

Overcoming or Reinforcing Coethnic Preferences? An Experiment on Information and Ethnic Voting *

Claire Adida[†] Jessica Gottlieb[‡] Eric Kramon[§] Gwyneth McClendon[¶]

January 25, 2017

Abstract

Social scientists often characterize ethnic politics as a threat to democracy and growth, and recent scholarship investigates factors that could exacerbate or alleviate it. A dominant view—that ethnicity acts as a heuristic in low-information contexts—implies that information access could reduce ethnic voting. But this view contrasts with evidence that identity often conditions information processing, potentially in ways that amplify in-group preferences. We test these expectations with a field experiment around Benin’s 2015 legislative elections. Behavioral and attitudinal data reveal that voters reward good-performing incumbents only if they are coethnics, and punish bad performers only if they are non-coethnics. Coethnics are also more (less) likely to accurately recall performance information if it is positive (negative). These results are consistent with a theory of ethnically motivated reasoning whereby voters act on new information only when it allows them to reaffirm their social identity. These findings improve our understanding of comparative ethnic politics, identity and information processing, and information and accountability.

*We thank EGAP for generous funding. This study is part of the larger Metaketa initiative to accumulate knowledge about the relationship between information and accountability across country contexts. The registered pre-analysis plan for the Metaketa project can be found at: <http://egap.org/registration/736>. The registered pre-analysis plan for this particular study can be found at: <http://egap.org/registration/735>. We thank Amanda Pinkston for sharing 2011 legislative election data and Ana Quiroz for excellent research assistance. We thank participants at Princeton University, Harvard University, George Washington University, MPSA and WGAPE, and Leonard Wantchekon, Karen Ferree, Nahomi Ichino, Dan Posner, Thad Dunning, Kadir Yildirim, Guy Grossman, and Jake Bowers for helpful comments. This research was conducted in collaboration with the Centre de Promotion de la Démocratie et du Développement (CEPRODE), and we thank Adam Chabi Bouko for leading the implementation effort. Our project received ethics approval from the authors’ home institutions. We also obtained permission to conduct the study from the President of the National Assembly of Benin. In each study village, permission to conduct research was obtained from the chief and consent was obtained from each surveyed participant in the study.

[†]University of California, San Diego, 9500 Gilman Drive #0521, La Jolla, CA 92093. cadida@ucsd.edu.

[‡]Texas A&M University, 4220 TAMU, College Station, TX 77843. jgottlieb@tamu.edu.

[§]George Washington University. Monroe 440, 2115 G. St NW, Washington, DC 20052. ekramon@gwu.edu.

[¶]Harvard University, 1737 Cambridge Street, CGIS Knafl 207, Cambridge, MA 02138.
gwyneth_mcclendon@harvard.edu.

Scholars of comparative politics have long been concerned about the political and economic consequences of ethnic politics. The normative and empirical literature on democracy warns that politically salient ethnic divisions are likely to undermine democratic stability, accountability, and prospects for democratic consolidation (e.g. Dahl 1973; Horowitz 1985; Lijphart 1977; Rabushka and Shepsle 1972), while a substantial literature in political economy suggests that ethnic politics is associated with the under-provision of public goods and poor economic growth around the world (Alesina et al. 1999; Easterly and Levine 1997; Miguel 2004; Miguel and Gugerty 2005).

The growing literature on information and accountability offers a potential solution: increasing voter access to information about politics, or the cultivation of an informed electorate, may help to reduce the salience of ethnic divisions in democratic politics. Drawing on research that emphasizes the role of heuristics in guiding voter decision making (Popkin 1991), this literature highlights that ethnicity often becomes important in electoral politics where information is scarce (Birnir 2007; Chandra 2004; Conroy-Krutz 2013; Ferree 2006; Posner 2005). In environments where voters lack access to information about competing candidates or parties, ethnic labels provide voters with a cheap and relatively accessible source of information about either candidate quality — whether they can be trusted, are competent, and will work hard — or about the extent to which a candidate will pursue a policy or redistributive agenda that is likely to benefit specific individuals and groups (Carlson 2015; Posner 2005). This set of arguments implies that *increasing voter access to information about candidates should reduce ethnic voting* (Conroy-Krutz 2013).

But we argue that this expectation contrasts with the implications of research in social and political psychology that finds that identity often conditions how individuals process and act upon information about politics, potentially in ways that reinforce or exacerbate, rather than overcome, group-based preferences. Social identity theory emphasizes that people often derive self-esteem and other psychological benefits from observing that in-group members do well relative to other groups (e.g Tajfel 1974; Lieberman 2009). These benefits from in-group status (Kasara 2007) can provide a “directional goal” (Kunda 1999) that guides information processing. Indeed, the literature on “motivated reasoning” in American Politics argues that individuals motivated to affirm

their social (usually partisan) identities will process information in ways consistent with that goal — that is, in ways that allow them to maintain or enhance their positive view of ingroup members relative to outgroup members (Taber and Lodge 2006; Bolen et al. 2014; Kunda 1990, 1987). We argue that, where ethnicity is a highly salient social identity category, voters may engage in “ethnically motivated reasoning,” updating their beliefs about politicians only when presented with positive information about coethnics and with negative information about non-coethnics, and not otherwise. In other words, *increasing voter access to information may reinforce or amplify ethnic voting.*

The empirical and normative implications of these two theoretical approaches are quite different. According to the first, improved access to information should supplant identity-based politics, freeing voters from the need to rely on ethnic markers when making their voting decisions and reducing the salience of ethnicity in electoral politics. According to the second, voters’ group identities should condition how they process and act upon political information in ways that reinforce or exacerbate ethnic-based divisions in the electorate. We test these competing expectations in a context in which ethnicity is a highly salient social identity. In this context, *does increasing voter access to information reduce or reinforce ethnic voting?*

We address this question by analyzing data collected as part of a large-scale field experiment conducted around the 2015 National Assembly elections in Benin.¹ Benin is a democratic West

¹This paper is part of an experimental study on information and accountability in Benin and part of the EGAP Metaketa initiative. In a companion paper, we present other results from the full experimental design, while here we focus on information and ethnic voting. In Appendix C.4-C.4.4, we provide full details on all treatment conditions and discuss and provide tables demonstrating the consistency of results across papers. In designing the experiment, we followed the ethical principles agreed upon by the Metaketa initiative, as outlined in the joint metaketa pre-analysis plan: that the intervention consist of information that existed in the political system, be provided with consent, in a non-partisan way, without deception, and in cooperation with a local group. For further notes, see Appendix C.3.

African country where ethnicity is a highly salient social identity (Adida 2015; Wantchekon 2003) and where the political information environment is poor. In the experiment, villages and urban neighborhoods were randomly assigned to receive information about the legislative performance of incumbent politicians running in the election.² Individuals and villages in our sample vary in their ethnic connection to the incumbents in their area — some are coethnics with the incumbent while others are not — which allows us to leverage our experimental data to examine how ethnic ties condition the impact of information access, while controlling for factors that may correlate with coethnicity, such as prior beliefs about candidates.³

Using both behavioral data from official village- and neighborhood-level election outcomes and panel survey data, we find results that are most consistent with ethnically motivated reasoning. That is, we find that access to information amplified voter preferences for coethnic politicians and reinforced voter disapproval of non-coethnic candidates. More specifically, voters rewarded the good performance of their coethnics but did not reward the good performance of non-coethnics. By contrast, they punished the poor performance of non-coethnics but did not punish the poor performance of coethnics. We uncover these patterns in both panel survey data and in official administrative data, alleviating potential concerns about social desirability bias and differential attrition in the survey.

Further bolstering the ethnically motivated reasoning argument, we find additional evidence that ethnic identity conditioned information processing in a comprehension survey that was con-

²The informational intervention focuses specifically on the tasks for which deputies are formally charged. We discuss details of the information treatment in subsequent sections.

³The hypothesis about coethnicity as a moderator was specified in our pre-analysis plan. Nevertheless, because coethnicity is not randomly assigned, it is possible that it is correlated with other omitted variables that are driving the results. We thus provide evidence that our results are robust to controlling for a number of variables that might correlate with coethnicity, such as respondent priors on incumbent performance, partisanship, education, and poverty.

ducted immediately following the provision of performance information. The results show that the incumbent's coethnics were more likely than non-coethnics to accurately recall the information provided in the experiment if the information was positive. On the other hand, when the information suggested that the incumbent was a poor performer, coethnics of the incumbent were substantially *less* likely to recall the information accurately. These results are consistent with the notion that voters took up new information when it allowed them to maintain or enhance a positive view of coethnics, and discarded it when it did not.

The study makes a number of contributions. First, we advance theories of ethnic voting in comparative politics by highlighting an additional mechanism that has received less attention: the possible role of ethnic identity in conditioning how voters process and ultimately act on good and bad information about political performance. Although this mechanism has informed how scholars of American politics understand partisan attachment and information processing in particular, it has been surprisingly absent from the literature on comparative ethnic politics. Our claim is not that alternative theories of ethnic voting are necessarily incorrect in all cases. Rather, we show that ethnically motivated reasoning best explains our results, and conclude that ethnic attachments *can* condition how voters process new political information. This finding raises questions about the extent to which information can overcome social-identity based voting.

Second, our study contributes to the literatures on motivated reasoning and biased information processing in politics. Scholars of American politics have recognized that partisan identities can play a profound role in shaping how individuals respond to political information about politics (e.g. Bartels 2002; Campbell et al. 1960; Taber and Lodge 2006; Zaller 1992). We extend this literature by providing evidence that motivated reasoning, driven by ethnic rather than partisan attachments, can have similar effects on information processing in a different context. We thus also advance a nascent comparative politics literature that explores the implications of motivated reasoning outside of the United States (Conroy-Krutz and Moehler 2015; Carlson 2016; Horowitz and Long 2016).⁴ We go beyond these studies by examining more explicitly how ethnic identity

⁴Conroy-Krutz and Moehler (2015) employ a field experimental approach to test whether parti-

shapes information processing.

Finally, we contribute to the literature on information and accountability. Theories of political accountability generally posit that increased access to information about politician performance shapes voting behavior, helping voters to distinguish between strong and weak performers (Fearon 1999; Pande 2011). Yet evidence to date suggests that the relationship between access to performance information and voting behavior is not straightforward. Some field experiments have found that providing performance information to voters does indeed result in the punishment of poorly performing politicians (Banerjee et al. 2011; Ferraz and Finan 2008). But other experimental work has been unable to reject the null hypothesis that increased access to performance information has no effect on citizen behavior (Humphreys and Weinstein 2012; Lieberman et al. 2014). Some experimental studies have even found that the provision of performance information resulted in the punishment of challengers as well as incumbents (Chong et al. 2015). Our results shed light on these puzzling findings by highlighting a key moderating factor, ethnic attachments, that conditions how voters respond to interventions designed to enhance information access.

Ethnicity, Information, and Voting

Ethnic voting is prevalent in many democracies.⁵ A significant body of research documents that voters in many parts of the world, including in Benin, are more likely (though not determined) to vote for coethnic candidates than for non-coethnic candidates (e.g., Adida 2015; Bratton and san media moderates political attitudes in urban Ghana but find little evidence that it does. Carlson (2016) finds evidence of partisan bias in evaluations of government service provision in Uganda. Horowitz and Long (2016) argue that motivated reasoning might explain why Kenyan voters support coethnic presidential candidates that have no chance of winning.

⁵Ethnicity is understood here to mean a politically and socially constructed group identity, based on real or perceived descent (Chandra 2006).

Kimenyi 2008; Ferree 2006; Heath et al. 2015; Hutchings and Valentino 2004; Horowitz 1985; Posner 2005).⁶ However, there is considerable debate about why this is so, with important implications for how new political information might interact with ethnic voting. In what follows, we first discuss the most prominent expectation from the comparative politics literature – that ethnicity acts as a heuristic for other features of the candidate, and how new information should thus condition ethnic voting. We then introduce an alternative expectation generated by insights from the American politics literature – that ethnicity conditions how individuals cognitively process new information, and the distinct predictions this implies for the effect of information on ethnic voting.

One explanation for ethnic voting emphasizes ethnic labels as heuristics used in the face of information scarcity. In low information environments, voters may use ethnicity as a short-cut for evaluating candidates (Birnir 2007; Chandra 2004; Conroy-Krutz 2013; Ferree 2006; Posner 2005). In contrast to other candidate attributes on which it might be difficult to gather information, coethnicity is often easily observable, thus providing a cheap source of information where other sources are unavailable or difficult to obtain (Chandra 2004).⁷

The comparative politics literature emphasizes two different factors for which coethnicity can serve as a short-cut. One strand of the ethnicity-as-heuristic literature suggests that ethnic labels convey information about candidate quality: the extent to which the candidate is competent, hard-working, trustworthy, or will not shirk on the job (Conroy-Krutz 2013; Fiske et al. 2007). In the absence of information to the contrary, people may tend to assume that members of their in-groups are more highly capable (Fiske et al. 2007). Here, coethnicity serves as a means for deciding which

⁶Chauchard (2016) points out that ethnic cues may also affect voting when choosing among non-coethnics.

