Learning, Salience, and Voting: Evidence from Criminal Politicians in India

Siddharth George, Yusuf Neggers, and Sarika Gupta*

May 2026

Abstract

We study how voters process information through two experiments around Indian elections. In a large-scale experiment, we show that providing voters information about candidates' criminal charges increases votes for clean candidates and reduces votes for criminal politicians, with larger penalties for candidates facing more and serious charges. A follow-up experiment replicates these results and identifies two mechanisms. First, information facilitates learning: voters form more accurate beliefs and evaluate criminal candidates less favourably. Second, using direct measures of voter attention, we show that information makes criminality more salient, and increases its weight in voting decisions. Salience effects are larger when information is surprising or highlights contrast, but do not vary with decision relevance, consistent with bottom-up attention. Causal forest estimates provide further evidence that learning and salience are both important drivers of changes in voting behaviour. We develop a simple model that integrates salience theory into a standard probabilistic voting framework to explain our results.

*Siddharth Eapen George, segeorge@nus.edu.sg, National University of Singapore; Yusuf Neggers, yneggers@umich.edu, University of Michigan; Sarika Gupta, sringwala@worldbank.org, World Bank. We thank our research partner Manoj Kumar at the Center for Governance and Development in India for deep institutional knowledge about election campaigns in Uttar Pradesh, and for invaluable support in implementing the experiment. We are indebted to Rohini Pande for her extensive support and advice throughout this project. We also thank Abhijit Banerjee, Robert Bates, Pedro Bordalo, Kirill Borusyak, Emily Breza, Kevin Carney, Melissa Dell, Ruben Durante, Ryan Enos, Michael Kremer, Judd Kessler, Horacio Larreguy, Michael Lerner, Filip Matějka, Martin Mattsson, Eduardo Montero, Brendan Nyhan, Vincent Pons, Nishith Prakash, Pablo Querubín, Pia Raffler, Collin Raymond, Gautam Rao, Chris Roth, Deepak Saraswat, Peter Schwardmann, M.R. Sharan, Andrei Shleifer, Henrik Sigstad, Tavneet Suri, Raphael Susewind, Edoardo Teso, and Egon Tripodi for useful feedback and helpful conversations. We thank seminar participants at Boston University, Brown, Columbia, Harvard, Michigan, National University of Singapore, the NEUDC, NYU, and UC Berkeley for comments and suggestions. Gaurav Peter Mathai provided outstanding research assistance. In addition, we thank Sawal Choudhary, Riddhimaa Gupta, Sameer Nair-Desai, Shashank Sreedharan, and Niaoniao You for their contributions to this project. AEA RCT Registry is AEARCTR-0017542. Financial support from the J-PAL Governance Initiative, the Harvard Lab for Economic Applications and Policy, and the Watson Institute for International and Public Affairs at Brown University is gratefully acknowledged.

1 Introduction

In principle, democratic elections empower citizens to select good leaders and hold them accountable. In practice, however, multiple constraints can prevent voters from electing high-quality politicians. Voters may lack the knowledge required to effectively screen candidates, or their attention may be captured by salient ideological or identity-based divisions. Better information on candidates could address both challenges and ultimately improve political selection—by helping voters learn about candidates and drawing their attention to valence attributes like competence and honesty (Besley 2005; Dal Bó and Finan 2018).

In this paper, we study how information about candidate quality shapes voters' beliefs, attention, and electoral choices. We focus on an egregious yet common form of adverse selection into politics: the election of criminal politicians—individuals facing criminal charges—to public office. Since 2000, more than 30 democratic countries have had criminally charged heads of state or government, raising the question of why voters elect such politicians.[1]

India offers a compelling setting to study criminal politicians. Following a Supreme Court ruling in 2003, all candidates contesting national and state legislative elections have been required to disclose their criminal charges, making this information accessible to researchers.[2] Nearly 35% of national legislators elected since 2003 have criminal charges—including for violent crimes like murder, kidnapping, rape, extortion, and armed robbery—and this share has risen steadily over time.[3] Criminal politicians are regularly elected despite causal evidence that they increase violent crime and crimes against women, reduce household consumption and female labour force participation, and lower local economic development (Chemin 2012; Prakash et al. 2019, 2024; Asher and Novosad 2023).[4] Although some voters may value criminal politicians' ability to violently defend group interests (Vaishnav 2017), many express a preference for non-criminal candidates (Banerjee et al. 2014). We hypothesise that voters' limited information and inattention to candidate quality contribute to the electoral success of criminal politicians.

We experimentally evaluate how voters process information about politicians' criminal charges. To guide our analysis, we develop a simple model that integrates salience theory into a probabilistic voting framework. The core idea is that information shapes both voters' beliefs and what is salient to them. Voters

[1] In ten of these countries, charges occurred before or during time in office. Democracy designation based on Polity5 scores for 2000 (Marshall and Gurr 2020). Appendix Table A1 gives additional details.

[2] As in other research studying criminal politicians in India (e.g., Vaishnav 2017; Prakash et al. 2019, 2024; Magesan et al. 2025), we focus on declared criminal charges rather than convictions. Charges appear on a candidate's record only if a court has taken cognizance of the crime or framed charges, requiring evidence of an offense. In addition, candidates can be disqualified, and face imprisonment and/or fines, for not declaring charges, mitigating non-reporting risk (Asher and Novosad 2023).

[3] From 24% in 2004 to 46% in 2024 (Association for Democratic Reforms 2025).

[4] Similarly, Britto et al. (2025) find in Brazil that political candidates and winners are adversely selected relative to the general population in terms of criminality, and their election at the local level worsens health outcomes.

value high-quality (e.g., non-criminal) and in-group (e.g., same party or caste) candidates, but weight each attribute according to its salience. In this framework, information about candidates affects voting decisions through two channels. First, *learning*: voters update beliefs about candidate quality. Second, *salience*: information about candidate quality draws voters' attention to the quality dimension, increasing the weight they place on quality when evaluating candidates. Salience effects may be particularly pronounced when the information is surprising or highlights contrast among candidates (Bordalo et al. 2013, 2022).

We test the predictions of the model by conducting two voter information experiments around successive elections in India's most populous state, Uttar Pradesh (UP), whose 240 million residents represent more than 6% of the global population living in a democracy (Government of India 2020; Economist Intelligence Unit 2025). Both experiments provided randomly selected voters with information about candidates' criminal charges via mobile phone shortly before election day.

Our first experiment was conducted in 38 constituencies during the 2017 state assembly elections. Out of 3,800 sample villages encompassing 4.2 million registered voters, 1,400 villages were randomly assigned to treatment. Citizens in treated villages received text messages and automated voice calls describing whether each major party candidate in their constituency had criminal cases against them and the nature of their charges—facts drawn from candidates' official sworn affidavits.[5] Two major telecom companies—whose mobile subscribers jointly comprised 37% of the electorate—delivered messages two days before the election. We randomised treatment across villages within each constituency and observe official electoral outcomes at the polling station level, enabling us to compare each candidate's electoral performance in treated villages (where voters received information) and control villages (where voters did not).

We find that information about candidate criminality affected voting behaviour. On average, "clean" candidates (those facing no charges) earned 3.4% more votes in treated villages, while criminal candidates received 2.7% fewer votes. Moreover, voters' responses were sophisticated. Electoral costs scale with crime severity: candidates received 1.3% fewer votes per criminal case, and incurred especially large penalties for murder-related charges. Voters' reactions also depend on the electoral context: clean candidates gained votes primarily when facing opponents with violent charges, indicating that voters compared candidates. Overall turnout remained unchanged.

Our treatment effects imply that 1–3% of message recipients were persuaded to vote for a clean candidate—a slightly higher persuasion rate than in get-out-the-vote mailer experiments in the US (Gerber and Green 2000; DellaVigna and Gentzkow 2010). Counterfactual calculations suggest that informing all voters about candidates' charges could have changed the outcome of 8–23% of recent UP elections where

[5] Candidate affidavits are publicly available online and also include information on assets, liabilities, and education.

criminal politicians defeated clean runners-up.[6] Our first experiment thus demonstrates that information has the potential to improve political selection. However, it does not allow us to unpack mechanisms.

To study how voters process information, we ran a second experiment during UP's subsequent assembly elections in 2022. We surveyed 3,700 voters in 23 constituencies—distinct from those in the 2017 experiment—several weeks before election day to collect data on demographics, beliefs, attention, and preferences. Randomising treatment assignment at the individual level within each constituency, we replicated the first experiment's intervention by providing treated voters information about candidates' criminal cases and charges. Messages had the same content and structure as in the first experiment. We surveyed treated and control voters again shortly after election day to collect endline measures of beliefs, attention, decision weights, candidate evaluations, and vote choice.

Most voters held inaccurate prior beliefs about politician criminality. At baseline, only 34% could correctly identify whether their incumbent legislator had criminal charges. At endline, control group voters had accurate beliefs about candidates' criminal status only 32% of the time. These errors were systematic, with the majority of voters underestimating the prevalence of criminal politicians.

The information intervention helped voters form more accurate beliefs, especially about criminal candidates. Belief accuracy increased by 4.7 percentage points (or 8.4% relative to control) for clean candidates and by 13 percentage points for criminal candidates, nearly doubling voters' ability to identify criminal politicians. These belief changes shaped electoral choices: treated voters rate criminal candidates 0.12 standard deviations (SD) worse and are 2 percentage points (or 7.3%) more likely to indicate voting for clean candidates. Turnout again remained unchanged.

Our 2022 experiment shows that information facilitates voter learning and increases support for clean candidates, echoing the findings of our first experiment in 2017. Although the two experiments implemented the same core treatment, they differed substantially in scale, randomisation level, constituency sample and electoral context, and how outcomes were measured (official electoral records in 2017; survey responses in 2022). The consistency of results across both designs strengthens confidence in the robustness and external validity of our findings. It also helps alleviate concerns about experimenter demand in the 2022 survey, since the 2017 results—which are based on administrative data—cannot plausibly be driven by social desirability bias (Bursztyn et al. 2025).[7]

Beyond documenting learning and external validity, our second experiment allows us to examine how salience shapes voter behaviour. We measure salience directly by asking voters what is "top of mind"

[6] Our intervention involved information provision to less than 0.7% of the population per sample assembly constituency, indicating in combination with our estimated treatment effects that it could not have affected election outcomes as implemented.

[7] Indeed, very few respondents (baseline, 0.6%; endline, 0.8%) believed the 2022 survey focus was on politician criminality.

when they think about the election and the candidates in their constituency, an approach aligned with the recommendations of Haaland et al. (2025). To validate our salience measure, we embedded a novel vignette-based recall task in the baseline survey: early in the survey, respondents heard a biographical profile of a candidate in a prior election, and at the end of the survey were asked to recall what they remembered. Across multiple candidate attributes (e.g., criminality, education, religion), we find that voters who report an attribute as top of mind are substantially more likely to recall it from the vignette—indicating that our top-of-mind measure captures attributes that are genuinely cognitively salient.[8]

We show, using our salience measure, that the information treatment affected voter attention. Criminality was not salient to a large portion of the electorate, with only 37% of control group voters attentive to it at endline. However, criminality was 3.6 percentage points (10% or 0.08 SD) more likely to be top of mind for treated voters. Information also heightened the salience of policy issues associated with criminal politicians—such as crime and corruption—even though our messages did not directly mention those subjects, consistent with similarity-based retrieval (Kahana 2012; Bordalo et al. 2020, 2023, 2025; Enke et al. 2024; Graeber et al. 2024). Combining these outcomes into a broader index, treatment increased the salience of criminality by 0.1 SD. By contrast, the salience of other candidate attributes and policy topics was unaffected, implying that information raised the relative salience of criminality.

These attention changes are driven by salient features of information attracting voters' focus. Information raises salience most when it is *surprising* (i.e., deviates more from voters' priors) or highlights *contrast* between candidates (i.e., reveals bigger differences in criminality among candidates). Notably, surprising information about electorally irrelevant candidates also boosts salience. By contrast, factors that raise the decision value of information, such as close elections or voter indecision, do not draw attention to criminality. Taken together, our findings are better explained by salient features of information drawing voters' attention bottom-up (Bordalo et al. 2022) than by models of top-down rational inattention (Matějka and Tabellini 2021).

Next, we elicit the decision weight voters place on different candidate attributes and policy topics, and test whether shifts in voter attention affect these weights, as predicted by salience theory (Bordalo et al. 2013). We find that treated voters place 0.07 SD greater decision weight on criminality, while weights on other attributes and policy issues remain unchanged. Overall, our 2022 experiment shows that information changes voter behaviour through two channels: belief updating (*learning*) and attention shifts (*salience*).

We conduct three additional analyses to shed light on the mechanisms underlying the salience channel.

[8] In addition, voters place greater weight on salient candidate attributes when making electoral choices. This is consistent with the predictions of both salience theory and rational inattention (Kőszegi and Szeidl 2013; Matějka and Tabellini 2021; Bordalo et al. 2022) and further validates our attention measure.

First, a simple mediation analysis (Imai et al. 2010; Andries et al. 2024) shows that treatment has a small, insignificant effect on decision weights once we control for salience. Second, applying the assumption-light test of Kwon and Roth (2026), we cannot reject the sharp null of full mediation—that the information treatment affects decision weights only through salience. Third, using the causal forest method of Wager and Athey (2018), we estimate each voter's conditional average treatment effect (CATE) for our key outcomes. We find that voters' salience responses are the strongest predictor of their decision weight responses. Together, these results support the causal chain in our conceptual framework and in salience models more broadly: information redirects voters' attention, influencing their decision weights (Bordalo et al. 2022).

To compare the relative importance of the learning and salience channels, we classify voters into types based on their predicted CATE response to information: *Learners* update beliefs but do not pay more attention to criminality, while *Engagers* increase attention but do not change beliefs. We find that *Learners* show slightly weaker treatment effects on candidate ratings and voting behaviour than *Engagers*. Information has the strongest treatment effects on *Responders*, for whom both channels operate. These results reinforce that salience is an important channel in its own right and suggest that it complements learning.

Finally, we examine whether group identities blunt the effectiveness of information interventions. A central concern in polarised polities is that voters may reject unfavourable information about in-group candidates. We find limited evidence of such motivated reasoning.[9] Belief updating is similar for in- and out-group candidates, indicating that learning occurs regardless of partisan or religious/caste alignments. While we find suggestive evidence that voters reward own-group candidates more for being clean, they impose similar electoral punishments on in- and out-group candidates with criminal charges. Information thus facilitates learning and influences voters' decisions even when it reflects negatively on in-group candidates, suggesting that it can improve political selection even in polarised contexts.

Our paper contributes to three literatures. First, we add to an experimental literature on voter information interventions (see Pande 2011; Dunning et al. 2019; Olken and Pande 2019; Haaland et al. 2023 for reviews). Prior studies have examined the impacts of informing voters about the electoral process (Aker et al. 2017; Marx et al. 2021; Schechter and Vasudevan 2023), leaders' powers and responsibilities (Cruz et al. 2021; Baysan 2022; Banerjee et al. 2026), candidate characteristics and policy positions (Kendall et al. 2015; Pons 2018; Cruz et al. 2024), and incumbent performance and corruption (Ferraz and Finan 2008; Boas et al. 2019; Chauchard et al. 2019; Arias et al. 2022; Bhandari et al. 2023; Rivera et al. 2025). However, the underlying mechanisms through which information influences voter behaviour are often difficult

[9]For work on motivated reasoning in political contexts see, e.g., Bénabou and Tirole (2016); Adida et al. (2017); Bolsen and Palm (2019); Nyhan (2020); D'Amico and Tabellini (2024); and Thaler (2024).

to determine, as most interventions plausibly affect beliefs, attention, and coordination.[10] We contribute by demonstrating the importance of two complementary channels—learning and salience. In addition, we provide novel evidence on the electoral impacts of informing voters about candidates' criminal charges.[11]

Second, we contribute to a growing empirical literature on salience. Influential theoretical work shows that salience can provide a unified explanation for a range of behavioural biases and choice anomalies (Bordalo et al. 2013, 2020, 2022; Loewenstein and Wojtowicz 2025). Empirical work demonstrates that salience shapes how people learn about climate change (Patel 2025), choose occupations (Conlon and Patel 2025), form macroeconomic expectations (Link et al. 2025), interpret statistics (Graeber et al. 2024; Bordalo et al. 2026), react to taxes (Chetty et al. 2009), and make education investments (Lichand et al. 2026). We apply salience theory to understand real-world voter behaviour. While prior work infers salience effects from responses to unrelated events (Bonomi et al. 2021; Colussi et al. 2021; Ajzenman and Durante 2023; Arya and Bhatiya 2025; Go et al. 2026) or using structural estimation (Cruz et al. 2024), we directly measure what is top-of-mind for voters, their decision weights on different candidate attributes, and their candidate ratings and electoral choices. These direct measures enable us to trace the full causal chain described in salience models: information shifts attention, raises the weight placed on salient attributes, and ultimately changes choices.

Third, we contribute to the literature on political selection (Besley 2005; Dal Bó and Finan 2018; Gulzar 2021). While prior work documents that politicians are often positively selected in terms of competence and ability in mature, well-functioning democracies (Dal Bó et al. 2017; Casey et al. 2021; Jokela et al. 2025), we study a widespread form of adverse selection into politics: the election of criminal politicians.[12] Existing explanations emphasise the defense of identity group interests (Vaishnav 2017), clientelism (Khemka 2025), illicit rents (Asher and Novosad 2023), and party strategy (Kapur and Vaishnav 2018; Magesan et al. 2025). We provide a complementary explanation rooted in voters' limited knowledge and attention. From a policy perspective, our mobile-based intervention highlights the potential of scalable, light-touch approaches that inform or draw voters' attention to valence attributes to improve political selection (Bandiera et al. 2024).[13]

[10] Related work documents the persuasive effects of political advertising (Spenkuch and Toniatti 2018), debates (Bidwell et al. 2020; Le Pennec and Pons 2023), television (DellaVigna and Kaplan 2007; Enikolopov et al. 2011; Durante et al. 2019; Canen and Martin 2023), social media (Fujiwara et al. 2024; Carney 2025), and the internet (Guriev et al. 2021), and highlights how information influences emotions, which in turn shape policy preferences (Manzoni et al. 2026).

[11] Banerjee et al. (2011) organized discussion groups for and provided Delhi voters with newspapers that contained report cards with information on incumbent public spending and candidate characteristics such as criminality, education, and wealth. Their working paper finds that voters did not react to criminality, likely due in part to the many categories of information provided.

[12] Our results parallel findings on the selection of dishonest individuals into the bureaucracy (Hanna and Wang 2017), though bureaucrats and politicians are screened differently.

[13] While a growing literature studies mobile-based interventions to increase voter registration and turnout in wealthy

The paper is structured as follows. Section 2 presents our conceptual framework. Sections 3 and 4 describe our context, experimental design, and data. Sections 5 and 6 present the findings of our 2017 and 2022 election experiments. Section 7 discusses mechanisms and additional results, and Section 8 concludes.

2 Conceptual Framework

We develop a simple theory of voter behaviour to guide our empirical analysis. Our model applies insights from salience theory (Bordalo et al. 2013, 2020, 2022) to electoral choices, by integrating salience thinking into a probabilistic voting framework à la Persson and Tabellini (2002). We then study and compare how voters process information under the benchmark model and in the model with salience thinking.

2.1 Benchmark Model

2.1.1 *Setup*

Citizens. A polity consists of a unit mass of citizens. Each citizen has two attributes: (i) quality $q \in \{0,1\}$, where $q=1$ denotes a high-quality citizen; and (ii) social identity $g \in \{L,R\}$, where group L is the majority ($f_L > \frac{1}{2}$). Both groups have the same share π of high-quality citizens.

Politicians. Each social group randomly nominates a candidate for a public office whose role is to provide a public good. A high-quality candidate ($q=1$) successfully delivers the public good, generating utility $B>0$ for all voters, whereas a low-quality candidate ($q=0$) does not deliver the public good. A candidate's identity is publicly observable but their quality is not.

Voting. Each citizen i holds subjective belief $\pi_{ij} = \mathbb{P}(q_j = 1)$ about the quality of candidate j. Absent information, citizens share a common rational prior π about each candidate's quality. When voting, citizens derive utility from two distinct components: the instrumental (expected) benefit from the public good and the expressive utility of supporting a candidate from their own group. Accordingly, citizen i's expected utility from voting for candidate j is:

$$U_{ij} = \theta_i \pi_{ij} B + (1-\theta_i)\gamma \cdot \mathbb{I}(g_j = g_i),$$

where θ_i is voter i's decision weight on quality which reflects her underlying "deep" preferences, and $\gamma > 0$ captures the expressive benefit of voting for an own-group candidate.

democracies (Dale and Strauss 2009; Malhotra et al. 2011; Cheng-Matsuno et al. 2023; Lindgren et al. 2025), existing studies in low-income democracies are concentrated in Sub-Saharan Africa and do not involve provision of candidate-specific information (Aker et al. 2017; Buntaine et al. 2018; Harris et al. 2021; Marx et al. 2021).

As in a standard probabilistic voting model, there are two sources of electoral uncertainty: (i) each voter receives an idiosyncratic shock $\sigma_i \sim U\left[-\frac{1}{2},\frac{1}{2}\right]$, which affects their support for the L party; and (ii) party L receives an aggregate popularity shock δ, which determines the outcome of the election. We assume that δ follows a general distribution function $h(\cdot)$ that is unimodal, non-uniform and symmetric around 0.

2.1.2 *Equilibrium*

No Information. We first solve the model for the benchmark case where citizens are uninformed about candidate quality. Consider the voting decision of citizen i from group L. This citizen chooses to support the L candidate if $U_{iL} > U_{iR} \iff \theta_i B \pi_L + (1-\theta_i)\gamma + \sigma_i + \delta > \theta_i B \pi_R$. Since citizens share a common prior π about each candidate's quality, this condition reduces to $\sigma_i > -\gamma(1-\theta_i) - \delta$, which occurs with probability $p_L = \frac{1}{2} + \gamma(1-\theta_i) + \delta$. By a similar logic, an R group citizen supports the L candidate with probability $p_R = \frac{1}{2} - \gamma(1-\theta_i) + \delta$. Thus, the L candidate's vote share is $V_L = f_L \cdot p_L + (1-f_L) \cdot p_R$, which does not depend on either candidate's quality, i.e., $\frac{\partial V_L}{\partial q_L} = \frac{\partial V_L}{\partial q_R} = 0$.

Full Information. Next, we consider the equilibrium when citizens are informed about candidate quality. Informed L group citizens supports the L candidate if $\sigma_i > -\theta_i B(q_L - q_R) - \gamma(1-\theta_i) - \delta$, i.e., with probability $p_L = \frac{1}{2} + \theta_i B(q_L - q_R) + \gamma(1-\theta_i) + \delta$. Similarly, informed R group citizens support the L candidate with probability $p_R = \frac{1}{2} + \theta_i B(q_L - q_R) - \gamma(1-\theta_i) + \delta$. Now, we see that the L candidate's vote share depends on candidate quality: $\frac{\partial V_L}{\partial q_L} = \theta_i B$ and $\frac{\partial V_L}{\partial q_R} = -\theta_i B$. Moreover, voting behaviour responds more strongly to information about candidate quality when citizens care more about quality (i.e., θ is higher) and when the return to quality (i.e., the public good utility B) is larger.

2.2 Incorporating Salience

2.2.1 *Setup*

In the benchmark model, voters' decision weights on quality and group identity reflect their underlying preferences. However, a growing body of research shows that individuals often pay more attention to salient attributes, and place greater decision weight on these attributes (Gennaioli and Shleifer 2010; Kőszegi and Szeidl 2013; Bordalo et al. 2013, 2022). Reviewing the literature, Bordalo et al. (2022) note that an attribute's salience depends on three factors: contrast, surprise, and prominence. Applying these ideas to our setting, we model the salience of each candidate attribute $k \in \{q, g\}$, denoted ϕ^k, as a function of these components.

- **Contrast** captures the extent to which the two candidates differ on attribute k. Following Bordalo

et al. (2013), we define it as:
$$\frac{|a_L^k - a_R^k|}{|\bar{a}^k|},$$
where a_L^k and a_R^k denote the values of k for each candidate, and $\bar{a}^k = (a_L^k + a_R^k)/2$ is their average.[14]

- **Surprise** captures the extent to which the observed candidate attribute deviates from a citizen's prior belief. We model surprise as:
$$\frac{1}{2} \sum_{j \in \{L,R\}} \left| a_j^k - \mathbb{E}_i[a_j^k] \right|,$$
where $\mathbb{E}_i[a_j^k]$ denotes citizen i's prior about the value of attribute k for candidate j.

- **Prominence** reflects the idea that some attributes are more cognitively accessible or top-of-mind. An attribute's prominence $\psi_i^k > 0$ can be shaped by exogenous forces such as advertising, media coverage, or information interventions. We assume that group identity is visible and prominent for all voters, so $\psi_i^g = \psi^g$. By contrast, candidate quality is not visible, and must be observed through information signals. Thus, $\psi_i^q = \psi^q + \beta T_i$, where $T_i \in \{0,1\}$ indicates whether voter i received information about candidate quality and $\beta > 0$ captures the effect of information on prominence.

Moreover, these components are complementary: prominent attributes are especially salient if they are also surprising or highlight contrast (Bordalo et al. 2022). Combining these elements, we define the salience of attribute k to voter i as:
$$\phi_i^k = \underbrace{\psi_i^k}_{\text{prominence}} \left[\lambda + \underbrace{\frac{|a_L^k - a_R^k|}{|\bar{a}^k|}}_{\text{contrast}} + \underbrace{\frac{1}{2} \sum_{j \in \{L,R\}} \left| a_j^k - \mathbb{E}_i[a_j^k] \right|}_{\text{surprise}} \right],$$
where $\lambda > 0$. For simplicity, we assume that our intervention—which only provided information about candidate quality—did not affect the prominence of group identity. We also assume that the contrast and surprise of the identity attribute are constant over time, since our information intervention was conducted two days before the election and candidates had no time to adjust campaigning in response.

