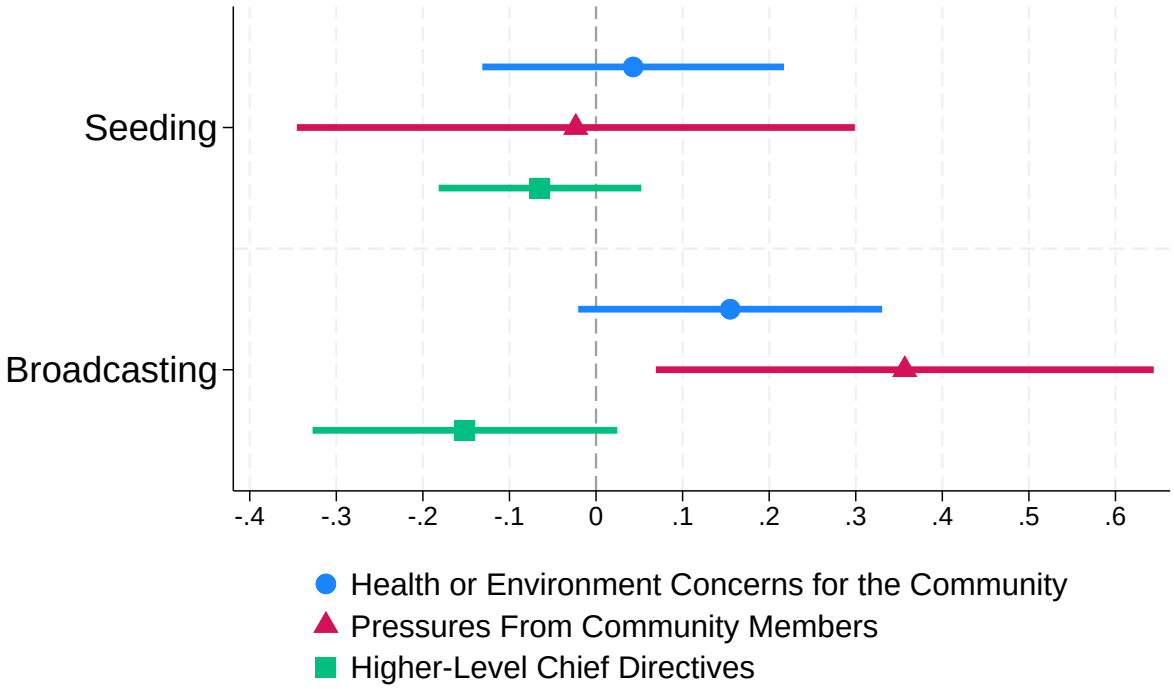


**Notes:** The unit of observation is a chief. All specifications district fixed effects, Conflicted Chief (above-median share of households perceiving the chief would accept money to allow *galamsey*), and High Mining Community (above-median share of mining households at baseline). Points denote coefficient estimates and bars indicate 95% confidence intervals. Standard errors are robust.

# Appendix Tables

Table 1.A1: Heterogeneous Treatment Effects on Community Members' Learning

| | Health Knowledge Index | Miscarriage | Birth Defects | Kidney/ Liver Diseases | Neuro. Disorder | Respiratory Cond. |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $\beta_1$ : Seeding | -0.049 | -0.065 | -0.031 | -0.007 | -0.109 | -0.109 |
| | (0.119) | (0.204) | (0.195) | (0.210) | (0.187) | (0.154) |
| $\beta_2$ : Broadcasting | 0.303*** | 0.288 | 0.323* | 0.535*** | 0.415** | 0.502*** |
| | (0.109) | (0.185) | (0.176) | (0.176) | (0.180) | (0.169) |
| $\gamma_1$ : Seeding × Conflicted Chief | 0.045 | 0.019 | 0.022 | -0.040 | 0.111 | 0.146 |
| | (0.187) | (0.344) | (0.327) | (0.311) | (0.287) | (0.239) |
| $\gamma_2$ : Broadcasting × Conflicted Chief | 0.156 | 0.179 | 0.146 | 0.164 | 0.200 | 0.342 |
| | (0.164) | (0.311) | (0.293) | (0.274) | (0.245) | (0.233) |
| Conflicted Chief | -0.126 | -0.204 | -0.125 | -0.176 | -0.163 | -0.212 |
| | (0.131) | (0.252) | (0.227) | (0.228) | (0.192) | (0.165) |
| Ctrl Mean | -0.085 | 3.599 | 3.601 | 3.614 | 3.371 | 3.545 |
| Ctrl SD | 0.817 | 1.495 | 1.518 | 1.403 | 1.476 | 1.418 |
| Obs. | 1367 | 1367 | 1367 | 1367 | 1367 | 1367 |
| R-squared | 0.154 | 0.112 | 0.123 | 0.140 | 0.096 | 0.090 |
| $\beta_1 + \gamma_1\,(est.)$ | -0.004 | -0.046 | -0.009 | -0.048 | 0.002 | 0.038 |
| $\beta_1 + \gamma_1 = 0\,(p\,value)$ | 0.975 | 0.868 | 0.974 | 0.840 | 0.992 | 0.830 |
| $\beta_2 + \gamma_2\,(est.)$ | 0.459 | 0.467 | 0.469 | 0.699 | 0.614 | 0.844 |
| $\beta_2 + \gamma_2 = 0\,(p\,value)$ | 0.001 | 0.085 | 0.083 | 0.003 | 0.001 | 0.000 |

**Notes:** The unit of observation is a household. The dependent variables include a standardized health knowledge index (Column 1) and Likert scales from 1 to 5 for recognizing specific long-term health risks associated with *galamsey* (Columns 2–6). Conflicted Chief is an indicator for communities where the share of households perceiving the chief would accept money to allow *galamsey* is above the sample median. All regressions include district fixed effects, voter gender, and Mining Household, and interactions between treatment indicators and the demeaned high-mining-community indicator (above-median share of mining households at baseline). Standard errors are clustered at the community level. The bottom panel reports estimated linear combinations and corresponding $p$-values for treatment effects in conflicted-chief communities, specifically $\beta_1 + \gamma_1$ for seeding and $\beta_2 + \gamma_2$ for broadcasting.

Table 1.A2: Voter Turnout at Town Halls by Treatment and Leader Type

| | All Households | | Mining Households | | Non-Mining Households | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $\beta_1$ : Seeding | 3.0 | -4.5 | -3.5 | -4.0 | 6.5 | -0.5 |
| | (4.2) | (6.5) | (3.0) | (5.0) | (4.0) | (6.9) |
| $\beta_2$ : Broadcasting | -9.8*** | -16.1*** | -8.3*** | -10.0** | -1.5 | -6.0 |
| | (3.1) | (5.2) | (2.5) | (4.4) | (3.2) | (5.4) |
| $\gamma_1$ : Seeding × Conflicted Chief | | 14.4 | | 0.4 | | 14.1* |
| | | (8.9) | | (5.9) | | (8.4) |
| $\gamma_2$ : Broadcasting × Conflicted Chief | | 11.6* | | 2.9 | | 8.7 |
| | | (6.9) | | (5.1) | | (7.3) |
| Conflicted Chief | -4.6 | -13.6** | -0.8 | -2.2 | -3.7 | -11.3** |
| | (3.5) | (5.7) | (2.0) | (4.5) | (3.5) | (5.4) |
| Ctrl Mean | 51.7 | 51.7 | 23.3 | 23.3 | 28.4 | 28.4 |
| Ctrl SD | 14.2 | 14.2 | 15.1 | 15.1 | 18.7 | 18.7 |
| Obs. | 95 | 95 | 95 | 95 | 95 | 95 |
| R-squared | 0.298 | 0.344 | 0.475 | 0.504 | 0.497 | 0.527 |
| $\beta_1 + \gamma_1\,(est.)$ | | 10.0 | | -3.6 | | 13.6 |
| $\beta_1 + \gamma_1 = 0\,(p\,value)$ | | 0.074 | | 0.265 | | 0.003 |
| $\beta_2 + \gamma_2\,(est.)$ | | -4.5 | | -7.1 | | 2.6 |
| $\beta_2 + \gamma_2 = 0\,(p\,value)$ | | 0.246 | | 0.006 | | 0.534 |

**Notes:** The unit of observation is a town hall. The dependent variable measures the number of ballots cast during community regulation voting. Columns 1–2 report overall turnout; Columns 3–4 restrict to mining households; and Columns 5–6 to non-mining households. *Conflicted Chief* is an indicator equal to one in communities where the share of households perceiving the chief would accept money to allow *galamsey* is above the sample median. All regressions include district fixed effects and interactions between treatment indicators and the demeaned high-mining-community indicator (equal to one for communities with an above-median share of mining households at baseline). Standard errors are robust. The bottom panel reports estimated linear combinations and corresponding $p$-values for treatment effects in conflicted-chief communities, specifically $\beta_1 + \gamma_1$ for seeding and $\beta_2 + \gamma_2$ for broadcasting.

Table 1.A3: Heterogeneous Treatment Effects on Community Mining Bylaw Measures

| | No. Measures | | | | |
|---|---|---|---|---|---|
| | Zero | One | Two | Three | No Change/ Full Ban |
| | (1) | (2) | (3) | (4) | (5) |
| $\beta_1$ : Seeding | -0.018 | -0.063 | 0.210*** | -0.129** | -0.147** |
| | (0.045) | (0.054) | (0.066) | (0.055) | (0.066) |
| $\beta_2$ : Broadcasting | -0.069* | -0.096** | 0.213*** | -0.048 | -0.117** |
| | (0.040) | (0.045) | (0.062) | (0.054) | (0.055) |
| $\gamma_1$ : Seeding × Conflicted Chief | 0.040 | 0.079 | -0.236** | 0.117 | 0.157* |
| | (0.053) | (0.069) | (0.090) | (0.085) | (0.090) |
| $\gamma_2$ : Broadcasting × Conflicted Chief | 0.114** | 0.022 | -0.311*** | 0.175** | 0.289*** |
| | (0.048) | (0.059) | (0.087) | (0.087) | (0.084) |
| Conflicted Chief | -0.072** | -0.026 | 0.197*** | -0.098 | -0.171** |
| | (0.033) | (0.046) | (0.060) | (0.060) | (0.065) |
| Ctrl Mean | 0.100 | 0.310 | 0.324 | 0.266 | 0.366 |
| Ctrl SD | 0.300 | 0.463 | 0.468 | 0.442 | 0.482 |
| Obs. | 4588 | 4588 | 4588 | 4588 | 4588 |
| R-squared | 0.067 | 0.048 | 0.044 | 0.155 | 0.080 |
| $\beta_1 + \gamma_1\,(est.)$ | 0.022 | 0.016 | -0.027 | -0.012 | 0.011 |
| $\beta_1 + \gamma_1 = 0\,(p\,value)$ | 0.420 | 0.724 | 0.692 | 0.867 | 0.873 |
| $\beta_2 + \gamma_2\,(est.)$ | 0.045 | -0.074 | -0.098 | 0.126 | 0.172 |
| $\beta_2 + \gamma_2 = 0\,(p\,value)$ | 0.131 | 0.056 | 0.114 | 0.073 | 0.010 |

**Notes:** The unit of observation is a household ballot from the town hall vote on community mining regulation. The dependent variable in Columns 1–4 is an indicator for voting for bylaws with the corresponding number of regulatory measures: 0 = no change (option C), 1 = ban mining in forests/rivers or amalgam burning near homes (option B1 or B2), 2 = ban both (option B3), and 3 = ban all mining (option A). Column 5 captures the most lenient and most stringent options, taking the value one if a town hall attendee voted for either no change or ban all mining (option A or C). Conflicted Chief is an indicator for communities where the share of households perceiving the chief would accept money to allow *galamsey* is above the sample median. All regressions include district fixed effects, voter gender, and Mining Household, and interactions between treatment indicators and the demeaned high-mining-community indicator (above-median share of mining households at baseline). Standard errors are clustered at the community level. The bottom panel reports estimated linear combinations and corresponding $p$-values for treatment effects in conflicted-chief communities, specifically $\beta_1 + \gamma_1$ for seeding and $\beta_2 + \gamma_2$ for broadcasting.

Table 1.A4: Heterogeneous Treatment Effects on Chiefs' Persuasive Efforts

| | No. Measures | | | | | |
|---|---|---|---|---|---|---|
| | Zero | One | Two | Three | No Change /Full Ban | Stringent Measure |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $\beta_1$ : Seeding | -0.175* | -0.054 | 0.389** | -0.160 | -0.335** | 0.229** |
| | (0.088) | (0.069) | (0.149) | (0.135) | (0.148) | (0.102) |
| $\beta_2$ : Broadcasting | -0.123 | 0.165 | 0.097 | -0.138 | -0.262 | -0.041 |
| | (0.109) | (0.129) | (0.198) | (0.147) | (0.164) | (0.162) |
| $\gamma_1$ : Seeding × Conflicted Chief | 0.157* | 0.043 | -0.433* | 0.233 | 0.390* | -0.200 |
| | (0.085) | (0.136) | (0.224) | (0.200) | (0.205) | (0.149) |
| $\gamma_2$ : Broadcasting × Conflicted Chief | 0.531*** | -0.192 | -0.514* | 0.175 | 0.706*** | -0.339 |
| | (0.180) | (0.165) | (0.273) | (0.207) | (0.239) | (0.236) |
| Conflicted Chief | -0.160* | 0.047 | 0.357** | -0.244 | -0.404** | 0.112 |
| | (0.091) | (0.109) | (0.174) | (0.162) | (0.163) | (0.133) |
| Ctrl Mean | 0.097 | 0.065 | 0.645 | 0.194 | 0.290 | 0.839 |
| Ctrl SD | 0.301 | 0.250 | 0.486 | 0.402 | 0.461 | 0.374 |
| Obs. | 94 | 94 | 94 | 94 | 94 | 94 |
| R-squared | 0.301 | 0.119 | 0.232 | 0.169 | 0.266 | 0.244 |
| $\beta_1 + \gamma_1$ $(est.)$ | -0.018 | -0.011 | -0.044 | 0.073 | 0.055 | 0.029 |
| $\beta_1 + \gamma_1 = 0$ $(p\,value)$ | 0.512 | 0.918 | 0.787 | 0.599 | 0.694 | 0.785 |
| $\beta_2 + \gamma_2$ $(est.)$ | 0.407 | -0.027 | -0.417 | 0.037 | 0.445 | -0.380 |
| $\beta_2 + \gamma_2 = 0$ $(p\,value)$ | 0.005 | 0.781 | 0.018 | 0.750 | 0.007 | 0.024 |

**Notes:** The unit of observation is a chief's (or delegate's) ballot from the town hall vote on community mining regulation. The dependent variable in Columns 1–4 is an indicator for voting for bylaws with the corresponding number of regulatory measures: 0 = no change (option C), 1 = ban mining in forests/rivers or amalgam burning near homes (option B1 or B2), 2 = ban both (option B3), and 3 = ban all mining (option A). Column 5 captures the most lenient and most stringent options, taking the value one if the chief voted for either no change or ban all mining (option A or C). Column 6 measures whether the chief voted for a stringent bylaw option of banning all mining or both mining in forest/river areas and amalgam burning near homes (option A or B3). Conflicted Chief is an indicator for communities where the share of households perceiving the chief would accept money to allow *galamsey* is above the sample median. All regressions include district fixed effects and interactions between treatment indicators and the demeaned high-mining-community indicator (equal to one for communities with an above-median share of mining households at baseline). Robust standard errors are shown in parentheses. The bottom panel reports estimated linear combinations and corresponding $p$-values for treatment effects in conflicted-chief communities, specifically $\beta_1 + \gamma_1$ for seeding and $\beta_2 + \gamma_2$ for broadcasting.

Table 1.A5: Heterogeneous Treatment Effects on Chiefs' Mobilization Efforts

| | Encourage Regular Meetings | Provide Meeting Logistics | Encourage Local WhatsApp Group | Support NGO Workshop | Bid NGO Workshop (GHS) |
|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) |
| $\beta_1$ : Seeding | 0.148 | 0.225* | 0.174 | 0.023 | 0.558 |
| | (0.173) | (0.123) | (0.173) | (0.127) | (4.760) |
| $\beta_2$ : Broadcasting | 0.004 | -0.026 | -0.221 | -0.023 | 4.204 |
| | (0.191) | (0.079) | (0.188) | (0.162) | (5.762) |
| $\gamma_1$ : Seeding × Conflicted Chief | -0.244 | -0.043 | -0.121 | 0.058 | 5.327 |
| | (0.225) | (0.209) | (0.249) | (0.184) | (6.689) |
| $\gamma_2$ : Broadcasting × Conflicted Chief | -0.539** | -0.294* | -0.491* | -0.393 | -5.362 |
| | (0.255) | (0.159) | (0.252) | (0.247) | (9.091) |
| Conflicted Chief | 0.197 | 0.112 | 0.166 | -0.041 | -3.108 |
| | (0.157) | (0.148) | (0.204) | (0.164) | (5.433) |
| Ctrl Mean | 0.774 | 0.161 | 0.613 | 0.742 | 34.839 |
| Ctrl SD | 0.425 | 0.374 | 0.495 | 0.445 | 16.906 |
| Obs. | 94 | 94 | 94 | 94 | 94 |
| R-squared | 0.152 | 0.337 | 0.352 | 0.241 | 0.274 |
| $\beta_1 + \gamma_1\,(est.)$ | -0.096 | 0.182 | 0.054 | 0.082 | 5.885 |
| $\beta_1 + \gamma_1 = 0\,(p\,value)$ | 0.473 | 0.296 | 0.745 | 0.556 | 0.211 |
| $\beta_2 + \gamma_2\,(est.)$ | -0.535 | -0.319 | -0.712 | -0.416 | -1.158 |
| $\beta_2 + \gamma_2 = 0\,(p\,value)$ | 0.001 | 0.020 | 0.000 | 0.023 | 0.870 |

**Notes:** The unit of observation is a town hall. Columns 1–4 are indicators coded from the speech checklist for whether the chief encouraged regular meetings (col. 1), offered meeting logistics (col. 2), encouraged joining the local WhatsApp group (col. 3), or expressed support for an NGO workshop (col. 4). Column 5 records the amount (GHS) the chief allocated in the lottery-based bidding game from a 50 GHS endowment to bring an NGO workshop to the community. Conflicted Chief is an indicator for communities where the share of households perceiving the chief would accept money to allow *galamsey* is above the sample median. All regressions include district fixed effects and interactions between treatment indicators and the demeaned high-mining-community indicator (equal to one for communities with an above-median share of mining households at baseline). Robust standard errors are shown in parentheses. The bottom panel reports estimated linear combinations and corresponding $p$-values for treatment effects in conflicted-chief communities, specifically $\beta_1 + \gamma_1$ for seeding and $\beta_2 + \gamma_2$ for broadcasting.

Table 1.A6: Effects on Community Mining Regulation Preferences (Split-Sample by Occupation)

| | Stringent Measure | | | | | |
|---|---|---|---|---|---|---|
| | All Households | | Mining Households | | Non-Mining Households | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $\beta_1$ : Seeding | 0.026 | 0.081 | 0.055 | 0.108 | 0.006 | 0.043 |
| | (0.043) | (0.061) | (0.056) | (0.074) | (0.046) | (0.070) |
| $\beta_2$ : Broadcasting | 0.100** | 0.165*** | 0.127** | 0.207** | 0.088** | 0.139*** |
| | (0.038) | (0.053) | (0.056) | (0.088) | (0.037) | (0.051) |
| $\gamma_1$ : Seeding × Conflicted Chief | | -0.119* | | -0.174* | | -0.062 |
| | | (0.069) | | (0.099) | | (0.075) |
| $\gamma_2$ : Broadcasting × Conflicted Chief | | -0.136** | | -0.197** | | -0.071 |
| | | (0.065) | | (0.099) | | (0.069) |
| Conflicted Chief | 0.013 | 0.098** | 0.031 | 0.145** | 0.002 | 0.053 |
| | (0.034) | (0.040) | (0.047) | (0.059) | (0.037) | (0.045) |
| Weighted | No | No | No | No | No | No |
| Ctrl Mean | 0.590 | 0.590 | 0.442 | 0.442 | 0.712 | 0.712 |
| Ctrl SD | 0.492 | 0.492 | 0.497 | 0.497 | 0.453 | 0.453 |
| Obs. | 4588 | 4588 | 1866 | 1866 | 2722 | 2722 |
| R-squared | 0.108 | 0.113 | 0.063 | 0.072 | 0.054 | 0.056 |
| $\beta_1 + \gamma_1\,(est.)$ | | -0.038 | | -0.067 | | -0.019 |
| $\beta_1 + \gamma_1 = 0\,(p\,value)$ | | 0.389 | | 0.417 | | 0.645 |
| $\beta_2 + \gamma_2\,(est.)$ | | 0.029 | | 0.010 | | 0.068 |
| $\beta_2 + \gamma_2 = 0\,(p\,value)$ | | 0.508 | | 0.886 | | 0.199 |

**Notes:** The unit of observation is a household ballot from the town hall vote on community mining regulation. The outcome is whether respondents voted for a stringent bylaw option (banning all mining or both mining in forest/river areas and amalgam burning near homes). Columns 1–2 report overall results; Columns 3–4 restrict to mining households; and Columns 5–6 to non-mining households. Conflicted Chief is an indicator for communities where the share of households perceiving the chief would accept money to allow *galamsey* is above the sample median. All regressions include district fixed effects, voter gender, and interactions between treatment indicators and the demeaned high-mining-community indicator (above-median share of mining households at baseline). Standard errors are clustered at the community level. Columns 1-2 also control for Mining Household. The bottom panel reports estimated linear combinations and corresponding $p$-values for treatment effects in conflicted-chief communities, specifically $\beta_1 + \gamma_1$ for seeding and $\beta_2 + \gamma_2$ for broadcasting.