⁷This does not mean that voters always identify coethnics without error (Habyarimana et al. 2009; Harris and Findley 2014). Instead, it means that relative to other politician characteristics, coethnicity is more easily surmised. Chandra (2004) points out that ethnicity is often determined on the basis of highly visible attributes such as name, physical features, speech and dress.

candidates are likely to be the most effective. If ethnicity acts as a heuristic for candidate quality, *then the introduction of new information about politician performance that conveys information about candidate quality should reduce the importance of ethnicity in shaping voting decisions overall.*

A second strand of this literature highlights instead that ethnic labels provide information about which groups of voters are likely to benefit from a candidate's policies or efforts to distribute patronage. This argument, which is dominant among scholars of African politics (e.g. Carlson 2015; Ferree 2010; Ichino and Nathan 2013; Posner 2005) but has been applied in other regions as well (Chandra 2004), posits that voters expect coethnic candidates to favor them when delivering goods and services and when deciding which legislative goals to pursue. Several studies have shown evidence of political ethnic favoritism when it comes to constituency service (e.g., Butler and Broockman 2011; McClendon 2016), the delivery of local public goods (Burgess et al. 2015; Kramon and Posner 2013, 2016; Franck and Rainer 2012) and the pursuit of policies that serve coethnics' interests (Pande 2003; Preuhs 2006), so these expectations may not be unreasonable (c.f., Kasara 2007). In fact, Chandra speaks of an equilibrium whereby voter expectations of coethnic favoritism are reinforced by the biased distribution of patronage and public goods (Chandra 2004). In places like sub-Saharan Africa where ethnic groups are regionally concentrated, legislation can also be used to direct resources to particular areas, and thus groups, within the country (e.g., Bates 1983; Hodler and Raschky 2014). Furthermore, national policies or institutional arrangements beyond resource transfers can have disparate impacts on groups with heterogeneous preferences, which can be leveraged in the name of political favoritism (Kramon and Posner 2011).⁸

⁸For example, broad-based decentralization reforms in Mali were undertaken to appease ethnic Tuaregs in the North who have greater preferences for autonomy than groups predominantly residing in the South (Seely 2001). In Benin, policies about cotton-pricing will be more important to the North, where cotton grows, while policies about cashew-growing will be more important to the South. In addition, policies about land-tenure and natural resource management will be more salient to localities with larger concentrations of Peuls, the traditional pastoral group in the Sahel.

If ethnicity is a heuristic for the favorable distribution of goods by a coethnic politician, then, as Carlson (2015) finds, *increased access to information about candidate quality may matter to voters only when the information is about members of their own ethnic group*. Believing they are unlikely to benefit from the efforts of non-coethnic politicians whether those politicians are competent and hardworking or not, voters may seek to select strong coethnic performers over weak coethnic performers but may care little about the performance records of non-coethnics.⁹ In this scenario, increasing access to information may serve to enhance accountability *within* ethnic groups (Carlson 2015) but do little to overcome ethnic-based divisions in the electorate as a whole.¹⁰

The ethnicity as heuristic arguments contrast with certain psychological approaches to ethnic voting that identify two possible, often competing, objectives guiding information processing. Individuals may be motivated by accuracy goals, e.g., "to seek out and carefully consider relevant evidence so as to reach a correct or otherwise best conclusion" (Taber and Lodge 2006, 756); alternatively, individuals may be motivated by social identity goals, that is, to confirm their prior beliefs or affirm their membership in social identity groups (Taber and Lodge 2006; Kahan 2016). In the context of our study, where ethnicity is a highly salient social and political identity category

⁹Voters living in a rural areas dominated by non-coethnics might be an exception (Ichino and Nathan 2013). These voters might consider the performance record of a non-coethnic candidate who shares the dominant ethnicity of the local area because they expect that a strong non-coethnic performer would deliver local public goods to their area. This argument applies to rural voters and the delivery of public goods. We therefore do not apply it to our empirical case, which is both rural and urban, and which considers legislative performance rather than public goods delivery.

¹⁰We note that this line of argument becomes more complicated if information about candidate quality also sends signals about expected targeting or favoritism (good coethnic legislators might be better *or* worse at providing coethnics with informal transfers), but the prediction remains that good information about coethnics should have an effect (whether positive or negative) on voter support, while information about non-coethnics should not affect voters' decisions.

(Adida 2015; Wantchekon 2003), individuals might therefore be motivated to process information in a way that affirms their ethnic identity, which we term ethnically motivated reasoning.¹¹

The potential benefits to the individual of ethnically motivated reasoning are twofold. First, as social identity theory emphasizes, people derive positive psychological benefits such as self-esteem from seeing their groups do well relative to other groups (e.g Tajfel 1974; Lieberman 2009). Individuals might therefore be motivated to process information in ways that maintain or enhance a positive view of their ethnic group relative to out-groups in order to reap these psychological benefits (Bolsen et al. 2014; Westen et al. 2006; Kunda 1990). Second, individuals can derive both psychological and social benefits from affirming their status within their identity group (Kahan 2016). They might therefore process and act upon information in ways that allow them to affirm their status as “good” members of their in-group (Kahan 2016).

In some instances, the two goals will not be in conflict. For example, if a voter receives new, positive factual information about a member of her ethnic group, updating her beliefs and behavior in light of that information can serve both goals. However, in other instances — when, for example, a voter receives new, negative factual information about a coethnic — there will be a tension between accuracy and identity group affirmation, and there are a number of reasons why ethnically motivated reasoning is likely to prevail over accuracy. First, some voters might simply value the benefits of identity group affirmation over the benefits of accuracy. Second, directionally motivated reasoning may be the result of unconscious cognitive processes (Kunda 1990; Taber and Lodge 2006). Finally, ethnically motivated reasoning might trump accuracy because of rational cost-benefit calculations (Kahan 2016). Since an individual voter will rarely if ever be pivotal to an election or policy outcome, the potential benefits to the individual of accurate information processing are minimal. On the other hand, the psychic and social benefits of social identity affir-

¹¹Research on motivated reasoning often focuses on partisan motivated reasoning in the context of the United States, e.g. (Bolsen et al. 2014; Kunda 1999; Taber and Lodge 2006). As far as we are aware, the comparative politics literature has paid some attention to partisan motivated reasoning but less attention to the implications of ethnicity for cognitive processes.

mation can be enjoyed regardless of an election outcome. Thus, the individual voter has incentive to engage in motivated reasoning, even though doing so can in some instances lead to collectively sub-optimal outcomes (Kahan 2016).

We adopt this view of motivated reasoning (Kahan 2016) in our present study: where ethnicity is highly salient (as in Benin) and where the benefits to accuracy are close to zero (as in any election where no individual voter is pivotal), we can expect individuals to engage in ethnically-motivated reasoning so long as they receive some greater-than-zero benefit from affirming their ethnic identity. When voters are coethnics with the incumbent and they receive new positive information about them, accuracy and ethnic motivations will not be in conflict: we can expect voters to update their beliefs and act upon this information (that is, they will be more likely to vote for the incumbent). By contrast, a conflict between accuracy and ethnic motivations is generated when the incumbent's coethnics receive new negative information. In this case, ethnically motivated reasoners will ultimately discard this information. As a result, the provision of negative information to coethnics will have no effect on voting behavior. Finally, ethnically motivated reasoning implies the opposite pattern with the incumbent's non-coethnics. For these voters, conflict between accuracy and ethnic motivations arises when new positive information about the incumbent is provided, and we expect this information to be more likely to be discarded and thus to have no effect on behavior. On the other hand, when voters receive new negative information about a non-coethnic, there is no conflict between accuracy and ethnicity and so these voters are likely to act upon the new information (that is, they will be less likely to vote for the incumbent). In sum, if the cognitive phenomenon of ethnically motivated reasoning is at work in places where ethnic identities are highly salient, then *voters will process performance information about coethnic politicians and about non-coethnic politicians differently (Bolsen et al. 2014; Taber and Lodge 2006) and potentially in ways that amplify voter preferences for coethnics and against non-coethnics.*¹²

¹² Alternatively, one could expect motivated reasoning to occur only for the strongest ethnic-identifiers, i.e., for those individuals that receive a higher utility from affirming their ethnic identity. Instead, drawing on Kahan (2016)'s logic, we argue that motivated reasoning will manifest so long

The above arguments about what drives coethnic favoritism yield different predictions about the impact of information on ethnic politics. If coethnicity functions primarily as a short-cut for evaluating candidate quality, then increased access to information about candidate quality should attenuate the link between ethnicity and vote choice. This is because coethnics and non-coethnics will similarly update beliefs in response to new information – and where some amount of ethnic voting is the status quo, vote choice across ethnicities should converge. If instead ethnic labels serve primarily as a short-cut for evaluating whether or not a voter is likely to benefit from policy or patronage, then voters should be influenced by information about candidate quality only when it is about coethnics, and such information will do little to change status quo ethnic voting behavior overall. Last, if voters engage in ethnically motivated reasoning, then increased access to politician performance information should affect vote choice only when it helps voters maintain or enhance a positive view of coethnics relative to out-group members (e.g., when positive performance information is provided to coethnics or when negative information is provided to non-coethnics). As a result, voters' relative preference for coethnic candidates could be amplified as new information about candidate quality is introduced. In sum, there are at least three hypotheses with distinct observable implications about the possible interaction between coethnicity, voter choice and information about candidate quality. Because our experiment, detailed below, manipulates voter access to performance information, we frame our hypotheses as expectations about how coethnicity, or lack thereof, should condition voter response to new performance information.

Hypothesis 1 *If coethnicity serves primarily as a heuristic for gauging candidate quality in the absence of other information about quality, then access to new information about candidate quality should have an impact on voter behavior, regardless of whether they are the candidate's coethnics or not.*

as individuals receive a positive benefit from affirming their ethnic identity. We unfortunately have no appropriate measure of strength of ethnic identity to test this claim, but expect to observe motivated reasoning on average in a context where ethnicity is highly politically salient, as in Benin.

- If the incumbent has performed well (“positive” information), access to information will make coethnics and non-coethnics more likely to vote for the incumbent.
- If the incumbent has performed poorly (“negative” information), access to information will make coethnics and non-coethnics less likely to vote for the incumbent.

Hypothesis 2 If coethnicity serves primarily as a heuristic for assessing who will be favored by policies or redistribution, then access to new information about candidate quality, good or bad, should have an impact on vote choice only among voters who are coethnics of the candidate.

- If the incumbent has performed well (“positive” information), access to information will make coethnics more likely to vote for the incumbent.
- If the incumbent has performed well (“positive” information), access to information will have no effect on the voting behavior of non-coethnics.
- If the incumbent has performed poorly (“negative” information), access to information will make coethnics less likely to vote for the incumbent.
- If the incumbent has performed poorly (“negative” information), access to information will have no effect on the voting behavior of non-coethnics.

Hypothesis 3 If voters engage in ethnically motivated reasoning, then increased access to new information about candidate quality should influence vote choice only when it is positive news about a coethnic or negative news about a non-coethnic.¹³

- If the incumbent has performed well (“positive” information), access to information will make coethnics more likely to vote for the incumbent.

¹³This was the hypothesis registered in our individual project’s pre-analysis plan, albeit without a full theoretical discussion.

- If the incumbent has performed well (“positive” information), access to information will have no effect on the voting behavior of non-coethnics.
- If the incumbent has performed poorly (“negative” information), access to information will have no effect on the voting behavior of coethnics.
- If the incumbent has performed poorly (“negative” information), access to information will make non-coethnics less likely to vote for the incumbent.

Table 1 summarizes the distinct observable implications of these hypotheses. The next section details our research design for testing these hypotheses.

Table 1: Summary of Expectations about the Impact of Information on Electoral Support for the Incumbent

	Quality Heuristic	Favoritism Heuristic	Motivated Reasoning
Coethnic/ Positive Information	Positive	Positive	Positive
Non-Coethnic/ Positive Information	Positive	No Effect	No Effect
Coethnic/ Negative Information	Negative	Negative	No Effect
Non-Coethnic/ Negative Information	Negative	No Effect	Negative

Empirical strategy

We adjudicate between the differing hypotheses above with a data collection effort in an African democracy, Benin, where the political salience of ethnicity has already been established (Adida 2015; Wantchekon 2003). Benin has been considered a stable democracy since it first transitioned to holding free and fair elections in 1990.

We focus our analysis on an electoral race about which voters have poor information — legislative elections — and thus where providing information about incumbent quality has potential to cause voters to update their beliefs about the candidates running. Few media outlets report on legislative activity and information about legislative performance is not readily available. Constituency service activities fall largely outside the realm of what legislators are formally tasked with doing — in large part because they are given no budget with which to make policy for, or provide services to, their specific constituency.

In a companion paper, we discuss the extent to which voters in Benin care about legislative performance compared to other activities in which legislators might informally engage, such as individual- or village-targeted transfers. We find that at baseline voters in Benin clearly valued transfers over legislative performance. But, in many of the treatment conditions in our experiment, voters were moved to care about legislative performance and to consider strong legislative performance to be “good news” for their own well-being.¹⁴ In this paper, we consider whether their response to legislative performance information was further conditioned by coethnicity (or lack thereof) with the incumbent. For the sake of transparency and in order to use all of our data, we show results when analyzing behavior in all treatment conditions combined, compared to control. But we also confirm that our results hold in the subset of treatment conditions in which we can be *sure* that voters viewed legislative performance as a salient legislator activity (see Tables C.2 and C.3 in Appendix B).

¹⁴The field experiment consisted of 4 variants of treatment described in Appendix C.4, as well as variation in how many villages within an incumbent’s assigned commune were given the information. Voters in each treatment condition received the same relative performance information; what varied was whether the performance information included an additional message about the importance of legislative activity to voters’ wellbeing (a civics message), whether the information was provided publicly or privately, and how widely the information was disseminated in the incumbent’s constituency. Where a civics message was widely disseminated, we are confident that voters viewed strong legislative performance as important and positive.

Deputies in the national assembly are technically elected in multi-member districts by proportional representation, but in practice the system operates largely like one with single-member districts. Voters elect an average of 3.5 deputies per constituency, and with 77 total communes distributed among the constituencies, there are 3.2 communes per constituency on average.¹⁵ This distribution makes feasible, as a rule of thumb for voters and legislators, a one-to-one mapping of communes to legislators. Indeed, in practice, each legislator focuses on and “takes care of” a particular commune within his constituency, thus facilitating a one-to-one correspondence of incumbent legislator to commune.¹⁶ We thus provide voters in treated villages with performance information about one incumbent. In order to be doubly confident that our approach is appropriate, we restrict our experimental sample to 30 communes in which our local partner organization firmly verified a one-to-one correspondence and in which the incumbent legislator was running again. We further verified the one-to-one correspondence in our baseline survey by asking respondents to identify pre-treatment the legislative deputy who is most responsible for their village. And in the Discussion section, we further explore whether relaxing our assumptions about a one-to-one incumbent mapping would have implications for the interpretation of our results.