Salience affects voting decisions by influencing the weight citizens place on each attribute. Specifically, following Bordalo et al. (2013), we assume that the salience-adjusted weight citizens place on quality is $w_i^q \theta_i$, where $w_i^q = \phi_i^q$. Incorporating these salience-driven decision weights, citizen i's expected

[14] Nunnari and Zápal (2025) formalize this idea differently. Building on Kőszegi and Szeidl (2013), they analyze how parties set platforms when voters place greater weight on issues where parties differ more.

utility from voting for candidate j is:

$$U_{ij} = w_i^q \theta_i \cdot \pi_{ij} B + w_i^g (1-\theta_i) \cdot \gamma \cdot \mathbb{I}(g_j = g_i).$$

2.2.2 *Equilibrium*

No Information. Now L voters support the L candidate if $U_{iL} = w_i^q \theta_i B\pi + w_i^g(1-\theta_i)\gamma + \sigma_i + \delta > w_i^q \theta_i B\pi = U_{iR}$, which occurs with probability $p_L = \frac{1}{2} + w_i^g \gamma(1-\theta_i) + \delta$. Similarly, R-group citizens support the L candidate with probability $p_R = \frac{1}{2} - w_i^g \gamma(1-\theta_i) + \delta$. As in the benchmark model, candidate quality does not affect vote shares. Observe that $\frac{\partial V_L}{\partial q_L} = 0$, since $\frac{\partial p_L}{\partial q_L} = \frac{\partial p_R}{\partial q_L} = 0$, as $w_i^g = \phi_i^g$ does not depend on q_L.

Full Information. Distinct from the benchmark model, information now impacts voting in two ways: in addition to changing citizens' beliefs about candidates, it affects the decision weight on candidate quality.

First, the prominence of quality, ψ_i^q, increases from ψ^q to $\psi^q + \beta$. Information increases the salience of quality, since

$$\phi_{i,info}^q = (\psi^q + \beta)[\lambda + \frac{2|q_L - q_R|}{|q_L + q_R|} + \frac{1}{2}(|q_L - \pi| + |q_R - \pi|)] > \lambda\psi^q = \phi_{i,noinfo}^q,$$

with larger changes if $q_L \neq q_R$ (candidates have contrasting quality) and $q_L, q_R \neq \pi$ (the information is surprising). As a result, information about candidate quality causes citizens to place greater decision weight on quality when voting: $w_{i,info}^q = \phi_{i,info}^q > \phi_{i,noinfo}^q = w_{i,noinfo}^q$.

Next, we consider how citizens respond electorally to a change in candidate quality. Observe that L voters now support the L candidate if $\sigma_i > -w_i^q \theta_i B(q_L - q_R) - w_i^g(1-\theta_i)\gamma - \delta$, which occurs with $p_L = \frac{1}{2} + \phi_i^q \theta_i B(q_L - q_R) + \phi_i^g \gamma(1-\theta_i) + \delta$.[15] Thus, we have

$$\frac{\partial V_L}{\partial q_L} = \underbrace{\phi_i^q \theta_i B}_{\text{learning}} + \underbrace{\frac{\partial \phi_i^q}{\partial q_L} \theta_i B(q_L - q_R)}_{\text{salience}},$$

The first term captures the *learning* channel: the effect of citizens updating beliefs about candidate quality, scaled by the salience-adjusted decision weight on quality. The second term captures the *salience* channel: the effect of information on voters' attention to quality, which magnifies the impact of quality differences between the two candidates. The overall electoral impacts of information reflect both learning and salience effects. However, the two channels need not be equally active for each voter, however. For some voters, information may shift beliefs without changing what is top of mind (i.e., $\frac{\partial \phi_i^q}{\partial q_L} \approx 0$). For others, information may redirect attention without meaningfully updating priors (e.g., if they are already accurate).[16]

[15]Analogously, $p_R = \frac{1}{2} + \phi_i^q \theta_i B(q_L - q_R) - \phi_i^g \gamma(1-\theta_i) + \delta$.

[16]In Section 7.2, we empirically examine the prevalence of different such subsets of voters, in order to compare the importance of the learning and salience channels in shaping voting behaviour.

2.3 Key Predictions

Our salience-augmented probabilistic voting model yields the following key testable predictions:

1. Informing voters about candidate quality increases (reduces) support for high-quality (low-quality) candidates, and these effects are increasing in the quality difference between candidates.
2. Informing voters about candidate quality can raise its salience, and these salience effects are stronger when information about quality is surprising or highlights contrast in quality among candidates.
3. Changes in the salience of candidate quality lead to changes in the decision weight on candidate quality.

This framework provides a unified account of how information shapes voters' beliefs, attention, and ultimately their electoral choices through both the learning and salience channels.[17]

3 Background

3.1 Politics in India

Role of State Legislators. India is a federal parliamentary republic where state legislators, called Members of the Legislative Assembly (MLAs), play an important role in governance. MLAs represent single-member constituencies and are elected through a first-past-the-post system to serve five-year terms. Beyond their legislative duties, MLAs exert significant influence over bureaucratic appointments and transfers (Iyer and Mani 2012), and shape the allocation of public resources within their constituencies (Lehne et al. 2018). Thus, as prior work documents, the quality of candidates elected to these offices has important consequences for economic and development outcomes (Clots-Figueras 2012; Iyer et al. 2012; Bhalotra and Clots-Figueras 2014; Prakash et al. 2019, 2024).

Criminality in Indian Politics. Criminal politicians are prevalent in India: as of March 2025, 45% of incumbent MLAs nationwide faced criminal charges (Singh 2025). Yet a growing body of evidence finds that governance by criminal politicians imposes significant economic and social costs. Studies using close-election regression discontinuity and instrumental variable designs show that criminal politicians increase crime rates and reduce female labour force participation (Prakash et al. 2024), decrease consumption by disadvantaged groups (Chemin 2012), and depress local economic activity, especially when

[17] In Section 7.1, we engage with alternative models, such as the possibility that voter attention is allocated in a top-down, goal-oriented manner, as in rational inattention frameworks (e.g., Matějka and Tabellini 2021).

these politicians have serious charges (Prakash et al. 2019).[18]

Why are criminal politicians elected despite such costs? Two prominent explanations focus on (i) the private benefits parties derive from fielding criminal politicians (Magesan et al. 2025) and (ii) deep-seated social forces, such as ethnic divisions, that may lead voters to value the "enforcer" credentials of criminal politicians who (violently) defend group interests (Vaishnav 2017). We examine a complementary explanation: that voters lack the information to screen candidates and are inattentive to candidate quality.

3.2 Uttar Pradesh: Social and Political Context

Social Context. Our study is set in Uttar Pradesh (UP), India's largest state and the world's most populous sub-national unit. UP has a predominantly rural population of about 240 million and is one of India's poorest states, with a Human Development Index level comparable to Angola (Smits and Permanyer 2019). Like many low-income democracies, UP has deep sectarian divisions along caste and religious lines. Communal issues are often salient during elections (Banerjee et al. 2014), and may capture voters' attention at the expense of candidate quality (Banerjee and Pande 2009).

UP has one of the highest concentrations of criminal politicians among Indian states, with 15–26% of candidates and 35–51% of elected winners facing criminal charges in the past four state elections.[19] At the same time, mobile phone penetration is high: over 86% of rural households owned a mobile phone by 2011 (Press Trust of India 2015).[20] This combination of weak political selection, entrenched identity politics, and widespread mobile access makes UP a good setting in which to evaluate mobile-based voter information interventions.

Elections in UP. We conducted experiments during UP's 2017 and 2022 state assembly elections. Due to their scale, both elections were held in seven phases between February and March, with 40–70 geographically contiguous constituencies (out of 403 total) voting in each phase, and results announced after the final phase. The same four main parties contested the 2017 and 2022 elections: (i) the Bharatiya Janata Party (BJP), a right-wing party that controlled the central government during both elections; (ii) the Samajwadi Party (SP), a center-left party supported by Other Backward Castes (OBCs) and Muslims that governed UP from 2012–17; (iii) the Indian National Congress (INC), India's largest opposition party; and (iv) the Bahujan Samaj Party (BSP), which primarily represents disadvantaged Scheduled

[18] The Supreme Court of India has viewed the election of criminally charged politicians as sufficiently detrimental to public well-being to state: "A time has come that Parliament must make law to ensure that persons facing serious criminal cases do not enter into the political stream" (Jain 2018).

[19] Authors' calculations from myneta.info.

[20] As more broadly in the developing world, India has seen a sharp rise in mobile phone usage over the last two decades (Waghmare 2024), making mobile technology a feasible means of spreading information widely even in rural areas.

Castes (SCs) and governed UP from 2007–12. The SP and INC formed an electoral alliance in the 2017 elections, but contested separately in 2022. The BJP won a majority of seats in both elections.

4 Experimental Design and Data

4.1 Intervention Structure

We conducted two voter information experiments in partnership with a local NGO during the 2017 and 2022 UP state assembly elections. Both experiments implemented the same core treatment—providing voters information about the criminal charges of candidates in their constituency—but differed in their scale, sample, unit of randomisation, and source of outcome data.

2017 Experiment. The 2017 election experiment was a large-scale intervention implemented with two major Indian telecom companies. Within our 2017 sample constituencies, treatment was randomised at the village level (Section 4.3 further discusses sampling and randomisation). In each treated village, all mobile subscribers of our partner telecom providers—whose combined market share averaged 37% of the electorate—received an automated voice call and text message.[21] The calls and messages were delivered two days before the election.[22]

The treatment voice and text messages contained information about the number of criminal cases and type of criminal charges faced by each major-party candidate in the recipient's constituency. If a candidate did not have criminal charges, the message stated that the candidate faced no charges. If a candidate had criminal charges, the message stated the number of criminal cases, and listed charges for any violent crimes.[23] If a candidate faced only non-violent criminal charges, the message included the number of cases and stated that none of the charges were for violent crimes. Appendix B.1 provides an example message for a constituency where the major-party candidates ranged from facing no criminal charges to violent charges.[24]

In a subset of treated villages, the information messages contained an additional one to two sentences,

[21] The corresponding 25th, 50th, and 75th percentiles of village-level market share are 32%, 37%, and 44%, respectively.

[22] In the days leading up to an election, citizens are commonly exposed to campaign content from different political parties, but the Election Commission of India requires parties to freeze all campaigning starting the day before the election. To increase their perceived objectivity and differentiation from party content, our calls and text messages stated that they were sent on behalf of an unbiased, non-partisan, non-governmental organization.

[23] We categorized charges as violent or non-violent following India's penal code. The following charges were categorized as violent: murder, attempt to murder, culpable homicide, attempt to commit culpable homicide, rape, attempt to commit rape, kidnapping and abduction, dacoity, preparation and assembly for committing dacoity, robbery, and arson.

[24] Text messages were sent in the Hindi language and Devanagari script, and voice messages were recorded in Hindi. The order of the candidates in the messages for each constituency was randomised.

either (i) notifying recipients that the message had been shared with many other voters in the constituency ("Information plus Coordination" sub-treatment) or (ii) urging them to "break the habit" of ethnic voting ("Information plus Ethnic Voting" sub-treatment). Additional details can be found in Appendix B.1.

2022 Experiment. The 2022 election experiment was a smaller-scale intervention that aimed to identify mechanisms. We worked with a local survey firm that had a sampling frame of mobile phone subscribers in our 2022 sample constituencies. After enrolling participants via a baseline survey, treatment was randomised at the individual level within each constituency. Two days before the election, treated voters received an information voice call and text message whose content and structure were identical to the 2017 messages. In this sense, voters in both experiments received the same treatment—information about the criminal charges of candidates in their constituency. As in the 2017 experiment, messages for a subset of treated individuals included additional content mentioning that other voters in the constituency had received the same information ("Information plus Coordination" sub-treatment). Appendix B.1 provides additional details.

4.2 Data

Our study relies on several sources of administrative and survey data.

Politician Criminality Data. We use information on candidate criminality from publicly available affidavits that candidates file with the Election Commission of India. In 2003, the Indian Supreme Court made it mandatory for all candidates in national and state legislative elections to file sworn affidavits containing details of their education, assets and liabilities, and criminal charges. Candidates must report any charges that occurred at least six months before a given election and for which a conviction would result in two or more years of imprisonment. The affidavits are public, easily accessible online (for instance, via myneta.info), and subject to intense scrutiny by opposition parties and the media. Moreover, candidates can be disqualified and face fines or imprisonment for lying in their affidavits (Asher and Novosad 2023). Omission of criminal charges in affidavits is therefore not considered a widespread problem (Vaishnav 2017; Prakash et al. 2019). Furthermore, charges only appear on a politician's record once a court has taken cognisance of the case or framed charges—i.e., after examining case materials and determining that there is *prima facie* evidence that the accused has committed the offense.[25] Charges on candidates' af-

[25] In Indian criminal procedure, a court taking cognisance means that it has assessed the alleged offense and determined that there are sufficient grounds to initiate judicial proceedings. This requires examination of the police charge-sheet or complaint, supporting evidence, and witness statements, and is substantially more demanding than simply registering a police complaint. For grave offenses such as attempt to murder, courts apply more stringent scrutiny, assessing intent to cause death based on the nature of the weapon, injuries, and surrounding circumstances, often relying on medical and forensic evidence.

fidavits have therefore passed a meaningful level of judicial screening, especially for severe crimes where the evidentiary threshold is higher, reducing concerns about spurious charges. Consistent with the above, candidate affidavit data—and specifically criminal charge information—has been used numerous times in economics and political science research, including Fisman et al. (2014), Chauchard et al. (2019), Gehring et al. (2019), Fisman et al. (2021), Prakash et al. (2024), Magesan et al. (2025), and Goyal (2025).[26]

Polling Station and Village Data. To measure each candidate's electoral performance in the 2017 experiment, we compile publicly available electoral returns data from UP's Office of the Chief Electoral Officer. This polling-station-level data includes the votes received by each candidate and the total number of registered voters. We also collect this information for two prior elections—the 2014 national parliamentary elections and the 2012 state assembly elections. For village demographic information, we use data from the 2011 Census of India and the 2011 Socioeconomic and Caste Census, acquired through the SHRUG platform (Government of India 2011a,b; Asher et al. 2021). To match polling stations to villages, we use 2011 census village shapefiles acquired from ML InfoMap together with polling station GPS coordinates from the dataset of Susewind (2014). Finally, our partner telecom companies provided us with village-level subscriber count data for the 2017 experiment.

Survey Data. For the 2022 experiment, we enrolled voters in a baseline survey approximately one month before election day. In addition to gathering demographic information, the phone-based survey captured measures of attention, beliefs, and preferences, which Section 6 describes in more detail.[27] A few days after election day, we conducted a phone-based endline survey to measure beliefs, attention, decision weights, candidate evaluations, and vote choice.

4.3 Sample Selection and Randomisation

2017 Experiment. Our first experiment was conducted in 38 constituencies during phase 4 of the 2017 elections. Within this set of constituencies, we restricted attention to villages with one or two polling stations to ensure that each polling station primarily served voters from within the village. We further excluded (i) villages with outlier population sizes—those in the top and bottom 1% of their district population distribution; and (ii) villages where our partner telecom companies had very low coverage—those in the bottom 1% of the village subscriber share distribution. This procedure yielded an experimental

[26] Criminal convictions are a poor alternative in the Indian context due to severe judicial delays (Sukhtankar and Vaishnav 2015) and selective enforcement (Poblete-Cazenave 2025). Moreover, measures based on pending criminal charges robustly predict the quality of governance (Asher and Novosad 2023; Prakash et al. 2024).

[27] Parties' candidate rosters were finalized only two weeks before election day in each phase. Measuring baseline beliefs regarding candidate criminality in the 2022 election was thus infeasible; we collect beliefs regarding candidates in the previous 2017 election.

sample of approximately 3,800 villages with 4.2 million registered voters. Sample villages within each constituency were randomly assigned to treatment versus control, giving 1,400 treatment villages and 2,400 control villages.[28]

2022 Experiment. Our second experiment was conducted in 23 assembly constituencies during phases 5–7 of the 2022 elections. None of these constituencies were covered by our 2017 experiment (Appendix Figure A1). Drawing from our partner survey company's database of mobile subscribers for each constituency, we enrolled approximately 3,700 individuals in the baseline survey, which was conducted roughly one month prior to election day. Within each constituency, individuals in our baseline survey sample were randomly assigned to treatment or control, giving 2,000 treatment and 1,700 control individuals.[29] As in the 2017 experiment, treated individuals received messages containing information about the criminal charges of candidates in their constituency. Control and treatment individuals were recontacted for an endline survey a few days following the election. Our endline attrition rate was 20.4%, similar to the rates for other endline surveys conducted in election-related field experiments in lower-income democracies (Marshall 2019; Acemoglu et al. 2025; Bowles and Larreguy 2025), and did not differ between treatment and control ($p = 0.20$), yielding an endline sample of 3,000 voters.

Summary Statistics. As in the rest of UP, many candidates in our 2017 and 2022 sample constituencies faced criminal charges. In our 2017 sample, nearly 35% of major-party candidates had at least one criminal case (Table 1, Panel A), and the average candidate faced 0.76 cases, where a criminal case may involve multiple charges.[30] Approximately 16% of all major-party candidates faced charges for violent crimes, and 9% faced murder-related charges.[31] Criminal candidates were even more prevalent in our 2022 sample constituencies: 52% of major party candidates faced at least one criminal case, 35% faced charges for violent crimes, and 11% faced murder-related charges.

Criminal candidates were not confined to a small set of "bad" constituencies, but geographically dispersed. In 2017, nearly 80% of constituencies had at least one major-party candidate with a criminal charge,

[28]Randomisation within each constituency was stratified by number of polling stations, as well as partner subscriber count in total and as a share of village population. Approximately one-third each of treated villages were randomly assigned to receive the additional coordination priming or ethnic voting content. For expositional clarity and given insignificant impact differences by information treatment type (Appendix Table A5), we pool these groups into a single treatment going forward.

[29]Randomisation was stratified by constituency. Approximately 20% of treatment group individuals were randomly assigned to receive the additional coordination priming content. As in 2017, we observe insignificant differences by information treatment type (Appendix Table A24) and pool both groups into a single treatment in our subsequent analysis.

[30]Three of the 116 candidates in our 2017 sample were affiliated with a smaller party allied with the BJP—coalition partners typically do not field candidates in the same race—and two were non-major-party incumbents. For simplicity, we include these as "major party" candidates throughout the paper.

[31]We use "murder" to refer to formal charges of either murder or culpable homicide.

45% included a candidate with a violent charge, and 26% had a candidate with a murder-related charge (Table 1, Panel B). In 2022, criminal politicians were nearly ubiquitous: 96% of constituencies had at least one major-party criminal candidate, 83% had a violent criminal candidate, and 35% featured a candidate with a murder-related charge. At the same time, at least one clean major-party candidate was present in 97% of constituencies in 2017 and in 96% of constituencies in 2022. Voters therefore had a choice between clean and criminal major-party candidates in the large majority of races in both 2017 (76%) and 2022 (91%).

4.4 Balance Checks

We validate our randomisation in each experiment by checking for balance on pre-treatment characteristics.

2017 Experiment. Given the village-level randomisation in the 2017 experiment, we test whether treatment status is correlated with pre-treatment village demographic, infrastructural, and political characteristics. Appendix Table A2 reports results from regressions of each baseline characteristic on an information treatment indicator and randomisation strata fixed effects.[32] We consider village-level demographic and infrastructural characteristics in Panel A. We examine polling-station-level electoral characteristics in Panel B, clustering standard errors by village. Treated and control villages have similar geographic and demographic characteristics, levels of public good provision, and economic development. They are also balanced on pre-treatment political characteristics, including electorate size as well as prior election turnout and voting behaviour.

2022 Experiment. For the 2022 experiment, where treatment was randomised at the individual level, we test for balance on baseline voter characteristics.[33] Appendix Table A6 presents results from regressions of voter attributes on an individual-level treatment indicator and randomisation strata fixed effects. Treated and control voters are balanced on a wide array of demographic characteristics, as well as on baseline political variables including political participation and support for different parties. Relative to the general population, our sample contains a lower share of women and is more highly educated. This is expected, since our phone-based survey and treatment restricts the sample to mobile phone owners—a group that in India is disproportionately male and more educated, reflecting gaps in mobile access (GSM Association 2019).[34]

[32] Here, and in subsequent analysis, we drop 7% of polling stations for which the location data used for matching to villages was found to be inaccurate.

[33] We restrict attention to voters used in our subsequent analysis, i.e., those who were also present at endline.

[34] To the extent that men and generally more educated individuals are more politically engaged—and so more likely to pay attention to and have accurate beliefs about candidate criminality in the first place—any estimated treatment-driven increases in criminality salience and belief accuracy will most likely be attenuated relative to corresponding general population effects.

5 Electoral Impacts of Information: Evidence from Large-Scale Experiment in 2017

5.1 Impacts on Voting Behaviour

We evaluate the impacts of providing voters information about candidate criminality by estimating the following regression specification:

$$Y_{cs} = \beta Info_v + \phi(Info_v \times Crim_c) + \alpha_c + X'_{sv}\theta + \varepsilon_{cs}, \tag{1}$$

where Y_{cs} is log votes received by candidate c at polling station s in census village v in assembly constituency a. $Crim_c$ indicates whether candidate c in assembly constituency a has any criminal charges, while $Info_v$ indicates whether village v received our information treatment. X_{sv} denotes controls for randomisation strata and log registered voters at polling station s in village v. By including candidate fixed effects α_c, we rely on within-candidate variation, comparing a candidate's electoral performance in treated and control villages. Standard errors are clustered by village, the level of randomisation. The coefficient β gives the average treatment effect of information on votes for clean candidates, while ϕ gives the marginal effect for criminal candidates.

We find that informing voters about candidate criminality benefits clean candidates. Panel A of Figure 1 shows that clean candidates receive 3.4% more votes at polling stations in treated villages, and this effect is significant at the 1% level.[35] By contrast, criminal candidates lose votes when voters receive criminality information: the interaction between treatment and criminality is negative and significant (Table 2 , column (1)), giving an imprecisely estimated 2.7% total reduction in votes for criminal candidates at treated polling stations (p = 0.198).[36]

Voters processed information about candidate criminality in a sophisticated manner. Electoral penalties for criminal candidates scale with crime severity. In column (2) of Table 2, we examine heterogeneity by number of criminal cases, replacing the any-charge indicator with criminal case count in Equation (1). We find that candidates receive 1.3% fewer votes per criminal case in treated villages. Next, we consider heterogeneity by criminal charge type. Though interacting treatment with multiple sub-groups reduces precision, column (3) shows that treatment effects are more negative for candidates facing murder-related charges, the most severe charges in our sample.[37]

Voters' reactions also depend on the electoral context. In Panels B and C of Figure 1, we re-estimate

[35] Table 2 provides underlying underlying regression results.

[36] Restricting attention to races involving at least one clean and one criminal candidate, information increases votes for clean candidates by 3.9% (p = 0.024) and reduces votes for criminal candidates by 3.9% (p = 0.074).

[37] Our information treatment structure does not allow us to separately consider different types of nonviolent crime.

Equation (1) separately for races with and without violent criminal candidates. Consistent with our conceptual framework, we find that clean candidates gain votes primarily when facing violent-charged opponents. Information increases clean candidates' votes by an insignificant 1.8% when there are no violent opponents, but leads to a significant 6.5% vote gain in races with a violent-charged opponent.[38]

Finally, we examine the effect of information on turnout by estimating the specification:

$$Y_s = \beta Info_v + \alpha_a + X'_{sv}\theta + \varepsilon_s. \quad (2)$$

The outcome is log total votes cast and we include assembly constituency fixed effects α_a, thus comparing turnout in treated and control villages within a constituency. The information treatment has a small (0.6%) and statistically insignificant impact on turnout (Appendix Table A3). This null result is precisely estimated: with 95% confidence, we can rule out turnout increases larger than 1.6% and decreases of more than 0.3%.[39] Turnout remains unchanged for elections with and without violent-charged candidates. These findings are consistent with voters switching support to clean candidates, or with offsetting turnout responses among supporters of clean and criminal candidates. While we cannot distinguish between these two explanations in our 2017 experiment, we consider them further in Section 6.1, using the individual-level data from our 2022 experiment.

Our 2017 experiment shows that informing voters about candidates' criminal charges affected their electoral choices.[40] Clean candidates receive more votes, and these electoral gains are concentrated in races where there is a stark contrast in criminality among candidates. Furthermore, these impacts are captured using administrative electoral data and thus are not subject to survey response or social desirability bias (Bursztyn et al. 2025).

5.2 Voter Persuasion Rates and Electoral Importance

Voter Persuasion Rates. Our results show that providing voters information about criminal charges increased votes for clean candidates. Following the approach of DellaVigna and Gentzkow (2010), we calculate our treatment's persuasion rate—i.e., the share of treated individuals who voted for a clean candidate because they received information about candidates' charges.[41]

We calculate persuasion rates under two bounding assumptions about information diffusion. The first scenario assumes that information remained confined to households who received the voice calls and

[38] A test of the equality of clean candidate electoral benefits across race types yields $p = 0.100$.

[39] Appendix Table A3 additionally shows no impact on votes received by non-major-party candidates.

[40] During a later election phase, we sent mobile suscribers in randomly selected villages "placebo" messages that encouraged them to learn more about candidates, but included no candidate information (see Appendix B.1 for details). The placebo treatment did not affect voting behaviour (Appendix Table A4), reinforcing the relevance of messages' informational content.

[41] Appendix B.2 provides calculation details.

text messages (the "receiving household" scenario). Under this assumption, the implied persuasion rate is 3.3% (Table 3, column (1)). The second scenario assumes that information spreads to all households in treated villages (the "entire village" scenario), and yields a lower-bound persuasion rate of 1.2%. Because the (unobserved) actual pattern of information exposure in treated villages lies between these two extremes, we estimate that the true persuasion rate is between 1.2–3.3%. These estimates are slightly larger than those reported for encouragement mailer cards in U.S. get-out-the-vote experiments, and unsurprisingly much smaller than those reported for direct-contact interventions such as door-to-door canvassing (Gerber and Green 2000; DellaVigna and Gentzkow 2010).