Table 1.A7: Effects on Community Mining Regulation Preferences (Weighted)

| | Stringent Measure | | Polarization | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| $\beta_1$ : Seeding | 0.026 | 0.080 | 0.005 | -0.048 |
| | (0.043) | (0.060) | (0.043) | (0.062) |
| $\beta_2$ : Broadcasting | 0.101** | 0.165*** | 0.055 | -0.016 |
| | (0.039) | (0.055) | (0.039) | (0.048) |
| $\gamma_1$ : Seeding × Conflicted Chief | | -0.118* | | 0.116 |
| | | (0.069) | | (0.081) |
| $\gamma_2$ : Broadcasting × Conflicted Chief | | -0.136** | | 0.156** |
| | | (0.066) | | (0.076) |
| Conflicted Chief | 0.013 | 0.098** | 0.007 | -0.080 |
| | (0.035) | (0.040) | (0.043) | (0.057) |
| Weighted | Yes | Yes | Yes | Yes |
| Ctrl Mean | 0.590 | 0.590 | 0.643 | 0.643 |
| Ctrl SD | 0.492 | 0.492 | 0.471 | 0.471 |
| Obs. | 4588 | 4588 | 4588 | 4588 |
| R-squared | 0.107 | 0.112 | 0.026 | 0.031 |
| $\beta_1 + \gamma_1\,(est.)$ | | -0.038 | | 0.068 |
| $\beta_1 + \gamma_1 = 0\,(p\,value)$ | | 0.409 | | 0.219 |
| $\beta_2 + \gamma_2\,(est.)$ | | 0.029 | | 0.140 |
| $\beta_2 + \gamma_2 = 0\,(p\,value)$ | | 0.499 | | 0.020 |

**Notes:** The unit of observation is a household ballot from the town hall vote on community mining regulation. Observations in treatment communities are weighted to match the share of mining households in control communities. Columns 1–2 measure whether respondents voted for a stringent bylaw option (banning all mining or both mining in forest/river areas and amalgam burning near homes). Columns 3–4 measure preference dispersion, defined as the absolute distance between an individual's preferred bylaw stringency and the community mean. Conflicted Chief is an indicator for communities where the share of households perceiving the chief would accept money to allow *galamsey* is above the sample median. All regressions include district fixed effects, voter gender, and interactions between treatment indicators and the demeaned high-mining-community indicator (above-median share of mining households at baseline). Standard errors are clustered at the community level. The bottom panel reports estimated linear combinations and corresponding $p$-values for treatment effects in conflicted-chief communities, specifically $\beta_1 + \gamma_1$ for seeding and $\beta_2 + \gamma_2$ for broadcasting.

Table 1.A8: Effects on Community Mining Regulation Preferences (Robust to Misclassification)

| | Stringent Measure | | Polarization | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| $\beta_1$ : Seeding | 0.019 | 0.130 | 0.011 | -0.034 |
| | (0.060) | (0.081) | (0.052) | (0.078) |
| $\beta_2$ : Broadcasting | 0.108** | 0.217*** | 0.094* | 0.014 |
| | (0.053) | (0.073) | (0.049) | (0.069) |
| $\gamma_1$ : Seeding × Conflicted Chief | | -0.223** | | 0.085 |
| | | (0.103) | | (0.100) |
| $\gamma_2$ : Broadcasting × Conflicted Chief | | -0.218** | | 0.186** |
| | | (0.091) | | (0.087) |
| Conflicted Chief | 0.010 | 0.171** | -0.028 | -0.128* |
| | (0.053) | (0.066) | (0.056) | (0.074) |
| Ctrl Mean | 0.633 | 0.633 | 0.591 | 0.591 |
| Ctrl SD | 0.482 | 0.482 | 0.473 | 0.473 |
| Obs. | 2948 | 2948 | 2948 | 2948 |
| R-squared | 0.111 | 0.120 | 0.026 | 0.032 |
| $\beta_1 + \gamma_1\,(est.)$ | | -0.093 | | 0.051 |
| $\beta_1 + \gamma_1 = 0\,(p\,value)$ | | 0.164 | | 0.422 |
| $\beta_2 + \gamma_2\,(est.)$ | | -0.002 | | 0.199 |
| $\beta_2 + \gamma_2 = 0\,(p\,value)$ | | 0.980 | | 0.002 |

**Notes:** The unit of observation is a household ballot from the town hall vote on community mining regulation. The regression specifications are the same as Table 1.2, with the only change being the exclusion of communities in the middle tercile of the conflict-of-interest distribution. Columns 1–2 measure whether participants voted for a stringent bylaw option (banning all mining or both mining in forest/river areas and amalgam burning near homes). Columns 3–4 measure preference dispersion, defined as the absolute distance between an individual's preferred bylaw stringency and the community mean. Conflicted Chief is an indicator for communities where the share of households perceiving the chief would accept money to allow *galamsey* is above the sample median. All regressions include district fixed effects, voter gender, and interactions between treatment indicators and the demeaned high-mining-community indicator (above-median share of mining households at baseline). Standard errors are clustered at the community level. The bottom panel reports estimated linear combinations and corresponding $p$-values for treatment effects in conflicted-chief communities, specifically $\beta_1 + \gamma_1$ for seeding and $\beta_2 + \gamma_2$ for broadcasting.

Table 1.A9: Heterogeneous Treatment Effects on Commitment to Sustain Collective Action

| | Regular Meeting Index (1) | Attend Regular Meetings (2) | Organize Regular Meetings (3) | Join NGO WhatsApp Group (4) | Bid NGO Workshop (GHS) (5) |
|---|---|---|---|---|---|
| $\beta_1$ : Seeding | -0.278* | -0.133** | -0.054 | -0.139*** | -1.450 |
| | (0.140) | (0.062) | (0.057) | (0.045) | (1.869) |
| $\beta_2$ : Broadcasting | 0.397*** | 0.187*** | 0.081 | -0.038 | -0.757 |
| | (0.118) | (0.052) | (0.053) | (0.046) | (2.733) |
| $\gamma_1$ : Seeding × Conflicted Chief | 0.317 | 0.085 | 0.127 | 0.208*** | 4.651 |
| | (0.197) | (0.079) | (0.083) | (0.072) | (2.872) |
| $\gamma_2$ : Broadcasting × Conflicted Chief | -0.098 | -0.096 | 0.028 | 0.155** | -2.934 |
| | (0.181) | (0.072) | (0.078) | (0.064) | (4.337) |
| Conflicted Chief | 0.027 | 0.058 | -0.038 | -0.108** | 1.146 |
| | (0.144) | (0.057) | (0.063) | (0.049) | (2.638) |
| Ctrl Mean | 0.000 | 0.663 | 0.300 | 0.301 | 27.058 |
| Ctrl SD | 1.000 | 0.473 | 0.458 | 0.459 | 17.651 |
| Obs. | 4505 | 4505 | 4505 | 4505 | 4505 |
| R-squared | 0.129 | 0.119 | 0.048 | 0.070 | 0.124 |
| $\beta_1 + \gamma_1\,(est.)$ | 0.039 | -0.048 | 0.073 | 0.069 | 3.201 |
| $\beta_1 + \gamma_1 = 0\,(p\,value)$ | 0.773 | 0.339 | 0.252 | 0.207 | 0.174 |
| $\beta_2 + \gamma_2\,(est.)$ | 0.299 | 0.091 | 0.109 | 0.116 | -3.692 |
| $\beta_2 + \gamma_2 = 0\,(p\,value)$ | 0.046 | 0.110 | 0.067 | 0.015 | 0.258 |

**Notes:** The unit of observation is a household ballot from the town hall activities measuring revealed commitment to sustained collective action. Columns 1–3 capture time commitments: Column 1 reports a standardized index combining willingness to attend and organize regular community meetings on *galamsey*; Columns 2 and 3 report these components separately. Column 4 measures participants' enrollment in a local NGO-moderated WhatsApp group. Column 5 records participants' financial contributions in a crowd-bidding activity, where individuals allocated part of a potential 50 GHS prize between themselves and a communal fund to host an NGO workshop. Conflicted Chief is an indicator for communities where the share of households perceiving the chief would accept money to allow *galamsey* is above the sample median. All regressions include district fixed effects, voter gender, and Mining Household, and interactions between treatment indicators and the demeaned high-mining-community indicator (above-median share of mining households at baseline). Standard errors are clustered at the community level. The bottom panel reports estimated linear combinations and corresponding $p$-values for treatment effects in conflicted-chief communities, specifically $\beta_1 + \gamma_1$ for seeding and $\beta_2 + \gamma_2$ for broadcasting.

Table 1.A10: Heterogeneous Treatment Effects on Post-Town Hall Discussion

| | Freq. Discuss Galamsey w/ Other | | |
|---|---|---|---|
| | Non-mining Community Members (1) | Mining Community Members (2) | Chief Organized Gathering (3) |
| $\beta_1$ : Seeding | -0.458*** | -0.716** | -0.090 |
| | (0.163) | (0.316) | (0.056) |
| $\beta_2$ : Broadcasting | 0.441** | 0.595 | 0.059 |
| | (0.207) | (0.384) | (0.086) |
| $\gamma_1$ : Seeding × Conflicted Chief | 0.353 | 0.735 | 0.019 |
| | (0.274) | (0.474) | (0.075) |
| $\gamma_2$ : Broadcasting × Conflicted Chief | 0.555 | 0.627 | -0.018 |
| | (0.356) | (0.585) | (0.109) |
| Conflicted Chief | 0.002 | -0.064 | -0.025 |
| | (0.209) | (0.389) | (0.073) |
| Ctrl Mean | 1.070 | 1.460 | 0.136 |
| Ctrl SD | 1.800 | 2.448 | 0.343 |
| Obs. | 1366 | 1366 | 1368 |
| R-squared | 0.081 | 0.154 | 0.041 |
| $\beta_1 + \gamma_1\,(est.)$ | -0.105 | 0.019 | -0.071 |
| $\beta_1 + \gamma_1 = 0\,(p\,value)$ | 0.644 | 0.960 | 0.202 |
| $\beta_2 + \gamma_2\,(est.)$ | 0.996 | 1.222 | 0.040 |
| $\beta_2 + \gamma_2 = 0\,(p\,value)$ | 0.001 | 0.007 | 0.583 |

**Notes:** The unit of observation is a household. Columns 1 and 2 measure the frequency with which respondents report discussing *galamsey* with others after the town hall, separately for non-mining and mining community members. Column 3 reports whether the chief organized a follow-up community gathering on *galamsey*. All regressions include district fixed effects, and interactions between treatment indicators and the demeaned high-mining-community indicator (equal to one for communities with an above-median share of mining households at baseline). Standard errors are clustered at the community level. The bottom panel reports estimated linear combinations and corresponding $p$-values for treatment effects in conflicted-chief communities, specifically $\beta_1 + \gamma_1$ for seeding and $\beta_2 + \gamma_2$ for broadcasting.

Table 1.A11: Heterogeneous Treatment Effects on Pessimism About Collective Action

| | Perceived or Own Willingness | | | | | |
|---|---|---|---|---|---|---|
| | 3 Closest Friends & Neigh. | 10 Other Non-min. HHs | 10 Other Mining HHs | Chief | Self | Perceived Collective Action Efficacy |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $\beta_1$ : Seeding | 0.010 | -0.137 | -0.518** | 0.041 | -0.002 | -0.028 |
| | (0.081) | (0.240) | (0.250) | (0.034) | (0.039) | (0.058) |
| $\beta_2$ : Broadcasting | -0.438*** | -1.266*** | -0.544 | -0.106* | -0.137*** | -0.134** |
| | (0.126) | (0.300) | (0.469) | (0.054) | (0.050) | (0.060) |
| $\gamma_1$ : Seeding × Conflicted Chief | -0.009 | 0.076 | 0.061 | -0.057 | -0.051 | 0.059 |
| | (0.124) | (0.288) | (0.380) | (0.069) | (0.048) | (0.075) |
| $\gamma_2$ : Broadcasting × Conflicted Chief | 0.129 | 0.051 | -0.089 | -0.038 | 0.035 | 0.029 |
| | (0.161) | (0.362) | (0.553) | (0.079) | (0.059) | (0.080) |
| Conflicted Chief | -0.107 | -0.355 | -0.070 | -0.009 | -0.004 | -0.076 |
| | (0.100) | (0.217) | (0.303) | (0.050) | (0.031) | (0.049) |
| Ctrl Mean | 2.310 | 7.838 | 4.776 | 0.791 | 0.916 | 0.603 |
| Ctrl SD | 0.782 | 1.886 | 2.668 | 0.407 | 0.278 | 0.490 |
| Obs. | 1366 | 1366 | 1366 | 1368 | 1368 | 1368 |
| R-squared | 0.072 | 0.109 | 0.074 | 0.058 | 0.047 | 0.034 |
| $\beta_1 + \gamma_1\,(est.)$ | 0.002 | -0.062 | -0.456 | -0.016 | -0.053 | 0.030 |
| $\beta_1 + \gamma_1 = 0\,(p\,value)$ | 0.987 | 0.710 | 0.128 | 0.766 | 0.093 | 0.548 |
| $\beta_2 + \gamma_2\,(est.)$ | -0.309 | -1.215 | -0.633 | -0.144 | -0.102 | -0.105 |
| $\beta_2 + \gamma_2 = 0\,(p\,value)$ | 0.003 | 0.000 | 0.058 | 0.012 | 0.003 | 0.038 |

**Notes:** The unit of observation is a household. Column (1) reports responses on whether community collective action is perceived as more effective than government intervention. Columns (2)–(6) measure respondents' own willingness and their beliefs about others' willingness to participate in community efforts to address *galamsey*. Conflicted Chief is an indicator for communities where the share of households perceiving the chief would accept money to allow *galamsey* is above the sample median. All regressions include district fixed effects, and interactions between treatment indicators and the demeaned high-mining-community indicator (equal to one for communities with an above-median share of mining households at baseline). Standard errors are clustered at the community level. The bottom panel reports estimated linear combinations and corresponding $p$-values for treatment effects in conflicted-chief communities, specifically $\beta_1 + \gamma_1$ for seeding and $\beta_2 + \gamma_2$ for broadcasting.

Table 1.A12: Heterogeneous Treatment Effects on Attitudes Toward Other Community Members

| | Give To | | Expect From | |
|---|---|---|---|---|
| | Non-mining Households | Mining Households | Non-mining Households | Mining Households |
| | (1) | (2) | (3) | (4) |
| $\beta_1$ : Seeding | -1.0 | 3.6 | -2.3 | 0.8 |
| | (1.6) | (2.4) | (2.1) | (2.6) |
| $\beta_2$ : Broadcasting | 0.2 | 0.2 | 2.3 | -1.9 |
| | (1.7) | (2.5) | (2.3) | (2.7) |
| $\gamma_1$ : Seeding × Conflicted Chief | 2.5 | -2.5 | 4.3 | 2.1 |
| | (2.8) | (3.6) | (2.9) | (4.0) |
| $\gamma_2$ : Broadcasting × Conflicted Chief | -0.6 | 2.7 | 3.2 | 7.9* |
| | (2.9) | (3.7) | (3.0) | (4.2) |
| Conflicted Chief | -1.1 | -0.1 | -4.1* | -8.9*** |
| | (2.1) | (2.6) | (2.5) | (2.8) |
| Ctrl Mean | 41.9 | 17.8 | 26.4 | 37.6 |
| Ctrl SD | 19.7 | 21.3 | 17.6 | 32.5 |
| Obs. | 1366 | 1366 | 1366 | 1366 |
| R-squared | 0.029 | 0.112 | 0.041 | 0.125 |
| $\beta_1 + \gamma_1\,(est.)$ | 1.5 | 1.2 | 2.0 | 2.9 |
| $\beta_1 + \gamma_1 = 0\,(p\,value)$ | 0.492 | 0.652 | 0.308 | 0.274 |
| $\beta_2 + \gamma_2\,(est.)$ | -0.5 | 2.9 | 5.6 | 6.0 |
| $\beta_2 + \gamma_2 = 0\,(p\,value)$ | 0.837 | 0.289 | 0.004 | 0.045 |

**Notes:** The unit of observation is a household. Column (1)-(2) how much of 100 GHS endowment respondents were willing to give to a randomly selected non-mining and mining households from their community respectively. Column (3)-(4) how much of 100 GHS endowment respondents expected to receive from a randomly selected non-mining and mining households from their community respectively. Conflicted Chief is an indicator for communities where the share of households perceiving the chief would accept money to allow *galamsey* is above the sample median. All regressions include district fixed effects, and interactions between treatment indicators and the demeaned high-mining-community indicator (equal to one for communities with an above-median share of mining households at baseline). Standard errors are clustered at the community level. The bottom panel reports estimated linear combinations and corresponding $p$-values for treatment effects in conflicted-chief communities, specifically $\beta_1 + \gamma_1$ for seeding and $\beta_2 + \gamma_2$ for broadcasting.

Table 1.A13: Heterogeneous Treatment Effects on National Civic Action

| | Petition Letter | NGO Donation (GHS) |
|---|---|---|
| | (1) | (2) |
| $\beta_1$ : Seeding | 0.063 | 0.214 |
| | (0.050) | (0.456) |
| $\beta_2$ : Broadcasting | 0.076 | 0.395 |
| | (0.053) | (0.538) |
| $\gamma_1$ : Seeding × Conflicted Chief | -0.083 | -0.693 |
| | (0.080) | (0.724) |
| $\gamma_2$ : Broadcasting × Conflicted Chief | -0.069 | -0.741 |
| | (0.075) | (0.775) |
| Conflicted Chief | 0.067 | 0.093 |
| | (0.053) | (0.541) |
| Ctrl Mean | 0.787 | 6.675 |
| Ctrl SD | 0.410 | 5.394 |
| Obs. | 1368 | 1366 |
| R-squared | 0.024 | 0.015 |
| $\beta_1 + \gamma_1\,(est.)$ | -0.020 | -0.479 |
| $\beta_1 + \gamma_1 = 0\,(p\,value)$ | 0.745 | 0.349 |
| $\beta_2 + \gamma_2\,(est.)$ | 0.007 | -0.346 |
| $\beta_2 + \gamma_2 = 0\,(p\,value)$ | 0.906 | 0.508 |

**Notes:** The unit of observation is a household. Column (1) captures signing a petition opposing a pro-mining legislative instrument. Column (2) captures how much of a 20 GHS survey incentive respondents wished to donate to an NGO supporting grassroots anti-*galamsey* activism. Conflicted Chief is an indicator for communities where the share of households perceiving the chief would accept money to allow *galamsey* is above the sample median. All regressions include district fixed effects, and interactions between treatment indicators and the demeaned high-mining-community indicator (equal to one for communities with an above-median share of mining households at baseline). Standard errors are clustered at the community level. The bottom panel reports estimated linear combinations and corresponding $p$-values for treatment effects in conflicted-chief communities, specifically $\beta_1 + \gamma_1$ for seeding and $\beta_2 + \gamma_2$ for broadcasting.

# Chapter 2

# Political Economy of Information Disclosure: Evidence from Indian Rural Road Quality Audits

## 2.1 Introduction

Poor governance is widespread in low-income democratic countries (UNDP, 2002). Two dominant theories explain why democratic accountability often functions poorly in these settings. One emphasizes institutions: a country's institutional environment jointly shapes governance and political development (Acemoglu et al., 2001; Acemoglu, 2010). In this view, poor governance follows from weak institutions. Elections held under weak institutional constraints are vulnerable to electoral malfeasance, from fraud to elite capture, while institutions themselves are shaped by historical events that generate institutional reforms.

The second theory emphasizes voter information and political participation. A well-functioning democracy requires voters to be sufficiently informed and politically active to hold politicians accountable (Przeworski et al., eds, 1999; Besley, 2007). This view is more directly actionable because it points to clear policy levers. Empirically, it is supported by a growing literature that uses exogenous variation in information disclosure to study electoral accountability (Ferraz and Finan, 2008; Banerjee et al., 2011; Bobonis et al., 2016; Fisman et al., 2019). In practice, however, independent and impartial information disclosure systems are the exception rather than the norm, especially in low-capacity states.

This paper brings these two perspectives together by studying how political dynamics shape an information disclosure system itself. I show that the ruling party can use an audit system designed to inform voters about public good provision as a political tool, benefiting aligned

local politicians and disadvantaging non-aligned ones. The setting is a nationwide rural road construction program in India with a multi-tier quality monitoring system intended to ensure that completed roads meet program standards. State-level audits are managed by a state agency headed by a formally politically neutral civil service bureaucrat, but such bureaucrats may be vulnerable to pressure from the state ruling party because of career concerns. National-level audits, by contrast, are coordinated centrally and carried out by external engineers, making them more insulated from local political influence.

I estimate how road quality audit reporting changes when a constituency is represented by a politician aligned with the state ruling party. I refer to these constituencies as aligned constituencies.[1] The empirical design uses more than 7,000 state-level road quality audits conducted within one year before and after the 2015 Bihar election. The design differs from a standard difference-in-differences (DID) setting because there is no clearly defined control group during the pre-election period: the state ruling party has incentives to treat aligned constituencies favorably and non-aligned constituencies unfavorably. I define the positive treatment group as constituencies represented by a politician aligned with the state ruling party before the election and the negative treatment group as constituencies represented by a non-aligned politician. This classification is fixed by pre-election party affiliation, even though alignment after the election may change because of turnover in the state ruling party or in local politicians. The one-year post-election period, when the next election is remote and alignment may have changed, serves as the control period. I show that, under the standard parallel trends assumption, the interaction coefficient in a standard DID regression identifies the differential quality reporting effect by alignment status.[2] This approach requires that neither treatment group is treated during the control period, but it does not require the positive and negative treatment groups to have the same audit scores in that period.