One advantage of studying ethnic voting in a proportional representation system is that ethnicity cannot fully predict vote choice, even if candidate ethnicity were a voter’s only decision criterion. This is because, in all cases, multiple parties represent a single ethnicity (see Figure B.1)

¹⁵Administratively, Benin is divided into 12 departments with two legislative constituencies in each, for a total of 24 constituencies. The next administrative level down is the commune, and there are, on average, three communes per constituency. Villages (or their urban equivalent, quarters) then nest within communes.

¹⁶We note that this mapping in practice is consistent with expert evaluations of the party system in Benin as fragmented and weak. Parties are created and dismantled frequently, lack programmatic character, and reflect instead the personality of their founder(s) (Banégas 2003; Gazibo 2012). Pre-experiment focus groups also confirm that villagers can name and agree on a single legislator as their incumbent representative.

and party lists are often comprised of candidates of multiple ethnicities (see Figure B.2). Because coethnics of the incumbent have multiple coethnic choices, it is thus plausible, at baseline, that there is still room to move coethnic voters to support the incumbent. And, since party lists are diverse, it is plausible that non-coethnics of the incumbent could already support the incumbent's party at baseline because they are coethnics of another party member.

Experimental Design

Scholars face an inference problem when trying to identify the effect of information on voter behavior: certain types of people are more likely to be politically informed than, *and* are likely to vote differently relative to, others. To identify the effect of information on vote choice, an experimental manipulation is therefore advantageous. Furthermore, the level of voter information is a relatively simple construct to manipulate externally and non-deceptively, and to do so in a way that avoids spillovers or violations of the independence of treatment assumption. In our particular case, we cluster treatment assignment within villages which is the unit at which information is most likely to travel. Cross-village or cross-quarter information transmission is possible but less common, and would bias against our finding a treatment effect.

Our experiment randomly assigns villages either to receive information about the incumbent legislative representative's performance or not to receive such information. Details of the treatment are described below. Because strong and weak legislative performers are not randomly distributed across space, we conduct a within-legislator design in which villages within each of our 30 communes are randomly assigned to treatment and control conditions. When evaluating treatment effects, we thus take the across-commune average of within-commune effects. We explain our assignment strategy further below.

While our conditioning variable, coethnicity, is also not externally manipulated, we argue that the effects we identify are plausibly causal for the following reasons. First, our measure of coethnicity with the incumbent is self-reported prior to the experimental manipulation (pre-

treatment), meaning that coethnicity cannot be affected by our treatment. Second, our assumption that ethnicity remains fixed during the period of study, e.g. that no sorting across ethnicities occurs as a consequence of our treatment, is reasonable: our experimental manipulation occurred over the course of a single month, and ethnic identity in Benin is sticky. Third, we show that our results are robust to controlling for a variety of attributes that might be correlated with coethnicity, including prior beliefs about incumbent performance, partisanship, and prior vote choice in the 2011 election (see Appendix B.3).¹⁷

Treated villages in the study were given information about the incumbent legislator's relative performance in the National Assembly in the form of a video. This mode of delivery ensured consistency in the wording and tone of the message across the sample while at the same time making the information accessible to people of all education levels, literate and illiterate.¹⁸ The video also approximates how media outlets might deliver information about candidates in a real-world setting.¹⁹ In the video, a male actor reads a script in a neutral tone, as a news caster or radio host might, and graphics illustrate key points. The text was recorded in French and then dubbed in local languages as necessary.²⁰

The information provided was drawn from official reports of the Office of the President of the National Assembly that, while supposedly public, required extensive time and effort to obtain. From the reports, the authors produced a set of relative performance indices drawn from a set of indicators about an incumbent legislator's: 1) rate of attendance at legislative sessions, 2) rate of posing questions during legislative sessions, 3) rate of attendance in committees, and 4) productivity of committee work (the number of laws considered by the committee). While the

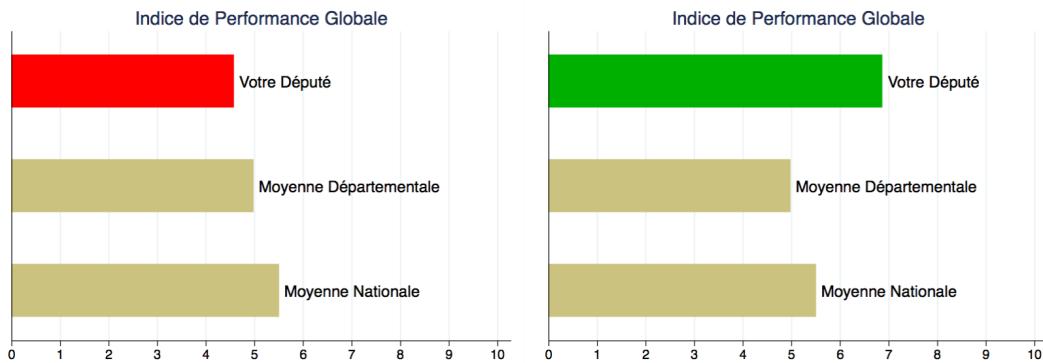
¹⁷In addition, we pre-specified this conditional effect in our pre-analysis plan.

¹⁸We presented the video to focus groups in rural villages prior to implementation, which confirmed that the information and images were accessible and comprehensible to villagers in Benin.

¹⁹See Bidwell et al. (2015) for the use of video-taped debates for similar reasons.

²⁰Full text of the video script in English is in Appendix C.4.1.

Figure 1: Two Examples of Intervention Bar Graphs



Note: Performance indices are constructed relative to both other legislators in the department (a local average) and the country (national average). Red bars are used when the incumbent's performance falls below the average and green bars when the incumbent's performance is above the average.

video provides raw data for each of these four indicators, it displayed graphics, like those in Figure 1 (which provides an example of a relatively poorly-performing politician on the left and a relatively well-performing politician on the right), of three key performance indices to increase comprehension and recall by participants: an index of plenary performance on a scale of 1-10 that takes the average of normalized scores on attendance and participation during full legislative plenary sessions, an index of committee performance also on a scale of 1-10 that takes an average of the normalized scores on attendance at committee meetings and productivity, and a global performance index which averages the first two indices.

Treatment was administered directly after a baseline survey. The baseline survey was also conducted (without the intervention) in control villages.²¹ Survey respondents were randomly selected in each village through a random-walk procedure (see Appendix C.2 for additional details). On average, 47 people per village received treatment. The treatment was conducted over the course of one day sometime within the month prior to the 2015 legislative elections.

After selecting our communes based on the one-to-one mapping confirmations described

²¹In the control condition, respondents received only a baseline survey. They did not receive any information about legislator performance or legislative responsibilities.

above, we drew our sample of villages (or their urban equivalent, quarters) and assigned them to treatment or control.²² To increase statistical efficiency, we sampled and randomized while stratifying on urban/rural status and electoral competitiveness of the village in the previous legislative election. Electorally non-competitive urban areas are rare, so we construct three blocks: urban, rural-competitive, and rural-non-competitive. Within communes, we randomly select five villages/quarters from each of the three blocks to form the sample.²³ We then randomly assign each of the five to an experimental condition.²⁴ In the Appendix, we provide a CONSORT diagram that illustrates the sampling and randomization process.

Data and Measurement

We use data from two sources – administrative election results and a panel survey, each with its own advantages and disadvantages. The administrative data allows us to test the effects of our intervention on actual behavioral outcomes. But for our theory of individual-level behavior, it presents an ecological inference problem. The panel survey data overcomes such inference problems by allowing us to test our argument directly at the individual-level. But, as a panel

²²We additionally vary whether communes receive a high dosage of treatment, e.g. a large number of treated villages, or a low dosage of treatment. The specific procedure for randomizing dosage treatment at the commune level is in Appendix C.1, but we do not disaggregate results by dosage in this paper as it is not directly related to the theory we are testing.

²³In low-dosage communes, only one treatment village is selected so blocking as this level does not apply. In some high-dosage communes, there were only five villages/quarters in the block, in which case all five were selected.

²⁴The experimental condition assigned to the single treated village in low-dosage communes is always the Civics/Public treatment to bias against finding an effect of the dosage treatment described in the appendix.

survey, these data present important challenges generated by response bias and attrition.²⁵ In this paper, we use both sources of data to leverage the advantages of each and to cross-validate our results. We describe each dataset in turn and then discuss how we construct our key variables.

To measure the effect of treatment on aggregate outcomes, we collected administrative data on election outcomes at the polling station level. We were able to match 2015 polling station data to all villages in our experimental sample except for one treated village and two surveyed control villages. These villages thus drop out of our analysis.

To measure the effect of treatment at the individual level, we collected panel survey data through a baseline in-person survey conducted 2 weeks to 1 month prior to the election and an endline phone survey conducted immediately after the election. The identities of the respondents were re-confirmed in the endline survey by calling the phone number provided in the baseline survey and asking for confirmation of respondents' first names and ages. To discourage attrition, one-third of total compensation per respondent was transferred as phone credit only after completion of the endline survey. In designing the study, we allowed for a possible 50% attrition rate between surveys and achieved a lower attrition rate (44%). A total of 3,419 individuals participated in the baseline *and* endline surveys (6,132 in the baseline).

Pre-treatment, we have village-level data on measures we used for blocking – urban/rural, incumbent legislative performance, and electoral competitiveness from the 2011 legislative elections. We also have village-level vote margin and the number of registered voters. In Appendix A, we use these data to provide evidence of balance across treatment groups.

Our key dependent variable is voting for the incumbent. At the aggregate level, we use administrative data to calculate vote share for the incumbent party at the polling station level in our experimental sample and then aggregate it to the village level ($N = 1,499$) in the case of multiple

²⁵In Appendix B.9, we discuss the extent to which inferences made from our survey data might be biased. Several tests to mitigate problems of attrition and response bias demonstrate that our conclusions are relatively robust.

polling stations per village. At the individual level, we use self-reports of voting for the incumbent party in the endline survey. The exact question, asked only of individuals who reported voting, is:

We would now like to know which political party you voted for in the legislative elections. Your response is entirely confidential and it will not be shared with anyone outside of the research team. We would like to know if you voted for the political party of [NAME OF PRINCIPAL DEPUTY]. The name of the party is [PARTY NAME] and its symbol is the [PARTY SYMBOL]. Just answer YES or NO. Did you vote for the party of [NAME OF PRINCIPAL DEPUTY]?²⁶

Our measure of coethnicity with the incumbent is self-reported on the baseline survey. The specific question is:

Thinking of the [NAME OF PRINCIPAL DEPUTY], would you say that you share the same ethnic group as this candidate?

We use a subjective measure of coethnicity because we believe the voter's own understanding of whether she has ethnic ties to the incumbent (pre-treatment) to be the theoretically relevant concept in the hypotheses we propose. Since individuals often have multiple, sometimes overlapping, group attachments that all could be considered ethnic identities (Chandra 2004; Posner 2005), the advantage to this approach to measurement is that it allows us to measure coethnicity without having to make assumptions about which ethnic categories are most relevant to specific voters.

In all analyses, we separate the effect of learning the incumbent was a good relative performer from the effect of learning the incumbent was a bad relative performer, as we pre-specified. We do so because we expect voters to respond differently to information about legislators that is positive

²⁶We deliberately asked this question in this format to protect respondent privacy and comfort level. The endline survey was conducted over the phone, and we preferred having respondents answer "yes" or "no" rather than openly voice their vote choice.

versus information that is negative. In our analysis of the official election results, we leverage the fact that the information provided in the intervention explicitly compares the incumbent legislator's performance to the performance of deputies in the surrounding area (those in the same department) and code positive and negative information relative to this local benchmark. More specifically, we define the information as positive if the incumbent's overall score is better than that of other deputies in the department. Poor legislative performers are those whose overall legislative score is worse than that of other legislators in this local area.²⁷ However, in our analysis of the survey data, we code positive and negative information relative to participant priors. In our baseline survey, we asked about the incumbent's relative performance, using the same scale that is provided in the intervention. We code the information as positive if the information provided in the intervention was better than the respondent's prior, and negative if it was worse. In instances where the information in the intervention is the same as the respondent's prior, or where the respondent reports that s/he does not know the incumbent's legislative performance, we follow the coding rule used with the official results data.²⁸ Findings using the survey data are robust to employing the same coding rule as the one we use with administrative data that does not account for priors.

²⁷This coding rule was pre-specified in our pre-analysis plan prior to project implementation.

²⁸For those whose priors match the information in the intervention, the logic is that the intervention should make them more confident in their assessment. For example, if their prior is that the incumbent is a bad performer and they receive information that validates that prior, they will become more confident in their beliefs. For those who have no priors (54% of baseline participants), the logic is that the intervention provides them with the only information they have. For both groups, it is thus reasonable to code the information they receive as good or bad based on the incumbent's objective performance relative to others in the local area. These coding rules were pre-specified in our pre-analysis plan prior to project implementation.

Does Ethnicity Condition the Impact of Information?