At-scale Election Impacts. Building on the voting impacts identified in our experiment, we consider the potential for such information campaigns to meaningfully impact election outcomes if conducted at scale.[42] Specifically, we simulate the impact of scaling our information campaign to all households and estimate the resulting change in the share of criminal winners in the 2017 and 2012 UP assembly elections. We assume that our experimental estimates of information effects in races featuring both clean and criminal candidates hold at scale and that each household owns a mobile phone.[43] In the "receiving household" scenario, only mobile subscribers receive information, so we scale up the experimental treatment effect on the clean-criminal vote share margin by the inverse of the average treated mobile subscriber share (0.37). In the "entire village" scenario, information diffuses to all households, so no scaling is necessary.[44]

Combining these implied vote shifts with observed victory margins in races where a criminal winner defeated a clean runner-up, we estimate potential changes in election outcomes. In the 2012 elections, a criminal candidate defeated a clean opponent in 103 constituencies. We estimate that 11 (32) of these races would flip to a clean winner under the entire-village (receiving-household) scenario (Table 3, column (2)). In 2017, criminal candidates defeated clean runners-up in 79 constituencies, and 3 (9) races are projected to flip in the entire-village (receiving-household) scenario (column (4)). These simulations imply that at-scale information provision could reverse outcomes in 8–23% of races where criminal candidates defeated clean opponents in recent UP elections.[45]

[42] To be clear, our own experiment was not conducted at sufficient scale to alter election outcomes, but our treatment effects guide this simulation.

[43] Given our experimental design, we are unable to consider potential nonlinearities in voting impacts that may emerge as the share of villages receiving information is increased—e.g., through a voter coordination channel (amplifying) or political party counter-reaction channel (attenuating). We thus view our analysis as a suggestive benchmarking exercise.

[44] The estimated vote share margin shift toward clean candidates is 1.1 (3.0) percentage points in the entire-village (receiving-household) scenario.

[45] Additionally, under the receiving-household scenario, one race in the 2012 elections (and none in 2017) had a vote share margin between a criminal winner and clean second-runner-up small enough (1.3pp) for at-scale information provision to result in a clean winner via a "leapfrogging" effect over a criminal runner-up.

Overall, our 2017 experiment demonstrates that providing voters information about candidate criminality can meaningfully impact real-world voting behaviour and has the potential to improve political selection. However, the experiment does not allow us to identify the mechanisms driving voter responses. In particular, we cannot determine whether voters learn about criminality or shift attention toward it. To examine mechanisms, we conducted a second experiment during the 2022 UP assembly elections that pairs the same information treatment with detailed survey measures of voter beliefs, attention, decision weights, candidate evaluations, and electoral choices.

6 How Voters Process Information: Evidence from Mechanism Experiment in 2022

6.1 Impacts on Candidate Knowledge and Evaluation

Voter Beliefs. We begin by describing voters' beliefs about politicians' criminal charges. At baseline, the majority of voters held inaccurate beliefs: only 34% could correctly identify whether their incumbent legislator faced criminal charges. Even at endline, control group voters had accurate beliefs about the criminal status of candidates in their constituency just 32% of the time (Appendix Figure A2). Moreover, these errors were systematic: most voters underestimated the prevalence of criminal candidates. Consequently, voters were four times more likely to have accurate beliefs about clean candidates (56% correct) than criminal candidates (14% correct).

We examine how information about candidates' charges affected voter beliefs by estimating the following regression, analogous to Equation (1), our baseline specification for the 2017 experiment analysis:

$$Y_{ci} = \beta Info_i + \phi(Info_i \times Crim_c) + \alpha_c + \varepsilon_{ci}. \tag{3}$$

Y_{ci} denotes an outcome for voter i in relation to candidate c in assembly constituency a. $Info_i$ denotes whether individual i received the information treatment and α_c are candidate fixed effects. This within-candidate variation enables us to compare how treated and control voters perceive and evaluate each candidate. Standard errors are clustered at the voter level.

Our information treatment improved belief accuracy. Relative to the control group, treated voters are 4.7 percentage points (or 8.4%) more likely to hold accurate beliefs about clean candidates (Panel A of Figure 2).[46] Beliefs about criminal candidates update more sharply: treatment increased accuracy by 12.9 percentage points (91%), nearly doubling voters' ability to identify criminal politicians.

[46]Table 4 provides underlying regression results.

Candidate Evaluations and Vote Choice. Having established that information improves belief accuracy, we examine whether it changes candidate evaluations and vote choice. We measure voter ratings of candidates on a five-point scale from "very bad" (one) to "very good" (five), standardized to a z-score. Information led voters to rate criminal candidates 0.12 SD lower, and improved clean candidates' ratings by 0.04 SD (Panel B of Figure 2), though the latter effect is marginally insignificant ($p = 0.12$). These evaluation changes translated into voting behaviour: treated voters are 2 percentage points (7.3%) more likely to indicate voting for clean candidates (Panel C of Figure 2), and 0.8 percentage points (5.3%) less likely to vote for criminal candidates, though this latter effect is imprecisely estimated ($p = 0.16$).

As in the 2017 experiment, we find no impacts on voter turnout (Appendix Table A7). Turnout remains unchanged for unaffiliated voters as well as for supporters of parties fielding clean or criminal candidates (Appendix Table A8). Taken together, our results across both experiments suggest that the electoral impacts of information are driven by voters updating beliefs and re-evaluating candidates, rather than turnout changes.

Stronger Effects in High-Contrast Constituencies. Our model predicts that information has stronger electoral effects when candidates differ more sharply in quality—even though voters update beliefs similarly across races. Consistent with this, Table 4 shows that the effects of information on belief accuracy do not vary with candidate composition (columns (2) and (3)). Treatment nevertheless increases clean candidates' ratings and vote shares only when they face violent-charged opponents (columns (6) and (9)), with no such impacts in races where candidates have more similar (low) levels of criminality (columns (5) and (8)).These patterns indicate that voters are comparing and contrasting candidates rather than responding mechanically to information about criminal charges.

Table 4 also shows that, within constituencies, violent-charged candidates incur larger rating penalties than non-violent criminal candidates (column (6)). This occurs even though belief updating is, if anything, slightly weaker for violent-charged candidates (column (3)).

Experimenter Demand. A natural concern in survey-based experiments is that estimated treatment effects may reflect experimenter demand or social desirability bias rather than genuine changes in beliefs or behaviour (Haaland et al. 2023; Bursztyn et al. 2025). We therefore take several steps, following best-practice recommendations in this literature, to assess and mitigate these concerns.

First, following Bursztyn et al. (2023) and Haaland et al. (2023), we directly elicit respondents' perceptions of the survey's purpose. Only 0.8% of respondents thought the focus of the survey was on politician criminality, while roughly 53% thought it related to views on candidates or parties in the election more broadly. Second, we test for treatment effect heterogeneity by perceived survey purpose. Voters who thought the survey focus related to candidate or party views respond no differently to the information treatment than other respondents (Appendix Figure A3), and we observe a similar pattern for the very small share who thought the survey focused on criminality, though this subgroup is too small for precise statistical inference.

Importantly, our survey-based findings from 2022 replicate the results of the 2017 experiment, which are based on administrative data from official election returns. These data capture a consequential behavioural outcome and are not susceptible to survey response or social desirability bias. The consistency of results across our two experiments is thus difficult to reconcile with experimenter demand being a key driver of our 2022 findings.

Replication and External Validity. Our 2017 and 2022 experiments both find that informing voters about candidate criminality raises support for clean candidates and reduces support for criminal candidates, especially when differences in criminality are large.[47] Beyond alleviating concerns about experimenter demand, this replication strengthens confidence in the robustness and external validity of our results. Although both experiments implemented the same core information treatment, they were conducted in different constituencies and election cycles, operated at different scales (4.2 million vs 3,700 voters), were randomised at different levels (village vs individual), and measured outcomes in different ways (administrative vs survey data). The consistency of our findings across these very different research designs suggests that our findings capture a genuine and generalisable behavioural response to information about candidate criminality at least within Uttar Pradesh, which comprises 6% of the global population living in a democracy (Government of India 2020; Economist Intelligence Unit 2025).[48]

6.2 Effects on Attention and Decision Weights

Our conceptual framework predicts that providing voters information about candidate criminality not only facilitates learning about that attribute, but also heightens its salience and the weight it receives in electoral choices. In this section, we describe and validate our measure of voter attention. Then, we assess whether our information treatment affected what voters paid attention to and placed weight on in their decisions.

[47] Also consistent with the 2017 experiment, Appendix Table A9 shows for the 2022 experiment that information leads to larger ratings and electoral penalties against candidates with more criminal cases.

[48] In addition, while we did not pre-register an analysis plan, replication also reduces false positive concerns (Coffman and Niederle 2015; Olken 2015).

Measuring Salience. We first ask voters open-ended questions eliciting what is top of mind when they think about the election and the candidates in their constituency.[49] Enumerators note down responses before asking a follow-up structured question about whether specific candidate attributes and policy issues come to mind.[50] Using responses to the structured question, we construct indicator variables for the salience of each candidate attribute and policy issue. The initial open-ended elicitation reduces the risk of priming voters toward specific topics, while the structured follow-up question reduces the measurement error that would arise from capturing and coding lengthy open-ended answers in a phone survey (Haaland et al. 2025).[51]

At baseline, we find that voters were most attentive to valence attributes (Appendix Figure A4), such as candidate education (81%), political experience (62%) and criminality (48%). By contrast, identity-based attributes like religion (8%) and caste (5%) were top of mind for comparatively fewer voters. Among policy topics, economic issues dominated: employment (71%) and farmer distress (60%) received more attention than the divisive Ayodhya temple controversy (30%).[52] Crime and corruption—policy issues closely associated with criminal politicians (Chemin 2012; Prakash et al. 2024)—were also salient to substantial shares of voters at baseline (42% and 50%, respectively).

Validating Salience Measure. To validate our salience measure, we embedded a vignette-based recall task in the baseline survey. Early in the survey—prior to any salience questions—respondents heard a profile of a politician who contested a prior election in another region of UP. The profile described the candidate's characteristics—including his or her education, wealth, and criminality—and the candidate's name also clearly signalled religious/ethnic identity.[53] At the end of the baseline survey, roughly 25 minutes later, we asked respondents to recall what they remembered from this profile.

We find that voters are substantially more likely to recall candidate attributes from the vignette that they reported as being top of mind during the salience elicitation procedure. For example, voters for whom criminality was salient were more than twice as likely to remember criminality from the vignette: recall rates increase from 16% to 37%, and this difference is robust to controlling for a broad

[49] Specifically, "*When you think of the election, what are some of the thoughts that come to mind?*" and "*Now think about the candidates who will stand in your constituency in the upcoming election. What are some thoughts that come to your mind?*".

[50] In particular, "*When you think about the election, do any of the following thoughts come to your mind?*"

[51] Appendix Table A10 shows that salience measures based on structured and open-ended responses are strongly correlated.

[52] The Ayodhya temple issue is a long-running socio-religious dispute over the site in UP where the 16th-century Babri Masjid mosque stood until its demolition by right-wing Hindu extremists in 1992, triggering nationwide communal riots and decades of legal and political conflict. During the 2022 UP elections, the incumbent BJP promised to complete construction of a Hindu temple on the disputed site—which occurred in 2024. For additional details see, e.g., Lodhi (2024) and Serhan (2024).

[53] Voters were randomly assigned one of three different vignettes. Appendix B.4 provides the full vignettes. After reading out the vignette, surveyors asked respondents whether they had previously heard of the candidate and then moved on to ask questions on other topics.

set of individual-level characteristics (Appendix Table A11). This pattern extends beyond criminality. Regressing recall of a given attribute from the vignette on the respondent's reported salience of that attribute, we find a strong positive association between salience and recall across candidate attributes (e.g., education, religion, wealth). These results provide validation that our salience measure identifies cognitively accessible attributes that are top of mind for voters (Bordalo et al. 2020).

Effect of Information on Salience. To evaluate whether informing voters about candidates' criminal charges influenced what they were attentive to, we estimate the specification:

$$Y_i = \beta Info_i + \alpha_a + \varepsilon_i. \tag{4}$$

Y_i denotes an indicator for whether a given attribute or issue is salient to voter i at endline, measured in the same way as at baseline, and α_a denotes constituency fixed effects.

At endline, criminality was salient to only 37% of control group voters. Treated voters are 3.6 percentage points (10% or 0.08 SD) more likely to pay attention to candidates' criminal charges (Appendix Table A12, Panel A).[54] Our information treatment also raised the salience of crime and corruption, two policy issues that are closely associated with criminal politicians (Asher and Novosad 2023; Prakash et al. 2024), even though our messages did not directly mention those subjects. This is consistent with similarity-based retrieval (Kahana 2012; Bordalo et al. 2020, 2023, 2025; Enke et al. 2024; Graeber et al. 2024). Combining these three outcomes into an index, constructed as the average of component z-scores (Kling et al. 2007), we find that treatment raises the salience of criminality by 0.10 SD (Figure 3, Panel A).[55]

By contrast, we find no evidence that treatment affected the salience of identity-related attributes or issues: information about candidates' charges does not affect voters' attention to candidate caste or religion, or impact the salience of the divisive Ayodhya temple issue (Appendix Table A12, Panel B). Similarly, the salience of important economic policy issues such as employment and COVID management remains unaffected (Panel C). Corresponding identity-related and economic policy salience indices are therefore unchanged (Figure 3, Panel A).

Impact of Information on Decision Weights. Having shown that our information treatment raises the salience of criminality, we evaluate whether it impacts voters' decision weights, as predicted by salience theory (Bordalo et al. 2022). We asked voters to assess the importance of different factors when deciding their vote, measured on a five-point scale from "not at all important" (one) to "most important" (five)

[54]While sizeable, these effects may have be even larger had criminality been less initially salient to voters—a prediction supported by our heterogeneity results in Section 7.2.

[55]The corresponding impact using the open-ended-response-based criminality salience index is 0.07 average SD (*p* = 0.022).

then standardized to a z-score.[56] Consistent with predictions of both salience and rational inattention models (Bordalo et al. 2022; Matějka and Tabellini 2021), voters placed greater decision weight on salient attributes at baseline (Appendix Table A13).[57]

To examine the effects of criminality information, we estimate Equation (4) with endline decision weights as the outcome. The information treatment increases voters' weight on candidate criminality by 0.07 SD (Figure 3, Panel B). By contrast, the intervention had no impact on decision weight indices for identity-related topics and economic policies, echoing the null results on the salience of those attributes and issues.[58] Together, our results show that information influences what voters pay attention to and place weight on in their electoral decisions.

7 Unpacking Mechanisms

7.1 Bottom-Up and Top-Down Determinants of Voter Attention

Two canonical traditions offer competing accounts of how attention is allocated. Our framework builds on bottom-up salience models, in which attention is captured by inherently salient stimuli (Bordalo et al. 2022); by contrast, top-down rational inattention models hold that attention is deployed strategically, in proportion to its decision relevance (Sims 2003; Matějka and McKay 2015; Matějka and Tabellini 2021). In this section, we examine the distinct predictions of these two traditions using our empirical evidence, finding that voters' responses in our context are better explained by bottom-up salience than by top-down rational inattention.

Surprise and Contrast. We examine the role of surprise and contrast, highlighted by Bordalo et al. (2022) and incorporated in our model as important determinants of salience. Information has a stronger salience impact when it is more surprising: treatment effects are larger for voters with inaccurate baseline beliefs about criminality (Appendix Table A15, columns (1) and (2)).[59] Salience effects are also larger in races with violent-charged candidates (columns (3) and (4)), where the contrast in criminality among candidates is more pronounced. Both relationships are robust to the inclusion of interactions between treatment and a host of voter characteristics. In addition, among control voters—for whom criminality is

[56] Specifically, at baseline, "*When you are deciding who to vote for in the upcoming election, how much importance will you give to each of the following factors?*" At endline, we ask the same question with regard to the recently concluded election.

[57] Decision weights were collected for a smaller set of attributes and policy issues as compared to salience. For criminality, we therefore consider candidate criminal record alone.

[58] Appendix Table A14 shows impacts separately for each component.

[59] Given the constraint we faced regarding the timing of parties finalizing candidate selection, we capture baseline informedness using belief accuracy about previous-cycle candidate criminality. Appendix Table A16 presents summary statistics for our measures of surprise and contrast.

less prominent—we see no relationship between salience and surprise or contrast (Appendix Table A17). This is consistent with the idea that prominence, surprise, and contrast are complementary: prominent features draw attention, especially if they are contrasting or surprising, while non-prominent features may be neglected (Bordalo et al. 2020).

Decision Value. Models of rational inattention predict that voters will pay more attention to decision-relevant information. We test this prediction using two factors that vary the decision value of criminality information to voters: election competitiveness and voter indecision between parties (Devdariani and Hirsch 2023; Bursztyn et al. 2024). Information about candidates' charges generates a similar increase in attention to criminality in close and less competitive elections (Appendix Table A18, columns (1) and (2)). We also find no evidence that treatment has a larger effect on voters who rate parties more closely at baseline (columns (3) and (4)), for whom information may be more decision-relevant. Our results suggest that the salience effects of information do not vary with either proxy of decision value.

Additionally, under top-down allocation of attention, voters would not be expected to devote costly attention to electorally irrelevant candidates. With bottom-up salience, however, such candidates may attract attention if they increase contrast or surprise. We therefore assess whether information about electorally irrelevant candidates affects salience (Appendix Table A19), classifying candidates as irrelevant if they received less than 10% of the vote in their constituency. We find that surprising information about irrelevant candidates draws attention to criminality (column (1)). Similarly, sharing information about electorally irrelevant violent-charged candidates—which has limited decision value but still heightens contrast among candidates—increases the salience of criminality (column (2)), though this effect is less precisely estimated ($p = 0.137$). While difficult to reconcile with rational inattention, these patterns are readily explained by bottom-up salience.

7.2 Interplay of Learning and Salience

Section 6 showed that information affects voter behaviour through two mechanisms: belief updating (*learning*) and attention shifts (*salience*). In this subsection, we provide evidence that salience is the primary channel through which information impacts decision weights—a key step in the causal chain underlying salience models. Then, we assess the relative importance of learning and salience in shaping ultimate voting behaviour.

Tracing Causal Chain of Salience Models. We conduct three analyses to test whether salience mediates the effect of information on decision weights. First, we run a standard mediation analysis (Imai et al. 2010; Andries et al. 2024). Information increases the decision weight on criminality, and this effect is robust to controlling for voters' demographic characteristics and baseline beliefs and salience (Appendix Table A20, columns (1)–(4)). However, the treatment effect becomes small and statistically insignificant once we control for endline salience (columns (5)–(6)), suggesting that salience is the key mediating channel (Andries et al. 2024). As with any mediation analysis, this interpretation rests on a sequential ignorability condition that salience is as-good-as-randomly assigned conditional on treatment and observed covariates (Imai et al. 2010). This is a strong assumption in our setting, since unobserved factors could plausibly affect both salience and decision weights.

Given this concern, we next apply the assumption-light test of Kwon and Roth (2026). This approach only requires that the information treatment is randomly assigned—satisfied by our experimental design—and weakly increases the salience of criminality, which is plausible since our messages explicitly highlight candidates' charges. We cannot reject the sharp null of full mediation—that information affects decision weights only through salience (p = 0.49). Consistent with this, the distribution of criminality decision weights is very similar among treated and control voters for whom criminality is not salient, suggesting no alternative mechanism by which information affects decision weights (Appendix Figure A5).

Third, we use the causal forest method of Wager and Athey (2018) to estimate conditional average treatment effects (CATEs) for each voter across our five key outcomes: belief accuracy, salience, decision weight, candidate ratings, and vote choice. The CATE estimates allow us to evaluate the model's predicted causal chain at the individual level.[60] We find that voters who update beliefs more substantially in response to information—those for whom information was more surprising—experience greater increases in the salience of criminality (Figure 4, Panel A). In turn, voters with stronger salience responses exhibit larger increases in their decision weight on criminality (Panel B). LASSO regressions corroborate this second link: voters' salience responses are the strongest predictor of their decision weight responses (Appendix Figure A7).

Together, these results support the causal chain in salience models: information draws voters' attention to candidate quality, leading them to place greater decision weight on it when voting.

Relative Importance of Learning and Salience. To consider the relative importance of the learning and salience channels for ultimate changes in electoral behaviour, we first identify the voters for whom only one channel operates. Using our CATE estimates, we classify voters based on their predicted belief and attention response to information. We categorise voters as *Learners* if information improved their

[60]The causal forest algorithm is a non-parametric, data-driven approach to estimating heterogeneous treatment effects. We implement the method using the generalized random forest package (Athey et al. 2019), with details in Appendix B.5. Appendix Figure A6 presents the CATE distributions for each outcome.

belief accuracy about criminal and clean candidates (i.e., they have positive and significant CATEs), but did not raise the salience of criminality. Conversely, we classify voters as *Engagers* if they respond to information by paying more attention to criminality, but do not update beliefs.[61] Finally, *Responders* show increases on both margins.[62]

Comparing candidate rating and vote choice responses across these voter types, we find that *Engagers* have slightly stronger treatment impacts in general than *Learners* (Appendix Figure A8, Panels A–D). This is driven by *Engagers*' much larger increase in the decision weight on criminality (Panel E). Our results thus indicate that learning and salience both matter for voter behavior and that salience is at least as important, even in a context where most voters were poorly informed at baseline. We also find that information has the strongest overall effects on *Responders*, indicating that salience and belief updating are complementary forces in changing voter behaviour. To our knowledge, ours is one of the first attempts to empirically compare the importance of the salience and learning channels.

7.3 Alternative Mechanisms: Motivated Reasoning and Coordination

Motivated Reasoning. Information may be ineffective at improving political selection if voters disregard unfavorable information about in-group candidates. Scholars are increasingly concerned that such motivated reasoning—where individuals gather and evaluate evidence with a bias toward maintaining their existing beliefs—could distort voter learning in polarised polities (Epley and Gilovich 2016; Bénabou and Tirole 2016; Adida et al. 2017; Nyhan 2020; Thaler 2024; D'Amico and Tabellini 2024). Thus, next we examine whether voters process information differently when it challenges their group identity.

To assess whether voters respond asymmetrically when presented with negative versus positive information about own-group versus out-group candidates, we add terms to Equation (3):

$$\begin{aligned} Y_{ci} = \beta Info_i + \phi(Info_i \times Crim_c) + \gamma(Info_i \times Own_{ci}) + \sigma(Info_i \times Crim_c \times Own_{ci}) \\ + \tau Own_{ci} + \omega(Crim_c \times Own_{ci}) + \alpha_c + \varepsilon_{ci}. \end{aligned} \tag{5}$$

Own_{ci} indicates whether candidate c and voter i share group identity. The coefficient γ captures voters' differential response to positive candidate information (i.e., that the candidate is clean) when it refers to in-group candidates, while $(\gamma+\sigma)$ captures voters' differential reaction to negative candidate information (i.e., that the candidate has criminal charges) when it refers to in-group candidates. In Figure 5, we consider two dimensions of group identity relevant in the Indian context: religion/caste group (top row)

[61] *Learners* and *Engagers* have similar demographic and baseline characteristics; the main observable difference is that the baseline salience of criminality is high for *Learners*, but low for *Engagers*. This is also consistent with our model predictions.

[62] *Learners*, *Engagers*, and *Responders* comprise 28%, 24%, and 25% of our sample, respectively.

and political party (bottom row).[63]

We find limited evidence of asymmetric belief updating. For both religion/caste group and party, voters update beliefs similarly when informed that an in- or out-group candidate is clean (Panel A). Furthermore, far from ignoring negative signals about own-group candidates, treated voters if anything update beliefs more strongly when criminal candidates are in-group. Treated voters are also no less likely to recall receiving information about candidates' charges, or to perceive the information provided as credible, when in-group candidates are revealed to be criminal (Appendix Table A22). These results suggest that voters engage meaningfully with, and learn from, both positive and negative information about in-group candidates.

Next, we test for motivated candidate evaluations. We find that voters reward in-group candidates more for positive signals: own-party candidates receive a larger rating boost when revealed as clean, and this pattern holds, albeit more noisily, for own-religion/caste candidates (Figure 5, Panel B). Similarly, information generates larger electoral gains for clean candidates when they are in-group, though the estimated difference is significant only for shared religion/caste group (Panel C).

By contrast, we do not find that voters punish own-group candidates less for negative signals. For both candidate ratings and vote choice (Panels B and C), the treatment-driven penalty to criminal candidates is statistically indistinguishable across in- and out-group candidates, defined based on party or religion/caste.

Therefore, at least in our context, group attachments do not appear to inhibit voter learning, and providing voters information about candidates led them to choose higher-quality candidates.

Coordination. Beyond enabling learning and shaping salience, information can also change electoral outcomes by coordinating voters (Arias et al. 2019; Bursztyn and Yang 2022). Information can facilitate both social coordination (i.e., discussion spurred by receiving information) or strategic coordination (i.e., updated second-order beliefs that can influence electoral choices). We consider both possibilities using data from our endline survey, regressing coordination-related outcomes on an individual-level treatment indicator and constituency fixed effects (Appendix Table A23, Panel A).

We find no evidence that information facilitated social coordination in our context. Information does not change the probability that voters discuss the election with their friends or neighbors, and it reduces the probability of discussing candidates' criminal charges by 3.9 percentage points (15%)—suggesting that information messaging partially substituted for social methods of information acquisition. By contrast, we find mixed evidence that information facilitated strategic coordination. Information does not change

[63] Appendix Table A21 reports underlying regression estimates. Religion/caste group match is based on religion for non-Hindus and on caste group (general, OBC, SC/ST) within Hindus. Partisan match is based on voter baseline-indicated preferred party and candidate party affiliation. Appendix B.3 gives additional details on variable construction.

voters' beliefs about the share of other voters with knowledge of candidates' criminal charges. However, treated voters are 4.6 percentage points (6%) more likely to believe that other voters place high importance on voting for a clean candidate.[64]

8 Conclusion

In this paper, we studied how voters process information about candidates for political office. We analyse the phenomenon of criminal politicians, a prevalent example of negative selection into politics. By conducting two experiments during the 2017 and 2022 Uttar Pradesh assembly elections, we show that informing voters about candidates' criminal charges meaningfully changes voting behaviour. Clean candidates gain support, criminal candidates lose support, and these effects are stronger when charges are more severe and there are stark criminality differences between candidates.