I find no differential extensive-margin effect of alignment: roads in aligned constituencies are not differentially more likely to be audited. The intensive margin, however, is economically meaningful. I estimate a 7.4 percentage point reported-quality differential by alignment status, relative to a sample average passing rate of 64.9%. The same pattern does not appear when I estimate the same specification using national quality audits. This contrast supports the interpretation that the state audit pattern reflects audit report manipulation rather than greater effort by aligned politicians or greater resource allocation to aligned constituencies. I then provide evidence on two potential mechanisms: auditor assignment and audit timing. To study auditor assignment, I use the auditor identifier associated with each quality audit to estimate auditor fixed effects, which I interpret as auditor leniency. The fixed effects are

[1] Blocks and constituencies are administrative and political regions within districts. Their boundaries might overlap. Due to data constraints, I consider the set of blocks that are entirely aligned or entirely non-aligned. Hence, blocks and constituencies are used interchangeably. The matching procedure is detailed in section 3.3.

[2] It is effectively the cross-alignment difference in the pre-post differences.

estimated using a sample that extends beyond the one-year window around the election, reducing the concern that they only reflect the endogenous assignment of auditors to roads of different quality during the electoral period. Consistent with the information manipulation interpretation, roads in aligned constituencies are differentially assigned more lenient auditors before the election. On audit timing, roads that eventually fail state audits are more likely to be inspected before the election in non-aligned constituencies than in aligned constituencies. This pattern is consistent with the state ruling party timing inspections to delay unfavorable information in aligned constituencies until after the election.

Although this paper does not take a stance on whether voters recognize manipulation, the information manipulation hypothesis requires audit outcomes to be electorally relevant. I therefore provide suggestive evidence that reported quality matters for voter behavior: in aligned blocks, the state ruling party is more likely to win when reported audit outcomes are better. This relationship appears for reports released before the election, but not after, which is consistent with audit information entering voters' electoral decisions.

The paper provides a cautionary perspective on empirical work showing that improved voter access to information enhances electoral accountability (Ferraz and Finan, 2008; Snyder and Strömberg, 2010; Banerjee et al., 2011; Bobonis et al., 2016; Fisman et al., 2019), reduces corruption, and improves public good provision (Olken, 2007; Djankov et al., 2010; Reinikka and Svensson, 2011; Avis et al., 2018; Banerjee et al., 2018, 2020). This literature typically exploits quasi-experimental or experimental variation in access to information. My setting-specific finding shows that audit programs themselves can be vulnerable to manipulation by local politicians. The closest related paper is Fisman and Wang (2017), which shows that local Chinese bureaucrats underreport workplace-related accidental deaths to satisfy promotion requirements set by higher-level government. My paper studies politically motivated misreporting in a democratic setting with weak local accountability and suggests that manipulation of official information is not limited to authoritarian regimes.

Relatedly, the paper contributes to work on policies designed to reduce misreporting. Duflo et al. (2013) finds that auditors chosen and paid by the firms they audit systematically underreport plant emissions, and shows that randomly assigning third-party auditors reduces both misreporting and actual emissions. Vannutelli (2022) finds that shifting control over municipal budget auditor assignments away from local politicians improves Italian municipalities' fiscal performance, demonstrating the importance of auditor independence. My paper differs from these settings because neither the auditors nor the politicians are private-sector actors.[3]

[3]In Duflo et al. (2013), both the auditors and the firms are private sector players. In Vannutelli (2022), the auditors are private practice Certified Public Accountants.

## 2.2 Institutional Background

### 2.2.1 PMGSY

According to the 2001 census, more than half of India's 600,000 rural villages were not connected to a road. Pradhan Mantri Gram Sadak Yojana (PMGSY), a nationwide program introduced in 2000 by then prime minister Atal Bihari Vajpayee, aims to provide paved "all-weather" rural roads to all previously unconnected villages. As of February 2021, 161,709 road projects had been completed under PMGSY, corresponding to 652,403 km of constructed roads. The program began as a 100% centrally sponsored scheme. From November 2015 onward, funding has been shared between the center and the states at a 60:40 ratio.[4]

The program's scale and geographic reach have attracted substantial interest from policymakers and researchers studying its welfare effects.[5] Rather than study another dimension of the program's welfare impact, this paper examines the internal operation of its quality monitoring system. In particular, I document politically shaped quality reporting. To study these political dynamics in detail, I focus on the eastern state of Bihar, the single largest PMGSY beneficiary, where 17,023 road works have been completed, connecting 53,678 km. Bihar is also one of the BIMARU states, an acronym for Bihar, Madhya Pradesh, Rajasthan, and Uttar Pradesh, which are often associated with corruption and state dysfunction.[6]

At the central level, the Ministry of Rural Development oversees PMGSY and has established the National Rural Roads Development Agency (NRRDA) to provide operational and management support to state-level implementation units.[7] PMGSY implementation is decentralized to the states. When the program was launched in 2000, each state government identified an existing agency with a record of executing road construction works to serve as the "Nodal Department," which assumes overall responsibility for state implementation. The Nodal Department then identifies a "State Rural Road Development Agency (SRRDA)" under its control as the agency dedicated to rural roads.[8] The Secretary in charge of the Nodal Department serves as the Chief Executive of the SRRDA. In Bihar, the Nodal Department is the Rural Works Department (RWD), and the SRRDA is the Bihar Rural Roads Development Agency (BRRDA). As with other state-level government departments in India, RWD is

[4]This took place during the sample period. However, due to the long time lag between road construction approval and quality inspection, none of the sample audits are carried out on roads approved under the new funding sharing scheme.

[5]Asher and Novosad (2020) shows that PMGSY rural road access has a small effect on the village economies relative to its cost. Adukia et al. (2020) found a positive effect on education attainment.

[6]The Hindi word "bimar" means "sick". Bihar has a population of 99 million and the lowest per capita income of INR 13,482 (USD 185) amongst 33 Indian states and union territories, one-third of the national average.

[7]Part II Section 14, PMGSY Programme Guidelines (April 2012)

[8]Part II Section 7, PMGSY Programme Guidelines (April 2012)

headed by a minister from the state ruling coalition. Below the minister is the Secretary of RWD, a formally politically neutral Indian Administrative Service (IAS) officer. The IAS is the elite layer of the government bureaucracy and has fewer than 5,000 members. Another IAS officer serves jointly as Additional Chief Executive of BRRDA and Additional Secretary of RWD.[9]

PMGSY has a three-tier quality monitoring system designed to ensure satisfactory road works. The NRRDA issues general quality control guidelines and compiles a quality control handbook followed by all three tiers. The first, most local tier concerns the quality of construction inputs, such as materials and workmanship. The second tier, at the state level, and the third tier, at the national level, focus on outcome quality. These two tiers are therefore directly comparable and form the basis of this study. At the state level, each state government appoints a senior engineer as the State Quality Coordinator (SQC). The SQC, under the SRRDA, coordinates State Quality Monitors (SQM), the auditors who conduct regular inspections and collect material samples for testing. National-level audits are coordinated centrally by the NRRDA, and roads are selected for inspection through randomization under program guidelines. These audits are carried out by National Quality Monitors (NQM), who are mainly retired engineers from other states.[10] In short, state quality audits are managed by state-level bureaucrats and may be exposed to political capture, while centrally managed national quality audits are less prone to such capture.

### 2.2.2 Roles of MLAs in PMGSY

Bihar has a bicameral legislature. Members of the Legislative Assembly (MLAs) are elected to the lower house of the state legislature. Electoral terms last five years, and elections use a first-past-the-post system. The three most recent MLA elections in Bihar were held in 2010, 2015, and 2020. Despite the parliamentary system, Bihar has closely followed the five-year electoral cycle since 1985, so I treat election timing as exogenous.

Under PMGSY program guidelines, MLAs have almost no formal role in program implementation. Their involvement is limited to the initial planning stage. In particular, they have no official role in either the state or national level of the quality monitoring system.[11] A systematic relationship between electoral cycles and reported outcomes under the quality control system can therefore be interpreted as evidence consistent with political intervention. In the main analyses, I document such a relationship for state quality monitoring, but not for national quality monitoring.

In practice, however, many accounts suggest that MLAs exercise influence beyond their

[9] https://rwdbihar.gov.in/ContactUS.aspx

[10] Part II Section 15, Programme Guidelines (April 2012)

[11] Part II Sections 4.1, 4.6, and 15.10.2, PMGSY Programme Guidelines (April 2012)

formal authority. Chopra (1996) shows that although MLAs' official role is to represent constituents in state legislative assemblies, only 3 percent of surveyed MLAs report assembly work as the activity on which they spend the most time. According to Jensenius (2013), state assemblies meet as few as 10 times per year, and "much more important to the MLAs are all their unofficial tasks of delivering pork, blessing occasions, and helping people out with their individual problems". In the PMGSY context, Lehne et al. (2018) document that MLAs intervene in contracting procedures on behalf of members of their private networks. On ruling-party alignment and road construction, Bohlken (2021) shows that alignment facilitates access to a PMGSY road. Yet road construction in constituencies represented by senior aligned politicians is rated lower quality than road construction in constituencies represented by ordinary legislators aligned with the ruling party and in non-aligned constituencies.[12]

As described above, IAS bureaucrats, rather than politicians, directly control state-level quality monitoring. Iyer and Mani (2012) proposes a mechanism through which bureaucratic and political incentives can converge: the chief minister, the de facto head of the state government formed by the state ruling coalition, can control bureaucrats by reassigning them to positions of varying importance. Because transfers across states are infrequent, bureaucrats with career concerns may have incentives to be responsive to the chief minister's priorities. Iyer and Mani (2012) also documents that when a new chief minister takes office, transfers become more frequent in districts with less aligned politicians, consistent with an explanation based on tightening political control for electoral reasons rather than improving administrative efficiency.

## 2.3 Data

### 2.3.1 PMGSY Audit Data

I obtain administrative data on road length, road costs, and reported quality for each PMGSY audit carried out in Bihar between January 2014 and December 2017 from the official PMGSY website (http://omms.nic.in). To isolate electoral incentives around the election, the main analyses use audits conducted within one year before and one year after the polling date of the 2015 Bihar election, October 12. The final dataset covers 7,404 state quality audits and 914 national quality audits.

Table 2.1 reports summary statistics for the state-level audit sample. The table includes the binary audit outcome, the budgeted construction cost, and the budgeted length of audited roads. These budgets are determined by the NRRDA using predetermined rules and are

[12]Bohlken (2021) does not distinguish the two levels of audits and takes inspection outcome from either level of audits as impartial measures of road quality. In light of the evidence below that state-level audit results may be politically influenced, the results from Bohlken (2021) should be interpreted with caution.

not malleable by the state-level bureaucracy or politicians. A given road can be audited multiple times and can also be audited before construction is complete. I therefore code a first-inspection dummy indicating the first audit on a given road and a completion-upon-inspection dummy indicating that road construction was complete at the time of inspection.

The observation counts in Columns 2 and 3 indicate that state-level audits are more likely to occur right before elections. More importantly, roads are more likely to fail state quality audits in the pre-election period. Figure 2.1 presents this pattern graphically in a binscatter plot. The same pattern does not appear in national-level audits. Despite the narrow time window, however, these descriptive patterns cannot by themselves be interpreted as evidence of manipulation. For example, they could reflect a newly elected state government imposing different quality monitoring guidelines that affect aligned and non-aligned constituencies similarly. The main empirical analyses therefore further divide the data by pre-election alignment status and test for differential quality reporting.

### 2.3.2 Bihar State Assembly Election Data

I use Bihar state election data from 2010 and 2015 made available by the Trivedi Centre for Political Data. For each of the 243 assembly constituencies in each election, I observe the winning candidate's party affiliation and winning margin. I treat the 2010 election data as predetermined. Specifically, I use the party affiliation of the 2010 winning candidate to assign constituencies to treatment groups: if the winner belonged to Janata Dal-United (JDU), the chief minister's party, the constituency is assigned to the positive treatment group; otherwise, it is assigned to the negative treatment group.

### 2.3.3 Probabilistic Matching

The PMGSY audit dataset records the administrative block of each audit, but not the assembly constituency.[13] To compare audits in aligned and non-aligned constituencies, I probabilistically match blocks to constituencies. Bihar has 38 districts. Below the district level, it has 534 blocks and 243 legislative assembly constituencies, whose boundaries often overlap. Although a block is smaller than a constituency on average, it does not necessarily fall entirely within one constituency. Both blocks and constituencies are composed of villages. I match blocks and constituencies using a correspondence that links each village, the administrative unit below the block, to a constituency. If all villages in a given block belong to constituencies aligned with JDU, the chief minister's party, I define the block as aligned. If all villages in a given block belong to non-aligned constituencies, I define the block as non-aligned. To reduce measurement concerns and make interpretation cleaner, I restrict the sample to the

[13]Administrative block is the administrative level below district designated for the purpose of rural development.

315 blocks, out of 534, that are entirely aligned or entirely non-aligned.

## 2.4 Empirical Design

### 2.4.1 Potential Outcome Framework in the Absence of Pure Controls

I use a potential outcomes framework to clarify the difference-in-differences (DID) identification strategy. Under the standard parallel trends assumption, the coefficient on the interaction term identifies the differential reporting effect for aligned blocks relative to non-aligned blocks. Treatment groups are defined by pre-election party affiliation, so the classification is time-invariant even though post-election alignment may change because of turnover in either the state ruling party or local politicians.

This setting departs from a standard DID setting in two ways. First, the treatment period is the pre-election period, $t = 0$, rather than the post period, $t = 1$. Second, and more importantly, there is no clearly defined control group during the treatment period. Both aligned ($A$) and non-aligned ($NA$) blocks are treated before the election, but in opposite directions. Aligned blocks are treated favorably by the ruling party, while non-aligned blocks are treated unfavorably. Thus, for every unit $i$ in period $t$, there are three potential treatment states, $z_{it} \in \{-1, 0, 1\}$, denoting unfavorable treatment, control, and favorable treatment, respectively. Treatment status can be summarized as follows: $z_{i0}(A) = 1, z_{i1}(A) = 0, z_{i0}(NA) = -1, z_{i1}(NA) = 0$. Let $Y_{it}(z_{it})$ be the potential outcome for unit $i$ in period $t$ when receiving treatment status $z_{it}$.

Because there is no pure control group, I cannot separately identify $E[Y_{it}(1) - Y_{it}(0)|A]$ and $E[Y_{it}(-1) - Y_{it}(0)|NA]$, namely the favorable treatment effect on aligned blocks and the unfavorable treatment effect on non-aligned blocks.

I can identify a different parameter of interest: the differential reporting effect for aligned blocks relative to non-aligned blocks. It is given by $\tau = E[Y_{i0}(1) - Y_{i0}(0)|A] - E[Y_{i0}(-1) - Y_{i0}(0)|NA] = E[Y_{i0}(1)|A] - E[Y_{i0}(-1)|NA] + E[Y_{i0}(0)|NA] - E[Y_{i0}(0)|A]$, where the last two terms are missing potential outcomes unobserved by the econometrician. Under the standard parallel trends assumption, $E[Y_{i0}(0) - Y_{i1}(0)|A] = E[Y_{i0}(0) - Y_{i1}(0)|NA]$, I can replace the last two terms with population moments as follows. $\tau = E[Y_{i0}(1)|A] - E[Y_{i0}(-1)|NA] + E[Y_{i1}(0)|NA] - E[Y_{i1}(0)|A] = E[Y_{i0}(1) - Y_{i1}(0)|A] - E[Y_{i0}(-1) - Y_{i1}(0)|NA]$. The sample analog and its standard error can then be estimated directly using a standard DID regression framework.

In summary, although alignment status is not randomly assigned, the coefficient estimate from the standard DID specification can be interpreted as the differential reporting effect by

alignment status as long as there is no time-varying confounding.

### 2.4.2 Empirical Specification

If the audit manipulation hypothesis is correct, blocks represented by aligned politicians, namely politicians belonging to Janata Dal (United), the state ruling party of Bihar, should have higher average reported road quality before the election relative to after the election. Non-aligned blocks, by contrast, should have lower reported quality before the election. I test for these heterogeneous effects using the following regression specification.

$$Y_{ibt} = \beta_1 \,\text{Before}_t + \beta_2 \,\text{Alignment}_b \times \text{Before}_t + \lambda' X_{ibt} + \alpha_b + \varepsilon_{ibt}. \tag{2.1}$$

$Y_{ibt}$ is the binary outcome for audit $i$ in block $b$ and period $t$, equal to 1 if the audit outcome is satisfactory. $\text{Alignment}_b$ indicates whether the MLA who won the 2010 election in block $b$ shares the chief minister's party affiliation. $\text{Before}_t$ is an indicator for inspections conducted before the election. $X_{ibt}$ collects predetermined observable road characteristics. $\alpha_b$ captures block fixed effects, and $\varepsilon_{ibt}$ represents idiosyncratic error. The sample is restricted to audits conducted within 365 days before and after the 2015 MLA election in Bihar. Throughout, I report standard errors clustered at the block level, the level at which treatment is assigned.

If the chief minister manipulates audit outcomes to favor aligned blocks relative to non-aligned blocks, audit outcomes should be differentially better in aligned blocks, implying $\beta_2 > 0$.

## 2.5 Secondary Findings

To help interpret the main results on differential quality reporting before the election, I first describe two sets of secondary findings.

### 2.5.1 No Differential Extensive Margin

If reporting manipulation is electorally beneficial, the ruling party may have an incentive to increase audit frequency before the election. However, because the occurrence of an audit does not necessarily send voters a positive or negative signal, it is unclear ex ante whether the ruling party would allocate more audits to aligned or non-aligned blocks. Column 1 shows that audit frequency does increase before the election. To test for a differential effect, I estimate the main specification at the block-period level, replacing the dependent variable with $ln(N\_audit_{bt})$. Table 2.2 shows no differential targeting in audit occurrence. Two explanations are consistent with this null extensive-margin result: technological constraints in targeting by alignment status, and similar net gains from manipulating audits in aligned and non-aligned blocks.

### 2.5.2 No Differential Targeting on Observables

If some observable road characteristics are associated with lower quality, the ruling party could audit more roads with those characteristics in non-aligned blocks and fewer such roads in aligned blocks. To test for this type of differential targeting, I estimate the main specification after replacing the dependent variable with several observable road characteristics. Table 2.3 shows no differential targeting based on observables. Because these variables are predetermined, I include them as controls in alternative specifications when examining the main results. The main results are virtually unchanged.

## 2.6 Main Results: Differential Quality Reporting

Table 2.4 presents the main findings. The statistic of interest is the coefficient on the interaction term, which represents differential quality reporting between aligned and non-aligned blocks. Columns 1 and 2 imply a differential reporting effect of 7.5 percentage points. This corresponds to an effect size of 10.7% relative to the sample average satisfactory rate of 69%, or 0.16 standard deviations given a sample standard deviation of 46.3%.

Two distinct hypotheses can rationalize this result: (i) politicians target resources to aligned blocks to improve actual road quality (resource targeting), and (ii) politicians influence the reporting of audit outcomes without improving actual quality (information manipulation).

To distinguish between these explanations, I use national audits, which are insulated from state-level political influence, as a more impartial measure of actual quality. If resource targeting explains the state audit results, the two tiers of quality audits should display differential reporting with the same sign. Column 3 suggests otherwise: the differential reporting estimate for national audits has the opposite sign from the state audit estimate and is statistically insignificant. This discrepancy between state and national quality reporting weighs against the resource targeting explanation and supports the information manipulation hypothesis. A remaining concern is that, by the same logic, the BJP, the national ruling party and a major opposition party in Bihar, might manipulate national audits. Column 4 is not consistent with this concern because the interaction term is null.[14] In sum, the pattern of differential reporting appears in state-level audits but not in national-level audits, supporting the interpretation that local politicians influence information disclosure.[15]

[14]The corresponding national-audit appendix exhibits appear at the end of the dissertation; see Table 2.A1 and Table 2.A2. Together, they show no clear differential pattern between blocks aligned with the BJP and those that are not.

[15]Ideally, I would compare results from the two levels of audits that are carried out on the same road, to precisely identify the within-road discrepancy in reporting. However, the small number of national audits in the sample period results in a lack of overlap and insufficient power to conduct a within-road analysis.

### 2.6.1 Mechanism #1: Auditor Assignment

The main results are estimated using 7,404 audits within a one-year bandwidth around the 2015 election. In the unrestricted 2011–2018 sample, there are 32,836 audits conducted by 171 unique auditors, 55 of whom audited at least one road during the main sample period. I residualize audit outcomes using block fixed effects and inspection-year fixed effects, then estimate a fixed effect for each auditor. I interpret this fixed effect as time-invariant auditor leniency. Estimating auditor fixed effects on the unrestricted sample, rather than only the narrow electoral window, reduces the concern that the predicted fixed effects reflect endogenous assignment to roads of different actual quality rather than auditor leniency.

The interaction in Column 1 of Table 2.5 suggests that, before the election, aligned blocks are assigned more lenient auditors who are on average 2.9 percentage points more likely to pass a state quality audit. To gauge the quantitative relevance of assigning auditors with different leniency, I add the predicted auditor fixed effects in Column 3. Comparing Columns 2 and 3 suggests that auditor assignment accounts for roughly one-third of the differential reporting in an accounting sense. Together with the null results for national quality audits, this pattern supports the interpretation that differential reporting reflects audit report manipulation rather than actual resource targeting.