How does ethnicity condition the impact of performance information? Figures 2 and 3 graphically summarize our main results using the official election results data and the survey data, respectively. The figures present the average treatment effect in four sub-groups of the sample: coethnic villages (or survey respondents) with good performing incumbents, non-coethnic villages (or survey respondents) with good performing incumbents, coethnic villages (or survey respondents) with poor performing incumbents, and non-coethnic villages (or survey respondents) with poor performing incumbents.²⁹ We estimate treatment effects using OLS with block fixed effects. As blocks are nested within communes, our treatment effect estimates are driven by differences between treated coethnics (non-coethnics) and control coethnics (non-coethnics) *that share the same incumbent*, which increases the precision of our estimates and controls for potential differences between communes and legislators.³⁰

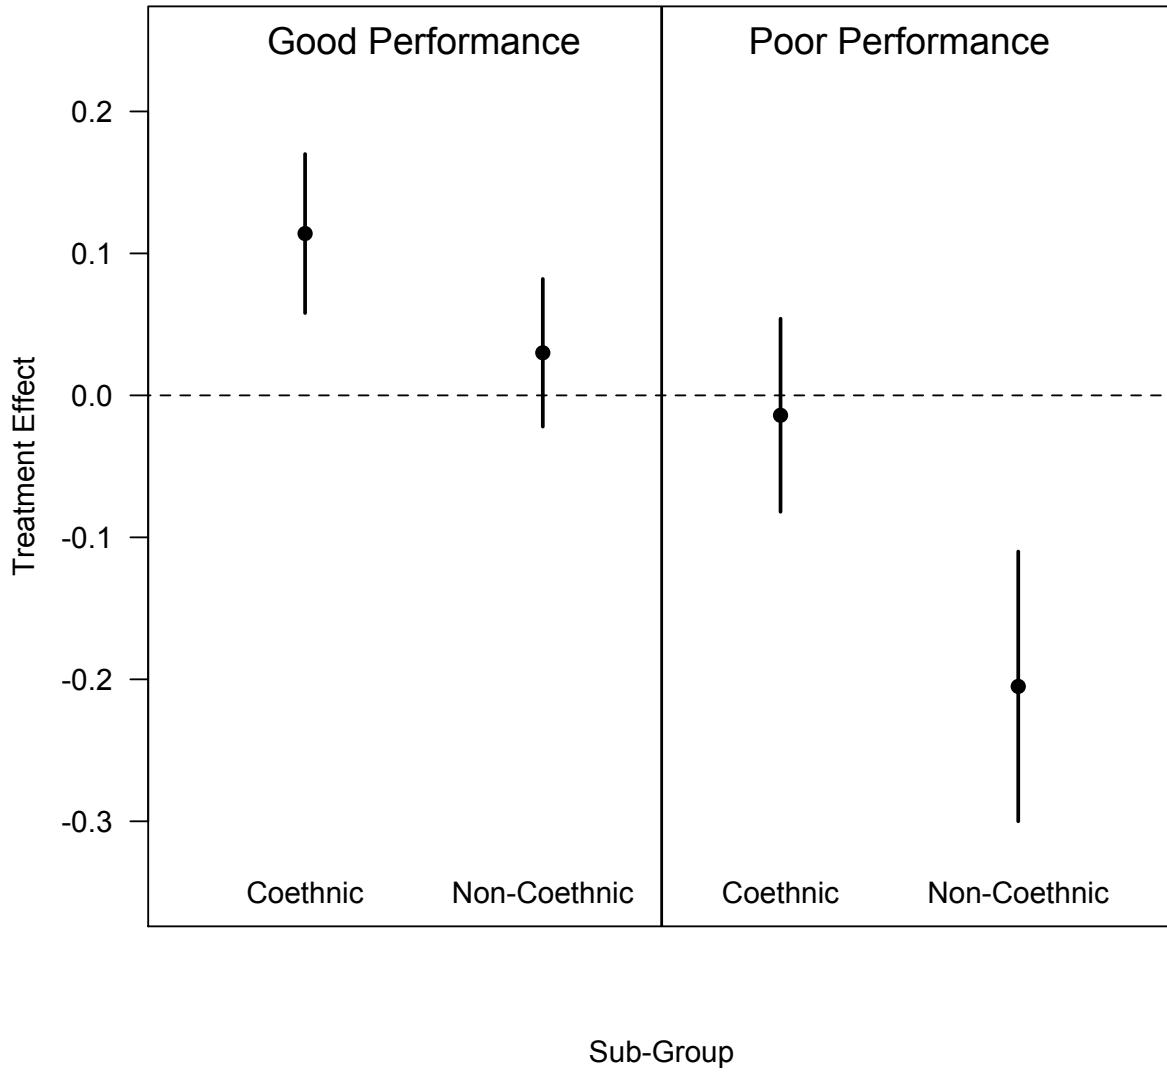
The results in both figures are substantively equivalent. Positive information increases support for the incumbent among coethnics; but it does nothing for support among non-coethnics. By contrast, negative information decreases support for the incumbent among non-coethnics; but it does nothing for support among coethnics.³¹

²⁹For simplicity, we use three thresholds for classifying a village as coethnic or not: 50%, 70%, and 90% of our village sample self-identified as coethnic with the incumbent in a pre-treatment survey question. While these were not pre-specified in our pre-analysis plan, our analysis is robust to any cutoff above 50% as demonstrated in Figure B.6.

³⁰In other words, no commune-level characteristic — whether it be urban/rural status, the level of economic development, the political competitiveness, etc... — is a plausible confounder in this analysis.

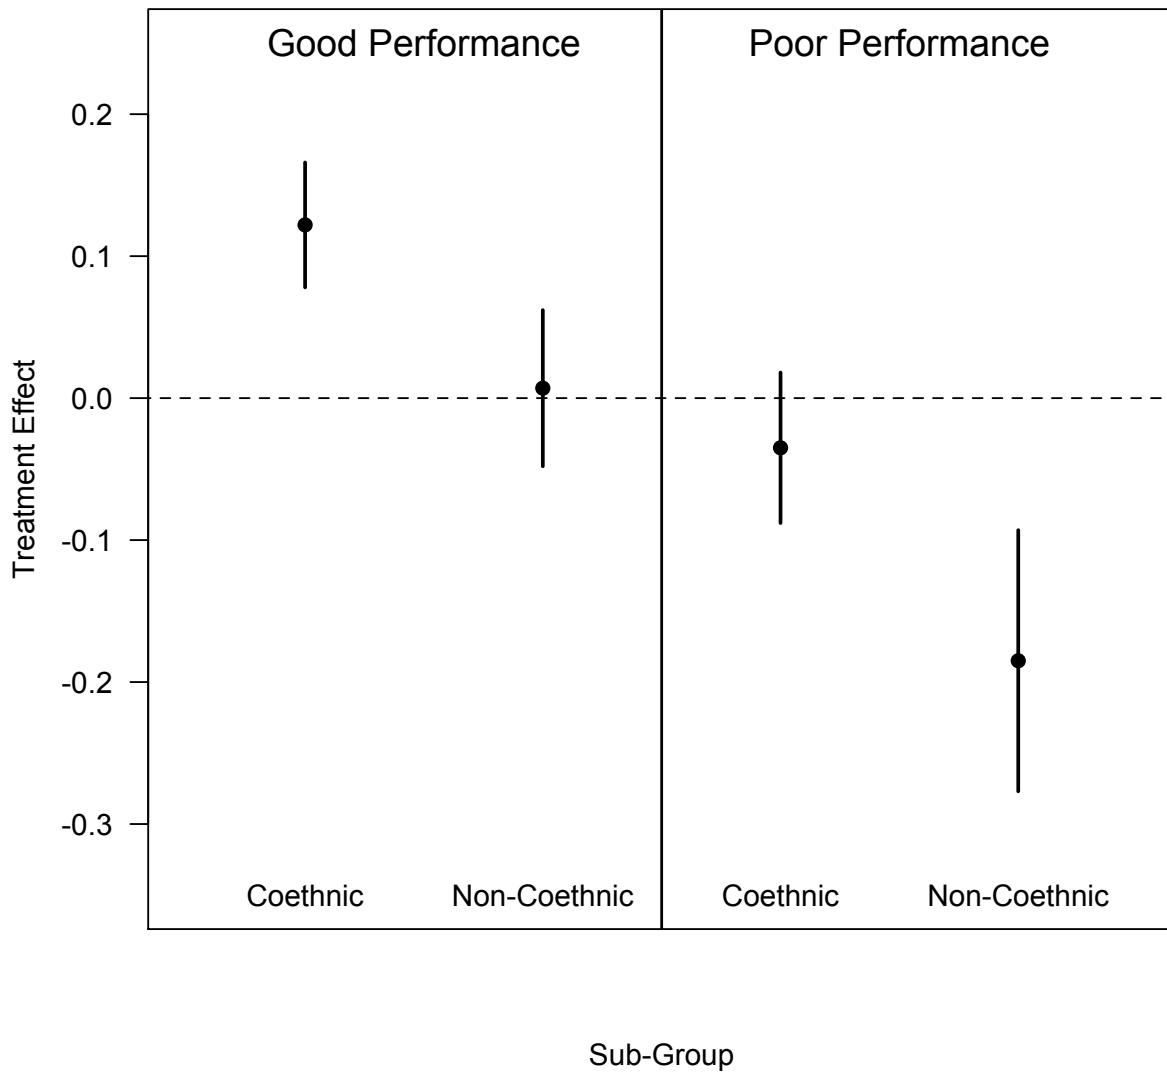
³¹Coethnicity was one of several moderators we were interested in testing in this project (and about which we have pre-specified hypotheses). We focus on ethnicity here because it is of greatest theoretical interest and because it allows us to speak to the ethnic politics literature in addition to the literature on information and accountability. However, in Appendix B.8 we show that other

Figure 2: The Effect of Information on Incumbent Voteshare, Conditional on Ethnic Connection and Level of Performance, Official Results Data



The dependent variable is incumbent vote share in the village, calculated from official election results. The figure presents the average treatment effect (and one standard error above and below the estimate) in four sub-groups: coethnic villages with good performing incumbents, non-coethnic villages with good performing incumbents, coethnic villages with poor performing incumbents, and non-coethnic villages with poor performing incumbents. In this figure, coethnic villages are coded using the 90% cutoff. Treatment effects are estimated using OLS with block fixed effects.

Figure 3: The Effect of Information on Incumbent Voteshare, Conditional on Ethnic Connection and Level of Performance, Survey Data



The dependent variable is an individual's vote for the incumbent, calculated using the survey data. The figure presents the average treatment effect (and one standard error above and below the estimate) in four sub-groups: coethnics with good performing incumbents, non-coethnics with good performing incumbents, coethnics with poor performing incumbents, and non-coethnics with poor performing incumbents. Treatment effects are estimated using OLS with block fixed effects.

Turning to a numeric presentation of the results, Table 2 displays the analysis using the official village/quarter-level election data. Columns 1 and 4 in each panel examine the unconditional treatment effect in good performance and poor performance communes, respectively. In good performance information communes, treatment has a positive but not statistically significant effect. In bad performance information communes, treatment has a negative but not statistically significant effect.

However, these null average treatment effects mask the heterogeneous treatment effects we find when conditioning on several village-level measures of politician-respondent coethnicity in columns (2)-(3) and (5)-(6). The pattern we uncover in each panel is remarkably consistent. In majority coethnic villages/quarters, positive performance information increases the vote share of the incumbent (column 2). The estimated treatment effects in column 2 are similar in magnitude and are statistically significant in the top and bottom panels. Although the coefficient in the middle panel is not statistically significant at conventional levels, the coefficient is qualitatively similar (and the p-value is 0.14). In non-coethnic villages/quarters, on the other hand, the impact of positive performance information is close to zero in all three panels and the effects are not significant at conventional levels (column 3). In coethnic areas, negative performance information has no impact on incumbent vote share (column 5). The coefficients in column 5 are substantively small and cannot be statistically distinguished from zero. By contrast, negative performance information reduces the vote share of incumbents in non-coethnic areas (column 6). In column 6, each of the coefficients are similar in magnitude, and estimates are statistically significant in the middle and bottom panel. In sum, positive performance information increases the vote share of the incumbent in coethnic areas but has no impact on the incumbent's vote share in non-coethnic areas. Negative performance information, on the other hand, has no impact on incumbent vote share in coethnic areas and a negative impact on vote share in non-coethnic areas.³²

moderators are less important empirically.

³²In some of the analyses of the administrative data, we can only report this result with 90% confidence. We find these results more convincing when paired with the survey results with which

Table 2: The Effect of Information on Incumbent Voteshare, Conditional on Ethnic Connection and Level of Performance, Official Results

VARIABLES	(1) Good Info Full Sample	(2) Good Info Coethnic (50)	(3) Good Info Non-Coethnic (50)	(4) Bad Info Full Sample	(5) Bad Info Coethnic (50)	(6) Bad Info Non-Coethnic (50)
Treatment	0.04 (0.03)	0.09** (0.04)	-0.03 (0.05)	-0.09 (0.07)	-0.02 (0.09)	-0.21 (0.17)
Constant	0.48*** (0.03)	0.47*** (0.03)	0.46*** (0.03)	0.57*** (0.07)	0.52*** (0.09)	0.62*** (0.14)
Observations	137	89	48	103	72	31
R-squared	0.59	0.70	0.56	0.60	0.67	0.58
VARIABLES	(1) Good Info Full Sample	(2) Good Info Coethnic (70)	(3) Good Info Non-Coethnic (70)	(4) Bad Info Full Sample	(5) Bad Info Coethnic (70)	(6) Bad Info Non-Coethnic (70)
Treatment	0.04 (0.03)	0.08 (0.05)	-0.02 (0.06)	-0.09 (0.07)	-0.01 (0.07)	-0.22* (0.12)
Constant	0.48*** (0.03)	0.43*** (0.05)	0.53*** (0.04)	0.57*** (0.07)	0.53*** (0.06)	0.62*** (0.10)
Observations	137	66	71	103	58	45
R-squared	0.59	0.72	0.59	0.60	0.76	0.63
VARIABLES	(1) Good Info Full Sample	(2) Good Info Coethnic (90)	(3) Good Info Non-Coethnic (90)	(4) Bad Info Full Sample	(5) Bad Info Coethnic (90)	(6) Bad Info Non-Coethnic (90)
Treatment	0.04 (0.03)	0.11* (0.06)	0.03 (0.05)	-0.09 (0.07)	-0.01 (0.07)	-0.21** (0.09)
Constant	0.48*** (0.03)	0.41*** (0.05)	0.48*** (0.04)	0.57*** (0.07)	0.55*** (0.06)	0.60*** (0.08)
Observations	137	44	93	103	53	50
R-squared	0.59	0.83	0.51	0.60	0.77	0.58

This table presents results using different cutpoints for defining a village as a coethnic village. In the top panel, villages are coded as coethnic if over 50 percent of survey respondents are coethnics of the incumbent. In the middle panel, villages are coded as coethnic if over 70 percent of survey respondents are coethnics of the incumbent. In the bottom panel, villages are coded as coethnic if over 90 percent of survey respondents are coethnics of the incumbent. Robust standard errors clustered by commune-treatment are in parentheses. *** p<0.01, ** p<0.05, * p<0.1

We also confirm in Tables C.2 and C.3 that these patterns hold when using only the treatment conditions in which we are confident that voters viewed legislative performance as an important and positive means of improving voter welfare. Because we consider only a subset of all treatment conditions here, we lose statistical power, but the patterns remain the same.