Beyond documenting electoral impacts, we identify two mechanisms through which information affects voting behaviour. First, voters form more accurate beliefs, re-evaluate candidates, and shift support toward clean candidates. Second, voters pay more attention to criminality.

Our direct measures of voters' attention, decision weights, candidate assessments, and electoral choices enable us to trace the causal chain underlying salience models: information captures attention, particularly when it is surprising or highlights contrast, and raises the decision weight voters place on criminality, amplifying electoral responses. Our causal forest estimates suggest that the learning and salience channels are each important on their own but also complementary, with the strongest effects on voting behaviour when both channels operate.

Our findings enrich our understanding of voter behaviour. Voters learn from information, but their responses are shaped by what captures their attention. They have strong ethnic and partisan identities, but respond to information about candidate quality. Our results further suggest that light-touch interventions that inform voters about candidates' valence attributes have the potential to improve political selection even in polarised environments—by helping voters learn about, and redirecting their attention towards, candidate attributes that matter for good governance and economic development.

Given our focus on how voters process information, we deliberately abstract from politicians' strategic responses in this study. However, political actors may adapt to voters becoming more informed or attentive to particular attributes (Bandiera et al. 2024). Politicians and parties could attempt to obfuscate or distort information, strategically shift voters' attention, or even adjust candidacy choices and in-office behaviour. Understanding such equilibrium responses is a valuable direction for future research.

[64] Results are similar when information is provided along with the coordination prime (Appendix Table A23, Panel B).

References

Acemoglu, D., C. G. Aksoy, C. Baysan, C. Molina, and G. Zeki (2025). Misperceptions and Demand for Democracy under Authoritarianism. Working Paper.

Adida, C., J. Gottlieb, E. Kramon, and G. McClendon (2017). Overcoming or Reinforcing Coethnic Preferences? An Experiment on Information and Ethnic Voting. *Quarterly Journal of Political Science 12*(4), 437–477.

Ajzenman, N. and R. Durante (2023). Salience and Accountability: School Infrastructure and Last-Minute Electoral Punishment. *Economic Journal 133*(649), 460–476.

Aker, J., P. Collier, and P. C. Vicente (2017). Is Information Power? Using Mobile Phones and Free Newspapers During an Election in Mozambique. *Review of Economics and Statistics* (2), 185–200.

Andries, M., L. Bursztyn, T. Chaney, M. Djourelova, and A. Imas (2024). In Their Shoes: Empathy Through Information. Working Paper.

Arias, E., P. Balán, H. Larreguy, J. Marshall, and P. Querubín (2019). Information Provision, Voter Coordination, and Electoral Accountability: Evidence from Mexican Social Networks. *American Political Science Review 113*(2), 475–498.

Arias, E., H. Larreguy, J. Marshall, and P. Querubin (2022). Priors Rule: When do Malfeasance Revelations Help or Hurt Incumbent Parties? *Journal of the European Economic Association 20*(4), 1433–1477.

Arya, Y. and A. Y. Bhatiya (2025). When do Voters 'Rally around the Flag'? The Salience of Political Messages. *Journal of Public Economics 245*, 105356.

Asher, S., T. Lunt, R. Matsuura, and P. Novosad (2021). Development Research at High Geographic Resolution: An Analysis of Night-Lights, Firms, and Poverty in India Using the SHRUG Open Data Platform. *World Bank Economic Review 35*(4).

Asher, S. and P. Novosad (2023). Rent-Seeking and Criminal Politicians: Evidence from Mining Booms. *Review of Economics and Statistics 105*(1), 20–39.

Association for Democratic Reforms (2025, May). Percentage of MPs with Declared Criminal Cases: Lok Sabha 2004, 2009, 2014, 2019, and 2024. https://web.archive.org/web/20250518073448/https://adrindia.org/content/lok-sabha-elections-2024. Accessed December 2025.

Athey, S., J. Tibshirani, and S. Wager (2019). Generalized Random Forests. *Annals of Statistics 47*(2), 1148–1178.

Bandiera, O., M. Callen, K. E. Casey, E. La Ferrara, A. Jensen, C. Landais, M. Teachout, and S. Vaziralli (2024). State Effectiveness. International Growth Centre Evidence Paper.

Banerjee, A. and R. Pande (2009). Parochial Politics: Ethnic Preferences and Politician Corruption. Working Paper.

Banerjee, A. V., E. Duflo, C. Imbert, and R. Pande (2026). Voters' Awareness of Social Insurance Programs Impacts Electoral Candidature: Experimental Evidence from India. Working Paper.

Banerjee, A. V., D. Green, J. McManus, and R. Pande (2014). Are Poor Voters Indifferent to Whether Elected

Leaders are Criminal or Corrupt? *Political Communication 31*(3), 391–407.

Banerjee, A. V., S. Kumar, R. Pande, and F. Su (2011). Do Informed Voters Make Better Choices? Experimental Evidence from Urban India. Working Paper.

Baysan, C. (2022). Persistent Polarizing Effects of Persuasion: Experimental Evidence from Turkey. *American Economic Review 112*(11), 3528–3546.

Bénabou, R. and J. Tirole (2016). Mindful Economics: The Production, Consumption, and Value of Beliefs. *Journal of Economic Perspectives 30*(3), 141–164.

Besley, T. (2005). Political Selection. *Journal of Economic Perspectives 19*(3), 43–60.

Bhalotra, S. and I. Clots-Figueras (2014). Health and the Political Agency of Women. *American Economic Journal: Economic Policy 6*(2), 164–197.

Bhandari, A., H. Larreguy, and J. Marshall (2023). Able and Mostly Willing: An Empirical Anatomy of Information's Effect on Voter-Driven Accountability in Senegal. *American Journal of Political Science 67*(4), 1040–1066.

Bidwell, K., K. Casey, and R. Glennerster (2020). Debates: Voting and Expenditure Responses to Political Communication. *Journal of Political Economy 128*(8), 2880–2924.

Boas, T. C., F. D. Hidalgo, and M. A. Melo (2019). Norms versus Action: Why Voters Fail to Sanction Malfeasance in Brazil. *American Journal of Political Science 63*(2), 385–400.

Bolsen, T. and R. Palm (2019). Motivated Reasoning and Political Decision Making. In D. P. Redlawsk (Ed.), *Oxford Encyclopedia of Political Decision Making*. Oxford University Press.

Bonomi, G., N. Gennaioli, and G. Tabellini (2021). Identity, Beliefs, and Political Conflict. *Quarterly Journal of Economics 136*(4), 2371–2411.

Bordalo, P., G. Burro, K. Coffman, N. Gennaioli, and A. Shleifer (2025). Imagining the Future: Memory, Simulation, and Beliefs. *Review of Economic Studies 92*(3), 1532–1563.

Bordalo, P., J. Conlon, N. Gennaioli, S. Kwon, and A. Shleifer (2026). How People Use Statistics. *Review of Economic Studies 93*(1), 250–285.

Bordalo, P., J. J. Conlon, N. Gennaioli, S. Y. Kwon, and A. Shleifer (2023). Memory and Probability. *Quarterly Journal of Economics 138*(1), 265–311.

Bordalo, P., N. Gennaioli, and A. Shleifer (2013). Salience and Consumer Choice. *Journal of Political Economy 121*(5), 803–843.

Bordalo, P., N. Gennaioli, and A. Shleifer (2020). Memory, Attention, and Choice. *Quarterly Journal of Economics 135*(3), 1399–1442.

Bordalo, P., N. Gennaioli, and A. Shleifer (2022). Salience. *Annual Review of Economics 14*(1), 521–544.

Bowles, J. and H. Larreguy (2025). Who Debates, Who Wins? At-Scale Experimental Evidence on the Supply of Policy Information in a Liberian Election. *American Political Science Review 119*(4), 1684–1703.

Britto, D. G., G. Daniele, M. Le Moglie, P. Pinotti, and B. Sampaio (2025). A Few Bad Apples? Criminal Charges, Political Careers, and Policy Outcomes. Working Paper.

Buntaine, M. T., R. Jablonski, D. L. Nielson, and P. M. Pickering (2018). SMS Texts on Corruption Help Ugandan Voters Hold Elected Councillors Accountable at the Polls. *Proceedings of the National Academy of Sciences 115*, 6668—6673.

Bursztyn, L., D. Cantoni, P. Funk, F. Schönenberger, and N. Yuchtman (2024, 04). Identifying the Effect of Election Closeness on Voter Turnout: Evidence from Swiss Referenda. *Journal of the European Economic Association 22*(2), 876–914.

Bursztyn, L., G. Egorov, I. Haaland, A. Rao, and C. Roth (2023). Justifying Dissent. *Quarterly Journal of Economics 138*(3), 1403–1451.

Bursztyn, L., I. K. Haaland, N. Röver, and C. Roth (2025). The Social Desirability Atlas. Working Paper.

Bursztyn, L. and D. Y. Yang (2022). Misperceptions About Others. *Annual Review of Economics 14*(1), 425–452.

Canen, N. and G. J. Martin (2023). How Campaign Ads Stimulate Political Interest. *Review of Economics and Statistics 105*(2), 292–310.

Carney, K. (2025). The Effect of Social Media on Voters: Experimental Evidence from an Indian Election. Working Paper.

Casey, K., A. B. Kamara, and N. Meriggi (2021). An Experiment in Candidate Selection. *American Economic Review 111*(5), 1575–1612.

Chauchard, S., M. Klašnja, and S. Harish (2019). Getting Rich Too Fast? Voters' Reactions to Politicians' Wealth Accumulation. *Journal of Politics 81*, 1197–1209.

Chemin, M. (2012). Welfare Effects of Criminal Politicians: A Discontinuity-Based Approach. *Journal of Law and Economics 55*(3), 667–690.

Cheng-Matsuno, V., F. Foos, P. John, and A. Unan (2023). Do Text Messages Increase Voter Registration? Evidence from RCTs with a Local Authority and an Advocacy Organisation in the UK. *British Journal of Political Science 81*(February), 102572.

Chetty, R., A. Looney, and K. Kroft (2009). Salience and Taxation: Theory and Evidence. *American economic review 99*(4), 1145–1177.

Clots-Figueras, I. (2012). Are Female Leaders Good for Education? Evidence from India. *American Economic Journal: Applied Economics 4*(1), 212–244.

Coffman, L. C. and M. Niederle (2015). Pre-Analysis Plans Have Limited Upside, Especially Where Replications Are Feasible. *Journal of Economic Perspectives 29*(3), 81–98.

Colussi, T., I. E. Isphording, and N. Pestel (2021). Minority Salience and Political Extremism. *American Economic Journal: Applied Economics 13*(3), 237–271.

Conlon, J. J. and D. Patel (2025). What Jobs Come to Mind? Stereotypes About Fields of Study. Working Paper.

Cruz, C., P. Keefer, and J. Labonne (2021). Buying Informed voters: New Effects of Information on Voters and Candidates. *Economic Journal 131*(635), 1105–1134.

Cruz, C., P. Keefer, J. Labonne, and F. Trebbi (2024). Making Policies Matter: Voter Responses to Campaign Promises. *Economic Journal 134*(661), 1875–1913.

Dal Bó, E. and F. Finan (2018). Progress and Perspectives in the Study of Political Selection. *Annual Review of Economics 10*, 541–575.

Dal Bó, E., F. Finan, O. Folke, T. Persson, and J. Rickne (2017). Who Becomes a Politician? *Quarterly Journal of Economics 132*(4), 1877–1914.

Dale, A. and A. Strauss (2009). Don't Forget to Vote: Text Message Reminders as a Mobilization Tool. *American Journal of Political Science 53*(4), 787–804.

D'Amico, L. and G. Tabellini (2024). Disengaging from Reality: Online Behavior and Unpleasant Political News. Working Paper.

DellaVigna, S. and M. Gentzkow (2010). Persuasion: Empirical Evidence. *Annual Review of Economics 2*(1), 643–669.

DellaVigna, S. and E. Kaplan (2007). The Fox News Effect: Media Bias and Voting. *Quarterly Journal of Economics 122*(3), 1187–1234.

Devdariani, S. and A. V. Hirsch (2023). Voter Attention and Electoral Accountability. *Journal of Public Economics 224*, 104918.

Dunning, T., G. Grossman, M. Humphreys, S. D. Hyde, C. McIntosh, G. Nellis, C. L. Adida, E. Arias, C. Bicalho, T. C. Boas, et al. (2019). Voter Information Campaigns and Political Accountability: Cumulative Findings from a Preregistered Meta-Analysis of Coordinated Trials. *Science Advances 5*(7), eaaw2612.

Durante, R., P. Pinotti, and A. Tesei (2019). The Political Legacy of Entertainment TV. *American Economic Review 109*(7), 2497–2530.

Economist Intelligence Unit (2025). Democracy Index 2024: What's Wrong with Representative Democracy?

Enikolopov, R., M. Petrova, and E. Zhuravskaya (2011). Media and Political Persuasion: Evidence from Russia. *American Economic Review 101*(7), 3253–3285.

Enke, B., F. Schwerter, and F. Zimmermann (2024). Associative Memory, Beliefs and Market Interactions. *Journal of Financial Economics 157*, 103853.

Epley, N. and T. Gilovich (2016). The Mechanics of Motivated Reasoning. *Journal of Economic Perspectives 30*(3), 133–140.

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

Fisman, R., F. Schulz, and V. Vig (2014). The Private Returns to Public Office. *Journal of Political Economy* (4), 806–862.

Fisman, R., F. Schulz, and V. Vig (2021). Financial Disclosure and Political Selection: Evidence from India. Working Paper.

Fujiwara, T., K. Müller, and C. Schwarz (2024). The Effect of Social Media on Elections: Evidence from the United States. *Journal of the European Economic Association 22*(3), 1495–1539.

Gehring, K., T. F. Kauffeldt, and K. C. Vadlamannati (2019). Crime, Incentives and Political Effort: Evidence from India. *European Journal of Political Economy 59*, 1–20.

Gennaioli, N. and A. Shleifer (2010). What Comes to Mind. *Quarterly Journal of Economics 125*(4), 1399–1433.

Gerber, A. S. and D. P. Green (2000). The Effects of Canvassing, Telephone Calls, and Direct Mail on Voter Turnout: A Field Experiment. *American Political Science Review 94*(3), 653–663.

Go, L., N. Fliess, and E. Østergaard-Nielsen (2026). Political Interdependence: Evidence from Emigrant Voter Turnout in 1,267 Elections Worldwide. *Journal of the European Economic Association*, jvag014.

Government of India (2011a). Census of India. Accessed May 2020 via SHRUG, https://www.devdatalab.org/shrug_download/, v1.1.samosa.

Government of India (2011b). Socioeconomic and Caste Census. Accessed May 2020 via SHRUG, https://www.devdatalab.org/shrug_download/, v1.1.samosa.

Government of India (2020). Population Projections for India and States 2011–2036. Report of the National Commission on Population.

Goyal, T. (2025). Local Political Representation as a Pathway to Power: A Natural Experiment in India. *American Journal of Political Science 69*(2), 515–530.

Graeber, T., C. Roth, and F. Zimmermann (2024). Stories, Statistics, and Memory. *Quarterly Journal of Economics 139*(4), 2181–2225.

GSM Association (2019). Connected Women – The Mobile Gender Gap Report 2019. Research Report.

Gulzar, S. (2021). Who Enters Politics and Why? *Annual Review of Political Science 24*, 253–275.

Guriev, S., N. Melnikov, and E. Zhuravskaya (2021). 3G Internet and Confidence in Government. *Quarterly Journal of Economics 136*(4), 2533–2613.

Haaland, I., C. Roth, S. Stantcheva, and J. Wohlfart (2025). Understanding Economic Behavior Using Open-Ended Survey Data. *Journal of Economic Literature 63*(4), 1244–1280.

Haaland, I., C. Roth, and J. Wohlfart (2023). Designing Information Provision Experiments. *Journal of Economic Literature 61*(1), 3–40.

Hanna, R. and S.-Y. Wang (2017). Dishonesty and Selection into Public Service: Evidence from India. *American Economic Journal: Economic Policy 9*(3), 262–290.

Harris, J. A., C. Kamindo, and P. van der Windt (2021). Electoral Administration in Fledgling Democracies: Experimental Evidence from Kenya. *Journal of Politics 83*(3), 947–960.

Imai, K., L. Keele, and T. Yamamoto (2010). Identification, Inference, and Sensitivity Analysis for Causal

Mediation Effects. *Statistical Science 25*(1), 51–71.

Iyer, L. and A. Mani (2012). Traveling Agents: Political Change and Bureaucratic Turnover in India. *Review of Economics and Statistics 94*(3), 723–739.

Iyer, L., A. Mani, P. Mishra, and P. Topalova (2012). The Power of Political Voice: Women's Political Representation and Crime in India. *American Economic Journal: Applied Economics 4*(4), 165–193.

Jain, R. (2018, September). Supreme Court Asks Parliament to Frame Law to Weed Out Crime from Politics. ThePrint.in. https://theprint.in/india/governance/supreme-court-asks-parliament-to-frame-law-to-weed-out-crime-from-politics/124004/. Accessed May 2025.

Jokela, M., J. Meriläinen, J. Tukiainen, and A. von Schoultz (2025). Personality Traits and Cognitive Ability in Political Selection. *Journal of the European Economic Association* (5), 1908–1949.

Kahana, M. J. (2012). *Foundations of Human Memory*. New York, NY: Oxford University Press.

Kapur, D. and M. Vaishnav (Eds.) (2018). *Costs of Democracy: Political Finance in India*. Oxford University Press.

Kendall, C., T. Nannicini, and F. Trebbi (2015). How Do Voters Respond to Information? Evidence from a Randomized Campaign. *American Economic Review 105*(1), 322–353.

Khemka, A. (2025). Manipulating the System: Clientelism and Criminality in Politics. Working Paper.

Kling, J. R., J. B. Liebman, and L. F. Katz (2007). Experimental Analysis of Neighborhood Effects. *Econometrica 75*(1), 83–119.

Kőszegi, B. and A. Szeidl (2013). A Model of Focusing in Economic Choice. *Quarterly Journal of Economics 128*(1), 53–104.

Kwon, S. and J. Roth (2026). Testing Mechanisms. *Review of Economic Studies*, rdag028.

Le Pennec, C. and V. Pons (2023). How do Campaigns Shape Vote Choice? Multicountry Evidence from 62 Elections and 56 TV Debates. *Quarterly Journal of Economics 138*(2), 703–767.

Lehne, J., J. N. Shapiro, and O. Vanden Eynde (2018). Building Connections: Political Corruption and Road Construction in India. *Journal of Development Economics 131*, 62–78.

Lichand, G., N. Cunha, R. A. Madeira, and E. Bettinger (2026). When the Effects of Informational Interventions Are Driven by Salience–Evidence from School Parents in Brazil. *American Economic Journal: Economic Policy*. Forthcoming.

Lindgren, K.-O. L., P. Nyman, and S. Oskarsson (2025). Examining Voting Spillover Effects of Text Message Reminders. *British Journal of Political Science 55*(e21), 1–9.

Link, S., A. Peichl, C. Roth, and J. Wohlfart (2025). Attention to the Macroeconomy. Working Paper.

Lodhi, A. (2024, January). Why is India's Ram Temple in Ayodhya Controversial? *Al Jazeera*. https://www.aljazeera.com/news/2024/1/22/why-is-indias-ram-temple-in-ayodhya-controversial.

Loewenstein, G. and Z. Wojtowicz (2025). The Economics of Attention. *Journal of Economic Literature 63*(3), 1038–1089.

Magesan, A., A. Szabó, and G. Ujhelyi (2025). Do Parties and Voters Disagree? An Equilibrium Analysis of Candidate Selection in India. Working Paper.

Malhotra, N., M. R. Michelson, T. Rogers, and A. A. Valenzuela (2011). Text Messages as Mobilization Tools: The Conditional Effect of Habitual Voting and Election Salience. *American Politics Research 39*(4), 664–681.

Manzoni, E., E. Murard, S. Quercia, and S. Tonini (2026). Emotions, Beliefs, and Policy Views. *Journal of the European Economic Association*, jvag016.

Marshall, J. (2019). Signaling Sophistication: How Social Expectations Can Increase Political Information Acquisition. *The Journal of Politics 81*(1), 167–186.

Marshall, M. G. and T. R. Gurr (2020). Polity5: Political Regime Characteristics and Transitions, 1800-2018. Dataset and User's Manual.

Marx, B., V. Pons, and T. Suri (2021). Voter Mobilisation and Trust in Electoral Institutions: Evidence from Kenya. *Economic Journal 131*(638), 2585–2612.

Matějka, F. and G. Tabellini (2021). Electoral Competition with Rationally Inattentive Voters. *Journal of the European Economic Association 19*(3), 1899–1935.

Matějka, F. and A. McKay (2015). Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model. *American Economic Review 105*(1), 272–298.

Nunnari, S. and J. Zápal (2025). A Model of Focusing in Political Choice. *Journal of Politics 87*(4), 1465–1481.

Nyhan, B. (2020). Facts and Myths about Misperceptions. *Journal of Economic Perspectives 34*(3), 220–236.

Olken, B. A. (2015). Promises and Perils of Pre-Analysis Plans. *Journal of Economic Perspectives 29*(3), 61–80.

Olken, B. A. and R. Pande (2019). Governance Initiative Review Paper. Jameel Poverty Action Lab.

Pande, R. (2011). Can Informed Voters Enforce Better Governance? Experiments in Low-Income Democracies. *Annual Review of Economics 3*, 215–237.

Patel, D. (2025). Environmental Beliefs and Adaptation to Climate Change. Working Paper.

Persson, T. and G. Tabellini (2002). *Political Economics: Explaining Economic Policy*. MIT Press.

Poblete-Cazenave, R. (2025). Do Politicians in Power Receive Special Treatment in Courts? Evidence from India. *American Journal of Political Science 69*, 78–95.

Pons, V. (2018). Will a Five-minute Discussion Change Your Mind? A Countrywide Experiment on Voter Choice in France. *American Economic Review 108*(6), 1322–1363.

Prakash, N., M. Rockmore, and Y. Uppal (2019). Do Criminally Accused Politicians Affect Economic Outcomes? Evidence from India. *Journal of Development Economics 141*, 102370.

Prakash, N., S. Sahoo, D. Saraswat, and R. Sindhi (2024). When Criminality Begets Crime: The Role of Elected Politicians in India. *Journal of Law, Economics, and Organization, forthcoming*.

Press Trust of India (2015, July). UP Tops List of States with Highest Mobile Penetration. The Times of India. https://timesofindia.indiatimes.com/uttar-pradesh/up-tops-list-of-states-with-highest-mobile-penetration/articleshow/47926636.cms. Accessed June 2025.

Rivera, E., E. Seira, and S. Jha (2025). Apex Corruption Erodes Democratic Values. Working Paper.

Schechter, L. and S. Vasudevan (2023). Persuading Voters to Punish Corrupt Vote-Buying Candidates: Experimental Evidence from a Large-Scale Radio Campaign in India. *Journal of Development Economics 160*, 102976.

Serhan, Y. (2024, January). India Is Unveiling Its Controversial Temple of Ram. Here's What to Know. *Time*. https://time.com/6564070/india-modi-temple-ram/. Accessed November 2025.

Sims, C. A. (2003). Implications of Rational Inattention. *Journal of Monetary Economics 50*(3), 665–690.

Singh, P. (2025, March). 45 Patralok Media Network. https://www.patralok.com/45-mlas-have-declared-criminal-cases-against-themselves-adr/. Accessed December 2025.

Smits, J. and I. Permanyer (2019). The Subnational Human Development Database. *Scientific Data 6*, 190038.

Spenkuch, J. L. and D. Toniatti (2018). Political Advertising and Election Results. *Quarterly Journal of Economics 133*(4), 1981–2036.

Sukhtankar, S. and M. Vaishnav (2015). Corruption in India: Bridging Research Evidence and Policy Options. *India Policy Forum, National Council of Applied Economic Research 11*(1), 193–276.

Susewind, R. (2014). GIS Shapefiles for India's Parliamentary and Assembly Constituencies Including Polling Booth Localities. doi: 10.4119/unibi/2674065 (Accessed January 2015).

Thaler, M. (2024). The Fake News Effect: Experimentally Identifying Motivated Reasoning using Trust in News. *American Economic Journal: Microeconomics 16*(2), 1–38.

Vaishnav, M. (2017). *When Crime Pays: Money and Muscle in Indian Politics*. New Haven: Yale University Press.

Wager, S. and S. Athey (2018). Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests. *Journal of the American Statistical Association 113*(523), 1228–1242.

Waghmare, A. (2024, February). Access to Phones and the Internet. Data for India. https://www.dataforindia.com/comm-tech/. Accessed June 2025.