### 2.6.2 Mechanism #2: Manipulation of Audit Timing

The auditor assignment evidence suggests one way in which reported outcomes may be shaped conditional on actual road quality. A second channel consistent with the main results is manipulation of audit timing conditional on expected audit outcomes. If the state ruling party has private information about road quality before formal inspections, it could strategically accelerate inspections of satisfactory roads and delay inspections of unsatisfactory roads in aligned blocks, while doing the reverse in non-aligned blocks. This would release favorable information before the election and unfavorable information after the election in aligned blocks, with the opposite pattern in non-aligned blocks. To examine this possibility, I use road construction completion dates to compute inspection lag, defined as the number of days between road construction completion and audit. Among roads completed before the election, and absent manipulation, roads of similar quality should be equally likely to be inspected before the election in aligned and non-aligned blocks. Figure 2.2 suggests a different pattern. For roads that turn out to be satisfactory, there is no discernible difference by alignment status in the share inspected before the election. For roads that turn out to be unsatisfactory, however, inspections are 4.2 percentage points more likely to occur before the election in non-aligned blocks than in aligned blocks. This is consistent with delayed audits of unsatisfactory roads in aligned blocks. The same pattern could also reflect accelerated audits of unsatisfactory roads in non-aligned blocks. Either interpretation is consistent with strategic audit timing, though the evidence does not directly observe intent or the scheduling

decision itself.

### 2.6.3 Robustness: Parallel Trend After Election

To better understand the dynamics around the election, I extend the sample to two years before and two years after the election and split it into eight six-month bins, with the first six months after the election serving as the benchmark period. Figure 2.3 presents the results of this event-study specification. The insignificant interaction terms in the post-period, which is the control period, are reassuring and consistent with the parallel trends assumption in the absence of treatment.

In addition, differential reporting is concentrated in the period 18 to 6 months before the election, while the interaction effect is insignificant during the 6 months immediately before the election.[16] This pattern suggests that the ruling party may time audit report release early enough for the information to circulate among voters during the re-election campaign.

### 2.6.4 Heterogeneity by Political Competitiveness

I explore heterogeneity in differential reporting along two dimensions of political competitiveness in Table 2.6. First, I consider whether JDU fielded a candidate in the block in the 2015 election (Columns 1 and 2). Consistent with the hypothesis that differential reporting is intended to gain political support, the effect is larger in blocks where JDU did field a candidate. This finding is consistent with the theory, but it cannot be interpreted causally because candidate entry is endogenous.[17]

Second, I study heterogeneity by splitting the sample according to block-level winning margins in the previous 2010 state election. Interestingly, the differential reporting effect appears in neither the least competitive nor the most competitive blocks. This pattern is not what one would expect if politicians were only trying to influence pivotal voters or consolidate electoral strongholds. It may reflect the incumbency disadvantage that characterizes Indian state-level politics. In my sample, only 70 of the 193 blocks aligned in 2010 were able to defend their seats in 2015. The ruling party may therefore concentrate limited manipulative capacity on blocks where its actions have the greatest potential to affect electoral outcomes.

[16]The general election, which elects members of the parliament and the president, took place in 2014 April. Thus, the insignificant term during the 24 to 18 month period leading up to 2015 December Bihar state election, can be explained by the different political dynamics before the general election.

[17]For example, the causal chain can go in the opposite direction: the JDU has less control over the bureaucracy in a given block and thus decides to opt out of running in this block. Proving the causality requires a study of policies that give rise to exogenous variation in candidate eligibility, such as the quota system by caste in India (Pande, 2003). This is a work in progress.

## 2.7 Voter Responses

I next present suggestive evidence in Table 2.7 that inspection outcomes may be politically relevant. In aligned blocks, the ruling party's probability of winning the election is positively correlated with audit outcomes released before the election, but not with audit outcomes released after the election (Column 1). This pattern is consistent with favorable inspection results being electorally useful.[18] Moreover, relative to road completion dates, inspections are delayed by a median of 316 days. Given the narrow one-year sample window, if electoral outcomes only reflected the completion of high-quality roads rather than the release of audit reports, favorable reports released after the election should also predict electoral success. The coefficient on post-election reporting is instead insignificant. This pattern is consistent with audit reports carrying politically relevant information, though the evidence remains correlational and does not directly identify voter updating.[19]

The results are robust to applying analytical weights corresponding to the number of audits at the block level (Column 2). Results for non-aligned blocks are insignificant (Columns 3 and 4).

## 2.8 Conclusions

Using a newly collected dataset covering the universe of state-level quality audits under PMGSY, a rural road-building program in India, I estimate a 7.4 percentage point reported-quality differential by alignment status in Bihar. The reduced-form results are consistent with two mechanisms: endogenous auditor assignment and strategic audit timing. Auditor assignment accounts for roughly one-third of the effect in an accounting sense, suggesting that the ruling party may influence quality reports by shaping which state-level auditors inspect which roads. In addition, roads completed before the election that turn out to be unsatisfactory are 4.2 percentage points more likely to be inspected before the election in non-aligned regions than in aligned regions. This pattern is consistent with audit timing manipulation that benefits aligned politicians electorally and harms non-aligned ones.

The same patterns do not appear in national audits, over which state-level politicians have no formal control and substantially less scope for informal influence. Taken together, the results

[18]This specification is in the spirit of Ferraz and Finan (2008), which compares the electoral outcomes of Brazilian municipalities audited before versus after the 2004 elections. However, my results cannot be interpreted as causal since the audit timings are not exogenously determined. A potential indirect test is to explore the interaction between report release and block-level media availability - voters should be more responsive to audit reports if such information is readily available (Besley and Burgess, 2002). This is a work in progress.

[19]This finding is related to experimental studies in India that explore the electoral response to information treatment on policy outcomes Banerjee et al. (2011, 2018) or candidates' criminal background (George et al., 2019).

from the two audit tiers favor the information manipulation hypothesis over the resource targeting hypothesis. The heterogeneity analysis shows that the effect is not concentrated in constituencies that were most competitive in the previous election, underscoring the value of the proposed DID strategy relative to a regression discontinuity design based only on close elections.

The literature studying the consequences of audits often leverages experimental variation or randomized variation in natural experiments. This paper extends that literature by examining the determinants of the audit process itself. The results suggest that audits do not necessarily take place in a political vacuum. Understanding information disclosure therefore requires understanding the political incentives that shape how information is produced and released.

Several issues remain outside the scope of the analysis. One concerns the complicated relationship between bureaucrats and politicians. Although I present anecdotal evidence of illicit exchange between bureaucrats and politicians in PMGSY and other contexts, the analyses effectively treat their interests as aligned. Future work could model this relationship as a principal-agent problem and examine when bureaucrats are more responsive to politicians' demands. A second issue concerns voter responses. This paper provides only suggestive correlational evidence on voter responses to information that may be biased in favor of the ruling party. Future work could study voters' belief updating, especially whether voters account for the possibility that audit information is politically biased. Doing so would require exogenous variation in information provision and is a possible direction for future research.

# Figures

Figure 2.1: State Audit Results Before and After the 2015 Bihar Election

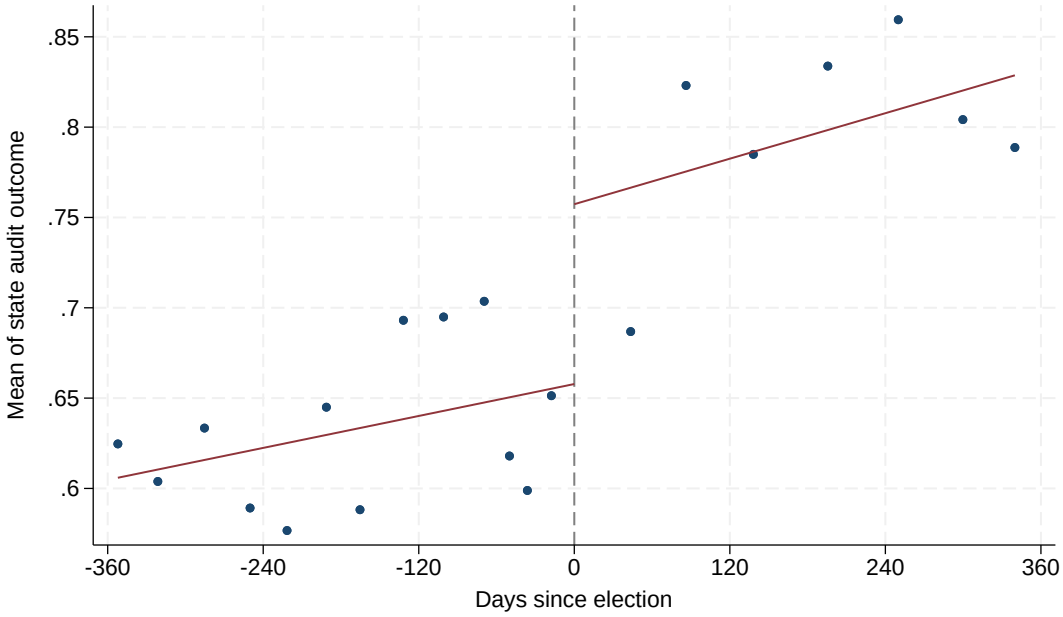


**Notes:** The figure plots binned means of the state-audit satisfactory indicator against days since the 2015 Bihar MLA election for audits conducted within the one-year window around the election. The fitted lines are estimated separately on each side of the election cutoff. Higher values indicate a higher share of satisfactory audit reports.

Figure 2.2: Share of Roads Inspected Before the Election by Inspection Outcome

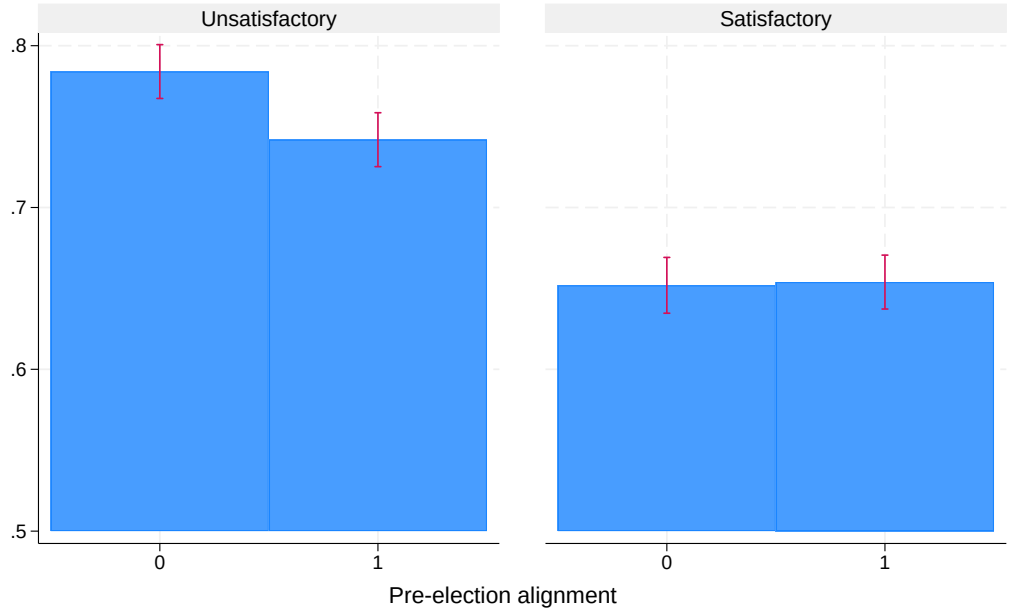


**Notes:** The figure reports the share of roads inspected before the 2015 Bihar MLA election by pre-election alignment with the chief minister, separately for roads that ultimately receive satisfactory and unsatisfactory state-audit outcomes. Bars show sample means and whiskers show 95 percent confidence intervals. The sample includes state quality audits of Bihar roads completed within five years before the election.

Figure 2.3: Event Study of Differential State Audit Reporting

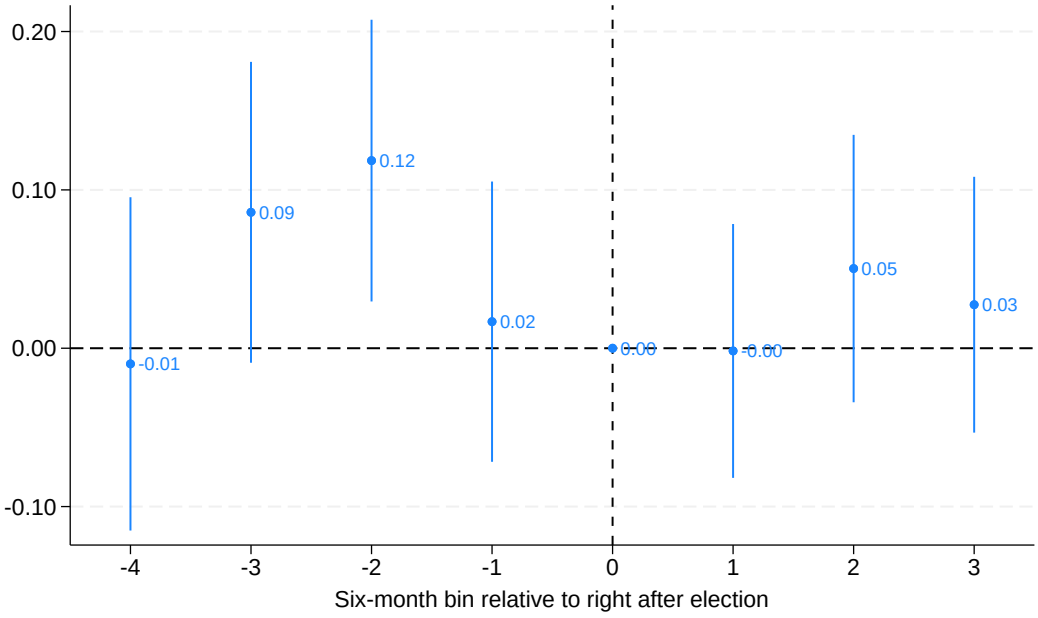


**Notes:** The figure plots coefficients on the interaction between 2010 chief-minister alignment and six-month bins relative to the election, estimated over a two-year window around the 2015 Bihar MLA election. The omitted period is the six months immediately after the election. Vertical bars show 95 percent confidence intervals based on standard errors clustered at the block level.

# Tables

Table 2.1: Summary Statistics Around Election: State Quality Audits

| Variable | (1) Full Sample | (2) Before | (3) After | (4) Difference |
|---|---|---|---|---|
| Reported quality (binary) | 0.69 | 0.63 | 0.80 | -0.17*** |
| | (0.46) | (0.48) | (0.40) | (0.00) |
| Road sanctioned cost (INR 100,000) | 160.83 | 162.12 | 158.38 | 3.75 |
| | (136.79) | (140.16) | (130.12) | (0.25) |
| ln(cost) | 4.83 | 4.83 | 4.84 | -0.01 |
| | (0.69) | (0.70) | (0.65) | (0.39) |
| Road sanctioned length (km) | 2.81 | 2.94 | 2.55 | 0.39*** |
| | (2.55) | (2.69) | (2.26) | (0.00) |
| ln(length) | 0.77 | 0.80 | 0.70 | 0.10*** |
| | (0.70) | (0.72) | (0.66) | (0.00) |
| Sanctioned cost per km (INR 100,000) | 60.41 | 58.22 | 64.59 | -6.37*** |
| | (15.28) | (15.91) | (13.01) | (0.00) |
| ln(cost per km) | 4.07 | 4.03 | 4.14 | -0.12*** |
| | (0.27) | (0.28) | (0.24) | (0.00) |
| First inspection of the road (binary) | 0.52 | 0.64 | 0.28 | 0.36*** |
| | (0.50) | (0.48) | (0.45) | (0.00) |
| Completion of the roads upon inspection (binary) | 0.49 | 0.52 | 0.42 | 0.10*** |
| | (0.50) | (0.50) | (0.49) | (0.00) |
| Observations | 7,404 | 4,856 | 2,548 | 7,404 |

**Notes:** The unit of observation is a state quality audit in Bihar conducted within 365 days before or after the 2015 Bihar MLA election. Columns (1)–(3) report sample means for the full sample, the pre-election subsample, and the post-election subsample. The Difference column reports the pre-election mean minus the post-election mean. Parentheses in Columns (1)–(3) report standard deviations; parentheses in Column (4) report p-values from robust standard errors. $^{*}p < 0.10$,$^{**}p < 0.05$,$^{***}p < 0.01$.

Table 2.2: Audit Frequency Around Election

| | ln(Audit Frequency) | |
|---|---|---|
| | (1) | (2) |
| Before election | 0.640*** | -0.419** |
| | (0.062) | (0.188) |
| Before election x CM alignment (2010) | -0.018 | |
| | (0.088) | |
| Before election x BJP (2010) | | 0.086 |
| | | (0.318) |
| Quality Monitor Level | State | National |
| Block F.E. | Yes | Yes |
| Mean dep. var. (untransformed) | 11.465 | 2.174 |
| S.D. | 9.745 | 1.328 |
| Obs. | 630 | 264 |
| R-squared | 0.804 | 0.570 |

**Notes:** The unit of observation is a block-period, where each block contributes one pre-election and one post-election observation within 365 days of the 2015 Bihar MLA election. Column (1) reports the log number of state quality audits in the block-period, and Column (2) reports the number of national quality audits in the block-period. The main regressors interact the pre-election indicator with 2010 chief-minister alignment in Column (1) and with 2010 BJP alignment in Column (2). All specifications include block fixed effects. Standard errors are clustered at the block level.

Table 2.3: State Audit Targeting Around Election

| | ln(cost) (1) | ln(length) (2) | ln(cost per km) (3) | First inspection (4) | Construction completed (5) |
|---|---|---|---|---|---|
| Before election | 0.011 | 0.119*** | -0.107*** | 0.399*** | 0.067** |
| | (0.024) | (0.027) | (0.013) | (0.030) | (0.030) |
| Before election x CM alignment (2010) | -0.042 | -0.048 | 0.006 | -0.012 | 0.026 |
| | (0.031) | (0.034) | (0.017) | (0.043) | (0.038) |
| Block F.E. | Yes | Yes | Yes | Yes | Yes |
| Obs. | 7403 | 7403 | 7403 | 7403 | 7403 |
| R-squared | 0.226 | 0.183 | 0.319 | 0.239 | 0.169 |

**Notes:** The unit of observation is a state quality audit in Bihar conducted within 365 days before or after the 2015 Bihar MLA election. Each column reports a separate regression with the indicated pre-determined road characteristic as the outcome. The main regressor is the interaction between the pre-election indicator and 2010 chief-minister alignment. Columns (1)–(3) use log sanctioned cost, log road length, and log sanctioned cost per kilometer; Columns (4)–(5) use indicators for whether the audit is the first inspection of the road and whether construction is complete at inspection. All specifications include block fixed effects. Standard errors are clustered at the block level.

Table 2.4: Audit Outcomes Around Election

| | Satisfactory Audit Outcome | | | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Before election | -0.207*** | -0.105*** | 0.078 | 0.027 |
| | (0.021) | (0.022) | (0.086) | (0.054) |
| Before election x CM alignment (2010) | 0.075** | 0.074** | -0.068 | |
| | (0.032) | (0.031) | (0.102) | |
| Before election x BJP (2010) | | | | 0.038 |
| | | | | (0.109) |
| Quality Monitor Level | State | State | National | National |
| Controls | No | Yes | Yes | Yes |
| Block F.E. | Yes | Yes | Yes | Yes |
| Mean dep. var. | 0.690 | 0.690 | 0.570 | 0.570 |
| S.D. | 0.463 | 0.463 | 0.495 | 0.495 |
| Obs. | 7403 | 7403 | 838 | 838 |
| R-squared | 0.165 | 0.224 | 0.365 | 0.364 |

**Notes:** The unit of observation is an audit in Bihar conducted within 365 days before or after the 2015 Bihar MLA election. In all columns, the outcome is an indicator equal to one if the audit outcome is satisfactory. Columns (1) and (2) use state quality audits and estimate the interaction between the pre-election indicator and 2010 chief-minister alignment; Column (2) adds controls for first inspection, completion upon inspection, and log sanctioned cost per kilometer. Column (3) applies the same interaction with 2010 chief-minister alignment to national quality audits, while Column (4) replaces that interaction with 2010 BJP alignment. All specifications include block fixed effects. Standard errors are clustered at the block level.

Table 2.5: State Audit Manipulation Mechanism

| | Predicted monitor FE | State audit outcome | |
|---|---|---|---|
| | (1) | (2) | (3) |
| Before election | -0.004 | -0.105*** | -0.102*** |
| | (0.011) | (0.022) | (0.021) |
| Before election x CM alignment (2010) | 0.029** | 0.074** | 0.049* |
| | (0.014) | (0.031) | (0.029) |
| Predicted monitor FE | | | 0.872*** |
| | | | (0.046) |
| Controls | Yes | Yes | Yes |
| Block F.E. | Yes | Yes | Yes |
| Mean dep. var. | -0.000 | 0.690 | 0.690 |
| S.D. | 0.154 | 0.463 | 0.463 |
| Obs. | 7403 | 7403 | 7403 |
| R-squared | 0.204 | 0.224 | 0.292 |

**Notes:** The unit of observation is a state quality audit in Bihar conducted within 365 days before or after the 2015 Bihar MLA election. Column (1) uses the predicted monitor fixed effect as the outcome, which is an auditor-leniency proxy estimated from the broader sample of state audits beyond the one-year election window. Columns (2) and (3) use the satisfactory-audit indicator as the outcome. The main regressor is the interaction between the pre-election indicator and 2010 chief-minister alignment. Column (3) adds the predicted monitor fixed effect to assess how much of the differential reporting operates through auditor assignment. Controls include first inspection, completion upon inspection, and log sanctioned cost per kilometer. All specifications include block fixed effects. Standard errors are clustered at the block level.