Table 3 presents results from analyses using the survey data. The results are almost identical to those from the analysis of the official election data. Panel A presents results from specifications that include only block fixed effects. When the incumbent's coethnics receive positive information about performance, they are about 12 percentage points more likely to report having voted for the incumbent (column 2). When non-coethnics receive the same positive information, however, they are no more likely to report voting for the incumbent (column 3). When coethnics of the incumbent receive negative performance information, they are no more or less likely to report supporting the incumbent (column 5). Yet when non-coethnics receive negative information, they are about 19 percentage points less likely to report voting for the incumbent (column 6). Once again, the evidence supports the claim that voters in this context reward good performance only if the incumbent is their coethnic, while they punish bad performance only when the incumbent is from a different ethnic group.³³

Panel B of Table 3 presents results from specifications that include controls for each respondent's *prior* evaluation of the incumbent's legislative performance. These analyses allow us to control for the possibility that differences in the prior beliefs of coethnics and non-coethnics could be driving our results. The omitted reference category in each specification is the set of respondents who report that they do not know about the incumbent's legislative performance (those without a

there is remarkable consistency.

³³In Appendix B.7, we present similar analyses in which the dependent variable is voter turnout. We find some evidence that negative performance information increases voter turnout, while positive performance information has no effect. We do not, however, find evidence that the impact of information on voter turnout is conditioned by coethnicity.

prior). The treatment effect estimates in these specifications are thus driven by comparisons between coethnics (non-coethnics) in treatment versus coethnics (non-coethnics) in control who have *both* the same incumbent *and* the same prior beliefs about the performance of the incumbent. In other words, this test speaks directly to the concern that co-ethnicity, which cannot be randomly assigned in any field experiment, is proxying for different distributions in prior beliefs about the incumbent's legislative performance. Even after we control for the legislator herself and for the respondent's prior beliefs about that legislator's performance, the conditional effects of coethnicity hold: There is a positive treatment effect among coethnics who receive positive information and a null effect among coethnics who receive negative information. Conversely, there is a null effect among non-coethnics who receive positive information, and a negative treatment effect among non-coethnics who receive negative information.

While the tables presented here show clear patterns across subsamples of individuals by coethnicity, Appendix Tables B.2 and B.4 present interaction models that demonstrate these differences are often statistically different from each other as well.³⁴ Since coethnicity with the incumbent is not randomly assigned, we also check that our results are robust to the inclusion of control variables that might be correlated with ethnicity and responsiveness to performance information. The results in Appendix B.3 show that our results are robust to the inclusion of controls for age, years of education, and a number of individual and community level measures of poverty and socio-economic development.

³⁴Following our pre-analysis plan, we first present uncorrected p values for these interaction coefficients and then in Appendix B.8 run multiple comparison corrections for the related family of hypotheses we pre-registered.

Table 3: The Effect of Information on Incumbent Voteshare, Conditional on Ethnic Connection and Level of Performance, Survey Data

	(1) Good Info Full Sample	(2) Good Info Coethnic	(3) Good Info Non-Coethnic	(4) Bad Info Full Sample	(5) Bad Info Coethnic	(6) Bad Info Non-Coethnic
Panel A						
Treatment	0.07* (0.03)	0.12*** (0.04)	0.01 (0.05)	-0.09 (0.06)	-0.03 (0.05)	-0.19** (0.09)
Constant	0.45*** (0.03)	0.40*** (0.04)	0.49*** (0.05)	0.57*** (0.05)	0.55*** (0.05)	0.60*** (0.08)
Observations	1,672	1,030	627	1,358	948	408
R-squared	0.18	0.22	0.18	0.24	0.28	0.22

	Good Info Full Sample	Good Info Coethnic	Good Info Non-Coethnic	Bad Info Full Sample	Bad Info Coethnic	Bad Info Non-Coethnic
Panel B						
Treatment	0.07* (0.04)	0.12*** (0.04)	0.01 (0.06)	-0.09 (0.06)	-0.04 (0.05)	-0.18** (0.09)
Prior — Much better	-0.01 (0.07)	-0.01 (0.08)	-0.02 (0.12)	0.02 (0.08)	0.02 (0.09)	0.03 (0.19)
Prior — A little better	0.02 (0.04)	0.04 (0.04)	-0.02 (0.07)	-0.05 (0.04)	-0.01 (0.05)	-0.12* (0.07)
Prior — A little worse	0.04 (0.04)	0.02 (0.05)	0.07 (0.08)	0.03 (0.04)	0.03 (0.05)	-0.02 (0.07)
Prior — Much worse	0.02 (0.06)	0.05 (0.07)	-0.04 (0.09)	-0.03 (0.04)	-0.00 (0.05)	-0.06 (0.09)
Constant	0.44*** (0.03)	0.39*** (0.04)	0.49*** (0.05)	0.58*** (0.06)	0.55*** (0.05)	0.62*** (0.09)
Observations	1,664	1,028	626	1,357	948	407
R-squared	0.18	0.22	0.18	0.25	0.28	0.23

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Why Does Ethnicity Condition the Impact of Information?

The results presented above demonstrate that ethnic identity conditions the impact of information on voting behavior. Coethnics of the incumbent are responsive to performance information only when it is positive, while non-coethnics are responsive to the information only when it is negative. In this section, we consider potential explanations for these results, as well as their implications

for theories of ethnic voting.

Ethnicity-as-Heuristic about Politician Quality

Our evidence is not consistent with models that emphasize the primary role of ethnicity as a heuristic for evaluating candidate quality (competence, effort). According to this set of arguments, ethnicity acts as a heuristic about politician effectiveness and capacity under informational constraints and, as a result, increased access to actual performance information should weaken the link between coethnicity and incumbent support (Hypothesis 1). Instead we find that voters respond to information in ways that amplify, rather than diminish, the association between coethnicity and incumbent support.

Ethnicity-as-Heuristic about Politician Favoritism

Likewise, our evidence is not consistent with the claim that ethnicity acts as a heuristic about politician favoritism. That line of argument suggests that voters will change their behavior in light of information about candidate quality only when it is about a coethnic candidate, since the quality of non-coethnics should be irrelevant to voters who use coethnicity as a heuristic about favoritism (Carlson 2015). Indeed, when they are not coethnics of the incumbent, voters under this argument should not expect to benefit from his efforts in office and so should be indifferent to whether he is a strong or weak performer (Hypothesis 2). Although the coethnic favoritism hypothesis has been developed primarily in reference to public goods provision, we expect these patterns to hold for legislative performance as well. Given the often geographically-segregated composition of ethnic group identities in Africa (Bates 1983; Hodler and Raschky 2014) and the resulting heterogeneity in preferences over policy, it is entirely plausible that voters might expect that coethnic politicians would be more likely to pursue legislative policy of greater benefit to them. Yet, contrary to the observable implications of this hypothesis, we find that *non-coethnics also* change their behavior when they receive negative information and that coethnics are responsive

only to positive information about the incumbent.³⁵

Ethnicity-as-Social Identity

Finally, we examine the possibility that voters favor coethnics because they derive positive psychological benefits such as self-esteem from seeing their group do well relative to other groups, and therefore process new information in such a way as to maintain or enhance this positive view. In other words, are our results consistent with voters engaging in ethnically motivated reasoning when presented with new information about politician performance (Hypothesis 3)? Our main experimental findings are consistent with the ethnically motivated reasoning hypothesis, which implies that performance information should influence vote choice only when it is positive news about a coethnic or negative news about a non-coethnic.

To further probe the plausibility of ethnically motivated reasoning, we analyze data from a comprehension survey conducted *immediately* after the performance treatment was administered (in other words, before participants had any opportunity to deliberate) among a randomly selected sample of 30 percent of treated participants. In the comprehension survey, we asked respondents to report on the relative performance of their incumbent in plenary and committee work.³⁶ That is, we asked respondents to provide us with the information to which they had been exposed just minutes earlier. We leverage these survey responses to test the psychological mechanism of motivated reasoning.³⁷

³⁵These patterns persist even when we examine only those places where a widely-disseminated civics message underscored the potential importance of legislation to voters' individual wellbeing, as can be seen in Appendix Table C.2.

³⁶Unfortunately, we did not ask a question about the relative overall performance of the incumbent; this is why we present results specifically on committee and plenary work.

³⁷In contrast to the analysis in previous sections, these analyses of mechanisms were not pre-specified. While we would have liked to test an additional observable implication – the effect

First, we create a dummy variable that takes a value of 1 if the respondent provides the correct answer to the plenary and committee performance comprehension questions, respectively, and 0 otherwise. Table 4 examines whether coethnicity with the incumbent is associated with correct answers in both the good and bad performance information communes.³⁸ In all models, we control for whether or not the respondent's prior belief about the incumbent collected at baseline was correct. This helps mitigate the possibility that our test is simply picking up pre-existing beliefs rather than an effect of information processing.

The results are consistent with the motivated reasoning explanation. In good performance areas, coethnics of the incumbent are significantly *more* likely to accurately recall the plenary information. They are not significantly more likely to do so in the bad information areas. On the committee score, coethnics are not more likely to accurately recall the information if it was positive, but they were significantly and substantially *less* likely to accurately recall the information if it was negative. Thus, coethnics appear more likely to accurately report the performance information if the information is positive and less likely to accurately report it if it is negative. Voters appear to have updated their beliefs only when the information was consistent with being able to view coethnics in a positive light and non-coethnics in a negative light.

Second, we analyze patterns of “Don’t Know” responses to the same two comprehension questions about incumbent plenary and committee performance. We expect that voters will be more likely to process, and thus learn, new information if the information allows them to affirm their social identities (they receive positive information about a coethnic) and less likely to process new information if the information does not allow them to affirm their social identities (they receive of motivated reasoning on cognitive effort, the best available measure of the effect of identity on information processing, response latency or time spent on answering the questions (Huckfeldt et al. 1999), was unusable due to a glitch in our survey software.

³⁸Again, respondents to the comprehension survey are from the treatment group only and so this table is of a different format than the previous tables of results.

Table 4: The effect of coethnicity on respondent comprehension of our treatment information

VARIABLES	(1) Good Performance Plenary	(2) Bad Performance Plenary	(3) Good Performance Committee	(4) Bad Performance Committee
Coethnic with incumbent	0.110** (0.046)	0.073 (0.054)	0.012 (0.043)	-0.188*** (0.051)
Correct Prior on Plenary Performance	0.250*** (0.059)	0.224*** (0.063)		
Correct Prior on Committee Performance			0.423*** (0.065)	0.219*** (0.061)
Constant	0.357*** (0.037)	0.404*** (0.047)	0.312*** (0.035)	0.455*** (0.045)
Observations	474	403	475	401
R-squared	0.046	0.033	0.082	0.072

Dependent variable takes a value of 1 if the respondent accurately reported the incumbent's performance score in plenary and committee during the comprehension survey immediately following the treatment. The sample only includes those who were randomly selected to participate in the comprehension survey. Standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

negative information about a coethnic).³⁹ To test this conjecture, we create a dummy variable for each comprehension question that takes a value of 1 if the respondent provided any substantive response (a sign of learning); 0 if the respondent replied “Don’t Know” (a signal the respondent did not learn). We expect more “Don’t Know” responses when negative information about a coethnic is provided. Similarly, we expect fewer “Don’t Know” responses when positive information about a coethnic is provided.

The results in Table 5 are fairly consistent with the expectations generated by the motivated reasoning hypothesis. First, coethnics of the incumbent are more likely to learn when the information is positive, although the associations are not significant (columns 1 and 3). Second, and importantly, coethnics of the incumbent are significantly *less* likely to learn (and thus more likely to provide a “Don’t Know” response) when the information is negative (columns 2 and 4), and here, the effect achieves statistical significance. The results thus provide some evidence that coeth-

³⁹This should be particularly true in a context like ours where legislative performance information is scarce: more than 50% of baseline survey respondents said “Don’t Know” when asked about legislative performance.

Table 5: The effect of coethnicity on respondent learning from our treatment information

VARIABLES	(1) Good Performance Plenary	(2) Bad Performance Plenary	(3) Good Performance Committee	(4) Bad Performance Committee
Coethnic with incumbent	0.032 (0.039)	-0.061* (0.035)	0.048 (0.039)	-0.065* (0.035)
Correct Prior on Plenary Performance	0.209*** (0.050)	0.074* (0.041)		
Correct Prior on Committee Performance			0.262*** (0.058)	0.061 (0.042)
Constant	0.714*** (0.031)	0.909*** (0.031)	0.712*** (0.031)	0.919*** (0.031)
Observations	474	403	475	401
R-squared	0.037	0.017	0.044	0.016

Dependent variable takes a value of 1 if the respondent provided a response to the questions about the incumbent's performance score in plenary and committee during the comprehension survey immediately following the treatment. The dependent variable instead takes a value of 0 if the respondent replied "Don't Know." The sample only includes those who were randomly selected to participate in the comprehension survey. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

nics and non-coethnics of the incumbent process information differently and in ways that reinforce preferences and beliefs about coethnics.⁴⁰

⁴⁰We cannot rule out the possibility that respondents did process the information but were unwilling to report it. But we note that both mechanisms – not accepting the information, and accepting the information but down-weighting it entirely – are consistent with motivated reasoning.

Discussion

Taken together, our results are most consistent with the hypothesis that ethnically motivated reasoning describes how voters respond to information in Benin. This section considers potential challenges to such an interpretation.