Figures and Tables

Figure 1: Effects of Information Treatment on Voting—2017 Experiment

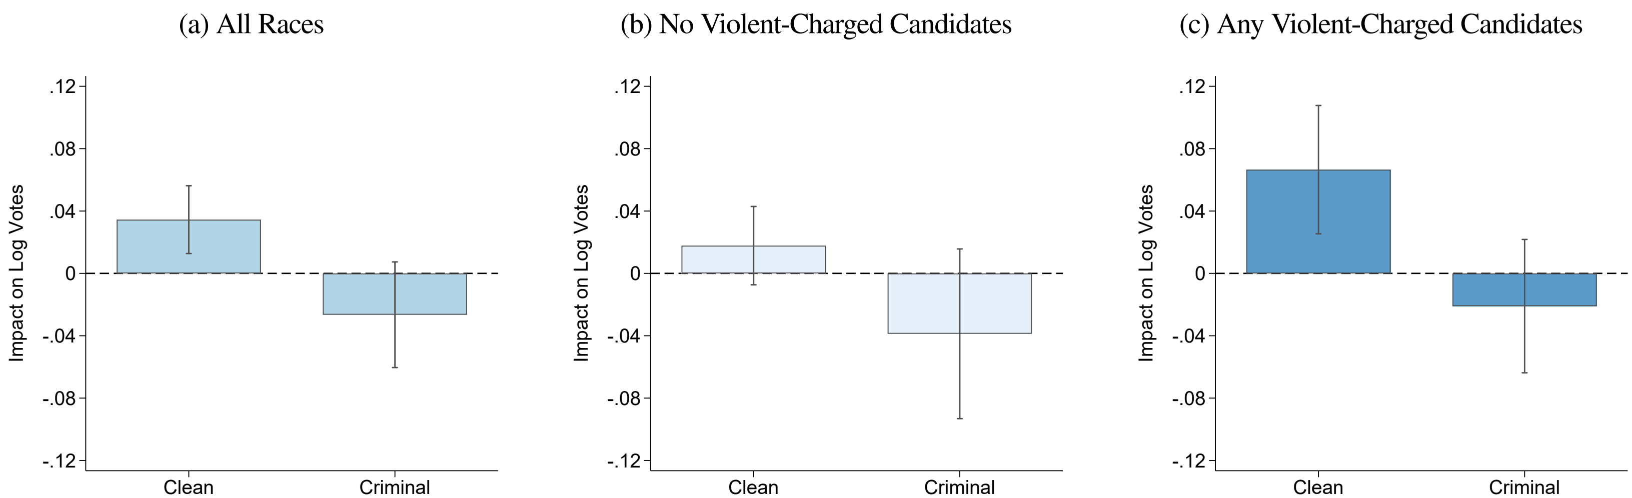

Notes: This figure presents results from estimating Equation (1). Each bar plot depicts the effect of the information treatment on polling-station-by-candidate-level log votes received by clean and criminally charged candidates. Panel A shows all races, while the sample in Panel B is restricted to races with no violent criminal candidates, and in Panel C to races with at least one violent criminal candidate. Standard errors in the underlying regressions are clustered at the village level, and error bars in each plot depict 90% confidence intervals. Results are discussed in Section 5.1.

Figure 2: Effects of Information Treatment on Voter Beliefs, Candidate Ratings, and Electoral Choices—2022 Experiment

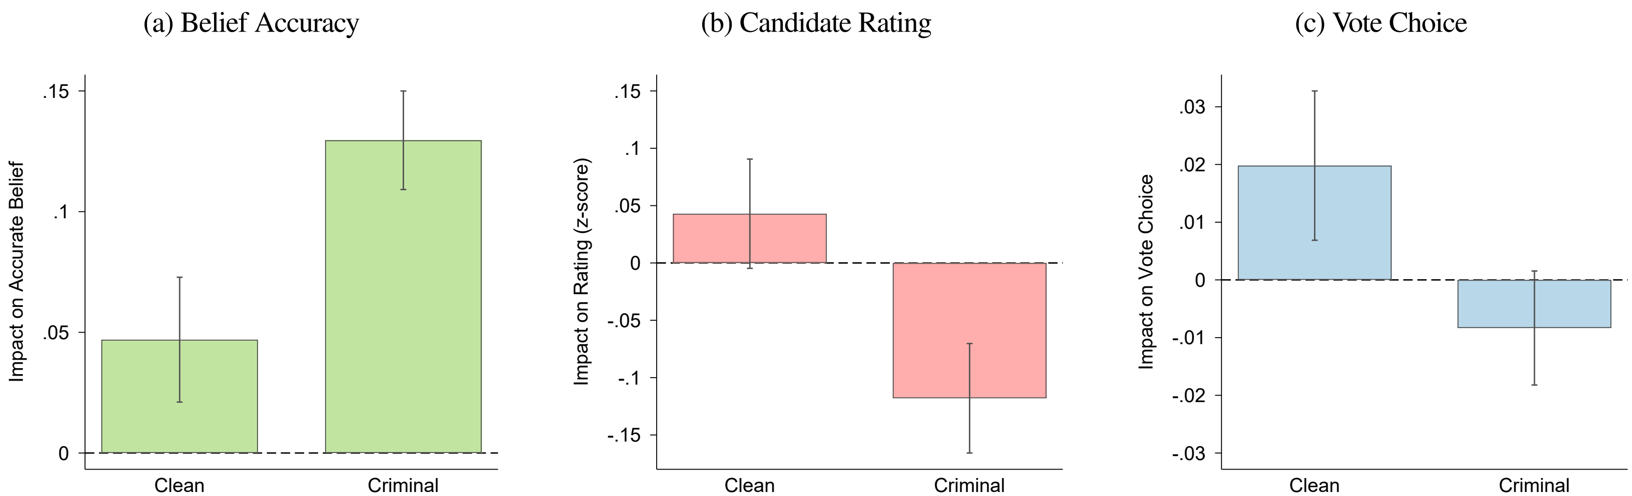

Notes: This figure presents results from estimating Equation (3). Each bar plot depicts the effect of the information treatment on a voter-by-candidate-level outcome for clean and criminally-charged candidates. The dependent variable in Panel A is an indicator for whether voter i had accurate beliefs about the criminal status of candidate c, while in Panel B it is voter i's rating of candidate c, converted from a 5-point Likert scale to a z-score. In Panel C, the dependent variable is an indicator for whether voter i stated that they voted for candidate c. Standard errors in the underlying regressions are clustered at the voter level, and error bars in each plot depict 90% confidence intervals. Results are discussed in Section 6.1.

Figure 3: Effects of Information Treatment on Salience and Decision Weights—2022 Experiment

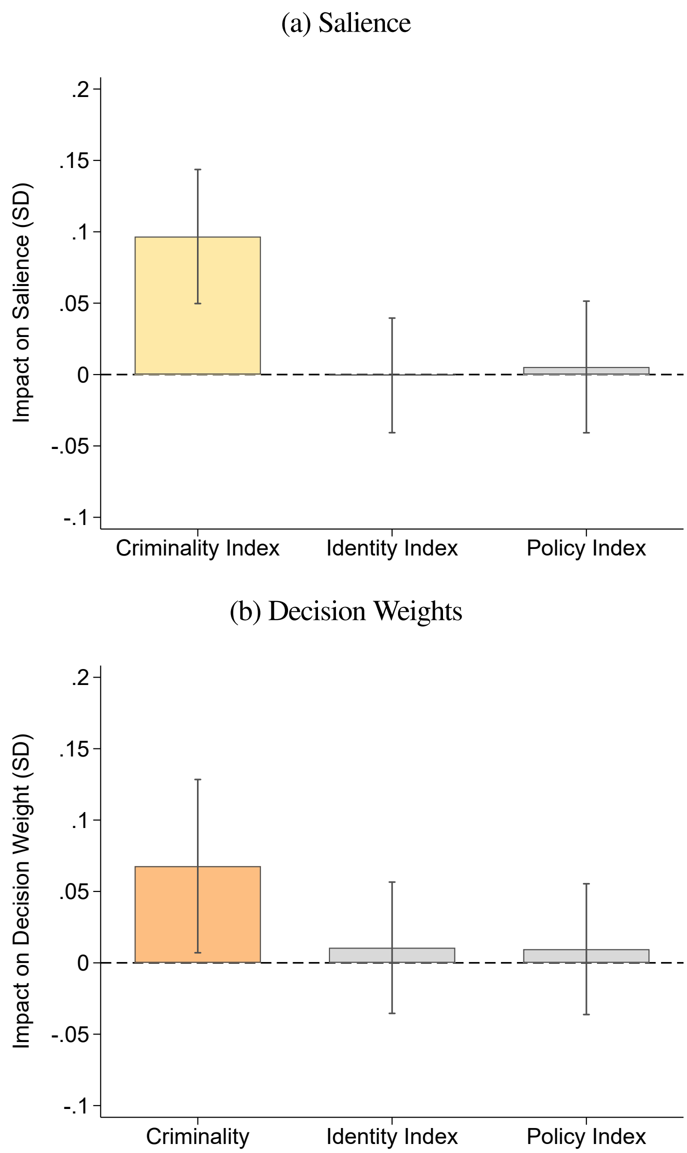

Notes: This figure presents the effects of the information treatment on salience (Panel A) and decision weights (Panel B) for voters, estimated using Equation (4). In Panel A, each outcome is a summary index constructed as the average of z-score-standardized binary indicators for whether specific attributes or issues are salient (top-of-mind) for a voter. The criminality index includes candidate criminal charges and the issues of crime and corruption. The identity index includes candidate religion and caste group, as well as party stance on the Ayodhya temple. The policy index includes farmer distress, employment, and COVID management. In Panel B, outcomes capture the importance voters place on attributes or issues when deciding their vote, based on 5-point Likert scale measures standardized to z-scores. Criminality corresponds to candidate criminal charges. The identity index is the average of candidate religion/caste group and party Ayodhya temple stance. The policy index is the average of farmer distress, employment, and COVID management. Error bars in each plot depict 90% confidence intervals. Results are discussed in Section 6.2. In addition, treatment effect estimates for each index component are available in Appendix Tables A12 (salience) and A14 (decision weights).

Figure 4: Associations of Conditional Average Treatment Effects on Voters' Beliefs, Salience, and Decision Weights—2022 Experiment

(a) Salience and Belief Updating

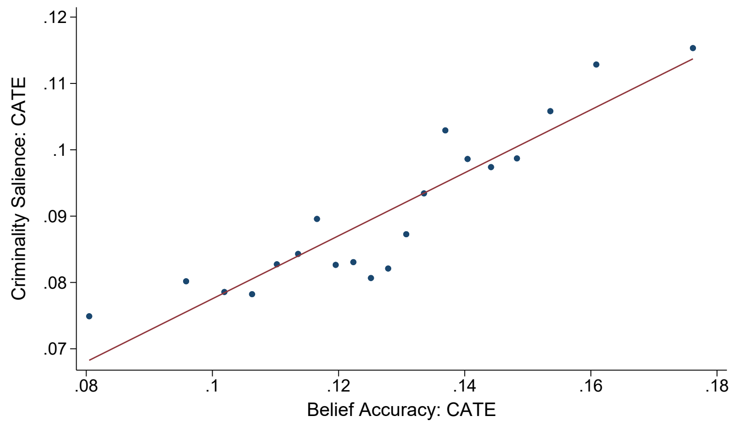

(b) Decision Weight and Salience

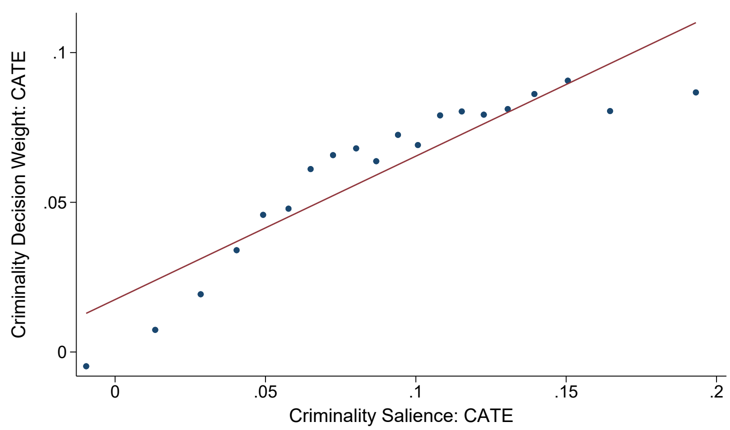

Notes: This figure presents binned scatterplots depicting the associations between predicted conditional average treatment effects (CATEs) on different outcomes. We estimate CATEs for each outcome using the causal forest procedure of Wager and Athey (2018), as described in Appendix B.5. Also plotted in each panel is the best linear fit line based on the underlying voter-level data. Panel A shows the association between each voter's belief accuracy CATE (i.e., the predicted effect of information on accuracy of beliefs about criminality) and their salience CATE (i.e., the predicted effect of information on whether criminality is top-of-mind). Panel B shows the association between each voter's salience CATE and their decision weight CATE (i.e., the predicted effect of information on the importance given to criminality when deciding how to vote). Results are discussed in Section 7.2.

Figure 5: Voter Responses to Information about In-Group and Out-Group Candidates—2022 experiment

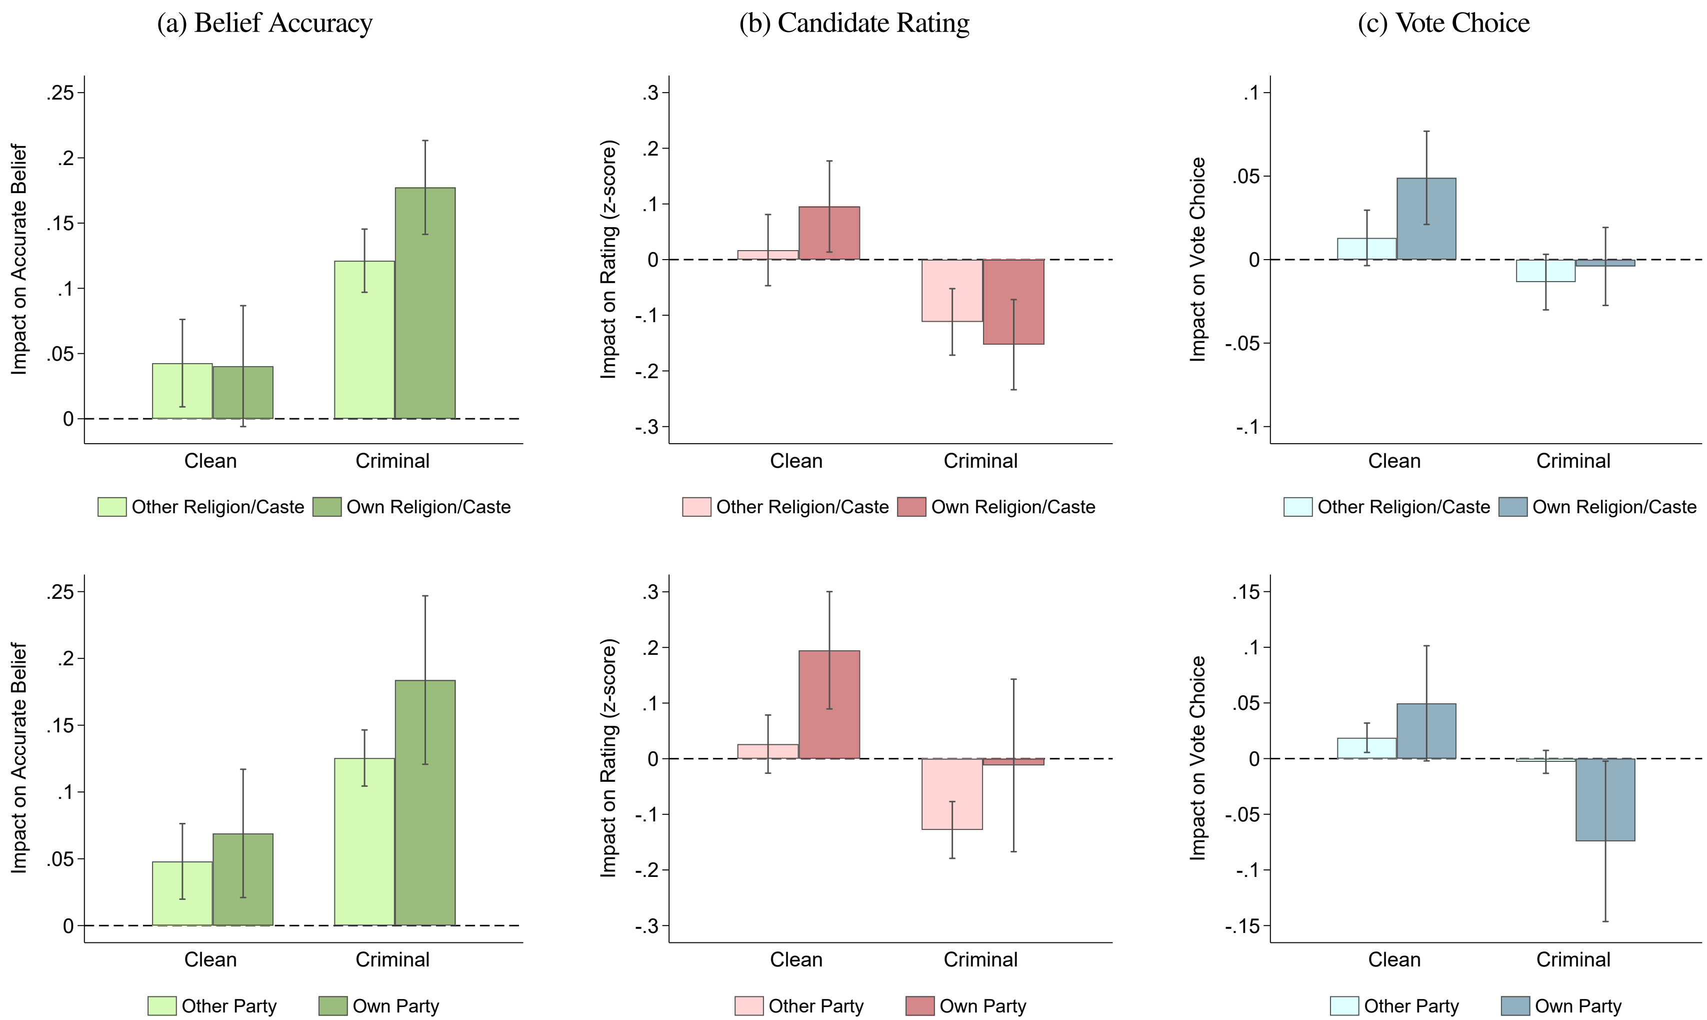

Notes: This figure presents results from estimating Equation (5). Each bar plot depicts the effect of the information treatment on voter beliefs, candidate ratings, and electoral choices for clean and criminally-charged candidates, as in Figure 2, but separately considering effects on in- and out-group candidates. Group identity is defined in terms of either religion/caste group (top row) or partisan alignment (bottom row). In Panel A, the outcome is an indicator for whether voter i had accurate beliefs about the criminal status of candidate c, while in Panel B it is voter i's rating of candidate c, converted from a 5-point Likert scale to a z-score. In Panel C, the outcome is an indicator for whether voter i stated that they voted for candidate c. Standard errors in the underlying regressions are clustered at the voter level, and error bars in each plot depict 90% confidence intervals. Results are discussed in Section 7.3.

Table 1: Prevalence of Criminal Candidates

	2017		2022	
	Mean	Obs.	Mean	Obs.
	(1)	(2)	(3)	(4)
Panel A: Candidate Level				
Any criminal charge	0.345	116	0.517	89
Any violent charge	0.155	116	0.348	89
Any murder-related charge	0.086	116	0.112	89
Criminal cases	0.759	116	1.753	89
Panel B: Constituency Level				
Any candidate with criminal charge	0.789	38	0.963	23
Any candidate with violent charge	0.447	38	0.826	23
Any candidate with murder-related charge	0.263	38	0.348	23
Any clean candidate	0.974	38	0.957	23

Notes: Panel A shows criminal charge and case statistics for major-party candidates in the 2017 and 2022 experimental sample constituencies. Panel B shows corresponding constituency-level statistics. Results are discussed in Section 4.3.

Table 2: Electoral Impacts of Information Treatment—2017 Experiment

	Log votes				
	(1)	(2)	(3)	(4)	(5)
Information	0.034***	0.024**	0.034***	0.018	0.065***
	(0.013)	(0.011)	(0.013)	(0.015)	(0.025)
Information					
× Criminal	-0.061**			-0.056	-0.086**
	(0.028)			(0.049)	(0.042)
× Number of cases		-0.013*			
		(0.007)			
× Nonviolent crime			-0.061*		
			(0.035)		
× Violent, non-murder			-0.015		
			(0.043)		
× Attempted murder			-0.064		
			(0.062)		
× Murder			-0.137**		
			(0.069)		
Candidate composition:					
No violent				X	
Any violent					X
Control mean	4.788	4.788	4.788	4.786	4.791
Observations	13,615	13,615	13,615	7,304	6,311

Notes: This table presents results on the electoral impacts of the information treatment during the 2017 elections. *Information* is an indicator for treated villages and the outcome is log votes received at the candidate-by-polling-station level (i.e., by candidate c at polling station s). Columns (1), (4), and (5) report estimates from Equation (1). Columns (2) and (3) report estimates based on a modified version of Equation (1) that allows the impacts of information to vary with a candidate's criminal case count (column 2) or most severe charge type (in column 3). Column (4) restricts the sample to races with no violent criminal candidates, and column (5) restricts the sample to races with at least one violent criminal candidate. Standard errors clustered at the village level are in parentheses. Results are discussed in Section 5.1. Significant at *10 percent, **5 percent,***1 percent.

Table 3: Implied Voter Persuasion Rates and Projected Election Impacts of At-Scale Information Provision

	Clean voting persuasion rate	2012		2017	
		Criminal to clean victory	Vote share margin range	Criminal to clean victory	Vote share margin range
	(1)	(2)	(3)	(4)	(5)
Treatment exposure:					
Receiving households	3.3%	32 (31.1%)	[0.001, 0.030]	9 (11.4%)	[0.001, 0.029]
Entire village	1.2%	11 (10.7%)	[0.001, 0.009]	3 (3.8%)	[0.001, 0.011]

Notes: Column (1) shows the estimated persuasion rate—i.e., the share of treated individuals who were persuaded to vote for a clean candidate because they received information about candidates' charges (DellaVigna and Gentzkow 2010). We consider two scenarios, which reflect two bounding assumptions about information diffusion. The "receiving households" scenario assumes that information remained confined to households who received the voice calls and text messages. The "entire village" scenario assumes that information spreads to all households in treated villages. Columns (2) through (5) present projected electoral impacts of an intervention providing information about candidates' criminal charges to voters at scale. Columns (2) and (4) show the number (share) of elections won by a criminal candidate over a clean runner-up that would have instead been won by a clean candidate in the 2012 and 2017 UP state assembly elections, respectively. Columns (3) and (5) give the range of the criminal winner vs clean runner-up vote share margins actually observed in those elections. Results are discussed in Section 5.2.

Table 4: Impacts of Information Treatment on Voter Beliefs, Candidate Evaluations, and Electoral Choices—2022 Experiment

	Accurate belief			Rating (z-score)			Vote choice		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Information	0.047***	0.065**	0.039**	0.043	-0.007	0.067**	0.020**	0.012	0.024**
	(0.016)	(0.030)	(0.018)	(0.029)	(0.056)	(0.033)	(0.008)	(0.013)	(0.010)
Information									
× Criminal	0.083***	0.126***		-0.161***	-0.111		-0.028**	-0.024	
	(0.019)	(0.044)		(0.039)	(0.093)		(0.012)	(0.033)	
× Nonviolent crime			0.112***			-0.064			-0.025
			(0.032)			(0.068)			(0.030)
× Violent crime			0.078***			-0.213***			-0.033**
			(0.021)			(0.046)			(0.015)
Candidate composition:									
No violent		X			X			X	
Any violent			X			X			X
Control mean	0.320	0.404	0.300	0.000	0.068	-0.017	0.209	0.217	0.207
Observations	11,623	2,472	9,151	9,895	2,108	7,787	9,224	2,008	7,216

Notes: This table presents results on the effects of the information treatment on voter beliefs, candidate ratings, and electoral choices during the 2022 elections. In columns (1)–(3), the dependent variable is a voter-by-candidate-level indicator for whether voter i had accurate beliefs about the criminal status of candidate c, while in columns (4)–(6) it is voter i's rating of candidate c, converted from a 5-point Likert scale to a z-score. In columns (7)–(9), the dependent variable is an indicator for whether voter i stated that they voted for candidate c. Columns (1), (2), (4), (5), (7), and (8) report estimates based on Equation (3). Columns (3), (6), and (9) report estimates based on a modified version of Equation (3) that allows the impacts of information to vary with a candidate's most severe criminal charge type. The sample is restricted in columns (2), (5), and (8) to races with no violent criminal candidates, and in columns (3), (6), and (9) to races with at least one violent criminal candidate. Standard errors clustered at the voter level are in parentheses. Results are discussed in Section 6.1. Significant at *10 percent, **5 percent, ***1 percent.

Appendix A: Additional Figures and Tables

Figure A1: Map of 2017 and 2022 Experiment Sample Constituencies in Uttar Pradesh

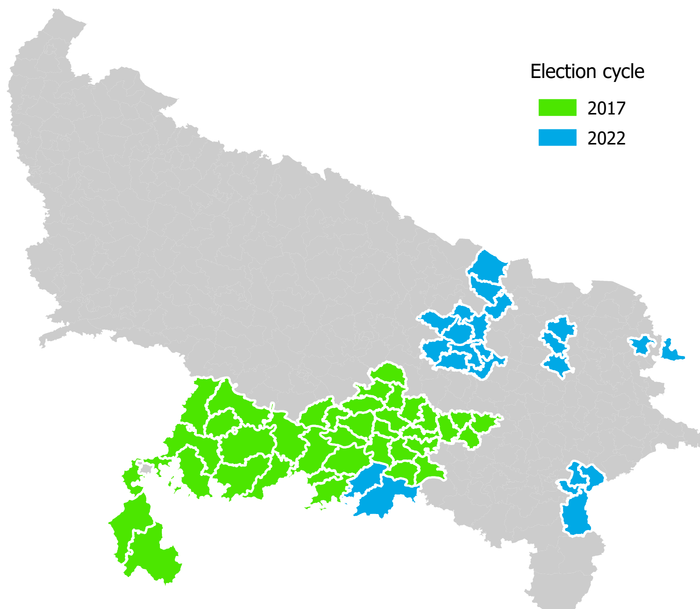

Notes: This map shows the constituencies in Uttar Pradesh included in the sample of the 2017 and 2022 election experiments. Discussed in Section 4.3.