Table 2.6: State Audit Outcome Heterogeneity

| | JDU candidate in 2015 | | 2010 winning margin | | |
|---|---|---|---|---|---|
| | No | Yes | Small | Middle | Large |
| | (1) | (2) | (3) | (4) | (5) |
| Before election | -0.106*** | -0.103*** | -0.078** | -0.186*** | -0.032 |
| | (0.029) | (0.035) | (0.034) | (0.034) | (0.045) |
| Before election x CM alignment (2010) | 0.043 | 0.082* | 0.035 | 0.183*** | -0.013 |
| | (0.051) | (0.042) | (0.053) | (0.055) | (0.053) |
| Controls | Yes | Yes | Yes | Yes | Yes |
| Block F.E. | Yes | Yes | Yes | Yes | Yes |
| Obs. | 3107 | 4296 | 2671 | 2419 | 2313 |
| R-squared | 0.231 | 0.218 | 0.243 | 0.214 | 0.225 |

**Notes:** The unit of observation is a state quality audit in Bihar conducted within 365 days before or after the 2015 Bihar MLA election. The outcome is an indicator equal to one if the audit outcome is satisfactory. Columns (1) and (2) split the sample by whether JDU fielded a candidate in the block in the 2015 election. Columns (3)–(5) split the sample by terciles of the 2010 winning margin. In every column, the main regressor is the interaction between the pre-election indicator and 2010 chief-minister alignment. Controls include first inspection, completion upon inspection, and log sanctioned cost per kilometer. All specifications include block fixed effects. Standard errors are clustered at the block level.

Table 2.7: Voter Response to Inspection Outcome

| | JDU winning in 2015 | | | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Avg audit outcome before election | 0.414*** | 0.455** | 0.054 | 0.005 |
| | (0.138) | (0.190) | (0.088) | (0.151) |
| Avg audit outcome after election | -0.001 | -0.076 | 0.029 | -0.016 |
| | (0.144) | (0.196) | (0.136) | (0.164) |
| Chief-minister alignment (2010) | Yes | Yes | No | No |
| Weights | No | Yes | No | Yes |
| Obs. | 168 | 168 | 138 | 138 |
| R-squared | 0.052 | 0.044 | 0.003 | 0.000 |

**Notes:** The unit of observation is a block. The outcome is an indicator equal to one if JDU won the block in the 2015 Bihar MLA election. The regressors are the block-level average state-audit outcome released before the election and the block-level average state-audit outcome released after the election, computed from audits within 365 days of the election. Columns (1) and (2) restrict the sample to blocks aligned with the chief minister in 2010; Columns (3) and (4) restrict the sample to non-aligned blocks. Columns (2) and (4) apply analytic weights equal to the total number of audits observed in the block. Robust standard errors are reported in parentheses.

# Appendix Tables

Table 2.A1: Summary Statistics Around Election: National Quality Audits

| Variable | (1) Full Sample | (2) Before | (3) After | (4) Difference |
|---|---|---|---|---|
| Reported quality (binary) | 0.56 | 0.57 | 0.56 | 0.02 |
| | (0.50) | (0.50) | (0.50) | (0.62) |
| Road sanctioned cost (INR 100,000) | 159.29 | 150.97 | 164.58 | -13.61 |
| | (133.33) | (130.17) | (135.14) | (0.13) |
| ln(cost) | 4.84 | 4.78 | 4.88 | -0.09** |
| | (0.65) | (0.66) | (0.65) | (0.04) |
| Road sanctioned length (km) | 2.53 | 2.51 | 2.55 | -0.04 |
| | (2.21) | (2.25) | (2.19) | (0.78) |
| ln(length) | 0.70 | 0.68 | 0.70 | -0.02 |
| | (0.65) | (0.65) | (0.65) | (0.63) |
| Sanctioned cost per km (INR 100,000) | 64.45 | 61.87 | 66.09 | -4.22*** |
| | (13.31) | (13.25) | (13.10) | (0.00) |
| ln(cost per km) | 4.15 | 4.10 | 4.17 | -0.07*** |
| | (0.21) | (0.22) | (0.19) | (0.00) |
| First inspection of the road (binary) | 0.97 | 0.98 | 0.97 | 0.00 |
| | (0.16) | (0.15) | (0.16) | (0.68) |
| Completion of the roads upon inspection (binary) | 0.23 | 0.17 | 0.27 | -0.11*** |
| | (0.42) | (0.37) | (0.45) | (0.00) |
| Observations | 914 | 355 | 559 | 914 |

**Notes:** The unit of observation is a national quality audit in Bihar conducted within 365 days before or after the 2015 Bihar MLA election. Columns (1)–(3) report sample means for the full sample, the pre-election subsample, and the post-election subsample. The Difference column reports the pre-election mean minus the post-election mean. Parentheses in Columns (1)–(3) report standard deviations; parentheses in Column (4) report p-values from robust standard errors. $^{*}p < 0.10,^{**}p < 0.05,^{***}p < 0.01$.

Table 2.A2: National Audit Targeting Around Election

| | ln(cost) (1) | ln(length) (2) | ln(cost per km) (3) | First inspection (4) | Construction completed (5) |
|---|---|---|---|---|---|
| Before election | -0.110* | -0.040 | -0.070*** | 0.018 | -0.136*** |
| | (0.065) | (0.070) | (0.020) | (0.012) | (0.050) |
| Before election x BJP (2010) | 0.240** | 0.152 | 0.088** | 0.012 | -0.048 |
| | (0.117) | (0.124) | (0.039) | (0.034) | (0.088) |
| Block F.E. | Yes | Yes | Yes | Yes | Yes |
| Obs. | 838 | 838 | 838 | 838 | 838 |
| R-squared | 0.420 | 0.381 | 0.493 | 0.314 | 0.357 |

**Notes:** The unit of observation is a national quality audit in Bihar conducted within 365 days before or after the 2015 Bihar MLA election. Each column reports a separate regression with the indicated pre-determined road characteristic as the outcome. The main regressor is the interaction between the pre-election indicator and 2010 BJP alignment. Columns (1)–(3) use log sanctioned cost, log road length, and log sanctioned cost per kilometer; Columns (4)–(5) use indicators for whether the audit is the first inspection of the road and whether construction is complete at inspection. All specifications include block fixed effects. Standard errors are clustered at the block level.

# Chapter 3

# Do Culture and Preferences Travel? Precautionary Behavior During COVID-19

## 3.1 Introduction

Many high-stakes health behaviors are private choices with both private and social consequences. The COVID-19 pandemic made this tradeoff unusually salient: decisions about whether to leave home, wear a mask, avoid gatherings, or practice hygiene involved private costs and benefits, but also affected the risks borne by others. The same tradeoff appears in other preventive health choices, such as vaccination, staying home when sick, and sexual protection, where private decisions can generate externalities. Whether people adopt such behaviors depends on more than medical information or formal policy. It may depend on how they weigh present costs against future benefits, how they perceive risk, how much they value others' welfare, whether they trust other people and institutions, and whether they view precaution as an individual choice or a social responsibility.

Following Guiso et al. (2006) and Bisin and Verdier (2011), I define culture as a group-level stock of transmitted preferences, beliefs, and values. By contrast, individual preferences are person-level characteristics that can vary within the same cultural group. The two are related but distinct. A more collectivist cultural environment, for example, may shape individual preferences by fostering greater concern for how one's actions affect others; a more individualistic environment may place greater weight on autonomy and personal choice. But the pathway from culture to preferences and then to behavior is not deterministic. Culture may also affect health behavior through channels not fully captured by measured preferences: a more collectivist cultural environment may shape beliefs about whether others

will comply, whether one's own compliance will be noticed, and how noncompliance will be perceived by others. At the same time, people from the same cultural background can differ in individualism and in other preferences, such as risk tolerance, patience, altruism, reciprocity, and trust, because of personal experience or idiosyncratic preference formation.

Distinguishing the behavioral effects of culture from individual preferences is not only a conceptual issue, but also an identification challenge. Culture may shape individual preferences through socialization, so group-level culture and individual-level preferences are empirically intertwined. At the same time, culture is usually measured at the county, regional, or country level, where it is bundled with institutions, public-health policies, disease severity, enforcement, media environments, and economic conditions (Alesina and Giuliano, 2015). A cross-country comparison between more collectivist and more individualistic societies therefore cannot tell whether differences in health behavior reflect culture itself, the individual preferences culture has produced, or the surrounding environment in which behavior is observed.

This paper addresses this challenge using an original online survey of international students at UC Berkeley, fielded with recruitment support from the Berkeley International Office. I focus on 899 respondents from 76 origin countries who did not return to their country of permanent residence during the same phase of the pandemic. I refer to these respondents as Berkeley stayers, but the term should be read as shorthand for non-returnees rather than verified physical residence in Berkeley for every respondent. Relative to respondents who returned home, this group was less likely to be re-exposed to home-country policy regimes, household environments, and local pandemic trajectories. The survey measures eight main respondent-level preference families: individualism and uncertainty avoidance using Hofstede-style modules (Hofstede, 2011); generalized trust using the World Values Survey (WVS) wording (Inglehart and Baker, 2000); and altruism, positive reciprocity, negative reciprocity, risk-taking, and patience using questions adapted from the Global Preference Survey (GPS) (Falk et al., 2018). Throughout the paper, respondent-level preferences refer to these survey-based individual measures. Origin-country norms refer to country-level Hofstede, World Values Survey, and Global Preference Survey proxies assigned by the respondent's country of permanent residence.

This setting allows me to separate two relationships that are usually bundled together. First, I compare students from the same origin country who differ in respondent-level preferences, using origin-country fixed effects to absorb shared cultural background and other origin-level factors. By holding origin-country factors fixed while focusing on non-returnees, this specification estimates a same-environment association between individual preferences and precautionary health behavior. Second, I ask whether origin-country norms predict precautionary behavior, both before and after controlling for the corresponding measured respondent-level preference. This comparison examines whether country-norm associations are fully accounted for by the corresponding measured individual preference, or whether they

remain predictive beyond that preference. The design does not randomly assign culture or preferences, but it substantially reduces one central confound in cross-area studies: cultural background is no longer mechanically tied to the policy and institutional environment of the origin country in which behavior is observed.

The preferred interpretation is deliberately narrow. The outcome is a standardized precautionary health behavior index of leaving home, expected need to leave home, face covering, avoiding gatherings, social distancing, and handwashing. The main individual-preference specification is a within-origin comparison among respondent stayers; the country-norm and joint specifications are cross-origin associations with observed controls. The estimates should therefore be read as conditional predictive relationships in the observed respondent-stayer sample, not as causal effects of culture or preferences.

The results point to both respondent-level preference and origin-country norm associations. In the origin-country fixed-effects specification, a one-standard-deviation increase in individualism and risk-taking is associated with a 0.10 and 0.09 standard-deviation decrease in the precautionary index, respectively ($p = 0.004$ and $p = 0.002$), while a one-standard-deviation increase in positive reciprocity and patience is associated with a 0.09 and 0.08 standard-deviation increase, respectively ($p < 0.001$ and $p = 0.006$).

At the origin-country level, a one-standard-deviation increase in origin-country uncertainty avoidance is associated with a 0.16 standard-deviation increase in precautionary behavior ($p = 0.039$), while a one-standard-deviation increase in origin-country individualism is associated with a 0.12 standard-deviation decrease ($p = 0.079$). The size and statistical significance of these country-norm associations remain unchanged after controlling for the matched individual preferences. This stability suggests that origin-country norm associations are not fully accounted for by the corresponding measured individual preference, although measurement error in the respondent-level preferences remains an important caveat.

Auxiliary analyses using beliefs about precaution effectiveness and COVID-related news attention help interpret these patterns. More individualistic and more risk-tolerant respondents pay less attention to COVID-related news, while more positively reciprocal and more patient respondents report stronger precaution-related beliefs and greater news attention. Adding these measures attenuates some individual-preference coefficients, especially patience, but the country-norm estimates remain broadly similar. These results are not formal mediation evidence, but they suggest that beliefs and information help organize the respondent-level preference patterns more than the origin-country norm patterns.

This paper contributes to three strands of literature. First, it adds to the literature on culture, trust, civic capital, and public-health behavior during crises. Recent COVID-19 studies relate area-level civic capital, civic culture, political trust, and individualism to mobility reductions, social distancing, mask use, and policy compliance using county, province, or regional panel

and event-study designs (Barrios et al., 2021; Durante et al., 2021; Bargain and Aminjonov, 2020; Bazzi et al., 2021; Bian et al., 2022). Related outbreak evidence connects public trust and social capital to Ebola and influenza protective behavior (Blair et al., 2017; Chuang et al., 2015). I complement this work by comparing individuals from different origin-country cultural backgrounds in the same environment. This difference matters because cross-area estimates can bundle culture with institutions, policy stringency, enforcement, epidemic severity, media environments, and other local conditions.

Second, the paper contributes to work linking individual preferences to preventive health behavior. Prior studies connect experimentally measured risk preference to health behaviors, time preference to health investment and vaccination, prosociality to COVID-19 health behaviors, and risk, time, trust, honesty, and social responsibility to pandemic compliance (Anderson, 2008; Fuchs, 1982; Chapman and Coups, 1999; Campos-Mercade et al., 2021; Müller and Rau, 2021); the Global Preference Survey provides validated measures of several of these preference dimensions across countries (Falk et al., 2018). I add direct survey measures of individualism, uncertainty avoidance, trust, altruism, reciprocity, risk-taking, and patience. This broader measurement matters because precautionary behavior combines immediate private costs, delayed and uncertain private benefits, and benefits to others, so no single preference fully captures the relevant decision problem.

Third, the paper bridges these culture and preference literatures by studying both origin-country norms and respondent-level preferences in the same empirical framework. I pair individual preferences with matched origin-country norms and estimate whether the country-level norm continues to predict precautionary behavior after controlling for the corresponding individual-level preference. This comparison asks whether country-norm associations are fully accounted for by the corresponding measured preference, or whether origin-country norms retain predictive power beyond preference transmission. This matters for interpretation: persistent origin-country associations would point to broader cultural channels, such as beliefs about others' compliance, perceived legitimacy of precaution, social expectations, unmeasured forms of socialization, or measurement error in the individual preference.

Finally, this paper contributes to the epidemiological approach in cultural economics. Fisman and Miguel (2007) compare diplomats from different countries in the same New York parking and enforcement environment, Fernandez and Fogli (2009) compare U.S.-born women with different ancestry under shared U.S. institutions, and Ichino and Maggi (2000) use workers in one Italian firm to separate birthplace, workplace environment, and sorting or peer effects (Fisman and Miguel, 2007; Fernandez and Fogli, 2009; Ichino and Maggi, 2000). I bring this same-environment logic to high-stakes health behavior and combine it with direct measures of individual preferences and origin-country norms. This matters because the design asks whether culture and traits travel with people and predict behavior when the immediate university and policy period is held closer to constant.

The rest of the paper proceeds as follows. Section 3.2 describes the Berkeley survey, the precaution index, the respondent-level preference measures, and the matched origin-country norms. Section 3.3 lays out the regression framework and clarifies the selection margins created by survey response and the decision not to return home. Section 3.4 presents the main individual-preference and country-norm results, together with auxiliary evidence on beliefs and information. Section 3.5 concludes by drawing out what the evidence implies for studying culture, preferences, and public-health behavior.

## 3.2 Data and Summary Statistics

The empirical design is anchored in an original online survey of international students at UC Berkeley. The survey was fielded on Qualtrics from May 4 to May 18, 2020, while local shelter-in-place restrictions remained in force. The Berkeley International Office distributed a recruitment email to the international-student population, and the invitation described a roughly ten-minute survey on behaviors and perceptions during the COVID-19 pandemic. Respondents who completed the questionnaire could separately enter a drawing for one of ten Amazon gift cards valued between $20 and $100. The survey was anonymous, with raffle emails collected separately from substantive responses. The final sample includes 1,137 completed responses from students originating in 87 countries. Relative to the closest complete administrative benchmark, namely the 6,833 international students in Berkeley International Office fall 2019 enrollment records, this corresponds to a response rate of about 16.6 percent. Among respondents, 899 reported not having returned to their country of permanent residence because of the coronavirus during the survey window. I refer to this group as Berkeley stayers.[1] The remaining 238 respondents had returned home. These respondents span 76 origin countries. Table 3.1 reports summary statistics for the Berkeley-stayer sample. The average respondent is 25.7 years old, 48.5 percent are female, 64.1 percent are graduate students, and average time in the United States and at Berkeley is 3.3 and 2.2 years, respectively.

Two selection margins matter for interpretation. The first is survey response: Table 3.2 shows that respondents are more likely than the Berkeley International Office population to be female and graduate students, are underrepresented among students from China and in Letters and Science, and are overrepresented in Engineering and several smaller schools.[2] The second is not returning to the country of permanent residence conditional on response: Table 3.3 compares Berkeley stayers with respondents who returned home. Relative to returnees,

[1]The survey directly asks whether the respondent went back to the country of permanent residence because of the coronavirus. Throughout the paper, "Berkeley stayers" is shorthand for respondents answering no, not a direct measurement of physical residence in Berkeley at the moment of survey completion.

[2]The gender comparison excludes survey respondents who selected "other gender" because the administrative benchmark is reported only for female and male students.

Berkeley stayers are older, more likely to be graduate students, and have spent more time in the United States and at Berkeley. Measured individual preferences are comparatively similar across the two groups, while the sharper differences arise in origin-country environments. Section 3.3 therefore interprets the main coefficients as relationships within the observed Berkeley-stayer sample, with broader extrapolation requiring additional assumptions about both selection margins.

### 3.2.1 Precautionary Behaviors

The first survey module records precautionary behaviors during the early pandemic. The main outcome is a standardized precaution index built from six behaviors: fewer days leaving home in the past seven days, a lower reported need to leave home over the next five days, covering one's nose or mouth during the most recent trip outside the home, avoiding social gatherings, practicing social distancing, and washing hands more frequently.[3] The last three components are recorded on 0–10 self-assessment scales. Each component is first standardized, and the aggregate index is then formed using inverse-covariance weighting (Anderson, 2008), so that higher values correspond to more precautionary behavior overall. In the Berkeley-stayer sample, respondents left home on 2.5 days on average during the prior week, 57.6 percent reported needing to leave home within the next five days, and 91.5 percent reported covering their nose or mouth during their most recent trip outside the home.

### 3.2.2 Individual Preferences

To measure respondent-level preferences, I closely follow the wording of established survey modules in the literature. The measures draw from three main sources: individualism and uncertainty avoidance from Hofstede (Hofstede, 2011), generalized trust from the World Values Survey (Inglehart and Baker, 2000), and altruism, positive and negative reciprocity, risk-taking, and patience from the Global Preference Survey (Falk et al., 2018). I briefly describe the construction of these measures below. Appendix Tables 3.A1 and 3.A2 report the exact question wording and response options.

Individualism is measured from the four-item ideal-job module and increases with greater weight on time for personal or home life and interesting work relative to job security and family respect. Uncertainty avoidance is constructed from four items covering nervousness, self-assessed health, managerial precision, and rule-following. Because nervousness and

[3]The precaution index is intended to capture broad defensive behavior rather than access to specific medical goods. During the May 2020 survey window, access to medical masks remained uneven. I therefore use the broad face-covering item, namely whether respondents covered their nose or mouth with any material when outside, rather than medical-mask ownership or willingness to pay for surgical masks. The same logic applies to the auxiliary belief index described below, which retains the broad face-covering item rather than the medical-mask-specific belief question.

self-assessed health were measured during the pandemic, the respondent-level uncertainty-avoidance index should be interpreted with more caution than the country-level Hofstede score. The survey contains two trust measures. The first follows the standard World Values Survey binary question on whether most people can be trusted or whether one needs to be very careful in dealing with people. The second follows the wording of the Global Preference Survey and is based on agreement with the statement that people have only the best intentions.[4]

Risk-taking and patience each come from the short 0–10 Global Preference Survey self-assessment items on willingness to take risks and willingness to give up something beneficial today in order to benefit more in the future.[5] Altruism combines a 0–10 self-assessment on willingness to share with others without expecting anything in return with a hypothetical donation amount out of 1,600 U.S. dollars. Positive reciprocity combines a 0–10 self-assessment on returning favors with a hypothetical thank-you gift choice for a helpful stranger, while negative reciprocity combines two 0–10 punishment questions with a revenge statement.

### 3.2.3 Country Norms

Each respondent is matched to home-country measures for the same nine dimensions. Individualism and uncertainty avoidance come from the Hofstede country scores. Trust is measured using both the World Values Survey country aggregate and the Global Preference Survey country aggregate so that each survey-side trust item is paired with its closest country-level counterpart. Patience, risk-taking, altruism, positive reciprocity, and negative reciprocity come from the Global Preference Survey country aggregates (Falk et al., 2018). This one-to-one mapping defines the matched preference-norm pairs used in Section 3.3: individualism, trust (WVS), trust (GPS), altruism, positive reciprocity, negative reciprocity, risk-taking, uncertainty avoidance, and patience.

I also analyze additional cultural factors that do not have a corresponding matched preference measure in the survey. Appendix Figure 3.A1 considers six such measures in separate regressions: power distance, masculinity, indulgence, long-term orientation, World Values Survey risk-taking, and World Values Survey altruism.

[4]In the Berkeley sample, the World Values Survey trust item and the GPS-style trust item are positively but only modestly correlated (raw correlation $r \approx 0.29$).

[5]This creates a measurement mismatch for risk-taking and patience between the Berkeley survey and the GPS country aggregates. In the Berkeley survey, I use the one-question short versions for both preferences because of survey-duration constraints. By contrast, the published GPS country measures combine those self-assessment items with additional staircase choice questions. In principle, the best fix would be to reconstruct country-level GPS measures from the corresponding short-form items only, but the question-by-question country data are not published in the paper's supplement, so I cannot readily recover directly comparable country aggregates here. This is therefore a measurement issue to revisit in future work.