First, although access to information in our experiment is randomly assigned, coethnicity with the incumbent is not. Coethnicity with the incumbent might therefore be correlated with omitted variables that are driving the results. For example, if the incumbent's coethnics are more educated or live in more economically developed areas, these differences may be responsible for the conditional treatment effects we identify. However, as Table B.1 in Appendix B shows, coethnics and non-coethnics in our sample look remarkably similar on a range of socio-economic indicators, including age, years of education, access to household electricity, and access to schools, clinics, and boreholes in their community. There are no statistically significant or substantively important differences between coethnics and non-coethnics on these measures.⁴¹ In addition, the results in Appendix B.3 show that our main results are robust to the inclusion of controls for these covariates in our main regression models. Finally, we note that potential omitted variables must account for the asymmetry in the conditional treatment effects we uncover in order to threaten our interpretation. For example, even if it were the case that coethnics were more educated and wealthier than non-coethnics, one would need an explanation for why education or wealth leads voters to reward good performance but *not* to punish poor performance.

We also consider the possibility that our results would also obtain if ethnicity were just a proxy for prior partisanship; in other words, voters might use the information to rationalize previous decisions such that prior partisans are more likely to vote for the incumbent when they hear good news and prior non-partisans are less likely to vote for the incumbent when they hear bad news.

⁴¹Where slight differences exist, they run in opposite directions. For example, coethnics are slightly more educated (3.5 versus 3.27 average years of schooling), but live in villages that are slightly less likely to have schools, boreholes, or health clinics.

Although partisanship and ethnic group membership are highly correlated in many contexts (Casey 2015; Michelitch 2015), we do not find that partisan rather than ethnically motivated reasoning accounts for our results. We created a dummy variable capturing whether each survey respondent voted for the political party of the incumbent in the previous legislative elections held in 2011. Table B.3 in Appendix B.3 shows that our main results are robust to controlling for this variable. Table B.2 in the same appendix further shows that our results are robust to controlling for this measure and for an interaction between treatment and voting for the incumbent in 2011. The results from this analysis also provide no evidence of partisan motivated reasoning. This pattern is consistent with the relatively weak partisan identification and party system fluidity that characterizes Benin, as well as with the results of Conroy-Krutz and Moehler (2015), which finds no evidence of partisan motivated reasoning in response to the media in another West African democracy, Ghana.

Another consideration relates to Benin’s multi-member-district electoral system, in which voters technically cast their ballot for the party list, rather than an individual candidate. However, we do not believe that this aspect of our context threatens our interpretation of the results for two reasons. First, as we note above, voting in legislative elections is quite personalistic in Benin (Hounkpe and Warren 2012), despite the electoral system, and we conducted significant pre-experiment research to identify the incumbents most associated with the communes in our sample. Second, if it were the case that voters thought more about parties than about candidates, we expect that, if anything, this feature would have made it difficult for us to detect *any* treatment effects, as voters should have been unlikely to act on performance information about one member of a party list. This is even more true given that party lists in our sample are more often than not multi-ethnic themselves. Additionally, the feature seems unlikely to account for the asymmetrical effects we found.

A third set of concerns relates to the ethnic composition of the other party lists competing in the election. In particular, voter response to performance information about a coethnic or non-coethnic might depend on the availability of other coethnic candidates competing in the election or on the ethnic composition of the incumbent’s party list. For example, voters may be willing

to punish coethnics for poor performance if there are other coethnics competing in the election, while they might not be willing to do so if there are no other coethnics competing. While party list leaders are more or less concentrated within a single ethnicity in our sample constituencies, in every case, there is at least one other party list head of the same ethnicity as our sample politician.

This line of argument suggests that ethnic demographics in a constituency should condition our treatment effects, as the ethnic composition of party lists is likely to be highly correlated with the ethnic composition of the constituency. To examine this possibility, we test whether treatment effects vary when voters are in the ethnic majority or ethnic minority in their commune.⁴² The results of these analyses, presented in Table B.6 of Appendix B.5, provide no evidence that ethnic demography in the commune moderates the impact of treatment. The incumbent's coethnics react similarly to the information when they are in the ethnic majority in the commune, and are thus more likely to have other coethnics competing in the elections, and when they are in the ethnic minority, and are thus less likely to have other coethnic options. The same patterns hold for non-coethnics, who are indifferent to good performance information and punish poor performance information regardless of whether they are a member of an ethnic majority or minority in the commune. While these results are inconsistent with the notion that voters respond strategically to performance information in ways that vary with local ethnic demography, they continue to be consistent with ethnically motivated reasoning.

Finally, neither the electoral system nor ethnic demography can explain the results of our comprehension survey, which show that coethnics and non-coethnics of the incumbent process information differently immediately after treatment. Those findings are consistent with the ethnically motivated reasoning mechanism and are unlikely to be explained by these alternative accounts.

⁴²Recall that inference is being made within constituencies, so we hold ethnic composition constant but can assess the relationship between information and ethnicity across different types of political environments.

Conclusion

Using attitudinal and behavioral data collected as part of a large-scale field experiment around Benin’s 2015 legislative elections, this paper investigated the relationship between information about candidate quality and ethnic voting. We find that positive information increases support for the incumbent among coethnics but has no impact on the voting behavior of non-coethnics, while negative information has no impact on the voting behavior of the incumbent’s coethnics, while non-coethnics punish incumbents for poor performance. In other words, in this context, voters respond to new information about incumbent performance only if they share an ethnic tie with the incumbent and the news is positive, or if they are not members of the incumbent’s ethnic group and the news is negative. We find remarkably consistent effects across both attitudinal measures of ethnic voting via a panel survey and official electoral results, and importantly, controlling for respondent priors about incumbent performance. We examine three sets of potential explanations for these results. Our evidence is most consistent with arguments that emphasize how social identity theory can shape information processing. In particular, the results suggest that many voters engage in ethnically motivated reasoning (Bolsen et al. 2014; Taber and Lodge 2006). *The overall pattern is that the introduction of information about candidate quality reinforced preferences for coethnic candidates.*

While we present field experimental evidence from Benin, we expect our results could extend to other democracies where ethnic identities are salient and where voter access to information about incumbents is limited. These conditions characterize many of the world’s democracies, including those in sub-Saharan Africa (Posner 2005), Asia (Chandra 2004), and Eastern Europe and Central Asia (Hale 2008). In addition, we do not believe that there is anything “special” about ethnic identity, relative to other social identities, that is producing our results. Rather, ethnic identity happens to be the salient social identity in our study context. Thus, to the extent that other types of identities — partisan, class, religious, and so on — are more salient in other contexts and may also impact how voters process information, our findings may also extend to contexts where non-

ethnic differences structure electoral competition. Indeed, as mentioned, our results resonate with a body of research in American politics that investigates how partisan identities shape information processing and knowledge about politics.

That said, it is also possible that ours is not the strongest test of the ethnicity-as-heuristic for favoritism hypothesis, since we provided information about the incumbent's efforts regarding legislation and executive oversight rather than the incumbent's performance in distributing patronage and material transfers. Had we provided information about the latter, it is possible that performance information would have interacted with ethnicity in ways more consistent with the ethnicity-as-heuristic for favoritism hypothesis. However, it is important to note that deputies can and do favor their coethnics through their activities in the legislature as well, for example by pursuing policies that favor the interests of particular ethnic groups or working to direct central government spending to coethnic constituents. Moreover, we emphasize that our goal is *not* to rule out the ethnicity-as-heuristic for favoritism completely. Rather, our claim is that this argument cannot account for our experimental results. Thus, we argue that the comparative politics literature should at the very least pay greater attention to the ways in which information might reinforce ethnic voting in some contexts. Indeed our findings raise the possibility that ethnically motivated reasoning might also condition how voters respond to information about patronage and material transfers, and we believe that this is an important area for future research.⁴³

Indeed, our paper makes a number of contributions. First, our results have implications for the literature on ethnic voting. We evaluate several theories about the relationship between co-ethnicity, performance information, and voting, and provide evidence that is most consistent with our argument that ethnic identity can shape information processing. Our results thus emphasize a mechanism that has garnered attention with regard to partisanship in American politics and more

⁴³Consistent with this idea, Carlson (2016) finds that supporters of Ugandan President Yoweri Museveni systematically overestimate the extent to which the government provides resources to them, while supporters of the opposition underestimate what they have received from government.

recent attention with regard to partisanship in developing democracies. In showing that ethnic identity also shapes how individuals process information about politics, and that it does so in ways that may serve to amplify voter preferences for coethnic candidates, our paper shows that ethnic voting may persist in part because identity conditions how voters respond to and process information about the political world. Thus, while the literature on ethnic politics has recently emphasized the rational and instrumental considerations that drive ethnic voting (e.g. Carlson 2015; Conroy-Krutz 2013; Ichino and Nathan 2013; Posner 2005), our results highlight important psychological dimensions of ethnic politics that reinforce ethnic voting. Our results also contribute to the literature on information and accountability. This literature has produced mixed results on the effects of information access and political behavior. Our results suggest that ethnic identity is an important moderator that should be investigated in future research.

Finally, there are policy and normative implications of our study. On the one hand, one could reasonably read our findings as being quite troublesome for the prospects of democracy in societies where ethnic or other social identities are salient in politics. Indeed our results show that improving voter access to information actually amplified, rather than diminished, voter preferences for coethnic candidates, echoing concerns about attitude polarization in other contexts (Taber and Lodge 2006). On the other hand, there is a more optimistic reading. To identify solutions, we must truly understand the challenge, and research in political science and related fields suggest a number of ways that ethnically motivated reasoning might be diminished. For example, the cultivation of overarching identities, such as national identities (Miguel and Gugerty 2005; Robinson 2016), could serve to mitigate ethnically motivated reasoning. Similarly, the priming of cross-cutting identities, such as cousinage over ethnicity in Mali (Dunning and Harrison 2010), could pivot politics toward more encompassing social categories. Finally, although the evidence for this remains nascent, research on discrimination suggests that individuals can overcome their biases when they are made aware of them. In one striking example, racial bias in the calls made by referees in the National Basketball Association seems to have disappeared in the aftermath of the media attention that was garnered by the original study demonstrating discrimination (Pope et al. 2013; Price and

Wolfers 2010). Widespread media attention, including front-page coverage in the New York Times and significant coverage in major news networks as well as in the sports media, worked to reduce bias in one context. Future work should investigate the effects of information dissemination that is accompanied by an awareness-raising campaign about the prevalence and implications of ethnic voting.

References

- Adida, C. L. (2015). Do african voters favor coethnics? Evidence from a survey experiment in benin. *Journal of Experimental Political Science* 2(1).
- Alesina, A., R. Baqir, and W. Easterly (1999). Public goods and ethnic divisions. *The Quarterly Journal of Economics* 114(4), 1243–1284.
- Banerjee, A., S. Kumar, R. Pande, and F. Su (2011). Do informed voters make better choices? experimental evidence from urban india. *Unpublished manuscript*.
- Bartels, L. M. (2002). Beyond the running tally: Partisan bias in political perceptions. *Political Behavior* 24(2), 117–150.
- Bates, R. H. (1983). Modernization, ethnic competition, and the rationality of politics in contemporary africa. In D. S. Rothchild and V. A. Olorunsola (Eds.), *State versus ethnic claims: African policy dilemmas*, pp. 171. Boulder, CO: Westview Press.
- Bidwell, K., K. Casey, and R. Glennerster (2015). Debates: Voting and expenditure responses to political communication. *Working Paper Stanford University*.
- Birnir, J. (2007). *Ethnicity and Electoral Politics*. New York: Cambridge University Press.
- Bolsen, T., J. N. Druckman, and F. L. Cook (2014). The influence of partisan motivated reasoning on public opinion. *Political Behavior* 36(2), 235–262.
- Bratton, M. and M. S. Kimenyi (2008). Voting in kenya: Putting ethnicity in perspective. *Journal of Eastern African Studies* 2(2), 272–289.
- Burgess, R., R. Jedwab, E. Miguel, A. Morjaria, and G. Padro i Miquel (2015). The value of democracy: Evidence from road building in kenya. *American Economic Review* 105(6), 1817–51.
- Butler, D. M. and D. E. Broockman (2011). Do politicians racially discriminate against constituents? a field experiment on state legislators. *American Journal of Political Science* 55(3), 463–477.
- Campbell, A., P. E. Converse, W. E. Miller, and D. E. Stokes (1960). *The American Voter*. New York: John Wiley.
- Carlson, E. (2015). Ethnic voting and accountability in africa: A choice experiment in uganda. *World Politics* 67(02), 353–385.
- Carlson, E. (2016). Finding partisanship where we least expect it: Evidence of partisan bias in a new African democracy. *Political Behavior* 38(1), 129–154.
- Casey, K. (2015). Crossing party lines: The effects of information on redistributive politics. *The American Economic Review* 105(8), 2410–2448.