Figure A2: Accuracy of Voter Beliefs about Candidate Criminality—2022 Experiment

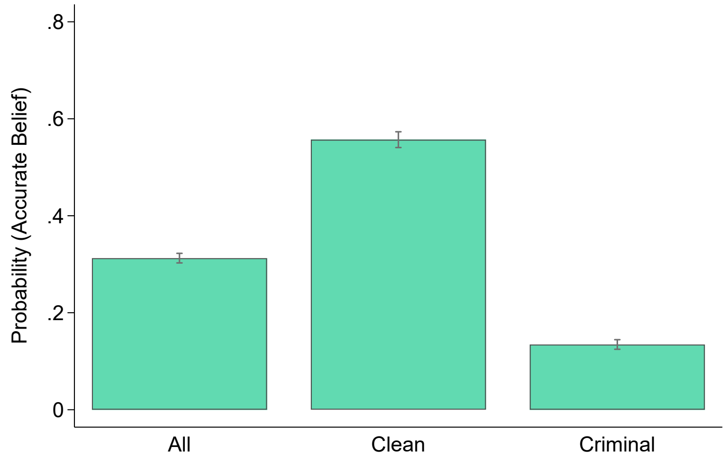

Notes: This figure shows the share of beliefs held by voters about the criminal status of candidates in their constituency that were accurate, for control group voters at endline. The bars present the accuracy of voters' beliefs about the criminal status of all candidates ("All"), non-criminal candidates ("Clean"), and criminal candidates ("Criminal"). Results are discussed in Section 6.1. Error bars depict 90% confidence intervals.

Figure A3: Impact Heterogeneity by Perceived Survey Focus—2022 Experiment

(a) Belief Accuracy (b) Candidate Rating (c) Vote Choice

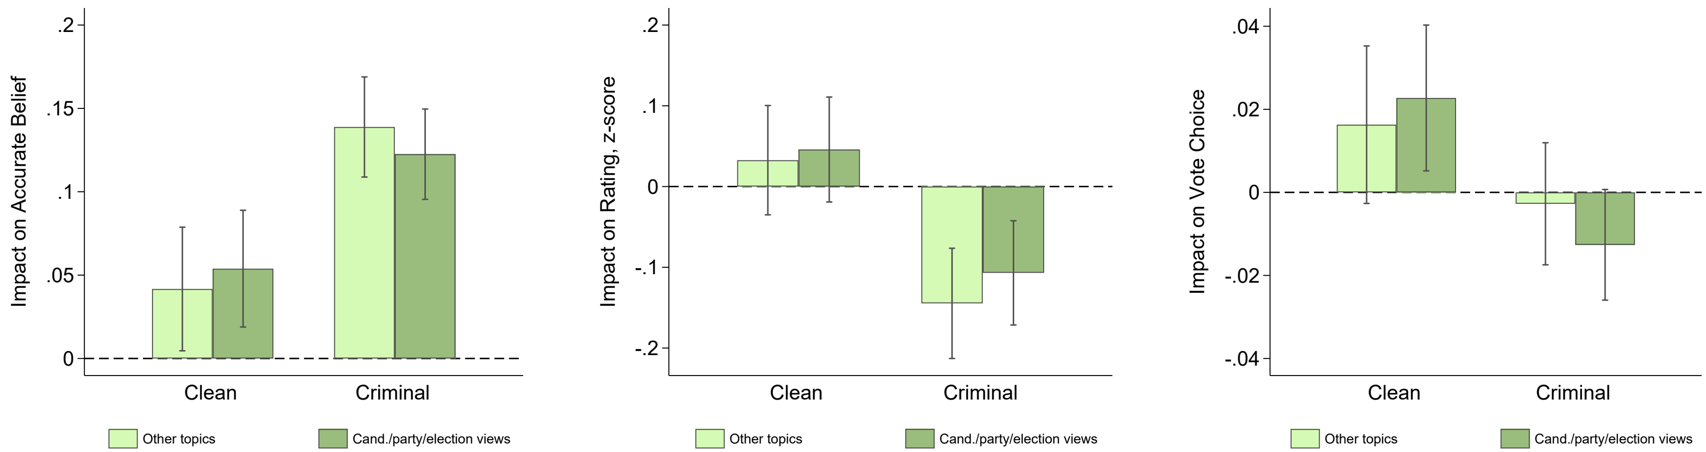

Notes: This figure presents results from estimating a version of Equation (5) that replaces the own-group indicator with a respondent-level indicator for perceiving the focus of the endline survey as related to views on candidates and parties in the election broadly. Each bar plot depicts the effect of the information treatment on voter beliefs, candidate ratings, and electoral choices for clean and criminally-charged candidates, as in Figure 2, but separately considering effects by the respondents' perceived focus of the endline survey. In Panel A, the outcome is an indicator for whether voter i had accurate beliefs about the criminal status of candidate c, while in Panel B it is voter i's rating of candidate c, converted from a 5-point Likert scale to a z-score. In Panel C, the outcome is an indicator for whether voter i stated that they voted for candidate c. Results are discussed in Section 6.1. Standard errors in the underlying regressions are clustered at the voter level, and error bars in each plot depict 90% confidence intervals.

Figure A4: Baseline Salience—2022 Experiment

(a) Candidate Characteristics

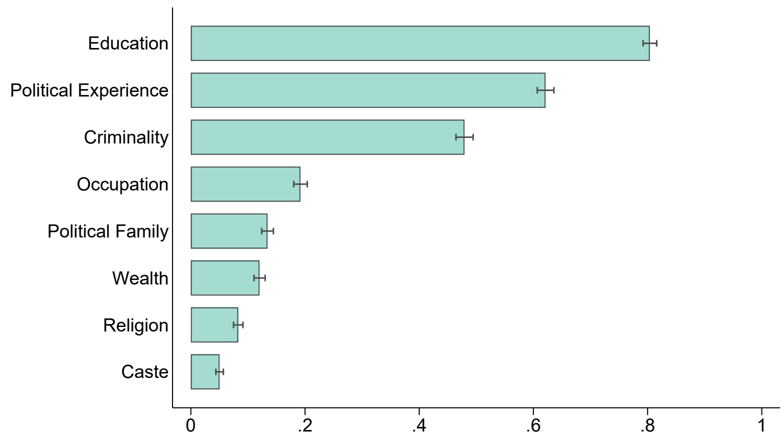

(b) Policy Issues

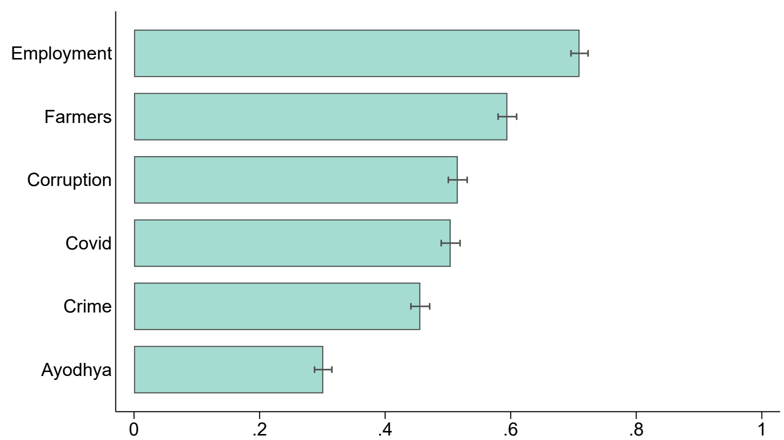

Notes: This figure presents descriptive statistics on voter attention using data from our baseline survey. Each bar shows the share of voters for whom a particular candidate attribute or policy issue was salient. Panel A presents the share of voters for whom each attribute was top of mind when prompted to think about the candidates in the election. Panel B presents the share of voters for whom each policy issue is top of mind when asked to think about the election. Section 6.2 describes our salience measure in more detail. Error bars in each plot depict 90% confidence intervals.

Figure A5: Test of Salience as Full Mediator of Treatment Effect on Decision Weights—2022 Experiment

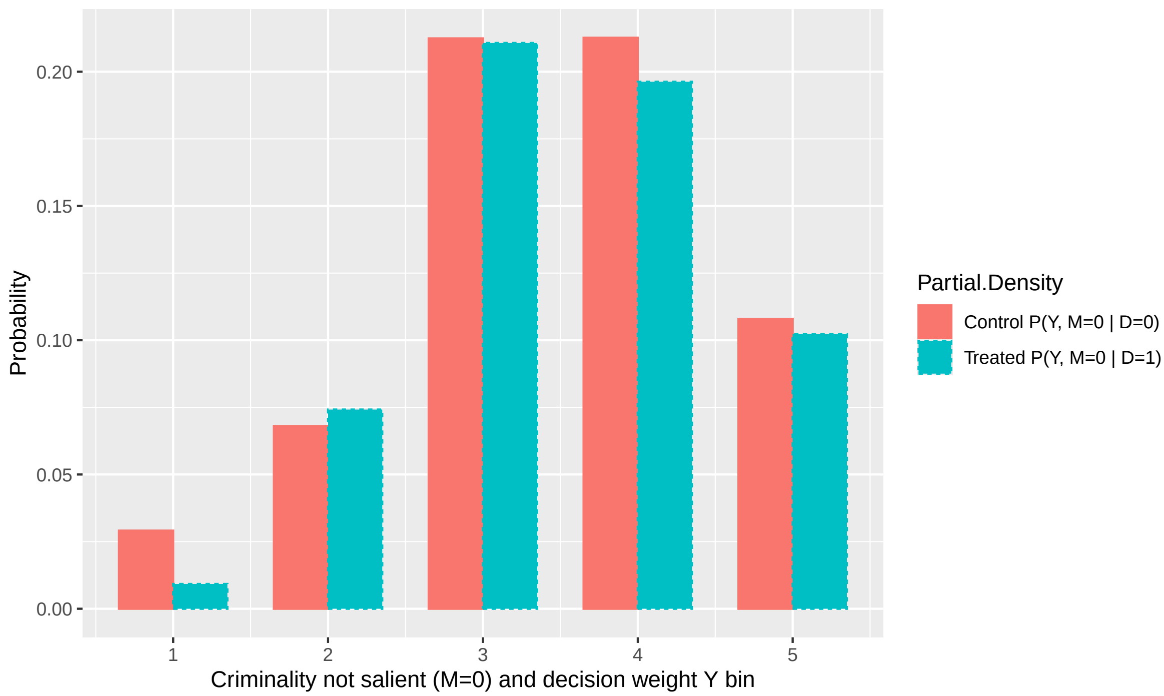

Notes: This figure presents results from the test for full mediation proposed by Kwon and Roth (2026), implemented using their *TestMechs* R package. The bars present the distribution of criminality decision weights for two groups of voters. The red bars depict the decision weights for control group voters for whom criminality was not salient, while the green bars present the decision weights for treated voters for whom criminality was not salient. Results are discussed in Section 7.2.

Figure A6: Histograms of Conditional Average Treatment Effects—2022 Experiment

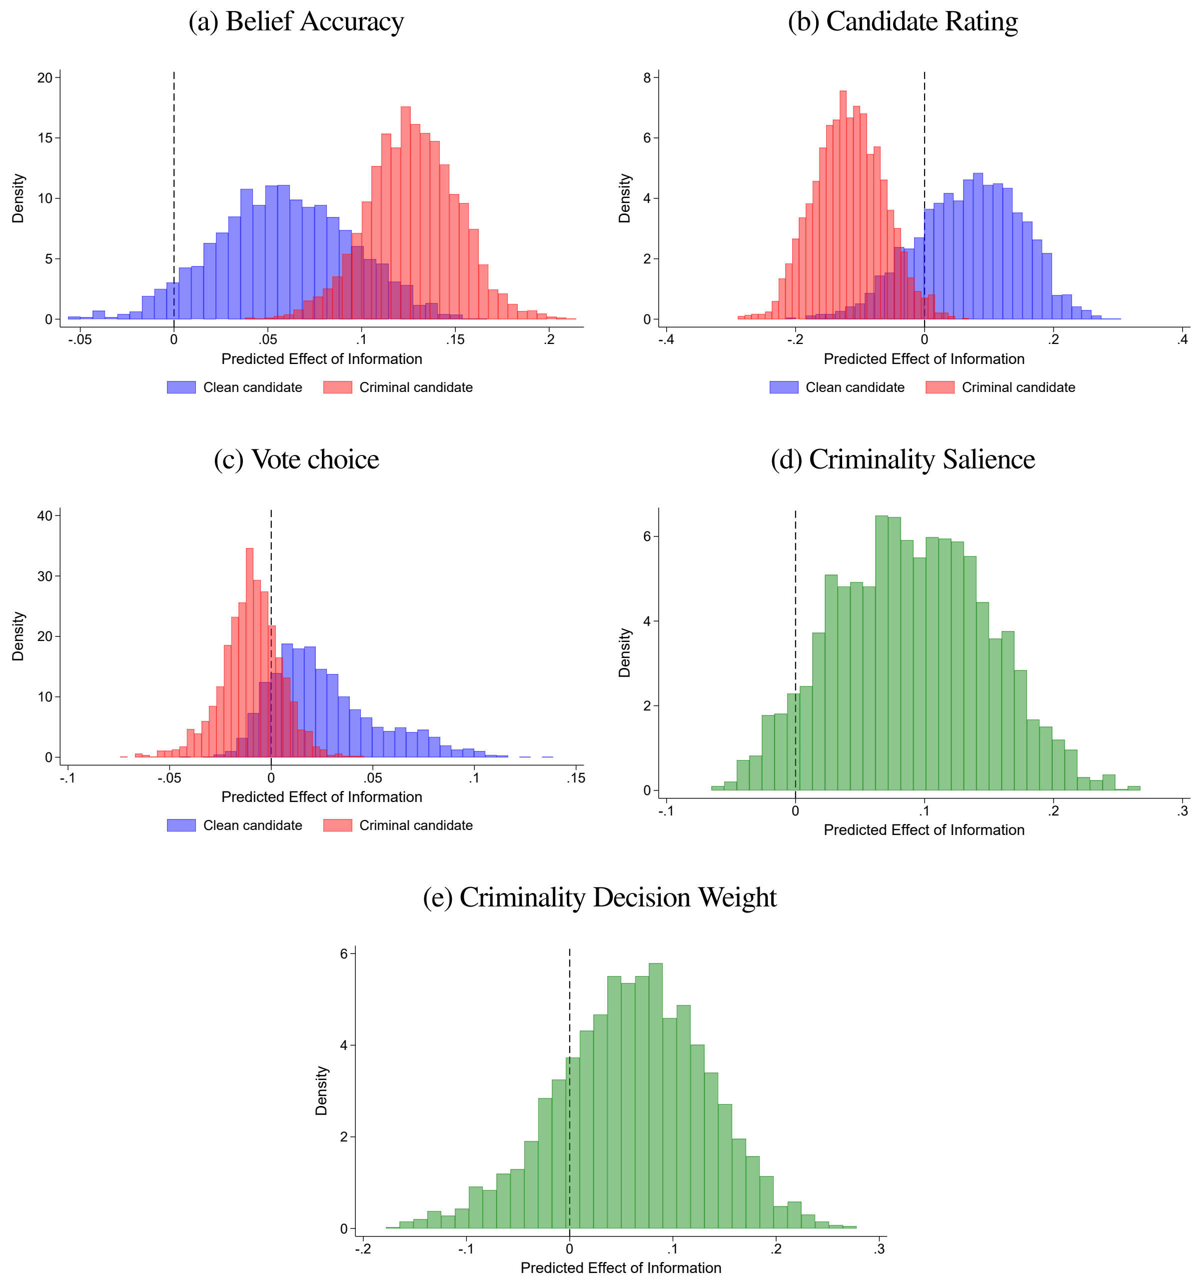

Notes: This figure plots histograms showing the distribution of predicted conditional average treatment effects (CATEs) of candidate criminality information, estimated using the causal forest procedure of Wager and Athey (2018) as described in Appendix B.5. Each panel presents the CATE distribution for a different outcome: belief accuracy (Panel A), candidate ratings (Panel B), vote choice (Panel C), salience of criminality (Panel D) and decision weight on criminality (Panel E). In Panels A–C, separate CATE distributions are shown for outcomes in relation to clean and criminal candidates. Results are discussed in Section 7.2.

Figure A7: Predictors of Decision Weight Treatment Effects—2022 Experiment

(a) OLS

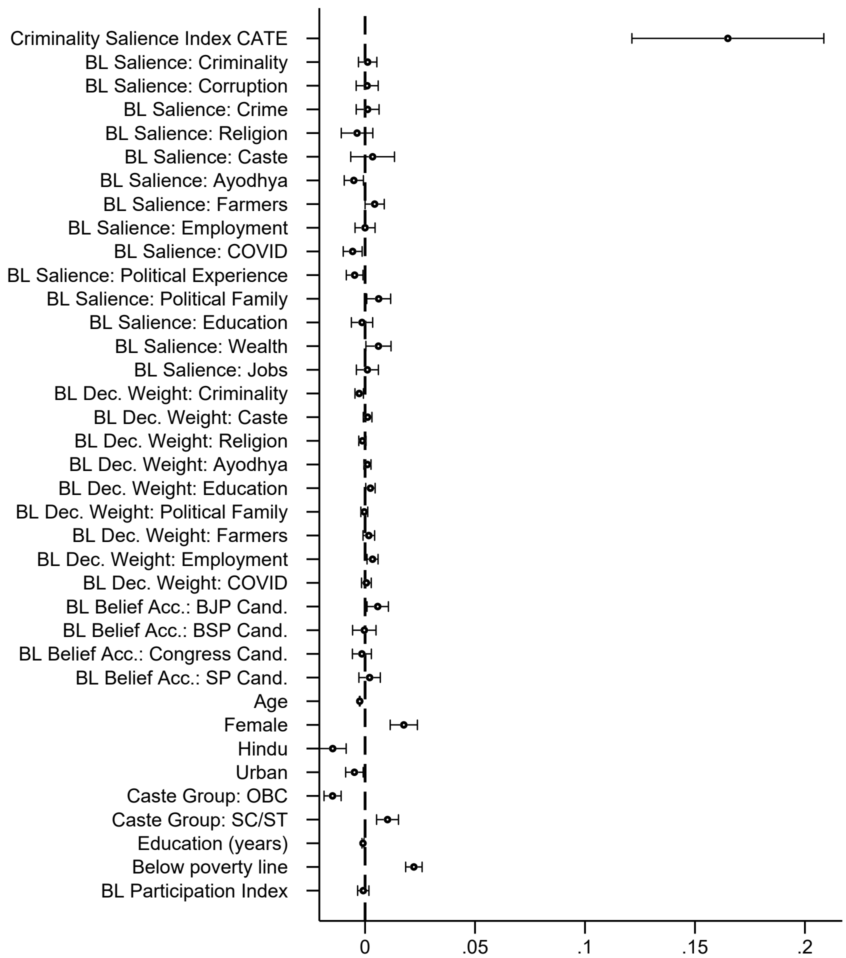

(b) LASSO

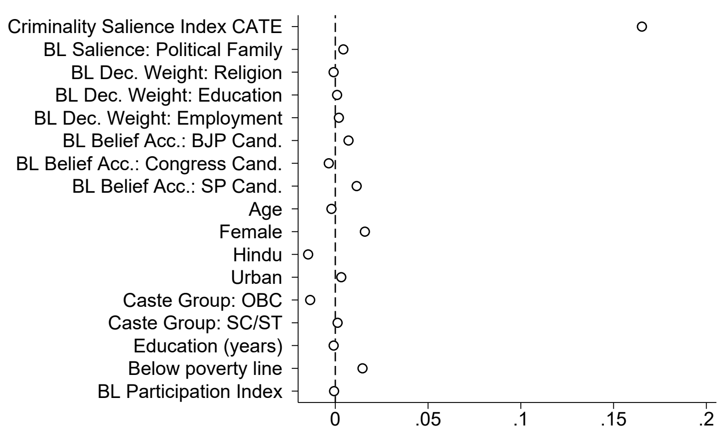

Notes: This figure shows the predictivity of different voter characteristics for the criminality decision weight CATE (conditional average treatment effect). Panel A reports coefficients from a voter-level OLS regression of criminality decision weight CATE on the listed voter characteristics as well as constituency fixed effects. Error bars depict 90% confidence intervals. Panel B reports coefficients for variables selected from the full set listed in Panel A by a LASSO model where the dependent variable is again the criminality decision weight CATE. Results are discussed in Section 7.2.

Figure A8: Conditional Average Treatment Effects by Voter Type—2022 Experiment

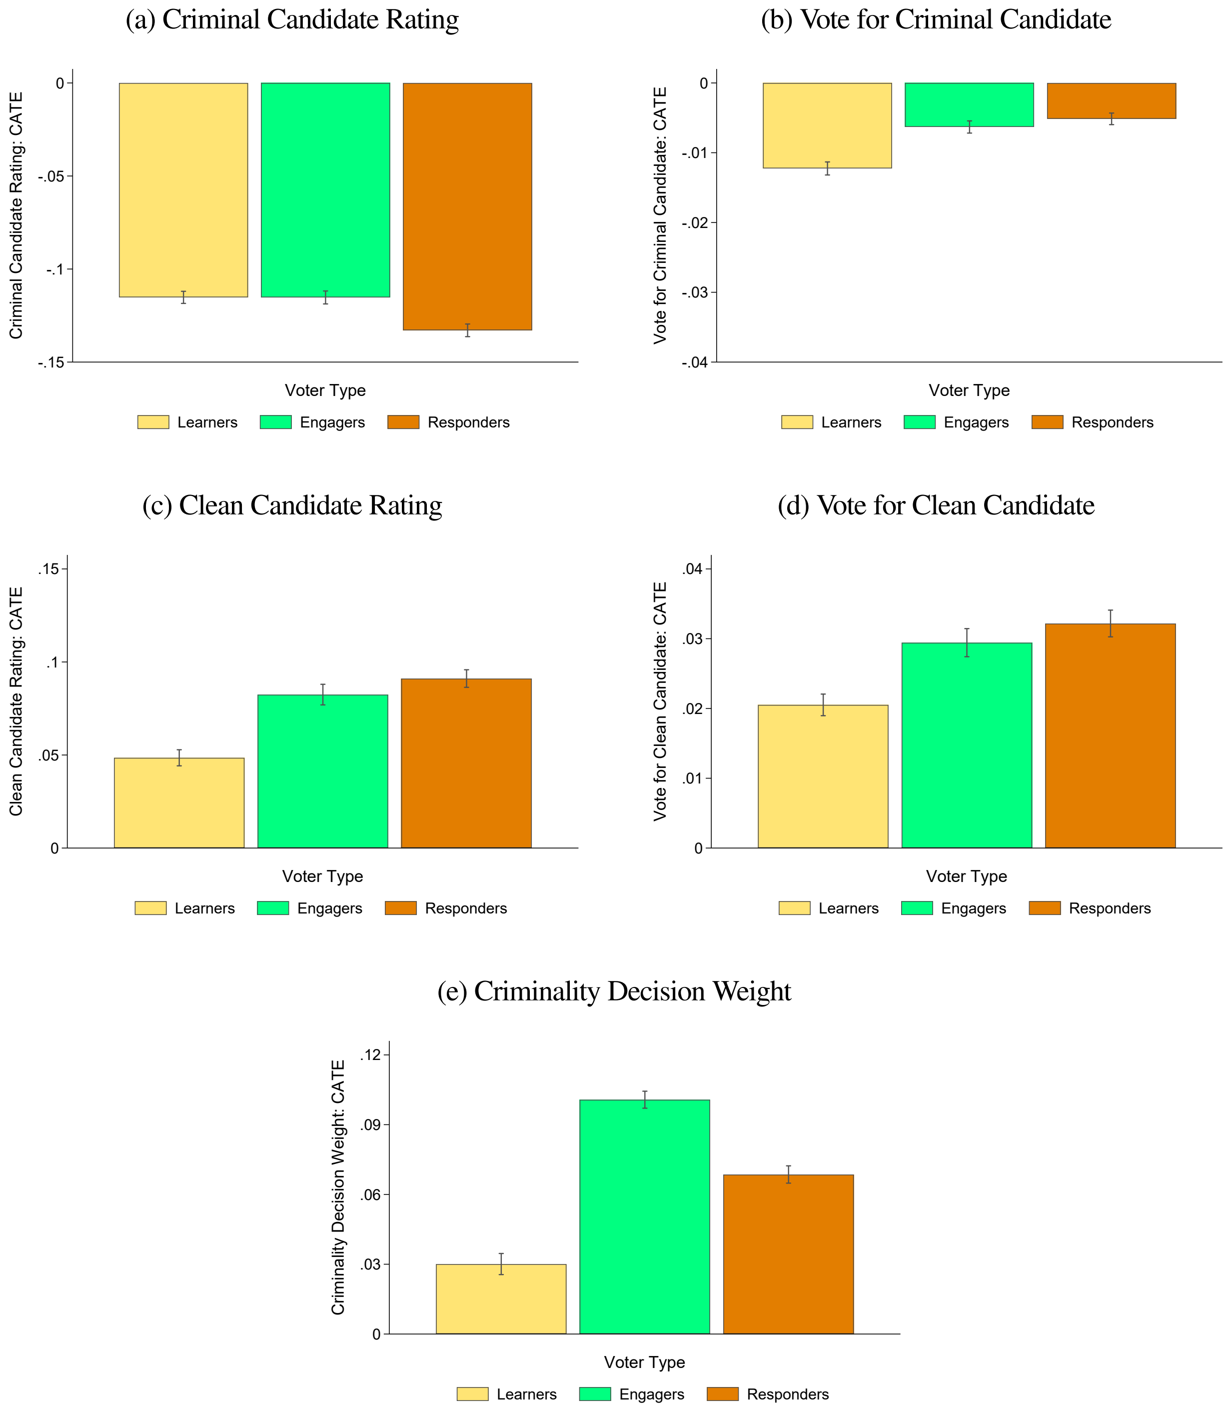

Notes: This figure presents predicted conditional average treatment effects (CATEs) by voter type. CATEs are estimated using the causal forest procedure of Wager and Athey (2018), as described in Appendix B.5. Each panel shows the average CATE for a particular outcome: ratings of criminal candidates (Panel A) and clean candidates (Panel C), measured as z-scores; intention to vote for a criminal candidate (Panel B) and clean candidate (Panel D), units of which are in percentage points; and criminality decision weight (Panel E), also a z-score. *Learners* are voters with statistically significant (10% level) positive CATEs on belief accuracy but insignificant CATEs on salience. *Engagers* are voters with statistically significant positive CATEs on salience but not on belief accuracy. *Responders* are voters who have statistically significant and positive CATEs on both belief accuracy and salience. Results are discussed in Section 7.2. Error bars in each panel depict 90% confidence intervals.