### 3.2.4 Individual Controls and Auxiliary Variables

The individual control vector includes age, gender, graduate-student status, religiosity, self-reported household income decile in the home country, years in the United States, years at Berkeley, and school indicators. I also include a short Socially Desirable Response Scale to account for differential response tendencies in self-reported measures. Beyond the controls, the survey contains two auxiliary indices. The belief index aggregates respondents' assessments of the effectiveness of each measure that composes the precautionary health behavior index, and the news index aggregates attention to COVID-related news in the United States and in the respondent's home country. These variables are used as additional outcome variables rather than as primary regressors.

One implementation detail is that for some dimensions, matched country norms are missing because a few students come from origin countries not covered for that dimension in the country-level dataset. For instance, Hofstede individualism is not available for some countries represented in the Berkeley sample, such as Nigeria. In those cases, I also restrict the country fixed-effects specification to students with non-missing country-norm data, so that all three specifications for a given dimension use the same sample. Because the country datasets have different country coverage, sample size may differ across dimensions.

### 3.2.5 Country-Level Controls

Country-level controls are matched to each respondent's survey date. Pandemic conditions are measured with Johns Hopkins CSSE country-day data and the Oxford stringency series, so the controls vary with the timing of the survey between May 4 and May 18, 2020. Specifically, the regressions include the log number of confirmed cases over the previous seven days per capita, the contemporaneous COVID-19 death rate, and the seven-day average policy stringency index. Because the main respondents did not return to their countries of permanent residence, these variables should not be interpreted as the local restrictions they personally faced. Instead, they adjust for contemporaneous home-country pandemic conditions that could affect news exposure, family concern, perceived return options, or attention to the pandemic. I also merge more persistent structural covariates: log GDP per capita, log inequality, log population density, a democracy indicator, and an indicator for countries with substantial SARS exposure.

## 3.3 Estimation Strategy

This section formalizes the estimating equations and clarifies the populations to which the coefficients can and cannot be interpreted to apply. The paper's main design remains the Berkeley-stayer sample, interpreted as respondents who did not return to their countries

of permanent residence, because this restriction keeps the analysis closest to the common environment that motivates the design.

The matched preference-norm analysis is estimated separately for each dimension $k$, where the matched set includes individualism, trust (WVS), trust (GPS), altruism, positive reciprocity, negative reciprocity, risk-taking, uncertainty avoidance, and patience. Let $i$ index individuals and $c$ index origin countries. For dimension $k$, let $Y_{ic}^k$ denote the precaution index standardized within the relevant estimation sample, $X_{ic}^k$ the respondent-level preference in dimension $k$, and $N_c^k$ the matched home-country norm for the same dimension. Let $W_i$ collect the individual controls described in Section 3.2.4, and let $Z_c$ collect the country-level pandemic and structural controls. In all regression analyses, the regressor attached to the reported coefficient is standardized within that same estimation sample, so the plotted magnitudes can be read as beta coefficients.

I begin with the within-origin-country specification

$$Y_{ic}^k = \beta_1^k X_{ic}^k + \Gamma' W_i + \alpha_c + \varepsilon_{ic}^k, \tag{3.1}$$

which absorbs all time-invariant origin-country heterogeneity through country fixed effects. The coefficient $\beta_1^k$ therefore comes from differences across Berkeley stayers from the same origin country who have different individual preferences in the observed respondent sample.

To study home-country norms, I estimate

$$Y_{ic}^k = \beta_2^k N_c^k + \Gamma' W_i + \Theta' Z_c + u_{ic}^k, \tag{3.2}$$

which replaces country fixed effects with observed country controls. Here $\beta_2^k$ captures whether Berkeley-stayer respondents from countries with different norm realizations behave differently after conditioning on observed individual characteristics and observed differences in home-country pandemic and structural conditions. This coefficient is therefore a cross-origin association rather than a within-origin comparison. Because identification now comes from cross-country variation, the interpretation of $\beta_2^k$ is inherently more exposed to remaining cross-country composition concerns than the individual-preference specification above.

The joint specification is

$$Y_{ic}^k = \beta_3^k X_{ic}^k + \beta_4^k N_c^k + \Gamma' W_i + \Theta' Z_c + v_{ic}^k. \tag{3.3}$$

This regression asks whether the individual-level preference and the matched home-country norm retain separate associations with precautionary behavior when entered together. Unlike the country fixed-effects regression, however, the joint specification does not identify $\beta_3^k$ from within-origin-country variation alone; once country fixed effects are removed, both coefficients rely on comparisons across origin countries after conditioning on the observed

controls in equation (2). Across all three specifications, the estimation sample is dimension-specific because observations must be non-missing for the outcome, the relevant regressor or regressors, and the control set. Each reported coefficient should therefore be read as the relevant selected-sample relationship in its own complete-case estimation sample. Standard errors are heteroskedasticity-robust in equation (1), where the focal regressor varies at the individual level within countries, and clustered by origin country in equations (2) and (3), where the matched norm and country controls vary only at the country level.

## 3.4 Results

### 3.4.1 Main Preference and Norm Results

Figure 3.1 summarizes the main reduced-form evidence for the matched preference-norm set using separate dimension-by-dimension regressions. Each point is a standardized coefficient, so the estimates can be read as beta coefficients. I begin with the country fixed-effects specification, which uses within-origin-country individual variation among respondents who did not return to their countries of permanent residence. I then turn to the country-norm specification with observed individual and country controls, which asks whether respondents from different home-country environments behaved differently in the common institutional and pandemic environment. I finally estimate a joint specification that places the individual preference and the matched country norm in the same regression. Because the figure reports many related estimates, I emphasize stable patterns by construct family rather than treating each $p$-value as a separate confirmatory test.

Three patterns stand out in the first specification. First, more individualistic respondents undertook significantly fewer precautions: a one-standard-deviation increase in individualism is associated with a 0.10 standard-deviation decrease in the precaution index ($p = 0.004$). Trust is more measurement-sensitive. Using the World Values Survey trust question, the coefficient is negative and precisely estimated ($\beta = -0.08$, $p = 0.003$), whereas the GPS-style trust measure enters with the opposite sign and is only marginally significant in the country fixed-effects specification ($\beta = 0.05$, $p = 0.053$).

Second, positive reciprocity and patience are positively associated with precautionary behavior, with coefficients of 0.09 and 0.08 standard deviations, respectively ($p < 0.001$ and $p = 0.006$), whereas risk-taking is negatively associated with precautionary behavior ($\beta = -0.09$, $p = 0.002$). Altruism is positive and marginally significant at the 10 percent level ($\beta = 0.07$, $p = 0.060$). These results speak to the broader health-economics view that preventive behavior depends on risk tolerance, time preference, and prosocial concerns.[6] Because the Berkeley

[6]Useful entry points include Anderson (2008) on risk preferences, Chapman and Coups (1999) on vaccination and time preferences, and Müller and Rau (2021) and Campos-Mercade et al. (2021) on COVID-19 prosociality

survey measures risk-taking and patience with short self-assessment items while the matched GPS country norms are published composites that also incorporate staircase choice questions, the risk-taking and patience comparisons should be read with somewhat more measurement caution than the other matched GPS dimensions.[7]

Third, negative reciprocity and the survey's uncertainty-avoidance measure are not precisely estimated in this within-country comparison. These results suggest that, conditional on origin fixed effects, precautionary behavior is most strongly related to a small set of individual-level preferences rather than to the full menu of measured characteristics. The side-by-side trust estimates are a useful reminder that some substantively plausible constructs are not measured cleanly enough to support strong conclusions without additional robustness work.

The second specification shifts the focus from individual heterogeneity to home-country norms. Here the clearest result is uncertainty avoidance: respondents from more uncertainty-avoidant countries take significantly more precautions, with a coefficient of 0.16 standard deviations after controlling for individual covariates and a rich set of country-level pandemic and structural controls ($p = 0.039$). The origin-country individualism norm is negative and marginally significant ($\beta = -0.12$, $p = 0.079$). The matched country norms for trust, reciprocity, risk-taking, and patience are not precisely estimated in the main figure. The result on individualism is consistent with other COVID-19 studies that connect individualism to distancing, mobility, and compliance across places.[8] Appendix Figure 3.A1 considers alternative country-only cultural proxies. The patience-related constructs do not line up cleanly across measures. The main matched patience norm is close to zero, indulgence is also negative but imprecise, while long-term orientation is positive and marginally significant ($\beta = 0.12$, $p = 0.051$). Taken together, those contrasts suggest that patience-like country traits are sensitive to how the underlying construct is measured. The same appendix figure also points to additional country-level patterns outside the matched set. Power distance is positively associated with precautionary behavior ($\beta = 0.17$, $p = 0.001$), the masculinity coefficient is close to zero and imprecise, and World Values Survey altruism enters with a significant negative coefficient ($\beta = -0.18$, $p = 0.018$). I view those country-only results as suggestive rather than central, but they are useful for showing that some origin-country cultural dimensions correlate with precaution in ways not fully captured by the matched preference-norm pairs alone.

The third specification puts the individual preference and the matched country norm in the same regression. The main message from this exercise is that the most prominent individual-

and preferences.

[7] A future extension should harmonize these constructs more tightly by rebuilding country-level risk-taking and patience norms from the short-form GPS items alone. The obstacle is data availability: the published GPS country files do not include the underlying question-by-question responses needed to reconstruct short-form country aggregates comparable to the Berkeley survey items.

[8] See Bazzi et al. (2021); Bian et al. (2022) on rugged individualism.

level associations survive. Individualism, positive reciprocity, risk-taking, and patience remain statistically significant with very similar signs and magnitudes to the country fixed-effects specification. On the country side, uncertainty avoidance remains positive and significant when entered jointly with the individual-level measure ($\beta = 0.16$, $p = 0.037$), while the negative individualism country-norm coefficient remains marginally significant ($p = 0.063$). Comparing the second and third specifications, the largely unchanged country-norm coefficients after adding the matched individual preference suggest that these origin-country norm associations are not fully accounted for by the corresponding measured individual preference. Rather, country norms, namely individualism and uncertainty avoidance, may matter through broader channels. For individualism, the relevant channel may be socially learned expectations about the proper balance between personal autonomy and responsibility to others, or norms about whether private behavior should internalize risks imposed on the broader community. For uncertainty avoidance, the channel may be culturally shaped responses to ambiguity and collective risk, including greater comfort with precautionary routines, stronger expectations of rule-following, or a lower threshold for adopting protective behavior when scientific and policy guidance is still evolving. Another possibility is measurement error: the individual survey measures may capture only part of the underlying trait transmitted from the home-country environment. Thus, the stability of the country-norm coefficients should not be interpreted as evidence that cultural transmission is unimportant. Instead, it suggests that the relevant transmission may operate through channels broader than the measured individual preference itself, or that the respondent-level preference measure captures the underlying trait with error.

Appendix Figure 3.A2 reports pooled within-family specifications. Panel A enters the six social and cultural orientation dimensions together, and Panel B does the same for the three economic preference dimensions. The motivation is to ask which results remain visible once conceptually similar dimensions are estimated jointly, while avoiding a single specification that combines all dimension families at once. The drawback is that the pooled coefficients are mechanically noisier and often harder to interpret, especially for country norms, because the dimensions remain correlated and the country-level variation is limited. I therefore treat this appendix figure as a robustness exercise rather than as the main design. Its broad takeaway nevertheless aligns with Figure 3.1: the most robust individual-level patterns continue to be lower precaution among more individualistic, World Values Survey-trusting, and risk-tolerant respondents, and higher precaution among more positively reciprocal and patient respondents. By contrast, the country-norm estimates are less stable in the pooled comparisons, reinforcing my emphasis on the same-period individual-preference evidence and on uncertainty avoidance as the most consistent norm-level result.

### 3.4.2 Auxiliary Belief and Information Evidence

Appendix Tables 3.A3 and 3.A4 report auxiliary analyses using beliefs about precaution effectiveness and COVID-related news attention. I treat these results as suggestive rather than as formal mediation evidence, since the belief and information measures are measured alongside the behavioral outcome rather than in a design that separates causal channels.

Appendix Table 3.A3 shows that the strongest individual-level patterns in the main figure line up in sensible, though not definitive, ways with these auxiliary measures. More individualistic and more risk-tolerant respondents pay less attention to COVID-related news, and their precaution coefficients attenuate somewhat once the belief and news indices are added to the precaution regression. Positive reciprocity and patience, by contrast, are associated with stronger precaution-related beliefs and greater news attention, and their precaution coefficients also attenuate somewhat when those auxiliary variables are included. The attenuation is especially noticeable for patience, whose coefficient becomes smaller and less precise in the augmented specification. I read these patterns as consistent with the idea that differences in beliefs and information exposure help organize the reduced-form preference results.

Appendix Table 3.A4 provides the parallel exercise for country norms. The most notable pattern there is that the home-country individualism norm positively predicts the belief index, whereas neither the individualism norm nor the uncertainty-avoidance norm strongly predicts the news index. When the belief and news indices are added to the precaution regression, the norm coefficients remain broadly similar, and uncertainty avoidance stays positive. This stability is useful descriptively because it suggests that the norm-level associations are not simply absorbed by these two auxiliary measures.

## 3.5 Conclusion

This paper uses a Berkeley international student setting to ask whether culture and individual preferences travel with people into a new public-health environment. The answer is partly yes, but the more useful lesson is about how they travel. Among international students who did not return to their countries of permanent residence during the same early-pandemic period, precautionary behavior is not explained by origin-country background alone. Respondents from the same origin country behaved differently depending on their own individualism, risk tolerance, reciprocity, and patience. At the same time, origin-country norms, especially uncertainty avoidance and individualism, continue to predict behavior after controlling for the matched respondent-level preferences. The evidence therefore points away from a simple model in which culture matters only because it shapes individual preferences.

This distinction matters for the interpretation of culture in empirical work. Same-environment designs are often used to reduce variation in institutions, policy, and local conditions so that

remaining differences can be interpreted as portable cultural influences. The Berkeley setting follows that logic, but it also shows that portable influences operate at more than one level. Some are visible as individual preferences: people who are more willing to bear risk, who put more weight on autonomy, or who discount future benefits more heavily appear less inclined to bear the private costs of precaution. Others reflect shared ways of interpreting situations rather than individual preferences: uncertainty-avoidant origin environments may shape how people respond to ambiguity, rules, and collective risk even when the measured individual uncertainty-avoidance item is held fixed. Likewise, individualism at the country level may proxy for norms about the boundary between private choice and social obligation that are broader than a respondent's own individualism.

The findings also speak to public-health policy. Communication that treats noncompliance as a homogeneous information problem will miss important heterogeneity. For some people, precaution may need to be framed around delayed private benefits or risk management; for others, around reciprocity and the external benefits of protecting classmates, co-workers, and neighbors. More broadly, culturally resonant policy need not mean tailoring rules by nationality. It means recognizing that the same rule can be interpreted through different expectations about autonomy, uncertainty, and mutual obligation. In a mobile university community, those expectations coexist in the same local environment.

The design has clear limits. The analysis sample consists of survey respondents who did not return to their countries of permanent residence, not a random sample of all international students or of all people from the origin countries represented in the data. The survey was fielded during a specific and unusual moment in May 2020, when information, mask norms, and policy guidance were changing quickly. The estimates are associational, and several constructs, especially trust, are sensitive to measurement choices.

The next step is to connect this kind of common-institution evidence with richer designs that observe behavior over time, measure social expectations directly, and distinguish information, beliefs about others' behavior, and intrinsic preferences. Doing so would help clarify when culture works through stable individual preferences and when it works through beliefs activated by a particular crisis. The COVID-19 pandemic made that distinction unusually visible, but the same issue applies to vaccination, staying home when sick, climate adaptation, and other behaviors whose private costs and social benefits are unevenly distributed. Understanding culture in those settings requires measuring both the person and the social scripts the person carries into the decision.

# Figures

Figure 3.1: Precautionary Behavior and Matched Traits

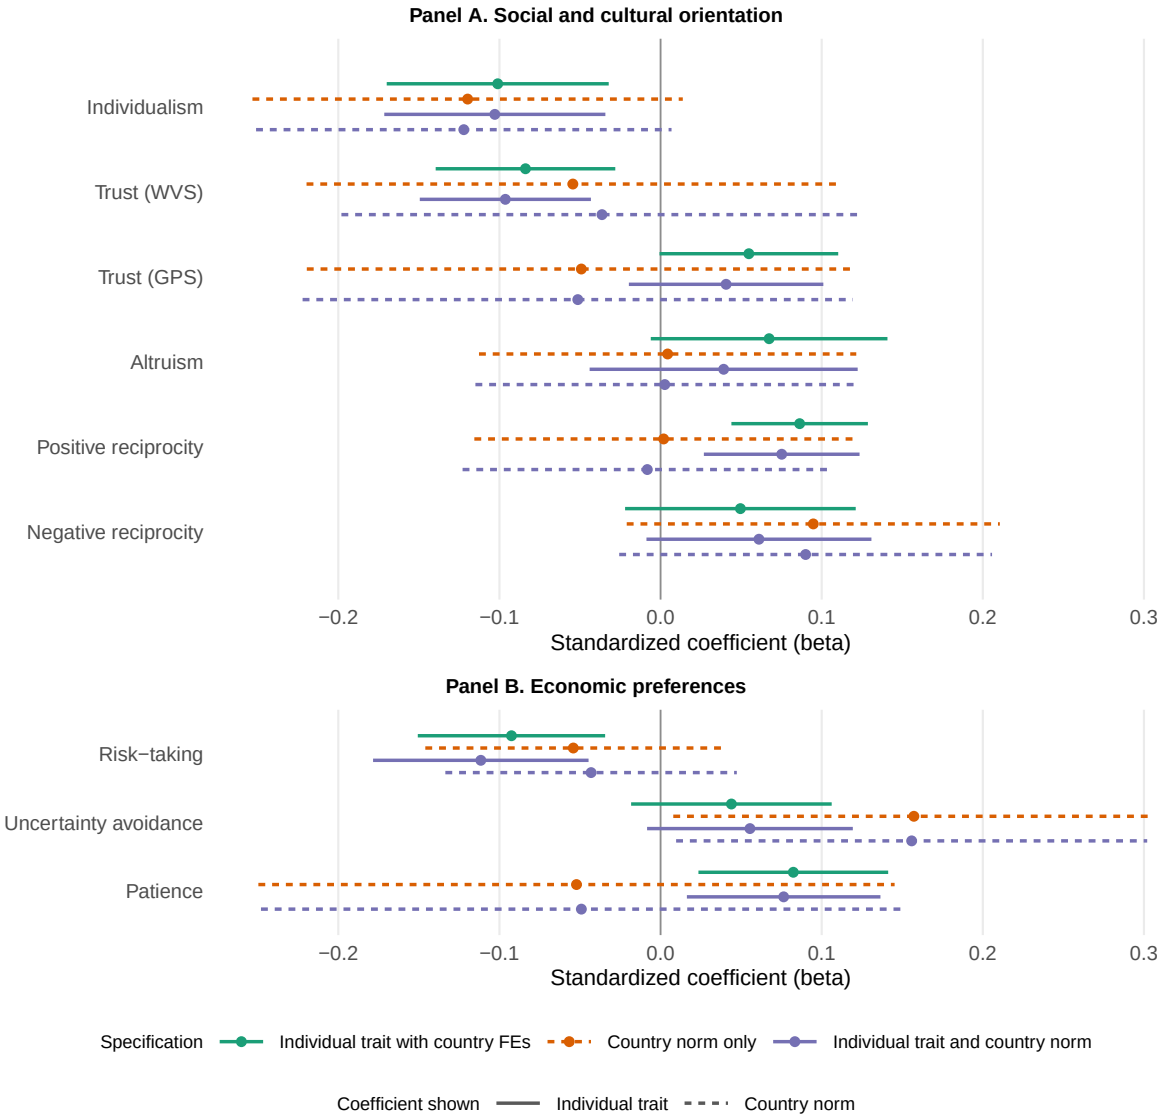

### Notes to Figure 3.1: Precautionary Behavior and Matched Traits

The unit of observation is an international student who did not return to the country of permanent residence during the survey period. The outcome is the standardized precaution index, where higher values indicate fewer days leaving home, a lower need to leave home, less participation in social gatherings, more social distancing, greater use of face covering, and more frequent handwashing. Panel A reports social and cultural orientation dimensions: individualism, trust (WVS), trust (GPS), altruism, positive reciprocity, and negative reciprocity. Panel B reports economic preference dimensions: risk-taking, uncertainty avoidance, and patience. For each dimension, the plotted coefficients come from three regressions estimated on a common dimension-specific sample: an individual-preference specification with country fixed effects and individual controls, a standalone country-norm specification with individual and country controls, and a joint specification that includes both the individual preference and the matched country norm. Solid confidence-interval lines denote individual-preference estimates, and dashed lines denote country-norm estimates. In every plotted regression, both the outcome and the plotted variable are standardized within the exact estimation sample, so the coefficients can be read as beta coefficients. Horizontal bars indicate 95 percent confidence intervals. Country controls include recent COVID cases per capita, the COVID-19 death rate, policy stringency, log GDP per capita, log Gini, log population density, democracy, and SARS exposure. Individual controls include age, gender, graduate status, religiosity, home-country income decile, years in the United States, years at Berkeley, school indicators, and the social desirability scale.