- Chandra, K. (2004). *Why ethnic parties succeed: Patronage and ethnic head counts in India*. New York: Cambridge University Press.
- Chandra, K. (2006). What is ethnic identity and does it matter? *Annu. Rev. Polit. Sci.* 9, 397–424.
- Chauchard, S. (2016). Unpacking ethnic preferences theory and micro-level evidence from north india. *Comparative Political Studies* 49(2), 253–284.
- Chong, A., L. Ana, D. Karlan, and L. Wantchekon (2015). Does corruption information inspire the fight or quash the hope? a field experiment in mexico on voter turnout, choice, and party identification. *The Journal of Politics* 77(1), 55–71.
- Conroy-Krutz, J. (2013). Information and ethnic politics in africa. *British Journal of Political Science* 43(02), 345–373.
- Conroy-Krutz, J. and D. C. Moehler (2015). Moderation from bias: A field experiment on partisan media in a new democracy. *Journal of Politics* 77(2), 575–587.
- Dahl, R. A. (1973). *Polyarchy: Participation and opposition*. Yale University Press.
- Dunning, T. and L. Harrison (2010). Cross-cutting cleavages and ethnic voting: an experimental study of cousinage in mali. *American Political Science Review* 104(1).
- Easterly, W. and R. Levine (1997). Africa's growth tragedy: Policies and ethnic divisions. *Quarterly Journal of Economics* 112(4), 1203–1250.
- Fearon, J. (1999). Electoral accountability and the control of politicians: Selecting good types versus sanctioning poor performance. In A. Przeworski, S. C. Stokes, and B. Manin (Eds.), *Democracy, accountability, and representation*, pp. 55–97. Cambridge: Cambridge University Press.
- Ferraz, C. and F. Finan (2008). Exposing corrupt politicians: The effects of brazil's publicly released audits on electoral outcomes. *The Quarterly Journal of Economics* 123(2), 703–745.
- Ferree, K. E. (2006). Explaining south africa's racial census. *Journal of Politics* 68(4), 803–815.
- Ferree, K. E. (2010). *Framing the race in South Africa: The political origins of racial census elections*. New York, NY: Cambridge University Press.
- Fiske, S. T., A. J. Cuddy, and P. Glick (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in cognitive sciences* 11(2), 77–83.
- Franck, R. and I. Rainer (2012). Does the leader's ethnicity matter? ethnic favoritism, education, and health in sub-saharan africa. *American Political Science Review* 106(02), 294–325.
- Habyarimana, J., M. Humphreys, D. N. Posner, and J. M. Weinstein (2009). *Coethnicity: Diversity and the dilemmas of collective action*. New York, NY: Russell Sage Foundation.
- Hale, H. E. (2008). *The Foundations of Ethnic Politics: Separatism of States and Nations in Eurasia and the World*. New York: Cambridge University Press.

- Harris, A. S. and M. G. Findley (2014). Is ethnicity identifiable? lessons from an rxperiment in south africa. *Journal of Conflict Resolution* 58(1), 4–33.
- Heath, O., G. Verniers, and S. Kumar (2015). Do muslim voters prefer muslim candidates? co-religiosity and voting behaviour in india. *Electoral Studies* 38, 10–18.
- Hodler, R. and P. A. Raschky (2014). Regional favoritism. *The Quarterly Journal of Economics* 129(2), 995–1033.
- Horowitz, D. L. (1985). *Ethnic groups in conflict*. Berkeley, CA: University of California Press.
- Horowitz, J. and J. Long (2016). Strategic voting, information, and ethnicity in emerging democracies: Evidence from Kenya. *Electoral Studies* 44(351-361).
- Hounkpe, M. and S. Warren (2012). Alp benin interim country report. Technical report, African Legislatures Project.
- Huckfeldt, R., J. Levine, W. Morgan, and J. Sprague (1999). Accessibility and the political utility of partisan and ideological orientations. *American Journal of Political Science*, 888–911.
- Humphreys, M. and J. Weinstein (2012). Policing politicians: Citizen empowerment and political accountability in uganda preliminary analysis. *International Growth Centre Working Paper S-5021-UGA-1*.
- Hutchings, V. L. and N. A. Valentino (2004). The centrality of race in american politics. *Annual Review of Political Science* 7, 383–408.
- Ichino, N. and N. L. Nathan (2013). Crossing the line: Local ethnic geography and voting in ghana. *American Political Science Review* 107(02), 344–361.
- Kahan, D. M. (2016). The politically motivated reasoning paradigm: What politically motivated reasoning is and how to measure it. In R. Scott and S. Kosslyn (Eds.), *Emerging Trends in the Social and Behavioral Sciences: An Interdisciplinary, Searchable, and Linkable Resource*. Wiley Online Library.
- Kasara, K. (2007). Tax me if you can: Ethnic geography, democracy, and the taxation of agriculture in africa. *American Political Science Review* 101(01), 159–172.
- Kramon, E. and D. N. Posner (2011). Kenya's new constitution. *Journal of Democracy* 22(2), 89–103.
- Kramon, E. and D. N. Posner (2013). Who benefits from distributive politics? how the outcome one studies affects the answer one gets. *Perspectives on Politics* 11(02), 461–474.
- Kramon, E. and D. N. Posner (2016). Ethnic favoritism in primary education in kenya. *Quarterly Journal of Political Science* 11, 1–58.
- Kunda, Z. (1987). Motivated inference: Self-serving generation and evaluation of causal theories. *Journal of personality and social psychology* 53(4), 636–47.

- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin* 108(3), 480–98.
- Kunda, Z. (1999). *Social cognition: Making sense of people*. Cambridge, MA: MIT press.
- Lieberman, E. S. (2009). *Boundaries of contagion: How ethnic politics have shaped government responses to AIDS*. Princeton, NJ: Princeton University Press.
- Lieberman, E. S., D. N. Posner, and L. L. Tsai (2014). Does information lead to more active citizenship? evidence from an education intervention in rural kenya. *World Development* 60, 69–38.
- Lijphart, A. (1977). *Democracy in Plural Societies: A Comparative Exploration*. Yale University Press.
- McClendon, G. H. (2016). Race and responsiveness: An experiment with south african politicians. *Journal of Experimental Political Science* 3, 60–74.
- Michelitch, K. (2015). Does electoral competition exacerbate interethnic or interpartisan economic discrimination? Evidence from a field experiment in market price bargaining. *American Political Science Review* 109(1), 43–61.
- Miguel, E. (2004). Tribe or nation?: Nation building and public goods in Kenya versus Tanzania. *World Politics* 56(3), 327–362.
- Miguel, E. and M. K. Gugerty (2005). Ethnic diversity, social sanctions, and public goods in Kenya. *Journal of Public Economics* 89(11-12), 2325–2368.
- Pande, R. (2003). Can mandated political representation increase policy influence for disadvantaged minorities? theory and evidence from india. *The American Economic Review* 93(4), 1132–1151.
- Pande, R. (2011). Can informed voters enforce better governance? experiments in low-income democracies. *Annual Review of Economics* 3(1), 215–237.
- Pope, D. G., J. Price, and J. Wolfers (2013). Awareness reduces racial bias. *NBER Working Paper 19765*.
- Popkin, S. L. (1991). *The reasoning voter: Communication and persuasion in presidential elections*. Chicago, IL: University of Chicago Press.
- Posner, D. N. (2005). *Institutions and ethnic politics in Africa*. New York, NY: Cambridge University Press.
- Preuhs, R. R. (2006). The conditional effects of minority descriptive representation: Black legislators and policy influence in the american states. *Journal of Politics* 68(3), 585–599.
- Price, J. and J. Wolfers (2010). Racial discrimination among nba referees. *The Quarterly Journal of Economics* 125(4), 1859–1887.
- Rabushka, A. and K. A. Shepsle (1972). *Politics in Plural Societies*. Citeseer.

- Robinson, A. (2016). Nationalism and ethnic-based trust: Evidence from an African border region. *Comparative Political Studies*.
- Seely, J. C. (2001). A political analysis of decentralisation: coopting the tuareg threat in mali. *The Journal of Modern African Studies* 39(03), 499–524.
- Taber, C. S. and M. Lodge (2006). Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science* 50(3), 755–69.
- Tajfel, H. (1974). *Social identity and intergroup behaviour*. Cambridge, UK: Cambridge University Press.
- Wantchekon, L. (2003). Clientelism and voting behavior: Evidence from a field experiment in benin. *World Politics* 55(3), 399–422.
- Westen, D., P. S. Blagov, K. Harenski, C. Kilts, and S. Hamann (2006). Neural bases of motivated reasoning: An fmri study of emotional constraints on partisan political judgment in the 2004 us presidential election. *Journal of cognitive neuroscience* 18(11), 1947–1958.
- Zaller, J. (1992). *The nature and origins of mass opinion*. Cambridge university press.

Supplemental Materials

A Balance Tests	50
B Additional Results and Robustness Checks	52
B.1 Socio-Economic Characteristics of Incumbent Coethnics and Non-Coethnics	52
B.2 Distribution of Candidate Ethnicities	53
B.3 Interaction Models and Control Variables	55
B.4 Treatment Effects Among Members of the Ethnic Minority in Villages	58
B.5 Treatment Effects Among Members of the Ethnic Minority in Communes	59
B.6 Robustness to Alternative Thresholds of Coethnicity	60
B.7 Voter Turnout Results	61
B.8 Multiple Comparisons Corrections	64
B.9 Addressing differential attrition and response bias in survey	68
B.9.1 Differential attrition	68
B.9.2 Social desirability bias	73
C Further Details on Experimental Design	76
C.1 Geographic Sampling and Randomization Procedures	76
C.2 Survey Sampling Procedures	78
C.3 Additional Notes on Ethical Considerations	79
C.4 Treatment Conditions	79
C.4.1 Information Only Condition Script (in English)	80
C.4.2 Civics Condition Script (in English)	83

C.4.3 Dosage Treatment	90
C.4.4 Linking the Results to the Companion Paper	92

A Balance Tests

Table A.1: Balance Between High and Low Dosage Communes

	Mean High Dosage	Mean Low Dosage	Difference	P-Value
Registered Voters (log)	6.3	6.52	.22	.18
Urban	.21	.25	.04	.41
Competitive (dichotomous)	.42	.5	.07	.25
Vote Margin	.28	.24	.05	.19
Overall Performance	4.97	5.35	.38	.67

P-values generated from tests in which we cluster on commune.

Table A.2: Balance in High Dosage Communes

	Control	Info Only/Private	Info Only/Public	Civics/Private	Civics/Public
Registered Voters (log)	687.5	828.26 (.08)	1066.25 (.06)	1110.73 (.00)	807.27 (.02)
Urban	.18	.03 (.00)	.34 (.00)	.27 (.07)	.29 (.01)
Competitive (dichotomous)	.41	.49 (.11)	.48 (.30)	.5 (.14)	.48 (.25)
Vote Margin	.29	.25 (.14)	.23 (.05)	.26 (.44)	.31 (.47)
Overall Performance	4.89	5.26 (.14)	5.19 (.21)	5.23 (.17)	5.23 (.17)

P-values in parentheses indicate significance of difference between the mean and each treatment group and the control group mean.

P-values generated from tests in which we cluster on commune.

5

Note: Because of our blocking and randomization process, there is a lack of balance in the raw means on urban and number of registered voters. This occurred because our rural blocks, where there are also fewer registered voters, contain larger numbers of units than our urban blocks. Since all non-treated units are used as controls, on average the proportion of rural areas in control is lower than in treatment. This lack of balance is not a problem as we use block fixed effects in all of our analyses, which controls for the urban/rural difference.

Table A.3: Balance in Low Dosage Communes

	Mean Treatment	Mean Control	Difference	P-Value
Registered Voters (log)	933.27	1024.24	90.97	.76
Urban	.27	.25	.01	.92
Competitive (dichotomous)	.53	.5	.04	.78
Vote Margin	.25	.24	.01	.86
Overall Performance	5.42	5.35	.07	.9

B Additional Results and Robustness Checks

B.1 Socio-Economic Characteristics of Incumbent Coethnics and Non-Coethnics

Table B.1: Socio-Economic Characteristics of Incumbent Coethnics and Non-Coethnics

	Coethnics	Non-Coethnics	Difference	P-Value
Age	36.47	36.01	.46	.41
Yrs of Education	3.5	3.27	.23	.27
Electricity in Household	.32	.34	.02	.47
Quality of Housing Material	2.1	2.07	.03	.22
School in Village	.9	.94	.05	.1
Borehole in Village	.89	.91	.02	.43
Health Clinic in Village	.47	.54	.06	.22

P-values generated from tests in which we cluster on village.