Table A1: Criminally Charged Elected Heads of Government and State

Country	Name	Position (Years in Office)
Argentina	Cristina Fernández de Kirchner	President (2007–2015)
Armenia	Robert Kocharyan	President (1998–2008)
Austria	Sebastian Kurz	Chancellor (2017–2019; 2020–2021)
Bangladesh	Khaleda Zia	Prime Minister (1991–1996; 2001–2006)
Bangladesh	Sheikh Hasina	Prime Minister (1996–2001; 2009–2024)
Bolivia	Jeanine Áñez	President (2019–2020)
Brazil	Luiz Inácio "Lula" da Silva	President (2003–2010; 2023–)
Brazil	Michel Temer	President (2016–2018)
Brazil	Jair Bolsonaro	President (2019–2023)
Bulgaria	Kiril Petkov	Prime Minister (2021–2022)
Colombia	Álvaro Uribe	President (2002–2010)
Costa Rica	Miguel Ángel Rodríguez	President (1998–2002)
Croatia	Ivo Sanader	Prime Minister (2003–2009)
Ecuador	Rafael Correa	President (2007–2017)
El Salvador	Antonio Saca	President (2004–2009)
Fiji	Laisenia Qarase	Prime Minister (2000–2001; 2001–2006)
Fiji	Frank Bainimarama	Prime Minister (2007–2022)
France	Jacques Chirac	President (1995–2007)
France	Nicolas Sarkozy	President (2007–2012)
Germany	Christian Wulff	President (2010–2012)
Guatemala	Alfonso Portillo	President (2000–2004)
Guatemala	Álvaro Colom	President (2008–2012)
Guatemala	Otto Pérez Molina	President (2012–2015)
Honduras	Juan Orlando Hernández	President (2014–2022)
Israel	Benjamin Netanyahu	Prime Minister (1996–1999; 2009–2021; 2022–)
Israel	Moshe Katsav	President (2000–2007)
Israel	Ehud Olmert	Prime Minister (2006–2009)
Italy	Silvio Berlusconi	Prime Minister (1994–1995; 2001–2006; 2008–11)
Lesotho	Tom Thabane	Prime Minister (2012–2015; 2017–2020)
Moldova	Vladimir Filat	Prime Minister (2009–2013)
Moldova	Igor Dodon	President (2016–2020)
Nicaragua	José Arnoldo Alemán Lacayo	President (1997–2002)
North Macedonia	Vlado Bučkovski	Prime Minister (2004–2006)
North Macedonia	Nikola Gruevski	Prime Minister (2006–2016)
Paraguay	Luis Ángel González Macchi	President (1999–2003)
Philippines	Joseph Estrada	President (1998–2001)
Philippines	Gloria Macapagal Arroyo	President (2001–2010)
Philippines	Rodrigo Duterte	President (2016–2022)
Portugal	José Sócrates	Prime Minister (2005–2011)
Romania	Adrian Năstase	Prime Minister (2000–2004)
South Africa	Jacob Zuma	President (2009–2018)
South Korea	Lee Myung-bak	President (2008–2013)
South Korea	Park Geun-hye	President (2013–2017)
South Korea	Yoon Suk Yeol	President (2022–2025)
Taiwan	Chen Shui-bian	President (2000–2008)
Thailand	Thaksin Shinawatra	Prime Minister (2001–2006)
Thailand	Yingluck Shinawatra	Prime Minister (2011–2014)
Trinidad and Tobago	Basdeo Panday	Prime Minister (1995–2001)
Ukraine	Viktor Yanukovych	Prime Minister (2002–2005; 2006–2007); President (2010–2014)
Ukraine	Yulia Tymoshenko	Prime Minister (2005; 2007–2010)
United States	Donald J. Trump	President (2017–2021; 2025–)

Note: This table lists heads of state and government in democratic countries (defined as having a Polity5 score for 2000 of 6 or above) since 2000 with criminal charges filed against them before, during, or following their time in office. List drawn from multiple media sources; additional details available upon request. Discussed in Section 1.

Table A2: Balance Check—2017 Experiment

	Control	Treatment difference	Obs.
	(1)	(2)	(3)
Panel A: Village Characteristics			
Log population	7.165	0.003	3,511
	[0.624]	(0.013)	
Population share:			
Female	0.478	-0.000	3,511
	[0.023]	(0.001)	
Literate	0.566	0.004*	3,511
	[0.087]	(0.002)	
SC/ST	0.284	0.007	3,511
	[0.157]	(0.005)	
Log partner mobile subscribers	5.400	0.004	3,511
	[0.638]	(0.013)	
School count:			
Primary	1.175	-0.037	3,511
	[0.767]	(0.024)	
Middle	0.668	-0.009	3,511
	[0.649]	(0.021)	
Secondary	0.116	-0.017	3,511
	[0.384]	(0.012)	
Agricultural household share	0.462	0.007	3,434
	[0.265]	(0.008)	
Log consumption per capita (INR)	9.632	0.003	3,434
	[0.192]	(0.006)	
Area (hectares)	391.7	7.218	3,511
	[348.9]	(9.389)	
All-weather road access	0.741	0.005	3,477
	[0.438]	(0.015)	
Power access:			
Domestic	0.947	0.006	3,481
	[0.223]	(0.008)	
Agriculture	0.758	-0.009	3,466
	[0.428]	(0.013)	
Distance (km) to nearest location:			
Above 10,000 population	9.787	-0.021	3,511
	[5.219]	(0.139)	
Above 50,000 population	27.417	0.420	3,511
	[12.574]	(0.263)	

Notes: This table presents balance tests for the 2017 election experiment, where randomly selected villages received the information treatment. For each pre-intervention characteristic, column (1) presents the control mean and standard deviation. Column (2) presents the coefficient and standard error from a regression of the characteristic on a village-level information treatment indicator and randomization strata. Column (3) gives the number of observations. In Panel A, observations are at village level and data for all variables come from the 2011 Census of India, with the exception of agricultural household share and log consumption per capita, which are from the Socioeconomic High-resolution Rural-Urban Geographic Platform for India (SHRUG). Standard errors in parentheses. Discussed in Section 4.4. Significant at *10 percent, **5 percent,***1 percent.

Table A2: Balance Check—2017 Experiment (continued)

	Control (1)	Treatment difference (2)	Obs. (3)
Panel B: Polling Station Characteristics			
Log registered voters	6.752	0.011	4,666
	[0.368]	(0.011)	
Log votes in 2014:			
Total	6.233	0.010	4,571
	[0.400]	(0.011)	
BJP	5.135	0.010	3,877
	[0.800]	(0.025)	
BSP	4.401	0.011	4,560
	[0.963]	(0.026)	
Congress	3.317	-0.001	4,539
	[1.409]	(0.028)	
SP	4.359	0.025	3,917
	[1.056]	(0.032)	
Log votes in 2012:			
Total	6.284	-0.001	4,591
	[0.368]	(0.011)	
BJP	3.648	-0.031	4,571
	[1.422]	(0.031)	
BSP	4.775	0.023	4,590
	[0.767]	(0.022)	
Congress	3.937	-0.004	4,576
	[1.254]	(0.030)	
SP	4.766	-0.003	4,461
	[0.883]	(0.026)	

Notes: This table presents balance tests for the 2017 election experiment, where randomly selected villages received the information treatment. For each pre-intervention characteristic, column (1) presents the control mean and standard deviation. Column (2) presents the coefficient and standard error from a regression of the characteristic on a village-level information treatment indicator and randomization strata. Column (3) gives the number of observations. In Panel B, observations are at the polling station level. Turnout and party voting outcomes are for the 2014 national parliamentary elections and 2012 state assembly elections, based on data from the Office of the Chief Electoral Officer, Uttar Pradesh. Observation counts can differ across parties in an election cycle due to their fielding candidates in different numbers of constituencies. Standard errors clustered at the village level are in parentheses. Discussed in Section 4.4. Significant at *10 percent, **5 percent,***1 percent.

Table A3: Impacts of Information Treatment on Turnout and Voting for Minor Parties—2017 Experiment

	Log total votes			Log total non-major-party votes		
	(1)	(2)	(3)	(4)	(5)	(6)
Information	0.006	0.007	0.005	-0.004	0.001	-0.008
	(0.005)	(0.006)	(0.007)	(0.018)	(0.025)	(0.026)
Candidate composition:						
No violent		X			X	
Any violent			X			X
Control mean	6.268	6.249	6.290	3.296	3.315	3.275
Observations	4,659	2,467	2,192	4,650	2,462	2,188

Notes: This table presents results from estimating Equation (2). Each observation is a polling-station-level electoral outcome. In columns (1)–(3), the outcome is log total votes cast, while in columns (4)–(6) it is log total votes for non-major-party candidates. In columns (2) and (5), the sample is restricted to races where a major-party candidate faces violent charges, while in columns (3) and (6) the sample is restricted to races where no major-party candidate faces violent charges. Standard errors clustered at the village level are in parentheses. Results are discussed in Section 5.1. Significant at *10 percent, **5 percent, ***1 percent.

Table A4: Electoral Impacts of Placebo Message—2017 Experiment

	Log votes	
	(1)	(2)
Placebo	0.014	-0.012
	(0.033)	(0.025)
Placebo		
$\times$ Criminal	-0.017	
	(0.055)	
$\times$ Number of cases		0.011
		(0.010)
Control mean	4.920	4.920
Observations	3,024	3,024

Notes: This table presents results from estimating a variant of Equation (1) where the village-level information treatment indicator is replaced by a placebo treatment indicator. The outcome is log votes received at the candidate-by-polling station level (i.e., by candidate c at polling station s). The experimental sample villages were located in constituencies that voted during the seventh phase of the 2017 UP elections. Standard errors clustered at the village level are in parentheses. Results are discussed in Section 5.1. Significant at *10 percent, **5 percent,***1 percent.

Table A5: Electoral Impacts of Information by Treatment Arm—2017 Experiment

	Log votes		
	(1)	(2)	(3)
Information only	0.037**	0.028*	0.037**
	(0.019)	(0.016)	(0.019)
Information and coordination	0.029	0.013	0.029
	(0.020)	(0.017)	(0.020)
Information and ethnic	0.038*	0.030*	0.038*
	(0.020)	(0.016)	(0.020)
Criminal			
× Information only	-0.060		
	(0.039)		
× Information and coordination	-0.074*		
	(0.040)		
× Information and ethnic	-0.050		
	(0.047)		
Number of cases			
× Information only		-0.015	
		(0.012)	
× Information and coordination		-0.012	
		(0.010)	
× Information and ethnic		-0.012	
		(0.009)	
Nonviolent crime			
× Information only			-0.057
			(0.048)
× Information and coordination			-0.077
			(0.053)
× Information and ethnic			-0.048
			(0.057)
Violent, non-murder			
× Information only			-0.005
			(0.059)
× Information and coordination			-0.004
			(0.070)
× Information and ethnic			-0.036
			(0.065)
Attempted murder			
× Information only			-0.097
			(0.103)
× Information and coordination			-0.151*
			(0.090)
× Information and ethnic			0.063
			(0.079)
Murder			
× Information only			-0.120
			(0.100)
× Information and coordination			-0.092
			(0.095)
× Information and ethnic			-0.203*
			(0.122)
Observations	13,615	13,615	13,615

Notes: This table presents results from estimating a modified version of Equation (1) where we replace the pooled information treatment indicator with separate indicators for the information only, information plus coordination, and information plus ethnic voting treatment arms. Message scripts for each treatment arm are provided in Appendix Section B.1. Columns (2) and (3) allow for information to have differential impacts by criminal case count or most severe charge type, respectively. Standard errors clustered at the village level are in parentheses. Results are discussed in Section 4.3. Significant at *10 percent, **5 percent,***1 percent.

Table A6: Balance Check—2022 Experiment

	Control (1)	Treatment difference (2)	Obs. (3)
Age	38.06	-0.901*	3,000
	[12.80]	(0.480)	
Female	0.075	0.004	3,000
	[0.263]	(0.010)	
Urban location	0.454	0.027	2,991
	[0.498]	(0.017)	
Below poverty line	0.328	-0.002	2,924
	[0.470]	(0.018)	
Hindu	0.913	-0.004	2,988
	[0.282]	(0.011)	
Caste group:			
General	0.461	0.018	2,925
	[0.499]	(0.019)	
OBC	0.415	-0.018	2,925
	[0.493]	(0.019)	
SC/ST	0.117	-0.005	2,925
	[0.321]	(0.012)	
Education level:			
Below secondary	0.229	-0.021	2,949
	[0.421]	(0.016)	
Secondary	0.186	-0.010	2,949
	[0.389]	(0.015)	
College	0.296	0.014	2,949
	[0.457]	(0.018)	
Post-graduate	0.262	0.022	2,949
	[0.440]	(0.017)	
Criminality belief accuracy	0.326	0.004	3,000
	[0.290]	(0.011)	
Political participation index	0.000	0.014	3,000
	[0.659]	(0.025)	
Party rating:			
BJP	7.447	-0.013	2,838
	[2.782]	(0.109)	
BSP	3.799	-0.127	2,736
	[2.488]	(0.096)	
Congress	3.343	-0.075	2,661
	[2.437]	(0.096)	
SP	4.948	-0.041	2,747
	[2.714]	(0.108)	

Notes: This table presents balance checks for the 2022 experiment, where the information treatment was randomized at the individual level. For each voter characteristic, column (1) presents the control mean and standard deviation. Column (2) presents the coefficient and standard error from a regression of the characteristic on an individual-level treatment indicator and randomization strata. Column (3) gives the number of observations. Criminality belief accuracy is the share of major-party candidates correctly identified as clean or criminal. The political participation index is constructed as the average z-scores of binary indicators for having voted in an election before, anyone in the household having ever contested in panchayat elections, and having attended the last Gram Sabha in the panchayat. The criminality salience index is constructed as described in Section 6.2. The party rating variables are on a 10-point scale from "very bad" (one) to "very good" (ten). Standard errors are in parentheses. Discussed in Section 4.4. Significant at *10 percent, **5 percent, ***1 percent.

Table A7: Impacts of Information Treatment on Turnout and Votes for Minor Parties—2022 Experiment

	Turnout			Vote for non-major party		
	(1)	(2)	(3)	(4)	(5)	(6)
Information	0.008	0.014	0.007	-0.003	-0.009	-0.001
	(0.013)	(0.028)	(0.014)	(0.005)	(0.012)	(0.005)
Candidate composition:						
No violent		X			X	
Any violent			X			X
Control mean	0.889	0.886	0.890	0.013	0.019	0.011
Observations	2,382	502	1,880	2,382	502	1,880

Notes: This table presents results on the effects of our 2022 information treatment on voter turnout and votes for minor parties. In columns (1)–(3), the dependent variable, turnout, is an indicator for whether voter i has voted. In columns (4)–(6), the dependent variable is an indictor for whether a voter i voted for a non-major party (i.e., not BJP, BSP, Congress, or SP). All regressions include constituency fixed effects. Columns (2) and (5) restrict the sample to races where no major-party candidate faces violent charges, while columns (3) and (6) restrict the sample to races where any major-party candidate faces violent charges. Standard errors are in parentheses. Results are discussed in Section 6.1. Significant at *10 percent, **5 percent, ***1 percent.

Table A8: Effects of Information Treatment on Turnout, by Baseline Support—2022 Experiment

	Turnout	
	(1)	(2)
Information	0.008	0.003
	(0.013)	(0.017)
Information		
× Criminal-party supporter		-0.012
		(0.032)
× Clean-party supporter		0.033
		(0.031)
Control mean	0.889	0.889
Observations	2,382	2,382

Notes: This table presents results on how the turnout effects of our 2022 information treatment vary with voters' baseline support for different parties. The outcome variable, turnout, is an indicator for whether voter i has voted. Column (2) additionally includes indicators for whether voter i supports a party that fielded a clean or criminal candidate and their interactions with the information treatment. All regressions include constituency fixed effects. Standard errors are in parentheses. Results are discussed in Section 6.1. Significant at *10 percent, **5 percent, ***1 percent.

Table A9: Information Impacts on Voter Learning and Candidate Evaluation by Number of Criminal Cases—2022 Experiment

	Accurate belief (1)	Rating (z-score) (2)	Vote choice (3)
Information	0.068***	0.026	0.011**
	(0.012)	(0.025)	(0.006)
Information × Number of cases	0.014***	-0.041***	-0.004*
	(0.003)	(0.008)	(0.002)
Control mean	0.320	0.000	0.209
Observations	11,623	9,895	9,224

Notes: This table presents results on the effects of the information treatment on voter beliefs, candidate ratings, and electoral choices during the 2022 elections. In column (1), the dependent variable is a voter-by-candidate-level indicator for whether voter i had accurate beliefs about the criminal status of candidate c, while in column (2) it is voter i's rating of candidate c, converted from a 5-point Likert scale to a z-score. In column (3), the dependent variable is an indicator for whether voter i stated that they voted for candidate c. All columns report estimates based on a modified version of Equation (3) that allows the impacts of information to vary with a candidate's criminal case count (top coded at the 99th percentile). Standard errors clustered at the voter level are in parentheses. Results are discussed in Section 6.1. Significant at *10 percent, **5 percent, ***1 percent.

Table A10: Relationship between Structured and Open-Ended Salience Measures—2022 Experiment

	Open-ended	Obs.	Outcome mean [SD]
Panel A: Criminality			
Candidate criminal record	0.284***	3,000	0.479
	(0.050)		[0.500]
Crime	0.177*	3,000	0.456
	(0.093)		[0.498]
Corruption	0.194	3,000	0.516
	(0.143)		[0.500]
Index	0.111***	3,000	0.000
	(0.024)		[0.820]
Panel B: Identity			
Candidate religion	0.133	3,000	0.083
	(0.106)		[0.276]
Candidate caste	0.010	3,000	0.050
	(0.018)		[0.219]
Ayodhya temple	-0.100	3,000	0.301
	(0.154)		[0.459]
Index	0.073**	3,000	0.000
	(0.029)		[0.681]
Panel C: Policy			
Farmers distress	0.131	3,000	0.595
	(0.115)		[0.491]
Employment situation	0.115***	3,000	0.710
	(0.022)		[0.454]
COVID management	0.380***	3,000	0.504
	(0.133)		[0.500]
Index	0.091***	3,000	0.000
	(0.023)		[0.793]
Panel D: Other Attributes			
Candidate political experience	0.204***	3,000	0.621
	(0.030)		[0.485]
Candidate dynastic political family	0.134	3,000	0.134
	(0.100)		[0.341]
Candidate education	0.156***	3,000	0.804
	(0.017)		[0.397]
Candidate wealth	0.056	3,000	0.120
	(0.074)		[0.325]
Candidate occupation	0.013	3,000	0.192
	(0.028)		[0.394]
Index	0.074***	3,000	0.000
	(0.022)		[0.667]

Notes: This table shows the association between the "structured" and "open-ended" salience measures in our 2022 experiment. Open-ended salience measures comes from analyzing open-ended text responses to the questions: "*When you think of the election, what are some of the thoughts that come to mind?*" and *"Now think about the candidates who will stand in your constituency in the upcoming election. What are some thoughts that come to your mind?"* Using a keyword dictionary, we generate indicators for mention of each candidate attribute and policy issue. Structured salience measures are based on responses to the structured question that immediately follows the open-ended question ("*When you think about the election, do any of the following thoughts come to your mind?*"). We generate indicators for responding "Yes" to each attribute/issue. We also construct four indices by averaging the standardized component indicators. Each row reports results from an individual-level regression of the baseline structured-question salience measure on the corresponding open-ended salience measure. All regressions include constituency fixed effects. Standard errors are in parentheses. Results are discussed in Section 6.2. Significant at *10 percent, **5 percent, ***1 percent.

Table A11: Validation of Salience Measure – Association between Salience and Recall—2022 Experiment

	Candidate attribute recall:							
	Criminality		Education		Religion		Wealth	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Corresponding attribute salience	0.209***	0.212***	0.054***	0.043***	0.136***	0.141***	0.112***	0.118***
	(0.014)	(0.014)	(0.013)	(0.014)	(0.030)	(0.031)	(0.024)	(0.025)
Individual controls		X		X		X		X
Omitted group mean	0.163	0.171	0.117	0.124	0.180	0.193	0.153	0.162
Observations	3,000	2,794	3,000	2,794	3,000	2,794	3,000	2,794

Notes: This table presents results on the relationship between voter attention to a candidate attribute and recall of that attribute from a vignette embedded in the baseline survey of the 2022 experiment. Early in the survey, voters were read a profile of a politician who contested the 2017 election in another region of UP. The profile included information about the candidate's education, wealth, and criminality, among other characteristics, with religious/ethnic identity also clearly inferrable from the candidate's name. At the end of the baseline survey, roughly 25 minutes later, respondents were asked to recall what they remembered from this profile. Full vignette text is provided in Appendix Section B.4. Each column presents the coefficient from a voter-level regression where the dependent variable is an indicator for voter recall of that specific attribute from the vignette. The independent variable is our baseline salience measure, and is an indicator for the voter mentioning that the attribute came to mind when prompted to think about the candidates in the election. Each regression includes constituency fixed effects. Columns (2), (4), (6), and (8) add individual-level controls for age, gender, religion (Hindu), urban residence, caste group (OBC, SC/ST), education (primary school, middle school, high school, pre-university, diploma, college, postgraduate, other), and economic status (below poverty line). Standard errors are in parentheses. Results are discussed in Section 6.2. Significant at *10 percent, **5 percent, ***1 percent.

Table A12: Effects of Information Treatment on Salience—2022 Experiment

	Information (1)	Obs. (2)	Control mean [SD] (3)
Panel A: Criminality			
Candidate criminal record	0.036**	3,000	0.377
	(0.018)		[0.485]
Crime	0.057***	3,000	0.420
	(0.019)		[0.494]
Corruption	0.049***	3,000	0.502
	(0.019)		[0.500]
Index	0.097***	3,000	0.000
	(0.029)		[0.778]
Panel B: Identity			
Candidate religion	-0.002	3,000	0.063
	(0.009)		[0.243]
Candidate caste group	-0.009	3,000	0.040
	(0.007)		[0.196]
Ayodhya temple	0.025	3,000	0.275
	(0.017)		[0.447]
Index	-0.001	3,000	0.000
	(0.024)		[0.654]
Panel C: Policy			
Farmers' distress	0.003	3,000	0.571
	(0.019)		[0.495]
Employment situation	0.025	3,000	0.698
	(0.017)		[0.459]
COVID management	-0.023	3,000	0.465
	(0.019)		[0.499]
Index	0.005	3,000	0.000
	(0.028)		[0.849]

Notes: This table presents estimates from Equation (4) that show the impact of our information treatment on salience. Each row contains the coefficient from a separate regression where the outcome variable is the salience of a particular candidate attribute or policy issue, i.e, an indicator for whether a voter mentioned that the attribute/issue came to mind when thinking about the election or the candidates contesting. For each family of attributes, we also consider as an outcome an index constructed by standardizing each component variable to a z-score, and then taking the average of standardized components. Standard errors are in parentheses. Results are discussed in Section 6.2. Significant at *10 percent, **5 percent, ***1 percent.

Table A13: Relationship between Salience and Decision Weights—2022 Experiment

	Decision weight:					
	Criminality		Identity index		Policy index	
	(1)	(2)	(3)	(4)	(5)	(6)
Corresponding salience index	0.342***	0.342***	0.261***	0.264***	0.090***	0.085***
	(0.022)	(0.023)	(0.021)	(0.022)	(0.019)	(0.020)
Individual controls		X		X		X
Control mean	0.000	-0.012	-0.002	0.010	-0.002	0.013
Observations	2,969	2,767	2,989	2,785	2,991	2,787

Notes: This tables presents results on the association between salience and decision weights using data from the baseline survey of the 2022 experiment. Each column presents the results from a voter-level regression where the dependent variable is the voter's decision weight on a particular attribute and the independent variable is the salience of that attribute. Decision weights are based on responses to the question, *When you are deciding who to vote for in the upcoming election, how much importance will you give to each of the following factors?* The responses are on a 5-point Likert scale from "Not at all important" (one) to "Most important" (five), which we standardize to a z-score. As with the salience indices, we construct decision weight indices by taking the average of the standardized component variables. Each regression includes constituency fixed effects. Columns (2), (4), and (6) add individual-level controls for age, gender, religion (Hindu), urban residence, caste group (OBC, SC/ST), education (primary school, middle school, high school, pre-university, diploma, college, postgraduate, other), and economic status (below poverty line). Standard errors are in parentheses. Results are discussed in Section 6.2. Significant at *10 percent, **5 percent, ***1 percent.

Table A14: Effects of Information Treatment on Decision Weights—2022 Experiment

	Information (1)	Obs. (2)	Control mean [SD] (3)
Panel A: Criminality			
Candidate criminal record	0.070*	2,939	3.760
	(0.038)		[1.036]
Panel B: Identity			
Candidate religion/caste group	0.039	2,978	1.824
	(0.038)		[0.995]
Ayodhya temple	-0.023	2,931	3.391
	(0.043)		[1.149]
Index	0.011	2,988	-0.005
	(0.028)		[0.763]
Panel C: Policy			
Farmers' distress	-0.008	2,984	4.317
	(0.029)		[0.782]
Employment situation	0.060**	2,984	4.310
	(0.027)		[0.762]
COVID management	-0.038	2,971	3.891
	(0.034)		[0.916]
Index	0.010	2,990	-0.003
	(0.028)		[0.772]

Notes: This table presents estimates of the effects of our 2022 information treatment on voters' decision weights on different attributes. Each row presents estimates from a separate regression based on a version of Equation (4) where the dependent variable is the decision weight for the listed attribute/issue. Decision weights are based on responses to the question, "*When you are deciding who to vote for in the upcoming election, how much importance will you give to each of the following factors?*" The responses are on a 5-point Likert scale from "Not at all important" (one) to "Most important" (five), which we standardize to a z-score. For each family of outcomes, we consider as an outcome a decision weight index constructed by taking the average of the standardized component variables. Standard errors are in parentheses. Results are discussed in Section 6.2. Significant at *10 percent, **5 percent, ***1 percent.