# Tables

Table 3.1: Summary Statistics for the Stayed-in-Berkeley Sample

| | Mean | SD | N |
|---|---|---|---|
| **Panel A. Precautionary behaviors** | | | |
| Precaution index | 0.061 | 0.939 | 899 |
| Days left home in past 7 days | 2.523 | 2.127 | 899 |
| Needs to leave home in next 5 days | 0.576 | 0.494 | 899 |
| Covered nose or mouth when outside | 0.915 | 0.278 | 899 |
| Avoided social gatherings | 9.018 | 2.317 | 899 |
| Practiced social distancing | 8.909 | 1.742 | 899 |
| Washed hands more frequently | 8.651 | 2.382 | 899 |
| **Panel B. Individual Preferences** | | | |
| Individualism | 0.010 | 0.967 | 899 |
| Trust (WVS) | -0.018 | 0.998 | 899 |
| Trust (GPS) | -0.002 | 1.013 | 899 |
| Altruism | -0.034 | 0.985 | 899 |
| Positive reciprocity | 0.009 | 0.990 | 899 |
| Negative reciprocity | 0.014 | 1.015 | 899 |
| Risk-taking | -0.040 | 1.012 | 899 |
| Uncertainty avoidance | -0.013 | 0.993 | 899 |
| Patience | 0.000 | 0.986 | 899 |
| **Panel C. Country norms** | | | |
| Individualism norm | -0.405 | 0.883 | 856 |
| Trust (WVS) norm | 0.838 | 1.062 | 870 |
| Trust (GPS) norm | 0.169 | 0.354 | 798 |
| Altruism norm | 0.164 | 0.367 | 798 |
| Positive reciprocity norm | 0.092 | 0.396 | 798 |
| Negative reciprocity norm | 0.094 | 0.224 | 798 |
| Risk-taking norm | -0.065 | 0.169 | 798 |
| Uncertainty-avoidance norm | -0.702 | 1.016 | 856 |
| Patience norm | 0.198 | 0.306 | 798 |
| **Panel D. Selected individual controls and auxiliary measures** | | | |
| Age | 25.661 | 4.705 | 899 |
| Female | 0.485 | 0.500 | 899 |
| Graduate student | 0.641 | 0.480 | 899 |
| Religious | 0.409 | 0.492 | 899 |
| Income decile in origin country | 6.522 | 1.902 | 899 |
| Years in the US | 3.299 | 2.588 | 899 |
| Years at Berkeley | 2.209 | 1.539 | 899 |
| Social desirability scale | 0.584 | 0.882 | 899 |
| Belief index | 0.016 | 1.013 | 899 |
| News-following index | -0.026 | 1.008 | 899 |
| **Panel E. Country-level controls** | | | |
| Log recent COVID cases per capita | 0.775 | 3.582 | 895 |
| COVID death rate | 0.050 | 0.034 | 899 |
| Policy stringency (7-day average) | 69.098 | 17.932 | 899 |
| Log GDP per capita | 9.599 | 0.770 | 899 |
| Democracy | 0.602 | 0.490 | 899 |

**Notes:** The sample is limited to the 899 Berkeley international students who reported not returning to their country of permanent residence during the May 4–18, 2020 survey window. The sample spans 76 origin countries. In Panel A, the last three precautionary-behavior items are reported on 0–10 scales. Panel B reports the standardized individual-preference measures used in the matched-trait analysis. Panel C reports the merged country-norm variables available in the Berkeley-stayer sample; sample sizes are smaller where matched country data are unavailable. Panel D reports selected individual controls and auxiliary measures. Panel E reports selected country-level controls matched to the respondent's survey date, with recent cases measured over the preceding seven days and policy stringency measured as the average over the preceding seven days.

Table 3.2: Representativeness of the Berkeley International Students Survey Sample

| | Population (%) | Survey Sample (%) | Difference (pp) |
|---|---|---|---|
| **Panel A. Gender and student status** | | | |
| Female | 44.3 | 50.1 | 5.8 |
| Graduate student | 42.9 | 56.7 | 13.9 |
| **Panel B. Top origin-country shares** | | | |
| China | 40.4 | 28.9 | -11.5 |
| India | 10.4 | 12.0 | 1.7 |
| South Korea | 8.1 | 6.1 | -2.0 |
| Canada | 5.8 | 5.0 | -0.8 |
| **Panel C. School shares** | | | |
| Letters and Science | 39.8 | 26.3 | -13.5 |
| Engineering | 24.2 | 27.1 | 2.9 |
| Graduate Division | 9.3 | 8.8 | -0.4 |
| Business | 5.9 | 7.3 | 1.5 |
| Chemistry | 4.9 | 5.0 | 0.0 |
| Environmental Design | 4.3 | 5.0 | 0.7 |
| Natural Resources | 4.2 | 5.9 | 1.7 |
| Boalt Law | 4.0 | 5.2 | 1.2 |
| Other schools | 3.3 | 9.3 | 6.0 |

**Notes:** Population shares come from Berkeley International Office Fall 2019 enrollment records for international students. The difference column reports survey share minus population share in percentage points. Panels A and B are calculated as shares of 6,833 international students. Panel C is calculated as shares of 6,529 degree-seeking undergraduate and graduate students, excluding 304 EAP exchange students. Survey size is 1130.

Table 3.3: Selection into Staying in Berkeley

| | Stayed | Returned | Raw diff. | p-value | Adj. diff. | p-value |
|---|---|---|---|---|---|---|
| **Demographics and Berkeley exposure** | | | | | | |
| Age | 25.661 | 22.597 | -3.064 | 0.000 | -0.949 | 0.000 |
| Female | 0.485 | 0.546 | 0.061 | 0.092 | 0.025 | 0.541 |
| Graduate student | 0.641 | 0.273 | -0.368 | 0.000 | -0.130 | 0.000 |
| Religious | 0.409 | 0.504 | 0.095 | 0.009 | 0.077 | 0.052 |
| Income decile in home country | 6.522 | 7.038 | 0.516 | 0.000 | 0.351 | 0.014 |
| Years in the US | 3.299 | 2.164 | -1.135 | 0.000 | -0.345 | 0.032 |
| Years at Berkeley | 2.209 | 1.521 | -0.688 | 0.000 | -0.300 | 0.000 |
| Survey day since May 4, 2020 | 1.941 | 2.462 | 0.521 | 0.032 | 0.541 | 0.050 |
| **Individual Preferences** | | | | | | |
| Individualism | 0.000 | -0.026 | -0.026 | 0.789 | -0.005 | 0.957 |
| Trust (WVS) | 0.000 | 0.011 | 0.011 | 0.896 | 0.014 | 0.879 |
| Trust (GPS) | 0.000 | -0.014 | -0.014 | 0.864 | -0.013 | 0.887 |
| Altruism | -0.000 | 0.199 | 0.199 | 0.025 | 0.128 | 0.169 |
| Positive reciprocity | -0.000 | -0.033 | -0.033 | 0.706 | -0.058 | 0.546 |
| Negative reciprocity | 0.000 | -0.037 | -0.037 | 0.631 | 0.100 | 0.255 |
| Risk-taking | -0.000 | 0.217 | 0.217 | 0.006 | -0.011 | 0.900 |
| Uncertainty avoidance | 0.000 | 0.056 | 0.056 | 0.521 | 0.108 | 0.267 |
| Patience | -0.000 | 0.063 | 0.063 | 0.473 | 0.032 | 0.752 |
| **Country norms** | | | | | | |
| Individualism | 0.000 | 0.659 | 0.659 | 0.000 | 0.139 | 0.009 |
| Trust (WVS) | 0.000 | -0.278 | -0.278 | 0.000 | 0.074 | 0.106 |
| Trust (GPS) | -0.000 | -0.277 | -0.277 | 0.000 | 0.092 | 0.060 |
| Altruism | -0.000 | -0.260 | -0.260 | 0.001 | 0.004 | 0.950 |
| Positive reciprocity | -0.000 | -0.262 | -0.262 | 0.000 | 0.026 | 0.649 |
| Negative reciprocity | -0.000 | 0.081 | 0.081 | 0.368 | -0.156 | 0.025 |
| Risk-taking | -0.000 | 0.140 | 0.140 | 0.147 | 0.101 | 0.196 |
| Uncertainty avoidance | -0.000 | 0.254 | 0.254 | 0.001 | -0.146 | 0.004 |
| Patience | 0.000 | 0.230 | 0.230 | 0.015 | 0.156 | 0.001 |

**Notes:** The table compares respondents who stayed in Berkeley during the survey window with respondents who returned to their home countries. Panel A is reported in natural units. Panels B and C are standardized using the stayer subsample, so stayers have zero means and unit variances by construction and returnees' means are reported in stayer-sample SD units. Raw differences come from univariate OLS regressions on a returned-home indicator. Adjusted differences use controls for age, gender, graduate status, religiosity, home-country income decile, years in the US, years at Berkeley, school indicators, the social desirability scale, recent COVID cases per capita, the death rate, policy stringency, log GDP per capita, log Gini, log population density, democracy, and SARS exposure. $p$-values are based on heteroskedasticity-robust standard errors.

# Appendix

## D Exact Wording for Individual Preference Measures

Table 3.A1: Survey Wording: Social Orientation Preferences

| Preference | Exact survey wording | Response options and construction |
|---|---|---|
| Individualism | *Q6:* "Please think of an ideal job, disregarding your present job, if you have one. In choosing an ideal job, how important would it be to you to . . . " The four items used in the main measure are: "Have sufficient time for your personal or home life"; "Have security of employment"; "Do work that is interesting"; "Have a job respected by your family and friends." | Five-point importance scale: "Of no importance"; "Of little importance"; "Of moderate importance"; "Very important"; "Of utmost importance." The individualism measure is (time/home life + interesting work)/(job security + family respect), standardized before estimation. |
| Trust (WVS) | *Q7:* "Generally speaking (disregarding the COVID-19 pandemic), would you say that most people can be trusted or that you need to be very careful in dealing with people?" | Two response options: "Most people can be trusted" and "Need to be very careful." The trust measure is an indicator for choosing "Most people can be trusted," standardized before estimation. |
| Trust (GPS) | *Q20:* "I assume that people have only the best intentions." | 0–10 scale from "Does not describe me at all" to "Describes me perfectly." The GPS-style trust measure is the standardized self-assessment. |
| Altruism | *Q21:* "How willing are you to share with others without expecting anything in return when it comes to charity?" *Q23:* "Imagine the following situation: Today you unexpectedly received 1,600 U.S. dollars. How much of this amount would you donate to a charity?" | *Q21* uses a 0–10 scale from "Completely unwilling to do so" to "Very willing to do so." *Q23* allows any amount from 0 to 1,600 U.S. dollars. The altruism measure is $0.3650 \times z(\text{self-assessment}) + 0.6350 \times z(\text{donation amount})$, standardized again after aggregation. |

*Continued on next page*

*Table 3.A1 continued*

| Preference | Exact survey wording | Response options and construction |
|---|---|---|
| Positive reciprocity | *Q20:* "When someone does me a favor, I am willing to return it."<br>*Q22:* "Do you give one of the presents to the stranger as a 'thank you' gift?" after a vignette in which a stranger helps the respondent at a cost of about 40 U.S. dollars.<br>*Q24:* "Which present do you give to the stranger?" | *Q20* uses a 0–10 scale from "Does not describe me at all" to "Describes me perfectly."<br>*Q22* offers "No, would not give present" and "Yes, would give a present."<br>If yes, *Q24* asks respondents to choose a present worth 10, 20, 30, 40, 50, or 60 U.S. dollars; respondents choosing no gift are coded as zero on the quantitative component.<br>The positive reciprocity measure is $0.4847 \times z(\text{favor-return item}) + 0.5153 \times z(\text{gift amount})$, standardized again after aggregation. |
| Negative reciprocity | *Q21:* "How willing are you to punish someone who treats you unfairly, even if there may be costs for you?" and "How willing are you to punish someone who treats others unfairly, even if there may be costs for you?"<br>*Q20:* "If I am treated very unjustly, I will take revenge at the first occasion, even if there is a cost to do so." | *Q21* uses a 0–10 scale from "Completely unwilling to do so" to "Very willing to do so."<br>*Q20* uses a 0–10 scale from "Does not describe me at all" to "Describes me perfectly."<br>The negative reciprocity measure is $0.3131 \times z(\text{punish unfairly}) + 0.3131 \times z(\text{punish others unfairly}) + 0.3738 \times z(\text{revenge item})$, standardized again after aggregation. |

Table 3.A2: Survey Wording: Preferences and Uncertainty

| Preference | Exact survey wording | Response options and construction |
|---|---|---|
| Risk-taking | *Q8:* "Please tell me, in general, how willing or unwilling you are to take risks." | 0–10 scale from "Completely unwilling to take risks" to "Very willing to take risks." The risk-taking measure is the standardized self-assessment. |
| Uncertainty avoidance | *Q13:* "To what extent do you agree or disagree with each of the following statements?" The two items used are: "One can be a good manager without having a precise answer to every question that a subordinate may raise about his or her work" and "A company's or organization's rules should not be broken – not even when the employee thinks breaking the rule would be in the organization's best interest." *Q11:* "How often do you feel nervous or tense?" *Q12:* "All in all, how would you describe your state of health these days?" | *Q13* uses a five-point agreement scale from "Strongly disagree" to "Strongly agree." *Q11* uses "Never," "Seldom," "Sometimes," "Usually," and "Always." *Q12* uses "Very poor," "Poor," "Fair," "Good," and "Very good." The uncertainty-avoidance measure is the cleaned-survey four-item index $40\times(\text{health}-\text{nervousness})+25\times$ (managerial precision item − rule-following item), standardized before estimation. |
| Patience | *Q21:* "How willing are you to give up something that is beneficial for you today in order to benefit more from that in the future?" | 0–10 scale from "Completely unwilling to do so" to "Very willing to do so." The patience measure is the standardized self-assessment. |

# Appendix Tables

Table 3.A3: Auxiliary Belief and Information Evidence: Individual Preferences

| | Individualism | Risk-taking | Positive reciprocity | Patience |
|---|---|---|---|---|
| **Panel A. Belief index** | | | | |
| Preference | -0.036 | -0.045 | 0.104*** | 0.148*** |
| | (0.030) | (0.055) | (0.039) | (0.039) |
| **Panel B. News index** | | | | |
| Preference | -0.106*** | -0.031 | 0.060** | 0.085* |
| | (0.033) | (0.035) | (0.025) | (0.049) |
| **Panel C. Precaution index** | | | | |
| *Baseline* | | | | |
| Preference | -0.087*** | -0.096*** | 0.074*** | 0.069** |
| | (0.033) | (0.030) | (0.019) | (0.027) |
| *With auxiliary controls* | | | | |
| Preference | -0.058* | -0.083*** | 0.048*** | 0.033 |
| | (0.032) | (0.024) | (0.018) | (0.028) |
| Belief index | 0.112*** | 0.110*** | 0.109*** | 0.109*** |
| | (0.034) | (0.034) | (0.034) | (0.035) |
| News index | 0.239*** | 0.243*** | 0.242*** | 0.243*** |
| | (0.023) | (0.024) | (0.024) | (0.024) |
| Individual controls | Yes | Yes | Yes | Yes |
| Country controls | No | No | No | No |
| Country FE | Yes | Yes | Yes | Yes |
| Observations | 899 | 899 | 899 | 899 |
| Countries | 76 | 76 | 76 | 76 |

**Notes:** The unit of observation is a Berkeley international student who remained in Berkeley during the survey period. All columns and panels use the same estimation sample. Panel A uses the standardized belief index, Panel B the standardized news index, and Panel C the standardized precaution index. The belief index is inverse-covariance weighted from the retained precaution-effectiveness items, and the news index is inverse-covariance weighted from following COVID-related news in the U.S. and in the respondent's home country. All specifications include individual controls and country fixed effects. Robust standard errors are in parentheses. The belief and news measures are auxiliary variables used to interpret the main reduced-form patterns.

Table 3.A4: Auxiliary Belief and Information Evidence: Country Norms

| | Individualism norm | Uncertainty-avoidance norm |
|---|---|---|
| **Panel A. Belief index** | | |
| Norm | 0.112** | -0.037 |
| | (0.055) | (0.059) |
| **Panel B. News index** | | |
| Norm | -0.047 | 0.098 |
| | (0.085) | (0.077) |
| **Panel C. Precaution index** | | |
| *Baseline* | | |
| Norm | -0.136* | 0.144** |
| | (0.072) | (0.071) |
| *With auxiliary controls* | | |
| Norm | -0.136** | 0.123* |
| | (0.066) | (0.062) |
| Belief index | 0.105*** | 0.103*** |
| | (0.034) | (0.033) |
| News index | 0.256*** | 0.255*** |
| | (0.024) | (0.023) |
| Individual controls | Yes | Yes |
| Country controls | Yes | Yes |
| Country FE | No | No |
| Observations | 853 | 853 |
| Countries | 50 | 50 |

**Notes:** The unit of observation is a Berkeley international student who remained in Berkeley during the survey period. Columns may use slightly different samples because matched norm data are unavailable for some origin countries. Panel A uses the standardized belief index, Panel B the standardized news index, and Panel C the standardized precaution index. The belief index is inverse-covariance weighted from the retained precaution-effectiveness items, and the news index is inverse-covariance weighted from following COVID-related news in the U.S. and in the respondent's home country. All specifications include individual and country controls. Standard errors clustered by country of origin are in parentheses. The belief and news measures are auxiliary variables used to interpret the main reduced-form patterns.

# Appendix Figures

Figure 3.A1: Precautionary Behavior and Country-Only Cultural Dimensions

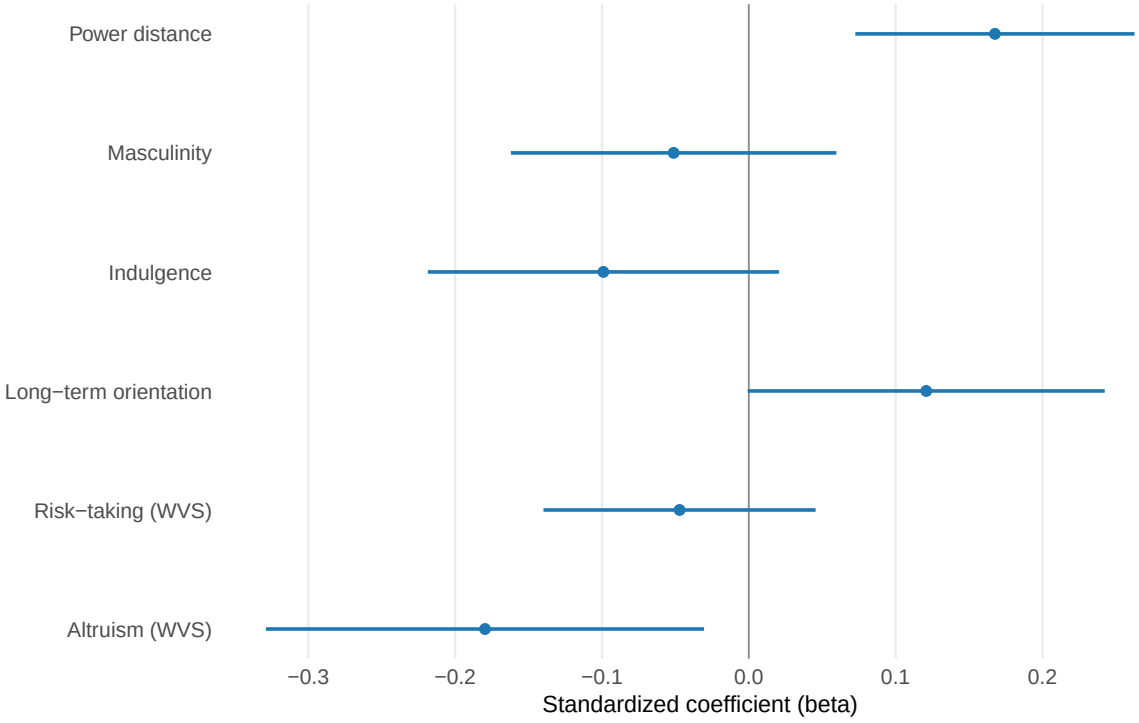

## Notes to Figure 3.A1: Precautionary Behavior and Country-Only Cultural Dimensions

The unit of observation is an international student who did not return to the country of permanent residence during the survey period. The outcome is the standardized precaution index, where higher values indicate fewer days leaving home, a lower need to leave home, less participation in social gatherings, more social distancing, greater use of face covering, and more frequent handwashing. Each point reports the coefficient on one country-level cultural measure from a separate regression that controls for individual covariates and country-level pandemic and structural controls. The country-only dimensions are power distance, masculinity, indulgence, long-term orientation, World Values Survey risk-taking, and World Values Survey altruism. For each plotted estimate, both the outcome and the country-level measure are standardized within the exact estimation sample, so the coefficients can be read as beta coefficients. Horizontal bars indicate 95 percent confidence intervals. The individual and country controls are the same as those described in Figure 3.1.