Table A15: Heterogeneity in Salience Effects of Information – Role of Surprise and Contrast—2022 Experiment

	Criminality salience index				
	(1)	(2)	(3)	(4)	(5)
Information					
× Above-median belief accuracy *(Low surprise)*	-0.110**	-0.110*			-0.106*
	(0.055)	(0.058)			(0.058)
× No violent-charged candidate *(Low contrast)*			-0.102	-0.124*	-0.119*
			(0.069)	(0.071)	(0.071)
Information	0.150***	-0.046	0.117***	-0.051	-0.018
	(0.040)	(0.214)	(0.032)	(0.214)	(0.214)
Treatment-interacted controls		X		X	X
Control mean	0.000	0.002	0.000	0.002	0.000
Observations	3,000	2,794	3,000	2,794	3,000

Notes: This table presents results showing how the effects of information on salience vary with measures of surprise and contrast. Each column contains estimates from a modified version of Equation (4) in which we interact the information treatment indicator with proxies for low surprise and/or low contrast. The outcome variable in all regressions is the criminality salience index, which is constructed by standardizing and averaging the salience indicators for candidate criminality, crime, and corruption. Columns (1), (2), and (5) interact the treatment indicator with a proxy for low surprise at the voter level—*Above-median belief accuracy*, which is constructed by taking the average accuracy of voters' beliefs about politician criminality from the baseline survey. The low surprise proxy is directly included as well. Columns (3), (4), and (5) interact the treatment dummy with a proxy for low contrast at the constituency level—*No violent-charged candidate*, an indicator for whether the election featured no violent criminal candidates. Columns (2), (4), and (5) further include interactions of the treatment indicator with individual-level controls for age, gender, religion (Hindu), urban residence, caste group (OBC, SC/ST), education (primary school, middle school, high school, pre-university, diploma, college, postgraduate, other), and economic status (below poverty line). Standard errors are in parentheses. Results are discussed in Section 7.1. Significant at *10 percent, **5 percent, ***1 percent.

Table A16: Surprise and Contrast Summary Statistics—2022 Experiment

	Mean	Obs.
	(1)	(2)
Surprise: Baseline average belief accuracy	0.337	3,000
	(0.300)	
Contrast: Election has any violent-charged candidate	0.794	3,000
	(0.404)	

Notes: This table presents voter-level summary statistics on our measures of surprise and contrast. *Surprise* is based on the average accuracy of each voter's baseline beliefs across all major-party candidates in the previous election cycle. *Contrast* captures whether there are large differences in criminality among candidates in the election. It is a constituency-level indicator variable that takes value one when at least one candidate contesting has a violent criminal charge. Column (1) presents the mean and standard deviation (in parentheses) for each variable. Column (2) reports the number of observations. Discussed in Section 7.1.

Table A17: Association of Salience with Surprise and Contrast among Control Group Voters—2022 Experiment

	Criminality salience index					
	(1)	(2)	(3)	(4)	(5)	(6)
Above-median belief accuracy	0.050	0.060			0.053	0.062
	(0.040)	(0.042)			(0.045)	(0.048)
No violent-charged candidate			0.030	0.058	0.040	0.065
			(0.050)	(0.051)	(0.074)	(0.076)
Above-median belief accuracy						
× No violent-charged candidate					-0.024	-0.018
					(0.101)	(0.103)
Individual controls		X		X		X
Control mean	0.000	0.002	0.000	0.002	0.000	0.002
Observations	1,481	1,368	1,481	1,368	1,481	1,368

Notes: This table presents results on the relationship between surprise, contrast, and salience among control group voters at endline in our 2022 experiment. Each column reports estimates from a regression of the criminality salience index on proxies for either surprise or contrast, restricting the sample to control group voters (who did not receive information). The criminality salience index is constructed using the salience indicators for candidate criminality, crime, and corruption. The index is the average of the standardized component variables. In columns (1) and (2), the key regressor is a voter-level proxy for surprise—*Above-median belief accuracy*, which is constructed using the average accuracy of voters' beliefs about politicians' criminality from the baseline survey. In columns (3) and (4), the key regressor is a constituency-level proxy for contrast, *No violent-charged candidate*, an indicator for the election featuring no violent criminal candidates. Columns (5) and (6) additionally include the interaction between the proxies for surprise and contrast. Columns (2), (4), and (6) further include individual-level controls for age, gender, religion (Hindu), urban residence, caste group (OBC, SC/ST), education (primary school, middle school, high school, pre-university, diploma, college, postgraduate, other), and economic status (below poverty line). All columns include constituency fixed effects. Standard errors are in parentheses. Results are discussed in Section 7.1. Significant at *10 percent, **5 percent, ***1 percent.

Table A18: Heterogeneity in Salience Effects of Information – Role of Decision Relevance—2022 Experiment

	Criminality salience index			
	(1)	(2)	(3)	(4)
Information				
× Close race (≤ 3pp)	-0.003			
	(0.093)			
× Margin of victory (pp)		0.007		
		(0.005)		
× Party rating gap (below-median)			-0.005	
			(0.056)	
× Party rating gap				0.006
				(0.010)
Information	0.097***	0.041	0.095**	0.071
	(0.030)	(0.046)	(0.042)	(0.046)
Control mean	0.000	0.000	0.006	0.000
Observations	3,000	3,000	2,932	2,932

Notes: This table presents results showing how the effects of information on salience vary with the decision relevance of the information to voters. Each column contains estimates from a modified version of Equation (4) in which we interact the information treatment indicator with a proxy of decision relevance. In columns (1) and (2), the decision relevance proxy is the competitiveness of the election contest. Column (1) includes an interaction of the treatment dummy with a constituency-level indicator for "close" elections, i.e., those with a margin of victory of ≤ 3 percentage points. Column (2) interacts the treatment dummy with a continuous version of the same proxy, the margin of victory in the constituency (i.e., the vote share gap between the top two candidates). In columns (3) and (4), the proxy of decision relevance is the voter's baseline indecision between parties. In column (3), the treatment dummy is interacted with a voter-level indicator for having a below-median ratings gap between parties at baseline. Specifically, we compute the gap in ratings (which are on a 10-point scale) between voters' two highest rated parties at baseline. In column (4), the treatment indicator is interacted with the ratings gap itself. In columns (3) and (4), the individual-level decision relevance measure is included directly as well. Standard errors are in parentheses. Results are discussed in Section 7.1. Significant at *10 percent, **5 percent, ***1 percent.

Table A19: Responsiveness of Salience to Information about Electorally Irrelevant Candidates—2022 Experiment

	Criminality salience index	
	(1)	(2)
Information		
× Irrelevant-party belief accuracy (*Low surprise*)	-0.153**	
	(0.071)	
× No violent-charged irrelevant candidate (*Low contrast*)		-0.087
		(0.059)
Information	0.152***	0.165***
	(0.043)	(0.045)
Control mean	0.000	0.000
Observations	3,000	3,000

Notes: This table presents evidence that information about electorally irrelevant candidates affects salience. Each column contains estimates from a version of Equation (4) that includes interactions of the treatment indicator with proxies for reduced surprise and contrast in relation to candidates from irrelevant parties. We classify a party as irrelevant for a voter if it received $<10\%$ total vote share in the voter's constituency. Column (1) interacts the treatment dummy with a proxy for low voter-level surprise about irrelevant-party candidates —accuracy of baseline candidate criminality beliefs, restricted to parties irrelevant in the 2022 election cycle. Also included is an interaction of the treatment indicator with the corresponding measure for relevant parties. Both average belief accuracy measures are directly included as well. Column (2) interacts the treatment indicator with a proxy for reduced constituency-level contrast due to irrelevant candidates—an indicator for the election featuring no violent-charged irrelevant candidates. Also included is an interaction of the treatment dummy with a corresponding indicator for relevant candidates. Standard errors are in parentheses. Results are discussed in Section 7.1. Significant at *10 percent, **5 percent, ***1 percent.

Table A20: Effects of Information on Decision Weights – Mediating Role of Salience—2022 experiment

	Criminality decision weight					
	(1)	(2)	(3)	(4)	(5)	(6)
Information	0.068*	0.078**	0.078**	0.076**	0.041	0.054
	(0.037)	(0.038)	(0.038)	(0.038)	(0.035)	(0.036)
Belief accuracy, baseline			-0.046	-0.054		-0.051
			(0.066)	(0.067)		(0.063)
Criminality salience, baseline				-0.049		-0.042
				(0.038)		(0.036)
Criminality salience, endline					0.709***	0.712***
					(0.033)	(0.035)
Individual controls		X	X	X		X
Control mean	0.000	-0.005	-0.005	-0.005	0.000	-0.005
Observations	2,939	2,737	2,737	2,737	2,939	2,737

Notes: This table reports estimates from a version of Equation (4) where the dependent variable is criminality decision weight at endline. Columns vary the individual-level controls included. Column (1) has no controls. Columns (2), (3), (4), and (6) include voter-level demographic controls: age, gender, religion (Hindu), urban residence, caste group (OBC, SC/ST), education (primary school, middle school, high school, pre-university, diploma, college, postgraduate, other), and economic status (below poverty line). A control for baseline accuracy of beliefs about politician criminality is also included in columns (3), (4), and (6), and columns (4) and (6) further include a control for baseline salience of candidate criminality. Columns (5) and (6) also include a control for endline salience of candidate criminality. Standard errors are in parentheses. Results are discussed in Section 7.2. Significant at *10 percent, **5 percent, ***1 percent.

Table A21: Voter Responses to Information about In- and Out-Group Candidates—2022 Experiment

	Accurate belief (1)	Rating (z-score) (2)	Vote choice (3)
Panel A: Shared Religion/Caste Group			
Information	0.043**	0.017	0.013
	(0.020)	(0.039)	(0.010)
Information			
× Shared group	-0.002	0.078	0.036*
	(0.034)	(0.062)	(0.020)
× Criminal	0.079***	-0.129**	-0.026
	(0.024)	(0.051)	(0.017)
× Criminal × Shared group	0.058	-0.119	-0.027
	(0.041)	(0.086)	(0.031)
Control mean	0.338	0.017	0.227
Observations	9,847	8,470	7,818
Panel B: Partisan Alignment			
Information	0.048***	0.026	0.019**
	(0.017)	(0.032)	(0.008)
Information			
× Preferred party	0.021	0.169**	0.031
	(0.032)	(0.070)	(0.033)
× Criminal	0.077***	-0.154***	-0.022*
	(0.020)	(0.041)	(0.012)
× Criminal × Preferred party	0.037	-0.052	-0.102*
	(0.051)	(0.120)	(0.055)
Control mean	0.321	0.000	0.209
Observations	11,283	9,596	8,958

Notes: This table reports estimates from Equation (5) showing how information about the criminality of in-group and out-group candidates affects voter beliefs, candidate ratings, and electoral choices. In Panel A, group identity is based on religion/caste group. *Shared group* is an indicator for voter i sharing the same religion/caste group (religion for non-Hindus; caste group within Hindus) as candidate c. In Panel B, group identity is based on partisan alignment. *Preferred party* is an indicator for candidate c belonging to voter i's preferred party as determined at baseline. Standard errors clustered at the respondent level are in parentheses. Results are discussed in Section 7.3. Significant at *10 percent, **5 percent, ***1 percent.

Table A22: Relationship of Information Recall and Credibility to In-Group Criminality—2022 Experiment

	Candidate criminality information			
	Recall receiving		Not credible, if recall	
	(1)	(2)	(3)	(4)
Panel A: Religion/Caste Group				
In-group criminal	0.001	-0.039	-0.045*	-0.036
	(0.031)	(0.050)	(0.025)	(0.043)
In-group violent crime		0.057		-0.012
		(0.057)		(0.046)
Outcome mean	0.542	0.542	0.078	0.078
Observations	1,488	1,488	807	807
Panel B: Political Party				
In-group criminal	0.122***	0.157***	0.010	0.013
	(0.037)	(0.050)	(0.025)	(0.039)
In-group violent crime		-0.069		-0.007
		(0.070)		(0.048)
Outcome mean	0.540	0.540	0.078	0.078
Observations	1,477	1,477	798	798

Notes: This table reports estimates based on individual-level regressions where the sample is restricted to voters in the treatment group (i.e., who were sent information about candidate criminality). The dependent variable is an endline measure of voters' perception of information and the regressor is an indicator for receiving negative information about in-group candidates—i.e., that a voter's own-group candidate faces criminal charges. In columns (1) and (2), the dependent variable is *Recall receiving*, an indicator which equals one if the respondent reports receiving any information about candidate criminality. In columns (3) and (4), the dependent variable is *Not credible, if recall*, an indicator which equals one if the respondent reports that the information received was not credible, conditional on recalling receipt of information. Columns (2) and (4) also include an indicator for own-group candidates facing violent criminal charges. Panels A and B define group match based on shared religious/caste group (religion for non-Hindus; caste group within Hindus) and partisan alignment, respectively. All regressions include constituency fixed effects. Standard errors are in parentheses. Results are discussed in Section 7.3. Significant at *10 percent, **5 percent, ***1 percent.

Table A23: Effects of Information Treatment on Voter Coordination—2022 Experiment

	Discussed with friends / neighbors		Second-order beliefs	
	Election	Candidate criminality	Candidate criminality knowledge	Vote for clean candidate, high importance
	(1)	(2)	(3)	(4)
Panel A: Pooled Treatment				
Information	-0.020	-0.039**	0.010	0.046***
	(0.018)	(0.016)	(0.010)	(0.017)
Control mean	0.346	0.254	0.363	0.711
Observations	3,000	3,000	3,000	2,880
Panel B: Treatment Arms				
Information only	-0.020	-0.033*	0.006	0.046**
	(0.019)	(0.017)	(0.010)	(0.018)
Information and coordination	-0.022	-0.069**	0.030	0.049
	(0.032)	(0.028)	(0.020)	(0.031)
Info only = Info & coord, p-value	0.965	0.206	0.283	0.905
Control mean	0.346	0.254	0.363	0.711
Observations	3,000	3,000	3,000	2,880

Notes: This table presents results on the effects of our information treatment on voter coordination. Each column in Panel A reports estimates from individual-level regressions of the listed endline survey measures on the information treatment indicator and constituency fixed effects. In Panel B, we include two separate treatment indicators, for the information-only treatment arm and information-plus-coordination treatment arm (where voters were additionally told that many other voters had received the information message). Message scripts of both treatment arms are provided in Appendix B.1. The outcome variables in columns (1) and (2) are indicator variables capturing whether the voter discussed the election and candidates' criminal charges, respectively, with their friends and neighbors. In column (3), the outcome variable is voters' perceived share of other voters in the same constituency who know about candidates' criminal charges. The outcome in column (4) is an indicator for believing that other voters in the constituency place high importance on voting for a clean candidate. Standard errors are in parentheses. Results are discussed in Section 7.3. Significant at *10 percent, **5 percent, ***1 percent.

Table A24: Effects of Information by Treatment Arm on Voter Beliefs, Candidate Ratings, and Electoral Choices—2022 Experiment

	Accurate belief		Rating (z-score)		Vote choice	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Criminality						
Information only	0.044***	0.044***	0.041	0.041	0.020**	0.020**
	(0.016)	(0.016)	(0.030)	(0.030)	(0.008)	(0.008)
Information and coordination	0.062**	0.062**	0.056	0.056	0.019	0.019
	(0.028)	(0.028)	(0.045)	(0.045)	(0.014)	(0.014)
Criminal						
× Information only	0.079***		-0.156***		-0.030**	
	(0.020)		(0.041)		(0.013)	
× Information and coordination	0.096***		-0.188***		-0.021	
	(0.034)		(0.065)		(0.021)	
Nonviolent crime						
× Information only		0.108***		-0.091		-0.030
		(0.027)		(0.057)		(0.023)
× Information and coordination		0.188***		-0.064		0.004
		(0.056)		(0.095)		(0.041)
Violent crime						
× Information only		0.068***		-0.181***		-0.030**
		(0.021)		(0.046)		(0.013)
× Information and coordination		0.069*		-0.224***		-0.027
		(0.036)		(0.073)		(0.021)
Control mean	0.320	0.320	0.000	0.000	0.209	0.209
Observations	11,623	11,623	9,895	9,895	9,224	9,224

Notes: This table presents results on the effects of each information treatment arm on voter beliefs, candidate ratings, and electoral choices during the 2022 elections. In columns (1) and (2), the dependent variable is a voter-by-candidate-level indicator for whether voter i had accurate beliefs about the criminal status of candidate c, while in columns (3) and (4) it is voter i's rating of candidate c, converted from a 5-point Likert scale to a z-score. In columns (5) and (6), the dependent variable is an indicator for whether voter i stated that they voted for candidate c. Columns (1), (3), and (5) present estimates from a modified version of Equation (3) where we replace the pooled information treatment indicator with separate indicators for the information only (voters received information about candidates' charges) and information plus coordination (voters were additionally told that many others in their constituency also received the information) treatment arms. Message scripts of both treatment arms are provided in Appendix Section B.1. In columns (2), (4), and (6), we use same appraoch, except allowing treatment effects to vary with the maximum severity of candidates' criminal charges. Standard errors clustered at the individual level are in parentheses. Results are discussed in Section 4.3. Significant at *10 percent, **5 percent, ***1 percent.

Appendix B: Additional Details

B.1 Messaging

B.1.1 *Example Information Message Content*

Example message content for a constituency in which the major-party candidates ranged from facing no criminal charges to serious violent charges:

"This message is from an unbiased, non-political, and non-governmental organization, the Center for Governance and Development. Get to know your candidates correctly, and on Election Day, give your vote only after thinking carefully! In your area: 1. Madhusudan Kushwaha from BSP (elephant party) has one criminal case, with attempt to murder charges. 2. Prakash Dwivedi from BJP (lotus party) has no criminal cases. 3. Vivek Kumar Singh from Congress (hand party) has three criminal cases, but has no violent charges."

B.1.2 *Additional Message Types—2017*

Information plus Coordination Message: This message contained the same content as the standard information message but also emphasized to recipients that the message had been shared widely to other individuals living in their area. This intervention was designed to test whether voters respond more strongly to information when they receive a public signal that may help them to coordinate their response with other voters. This message used the following format:

"This message is from an unbiased, non-political, and non-governmental organization, the Center for Governance and Development, and many people in your area have already received it. Get to know your candidates correctly, and on Election Day, give your vote only after thinking carefully! In your area: <<Insert Criminality Information of Candidates>>. Now you can elect the right candidate with the people in your area."

Information plus Ethnic Voting Message: This message also built upon the standard information message, adding content that urged recipients to break the habit of purely ethnic-based voting. Given the history of ethnic politics and voting in this context, this intervention was designed to test whether provision of criminality information has stronger impacts on voting behaviour when coupled with content that draws one's attention to flaws in default decision-making which may undermine the value of such information. This message contained the following:

"This message is from an unbiased, non-political, and non-governmental organization, the Center

for Governance and Development. Don't follow your old habits and vote only on the basis of caste or religion. Get to know your candidates correctly, and on Election Day, give your vote only after thinking carefully! In your area: <<Insert Criminality Information of Candidates>>."

Placebo Message: In phase 7 of the elections, 410 out of an experimental sample of 771 villages across 11 assembly constituencies were randomly assigned to receive "placebo" messages, using the same stratification approach as for the phase 4 experimental sample. The remaining 361 control villages received no messages. The placebo message included no candidate-specific information and only encouraged voters to get to know their candidates, as follows:

"This message is from an unbiased, non-political, and non-governmental organization, the Center for Governance and Development. Get to know your candidates correctly, and on Election Day, give your vote only after thinking carefully!"

B.1.3 *Additional Message Type—2022*

Information plus Coordination Message: As in the 2017 experiment, this message contained the same content as the standard information message but also emphasized to listeners that the message had been shared widely to other individuals living in their area. This message used the following format:

"This message is from the Center for Governance and Development, which is an unbiased, non-political, non-governmental organization, and has already been sent to many other people in your area. Informed voters prefer clean candidates over criminal candidates so you can now unite with other informed voters and make a difference to ensure that a clean candidate is chosen from your constituency. <<Insert Criminality Information of Candidates>>."

B.2 Persuasion Rates—2017 Experiment

Taking the expression for vote share persuasion effects from the supplemental appendix of DellaVigna and Gentzkow (2010), we have:

$$f=((v_T-v_C)/(e_T-e_C))\times(t_T/(1-v_C)), \tag{6}$$

where f is the persuasion rate, t is turnout, v is total vote share for clean candidates, and e is the treatment exposure rate. The T and C subscripts denote the treatment and control groups, respectively.

Using our data from the 2017 experiment, with an average of two clean candidates per race, we have $v_T = 0.571$, $v_C = 0.563$, and $t_T = 0.626$. For both treatment exposure scenarios, $e_C = 0$, while $e_T = 0.374$ for the "receiving households" scenario and $e_T = 1$ for the "entire village" scenario. Using these values together with Equation (6) gives clean candidate voting persuasion rates of 3.3% and 1.2% for the "receiving households" and "entire village" scenarios, respectively.

B.3 Shared Religion/Caste Group and Partisan Alignment—2022 Experiment

We consider two types of shared group identity: shared religion/caste and partisan alignment.

- **Religion**: Respondent religion is based on respondents' reported category in the baseline survey (Hindu, Muslim, Sikh, Christian, other).
- **Caste**: We consider broad caste category. Caste group is defined using standard Indian administrative groupings: General Caste, Other Backward Classes (OBC), and Scheduled Castes/Scheduled Tribes (SC/ST). Hindu survey respondents were assigned to one of these groups based on their baseline-reported caste category. For sitting MLAs and other contesting candidates, religion was generally straightforward to infer from candidate names and corroborated using public sources. For Hindu candidates, caste category was coded using publicly available information, including that available in candidate affidavits and on the Uttar Pradesh Vidhan Sabha website (upvs.neva.gov.in). Candidate caste category coding proceeded in two steps. First, for SC/ST-reserved constituencies, caste category assignment was direct by construction. Second, for non-reserved constituencies, caste category was determined through targeted verification using surnames, publicly available caste references, and cross-checks against official UP classifications of castes into the broader caste categories.
- **Partisan alignment**: For each of the four major parties (BJP, SP, Congress, and BSP). respondents are categorized as a supporter if they indicated that party as their most preferred in the baseline survey.

B.4 Candidate Vignettes—2022 Baseline Survey

Candidate vignettes included in the 2022 baseline survey:

- "*The candidate's name is Ajay Kumar Srivastava, he is 40 years old. He has a graduate degree and his profession is educationalist and businessman. He has one criminal charge against him, and has wealth of three crores. He was the BSP candidate from Lucknow North constituency.*"
- "*The candidate's name is Asha Maurya, she is 55 years old. She has a graduate degree and her profession is business. She has three criminal charges against her, and has wealth of four crores. She was the BJP candidate from Mohmoodabad constituency.*"

- "*The candidate's name is Dharmendra Singh, he is 42 years old. He has 8th pass education, and his profession is mentioned as rent. He has no criminal charges against him and has wealth of 68 lakhs. He was the RLD candidate from Kanpur Cantt. constituency.*"

B.5 Causal Forest Treatment Heterogeneity—2022 Experiment

To estimate each voter's predicted conditional average treatment effect (CATE) of receiving information, we use the causal forest approach of (Wager and Athey 2018). We implement the generalized random forest (grf) package in R (Athey et al. 2019), using the package's default parameter values except for the number of trees, which we increase from 2,000 to 10,000 to improve precision.

We estimate CATEs for each of our key outcomes. For belief accuracy, candidate ratings, and vote choice, we estimate voter CATEs separately in relation to clean and criminal candidates. For the salience of criminality and decision weight on criminality, we estimate CATEs for each voter.

For each outcome, we generate a forest composed of 10,000 trees. Each tree is generated using a random 50% subsample drawn without replacement from the underlying data as follows:

1. Randomly split the subsample in half; one half is used for training and the other for estimation.
2. Using the training portion of the subsample, begin with a single node. Randomly select a subset of available attributes as potential variables to split on.
3. Choose the single split (defined by some threshold value of one of the available attributes) which maximizes the difference in the treatment effect between the two resulting "child" nodes, while also yielding specified minimum numbers of treatment and control observations in each "child" node. Split the node and repeat this step for each of the "child" nodes. If no feasible split exists, cease splitting at this node, forming a terminal "leaf".
4. A tree is generated for the subsample once no more splits can be made in the previous step.
5. Using the estimation portion of the subsample, "repopulate" the leaf nodes by matching each observation to a leaf based on observed characteristics.

The following approach is then taken to obtain a predicted treatment effect for an observation with a given set of attribute values. For each tree, the observation is matched to a leaf based on its attribute values, and a set of neighbors (observations from the final step above belonging to the same leaf) is determined. Once this process has been completed for all trees in the forest, neighbors are given weights based on how many times they belonged to the same leaf as the observation of interest. Neighbors' weights are then used together with their outcome values and treatment assignments to estimate the observation of interest's treatment effect. Sections 2.1 and 6.2 of Athey et al. (2019) describe the calculation of weights and estimation of treatment effects in detail.

For outcomes measured at the voter level (salience and decision weight), the set of splitting variables include constituency, age, gender, urban location, economic status, religion, caste group, education level, occupation type, criminality belief accuracy, and political participation, interest, attitudes, and party support.

For outcomes measured at the voter-by-candidate level (belief accuracy, candidate ratings, and vote choice), the splitting variable set further includes candidate party and criminality type (if criminal). In addition, the procedure above and variance estimation are adjusted to account for voter-level clustering. For details of the adjustments, see the content in the GRF algorithm reference guide related to cluster-robust estimation at https://grf-labs.github.io/grf/REFERENCE.html.