Figure 3.A2: Precautionary Behavior and Matched Traits in Pooled Specifications

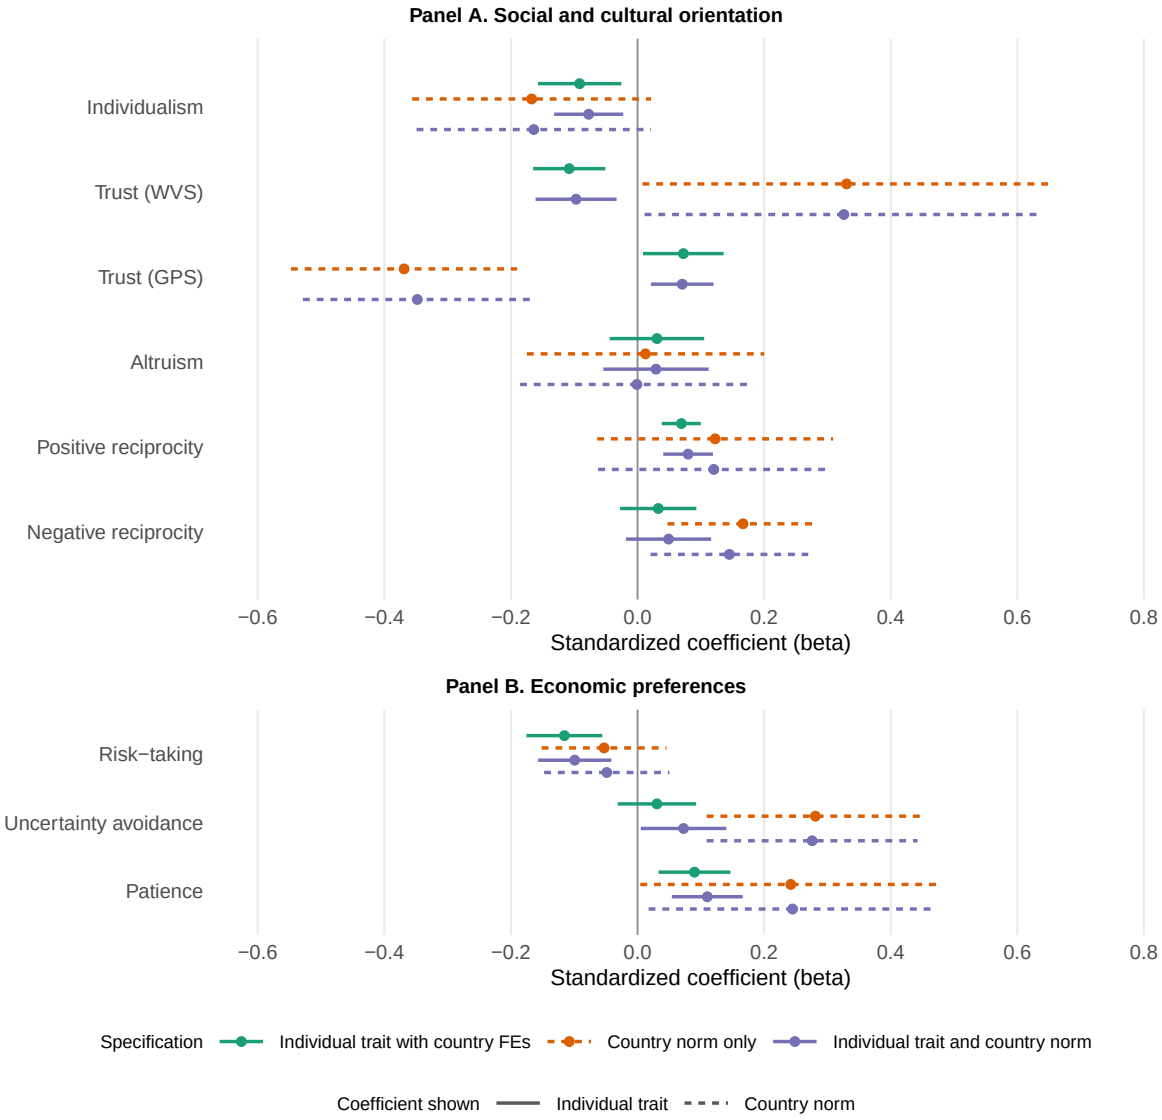

## Notes to Figure 3.A2: Precautionary Behavior and Matched Traits in Pooled Specifications

The unit of observation is an international student who did not return to the country of permanent residence during the survey period. The outcome is the standardized precaution index, where higher values indicate fewer days leaving home, a lower need to leave home, less participation in social gatherings, more social distancing, greater use of face covering, and more frequent handwashing. Panel A reports social and cultural orientation dimensions: individualism, trust (WVS), trust (GPS), altruism, positive reciprocity, and negative reciprocity. Panel B reports economic preference dimensions: risk-taking, uncertainty avoidance, and patience. The plotted coefficients come from pooled regressions estimated separately within each panel family rather than separate dimension-by-dimension models. Within Panel A, the pooled specifications include only the six social and cultural orientation dimensions and their matched country norms. Within Panel B, the pooled specifications include only the three economic preference dimensions and their matched country norms. For each family, the figure shows the individual-preference coefficient from the country fixed-effects specification, the country-norm coefficient from the standalone country-norm specification with controls, and both the individual-preference and country-norm coefficients from the joint specification. Solid confidence-interval lines denote individual-preference estimates, and dashed lines denote country-norm estimates. In every plotted regression, both the outcome and the plotted variable are standardized within that regression's exact estimation sample, so the coefficients can be read as beta coefficients. Horizontal bars indicate 95 percent confidence intervals. Country controls include recent COVID cases per capita, the COVID-19 death rate, policy stringency, log GDP per capita, log Gini, log population density, democracy, and SARS exposure. Individual controls include age, gender, graduate status, religiosity, home-country income decile, years in the United States, years at Berkeley, school indicators, and the social desirability scale.

# References


Acemoglu, Daron, "Theory, General Equilibrium, and Political Economy in Development Economics," *Journal of Economic Perspectives*, 2010, *24* (3), 17–32.

_ , Simon Johnson, and James A Robinson, "The Colonial Origins of Comparative Development: An Empirical Investigation," *American Economic Review*, dec 2001, *91* (5), 1369–1401.

_ , Tristan Reed, and James A. Robinson, "Chiefs: Economic Development and Elite Control of Civil Society in Sierra Leone," *Journal of Political Economy*, 2014, *122* (2), 319–368.

Adukia, Anjali, Sam Asher, and Paul Novosad, "Educational Investment Responses to Economic Opportunity: Evidence from Indian Road Construction," *American Economic Journal: Applied Economics*, 2020, *12* (1), 348–376.

Alesina, Alberto and Paola Giuliano, "Culture and Institutions," *Journal of Economic Literature*, 2015, *53* (4), 898–944.

Anderson, Michael L., "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects," *Journal of the American Statistical Association*, 2008, *103* (484), 1481–1495.

Arhin, Kwame., *Traditional Rule in Ghana: Past and Present*, Accra: Sedco, 1985.

Armand, Alex, Alexander Coutts, Pedro C Vicente, and Inês Vilela, "Does Information Break the Political Resource Curse? Experimental Evidence from Mozambique," *American Economic Review*, 2020, *110* (11), 3431–3453.

Asher, Sam and Paul Novosad, "Rural Roads and Local Economic Development," *American Economic Review*, 2020, *110* (3), 797–823.

Avis, Eric, Claudio Ferraz, and Frederico Finan, "Do Government Audits Reduce Corruption? Estimating the Impacts of Exposing Corrupt Politicians," *Journal of Political Economy*, oct 2018, *126* (5), 1912–1964.

Ayittey, George B.N., "Traditional Institutions and the State of Accountability in Africa," *Social Research*, 2010, *77* (4), 1183–1210.

Balán, Pablo, Augustin Bergeron, Gabriel Tourek, and Jonathan L. Weigel, "Local Elites as State Capacity: How City Chiefs Use Local Information to Increase Tax Compliance in the Democratic Republic of the Congo," *American Economic Review*, mar 2022, *112* (3), 762–797.

Baland, Jean-Marie. and J. P.. Platteau, *Halting Degradation of Natural Resources: Is There a Role for Rural Communities?*, Oxford University Press, 2000.

Baldwin, Kate, *The Paradox of Traditional Chiefs in Democratic Africa* Cambridge studies in comparative politics, first edit ed., Cambridge: Cambridge University Press, 2015.

Bamfo, Napoleon, "The Hidden Elements of Democracy among Akyem Chieftaincy: Enstoolment, Destoolment, and Other Limitations of Power," *Journal of Black Studies*, 2000, *31* (2), 149–173.

Banerjee, Abhijit, Nils Enevoldsen, Rohini Pande, and Michael Walton, "Public Information Is an Incentive for Politicians: Experimental Evidence from Delhi Elections," 2020, (1063693).

_ , Rema Hanna, Jordan Kyle, Benjamin A. Olken, and Sudarno Sumarto, "Tangible Information and Citizen Empowerment: Identification Cards and Food Subsidy Programs in Indonesia," *Journal of Political Economy*, apr 2018, *126* (2), 451–491.

Banerjee, Abhijit V., Eliana La Ferrara, and Victor H. Orozco-Olvera, "The Entertaining Way to Behavioral Change: Fighting HIV with MTV," 2019.

_ , Selvan Kumar, Rohini Pande, and Felix Su, "Do Informed Voters Make Better Choices? Experimental Evidence from Urban India," 2011.

Bargain, Olivier and Ulugbek Aminjonov, "Trust and Compliance to Public Health Policies in Times of COVID-19," *Journal of Public Economics*, 2020, *192*, 104316.

Barrios, John M, Efraim Benmelech, Yael V Hochberg, Paola Sapienza, and Luigi Zingales, "Civic Capital and Social Distancing During the COVID-19 Pandemic," *Journal of Public Economics*, 2021, *193*, 104310.

Barwick, Panle Jia, Shanjun Li, Liguo Lin, and Eric Yongchen Zou, "From Fog to Smog: The Value of Pollution Information," *American Economic Review*, may 2024, *114* (5), 1338–1381.

Basurto, Maria Pia, Pascaline Dupas, and Jonathan Robinson, "Decentralization and Efficiency of Subsidy Targeting: Evidence from Chiefs in Rural Malawi," *Journal of Public Economics*, may 2020, *185*.

Bazzi, Samuel, Martin Fiszbein, and Mesay Gebresilasse, "Rugged Individualism and Collective (In)action During the COVID-19 Pandemic," *Journal of Public Economics*, 2021, *195*, 104357.

Bennear, Lori, Alessandro Tarozzi, Alexander Pfaff, Soumya Balasubramanya, Kazi Matin Ahmed, and Alexander van Geen, "Impact of a Randomized Controlled Trial in Arsenic Risk Communication on Household Water-Source Choices in Bangladesh," *Journal of Environmental Economics and Management*, mar 2013, *65* (2), 225–240.

Besley, T. and R. Burgess, "The Political Economy of Government Responsiveness: Theory and Evidence from India," *The Quarterly Journal of Economics*, nov 2002, *117* (4), 1415–1451.

Besley, Timothy, *Principled Agents?: The Political Economy of Good Government*, Oxford University Press, aug 2007.

Bian, Bo, Jingjing Li, Ting Xu, and Natasha Z Foutz, "Individualism During Crises," *Review of Economics and Statistics*, 2022, *104* (2), 368–385.

Bisin, Alberto and Thierry Verdier, "The Economics of Cultural Transmission and Socialization," *Handbook of Social Economics*, 2011, *1A*, 339–416.

Blair, Robert A, Benjamin S Morse, and Lily L Tsai, "Public Health and Public Trust: Survey Evidence from the Ebola Virus Disease Epidemic in Liberia," *Social Science & Medicine*, 2017, *172*, 89–97.

Bobonis, Gustavo J., Luis R. Cámara Fuertes, and Rainer Schwabe, "Monitoring Corruptible Politicians," *American Economic Review*, 2016, *106* (8), 2371–2405.

Bohlken, Anjali Thomas, "Development or Rent-Seeking? How Political Influence Shapes Public Works Provision in India," *British Journal of Political Science*, 2021, *51* (1), 253–274.

Bursztyn, Leonardo and David Y. Yang, "Misperceptions about Others," *Annual Review of Economics*, 2022, *14*, 425–452.

Campos-Mercade, Pol, Armando N Meier, Florian H Schneider, and Erik Wengström, "Prosociality Predicts Health Behaviors During the COVID-19 Pandemic," *Journal of Public Economics*, 2021, *195*, 104367.

Chapman, Gretchen B and Elliot J Coups, "Time Preferences and Preventive Health Behavior: Acceptance of the Influenza Vaccine," *Medical Decision Making*, 1999, *19* (3), 307–314.

Chopra, Vir K., *Marginal Players in Marginal Assemblies: The Indian MLA*, Sangam Books, 1996.

Chuang, Yingchih, Ya-Lu Huang, Kuo-Chien Tseng, Chia-Hsin Yen, and Lin-Hui Yang, "Social Capital and Health-Protective Behavior Intentions in an Influenza Pandemic," *PLOS ONE*, 2015, *10* (4), e0122970.

Crawford, Gordon, Coleman Agyeyomah, Gabriel Botchwey, and Atinga Mba, "The Impact of Chinese Involvement in Small-Scale Gold Mining in Ghana," Technical Report May, International Growth Centre (IGC) 2015.

Dechezleprêtre, Antoine, Adrien Fabre, Tobias Kruse, Bluebery Planterose, Ana Sanchez Chico, and Stefanie Stantcheva, "Fighting Climate Change: International Attitudes toward Climate Policies," *American Economic Review*, apr 2025, *115* (4), 1258–1300.

Djankov, Simeon, Rafael La Porta, Florencio Lopez de Silanes, and Andrei Shleifer, "Disclosure by Politicians," *American Economic Journal: Applied Economics*, 2010, *2* (2), 179–209.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan, "Truth-Telling by Third-Party Auditors and the Response of Polluting Firms: Experimental Evidence from India," *Quarterly Journal of Economics*, 2013, *128* (4), 1499–1545.

Durante, Ruben, Luigi Guiso, and Giorgio Gulino, "Asocial Capital: Civic Culture and Social Distancing During COVID-19," *Journal of Public Economics*, 2021, *194*, 104342.

Esdaile, Louisa J. and Justin M. Chalker, "The Mercury Problem in Artisanal and Small-Scale Gold Mining," *Chemistry - A European Journal*, 2018, *24* (27), 6905–6916.

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde, "Global Evidence on Economic Preferences," *Quarterly Journal of Economics*, 2018, *133* (4), 1645–1692.

Fernandez, Raquel and Alessandra Fogli, "Culture: An Empirical Investigation of Beliefs, Work, and Fertility," *American Economic Journal: Macroeconomics*, 2009, *1* (1), 146–177.

Ferraz, Claudio and Frederico Finan, "Exposing Corrupt Politicians: The Effects of Brazil's Publicly Released Audits on Electoral Outcomes," *Quarterly Journal of Economics*, may 2008, *123* (2), 703–745.

Fisman, Raymond and Edward Miguel, "Corruption, Norms, and Legal Enforcement: Evidence from Diplomatic Parking Tickets," *Journal of Political Economy*, 2007, *115* (6), 1020–1048.

_ and Yongxiang Wang, "The Distortionary Effects of Incentives in Government: Evidence from China's "Death Ceiling" Program," *American Economic Journal: Applied Economics*, 2017, *9* (2), 202–218.

_ , Florian Schulz, and Vikrant Vig, "Is Sunlight (Plus Enforcement) the Best Disinfectant? Evidence of the Effects of Disclosure Laws and Demonetization on Political Selection," 2019.

Fuchs, Victor R, "Time Preference and Health: An Exploratory Study," *Economic Aspects of Health*, 1982, pp. 93–120.

George, Siddharth, Sarika Gupta, and Yusuf Neggers, "Can We Text Criminal Politicians Out of Office? Evidence from a Mobile Experiment in India," 2019.

Goldstein, Markus and Christopher Udry, "The Profits of Power: Land Rights and Agricultural Investment in Ghana," *Journal of Political Economy*, 2008, *116* (6), 981–1022.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, "Does Culture Affect Economic Outcomes?," *Journal of Economic Perspectives*, 2006, *20* (2), 23–48.

Henn, Soeren J., "Complements or Substitutes? How Institutional Arrangements Bind Traditional Authorities and the State in Africa," *American Political Science Review*, aug 2023, *117* (3), 871–890.

Hilson, Gavin, "Small-Scale Mining and Its Socio-Economic Impact in Developing Countries," *Natural Resources Forum*, 2002, *26* (1), 3–13.

_ and Roy Maconachie, "Formalising Artisanal and Small-Scale Mining: Insights, Contestations and Clarifications," *Area*, 2017, *49* (4), 443–451.

_ , Ekow Bartels, and Yanfei Hu, "Brick by Brick, Block by Block: Building a Sustainable Formalization Strategy for Small-Scale Gold Mining in Ghana," *Environmental Science and Policy*, 2022, *135* (March), 207–225.

Hofstede, Geert, "Dimensionalizing Cultures: The Hofstede Model in Context," *Online Readings in Psychology and Culture*, 2011, *2* (1), 1–26.

Hussam, Reshmaan, Kailash Pandey, Abu Shonchoy, and Chikako Yamauchi, "Translating Information into Action: A Public Health Experiment in Bangladesh," 2023.

Ichino, Andrea and Giovanni Maggi, "Work Environment and Individual Background: Explaining Regional Shirking Differentials in a Large Italian Firm," *Quarterly Journal of Economics*, 2000, *115* (3), 1057–1090.

Inglehart, Ronald and Wayne E Baker, "Modernization, Cultural Change, and the Persistence of Traditional Values," *American Sociological Review*, 2000, *65* (1), 19–51.

Iyer, Lakshmi and Anandi Mani, "Traveling Agents: Political Change and Bureaucratic Turnover in India," *Review of Economics and Statistics*, 2012.

Jack, B. Kelsey and María P. Recalde, "Leadership and the Voluntary Provision of Public Goods: Field Evidence from Bolivia," *Journal of Public Economics*, feb 2015, *122*, 80–93.

Jacobsen, Grant D., "The Al Gore Effect: An Inconvenient Truth and Voluntary Carbon Offsets," *Journal of Environmental Economics and Management*, jan 2011, *61* (1), 67–78.

Jensenius, Francesca Refsum, "Power, Performance and Bias: Evaluating the Electoral Quotas for Scheduled Castes in India." PhD dissertation, University of California, Berkeley 2013.

Kasanga, Kasim and Nii Ashie Kotey, *Land Management in Ghana: Building on Tradition and Modernity*, London: International Institute for Environment and Development, 2001.

Kosfeld, Michael and Devesh Rustagi, "Leader Punishment and Cooperation in Groups: Experimental Field Evidence from Commons Management in Ethiopia," *American Economic Review*, feb 2015, *105* (2), 747–783.

Lehne, Jonathan, Jacob N. Shapiro, and Oliver Vanden Eynde, "Building Connections: Political Corruption and Road Construction in India," *Journal of Development Economics*, 2018, *131* (November 2016), 62–78.

Madajewicz, Malgosia, Alexander Pfaff, Alexander van Geen, Joseph Graziano, Iftikhar Hussein, Hasina Momotaj, Roksana Sylvi, and Habibul Ahsan, "Can Information Alone Change Behavior? Response to Arsenic Contamination of Groundwater in Bangladesh," *Journal of Development Economics*, 2007, *84* (2), 731–754.

Mamdani, Mahmood, *Citizen and Subject: Contemporary Africa and the Legacy of Late Colonialism*, Princeton, NJ: Princeton University Press, 1996.

Müller, Stephan and Holger A Rau, "Economic Preferences and Compliance in the Social Stress Test of the COVID-19 Crisis," *Journal of Public Economics*, 2021, *194*, 104322.

Olken, Benjamin A., "Monitoring Corruption: Evidence from a Field Experiment in Indonesia," *Journal of Political Economy*, 2007, *115* (2), 200–249.

Olson, Mancur, *The Logic of Collective Action: Public Goods and the Theory of Groups*, Harvard University Press, 1965.

Ostrom, Elinor, *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge University Press, 1990.

Pande, Rohini, "Can Mandated Political Representation Provide Disadvantaged Minorities Policy Influence? Theory and Evidence from India," *American Economic Review*, 2003, *93* (4), 1132–1151.

Przeworski, Adam, Susan C. Stokes, and Bernard Manin, eds, *Democracy, Accountability, and Representation*, Cambridge: Cambridge University Press, 1999.

Rattray, Robert Sutherland, *Ashanti Law and Constitution* 1929.

Reinikka, Ritva and Jakob Svensson, "The Power of Information in Public Services: Evidence from Education in Uganda," *Journal of Public Economics*, 2011, *95* (7-8), 956–966.

Riley, Emma, "Role Models in Movies: The Impact of Queen of Katwe on Students' Educational Attainment," *Review of Economics and Statistics*, 2021, *106* (2), 334–351.

Schelling, Thomas, *The Strategy of Conflict*, Harvard University Press, 1960.

Siddique, Abu, Michael Vlassopoulos, and Yves Zenou, "Leveraging Edutainment and Social Networks to Foster Interethnic Harmony," 2024.

Snyder, James M and David Strömberg, "Press Coverage and Political Accountability," *Journal of Political Economy*, 2010, *118* (2), 355–408.

Szponar, Natalie, Claudia M Vega, Jacqueline Gerson, David Scott Mclagan, Martin Pillaca, Shamir Delgado, Domenica Lee, Nabila Rahman, Luis E Fernandez, Emily S Bernhardt, Adam M Kiefer, Carl P J Mitchell, Frank Wania, and Bridget A Bergquist, "Tracing Atmospheric Mercury from Artisanal and Small-Scale Gold Mining," *Environmental Science & Technology*, 2025, *59* (10), 5021–5033.

Tu, Meng, Bing Zhang, Jianhua Xu, and Fangwen Lu, "Mass Media, Information and Demand for Environmental Quality: Evidence from the "Under the Dome"," *Journal of Development Economics*, mar 2020, *143*.

Ubink, Janine M. and Kojo S. Amanor, *Contesting Land and Custom in Ghana: State, Chief and the Citizen*, Leiden University Press, 2008.

UNDP, "Human Development Report 2002: Deepening Democracy in a Fragmented World," Technical Report 2002.

UNEP, *Reducing Mercury Use in Artisanal and Small-Scale Gold Mining: A Practical Guide* 2012.

_ , "Minamata Convention on Mercury," 2013.

_ , *Global Mercury Assessment 2018*, Geneva, Switzerland: UN Environment Programme, Chemicals and Health Branch, 2019.

Vannutelli, Silvia, "From Lapdogs to Watchdogs: Random Auditor Assignment and Municipal Fiscal Performance in Italy," NBER Working Paper 30644, National Bureau of Economic Research 2022.