B.2 Distribution of Candidate Ethnicities

Figure B.1: Distribution of Candidate Ethnicities by Electoral Circonscription

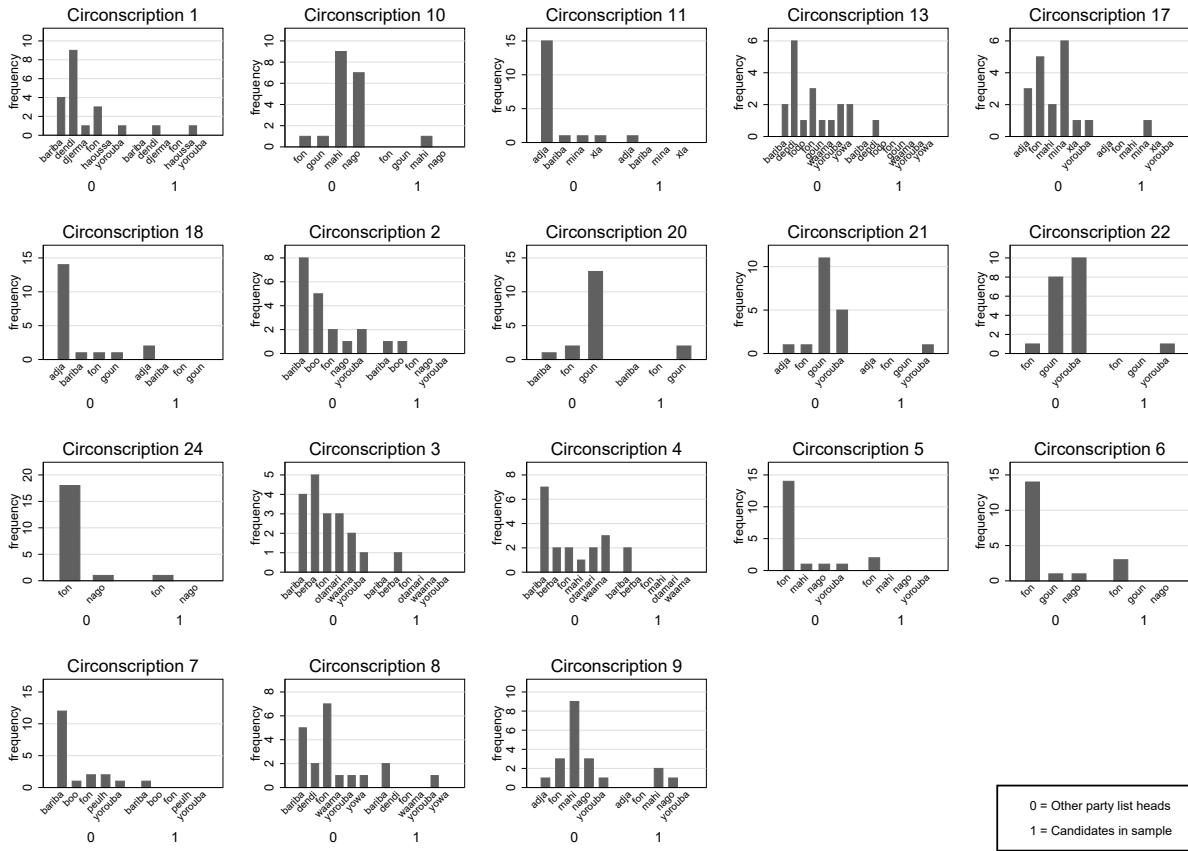
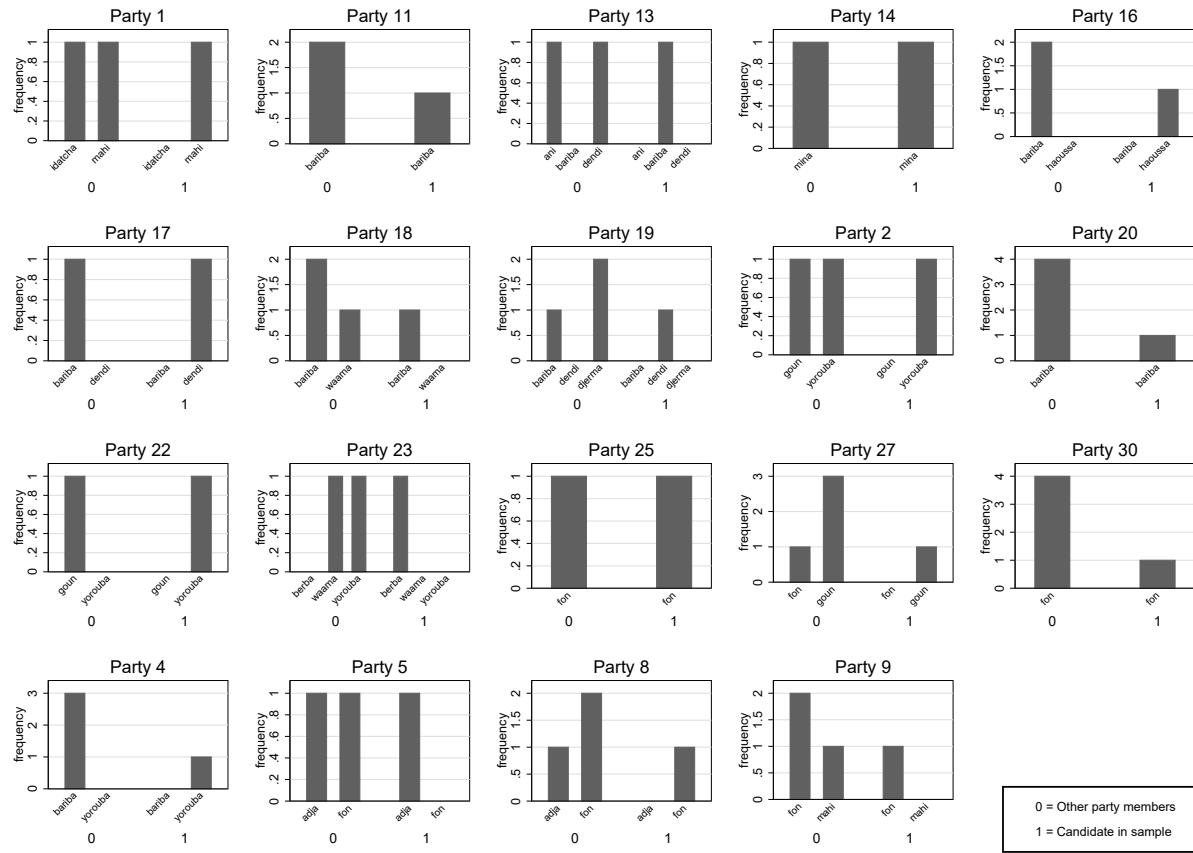


Figure B.2: Distribution of Candidate Ethnicities by Party List



Note: This figure shows 19 of the 30 party lists for which we were able to code candidate ethnicity for more than our sample politician.

B.3 Interaction Models and Control Variables

Table B.2: Interaction Models and Models with Controls, Survey Data

VARIABLES	(1) Good	(2) Good	(3) Bad	(4) Bad	(5) Good	(6) Bad
Treatment	0.007 (0.052)	-0.001 (0.059)	-0.169* (0.089)	-0.210** (0.090)	0.012 (0.070)	-0.244** (0.114)
Treatment x Coethnicity	0.110* (0.066)	0.132* (0.071)	0.124 (0.094)	0.122 (0.086)	0.111 (0.077)	0.204* (0.113)
Coethnic with incumbent	-0.124** (0.060)	-0.145** (0.066)	-0.107 (0.090)	-0.094 (0.077)	-0.132* (0.070)	-0.192* (0.112)
Age		0.001 (0.001)		-0.000 (0.001)		
Years of education			-0.000 (0.003)		-0.006 (0.004)	
electricity_type			-0.023 (0.029)		-0.016 (0.045)	
house_poverty			0.045* (0.024)		-0.008 (0.044)	
School in the community			-0.110 (0.071)		-0.074 (0.106)	
Borehole in the community			0.044 (0.081)		0.073 (0.095)	
Health clinic in the community			0.017 (0.036)		0.045 (0.069)	
Treatment x 2011 Incumbent Vote					0.013 (0.069)	-0.002 (0.104)
votedincumbent2011					0.000 (0.070)	0.132 (0.097)
Constant	0.517*** (0.044)	0.450*** (0.124)	0.642*** (0.081)	0.652*** (0.151)	0.519*** (0.066)	0.679*** (0.102)
Observations	1,657	1,444	1,356	812	1,342	1,100
R-squared	0.184	0.197	0.246	0.248	0.191	0.234

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table B.3: Treatment Effects by Ethnic Connection and Level of Performance Controlling for Prior Partisanship, Survey Data

VARIABLES	(1) Good Info Full Sample	(2) Good Info Coethnic	(3) Good Info Non-Coethnic	(4) Bad Info Full Sample	(5) Bad Info Coethnic	(6) Bad Info Non-Coethnic
Treatment	0.08** (0.04)	0.14*** (0.05)	0.02 (0.07)	-0.10* (0.06)	-0.04 (0.06)	-0.28*** (0.08)
2011 Incumbent Vote	0.02 (0.04)	0.04 (0.05)	-0.05 (0.06)	0.06* (0.04)	0.08* (0.04)	0.03 (0.07)
Constant	0.43*** (0.04)	0.37*** (0.05)	0.53*** (0.07)	0.55*** (0.06)	0.49*** (0.06)	0.69*** (0.08)
Observations	1,353	845	497	1,102	782	318
R-squared	0.19	0.23	0.19	0.23	0.25	0.26

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table B.4: Interaction Models, Official Results

VARIABLES	(1) Good Performance	(2) Bad Performance	(3) Good Performance	(4) Bad Performance	(5) Good Performance	(6) Bad Performance
Treatment	-0.060 (0.057)	-0.201 (0.123)	0.007 (0.056)	-0.223** (0.096)	0.019 (0.047)	-0.184** (0.086)
Treatment x Coethnic (50)	0.155** (0.068)	0.164 (0.155)				
Coethnic village (50 percent)	-0.111 (0.079)	-0.124 (0.145)				
Treatment x Coethnic (70)			0.082 (0.075)	0.232* (0.121)		
Coethnic village (70 percent)			-0.159** (0.073)	-0.274* (0.145)		
Treatment x Coethnic (90)					0.072 (0.076)	0.203* (0.110)
Coethnic village (90 percent)					-0.096 (0.067)	-0.293*** (0.098)
Constant	0.545*** (0.057)	0.645*** (0.106)	0.545*** (0.049)	0.717*** (0.094)	0.503*** (0.036)	0.703*** (0.063)
Observations	137	103	137	103	137	103
R-squared	0.604	0.615	0.610	0.637	0.596	0.639

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

B.4 Treatment Effects Among Members of the Ethnic Minority in Villages

Table B.5: Comparing Treatment Effect Among Ethnic Majority and Minority in Village

VARIABLES	(1) Good Info Coethnic	(2) Good Info Non-Coethnic	(3) Bad Info Coethnic	(4) Bad Info Non-Coethnic
Treatment	0.12*** (0.05)	-0.04 (0.07)	-0.04 (0.06)	-0.21* (0.12)
Treatment x Ethnic Minority in Village	-0.02 (0.19)	0.18 (0.13)	0.08 (0.15)	0.13 (0.22)
Ethnic Minority in Village	-0.04 (0.19)	-0.02 (0.12)	-0.06 (0.13)	-0.16 (0.16)
Constant	0.41*** (0.04)	0.49*** (0.06)	0.56*** (0.06)	0.63*** (0.10)
Observations	1,030	627	948	408
R-squared	0.22	0.19	0.28	0.23

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

B.5 Treatment Effects Among Members of the Ethnic Minority in Communes

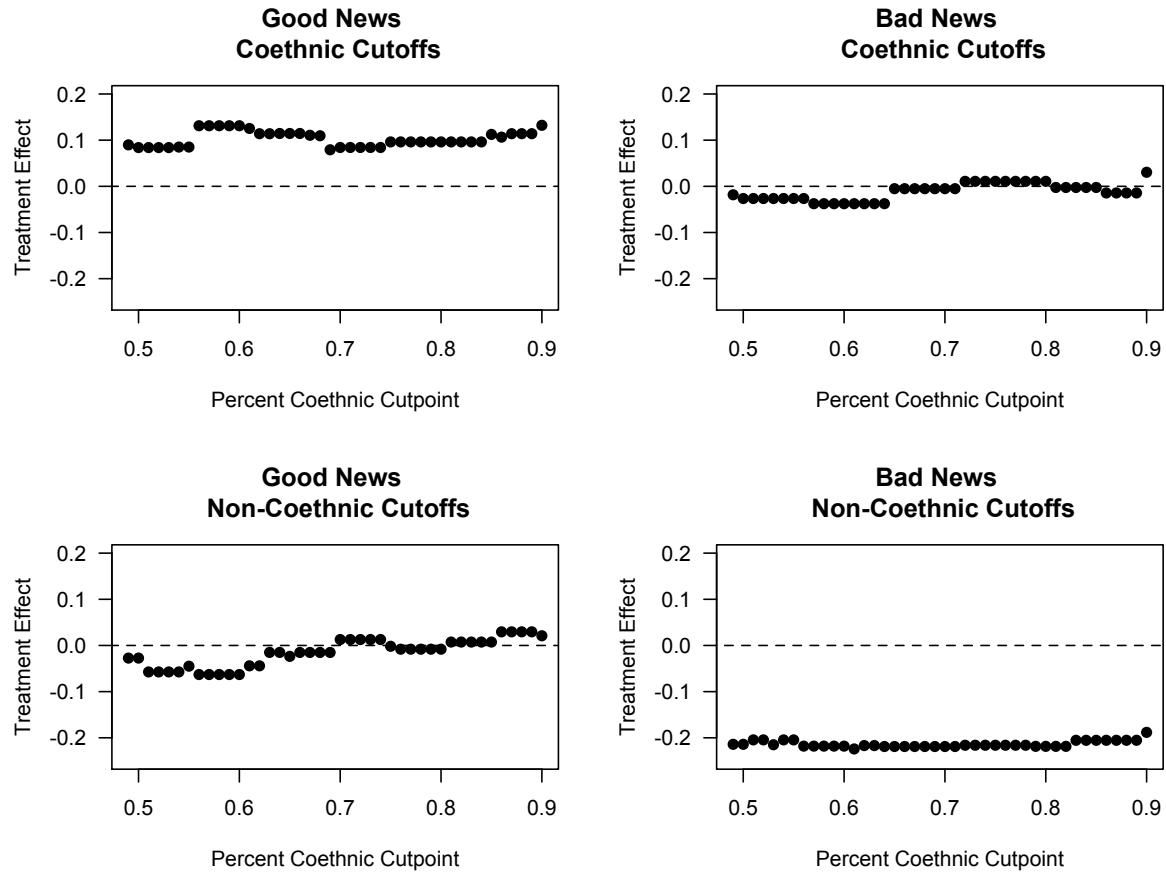
Table B.6: Comparing Treatment Effect Among Ethnic Majority and Minority in Commune

VARIABLES	(1) Good Info Coethnic	(2) Good Info Non-Coethnic	(3) Bad Info Coethnic	(4) Bad Info Non-Coethnic
Treatment	0.10** (0.05)	-0.06 (0.07)	-0.04 (0.06)	-0.08 (0.09)
Treatment x Ethnic Minority in Commune	0.15 (0.13)	0.16 (0.11)	0.06 (0.13)	-0.22 (0.16)
Constant	0.40*** (0.04)	0.50*** (0.05)	0.55*** (0.05)	0.63*** (0.08)
Observations	1,030	627	948	408
R-squared	0.22	0.19	0.28	0.23

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

B.6 Robustness to Alternative Thresholds of Coethnicity



B.7 Voter Turnout Results

This section presents the results of analyses in which the dependent variable is voter turnout rather than vote share. Table B.7 presents the results using official administrative election results. Positive performance information has no impact on turnout in the full sample. The impact of positive performance information is also not conditioned by coethnicity. Negative performance information, by contrast, increases voter turnout by about 2 percentage points in the full sample. There is some evidence that this effect is concentrated among non-coethnic villages, but none of the results is statistically significant.

Table B.8 presents results of the same tests using survey data. The evidence in this table should be interpreted with caution given well-known problems of social desirability bias in survey measures of voter turnout. Once again, we find no evidence that positive performance information has an effect on voter turnout, and the results do not vary by coethnicity. Negative performance information increases voter turnout, and the effect is similar for coethnics and non-coethnics.

In sum, we find evidence that negative information increases voter turnout, while positive information has no effect on participation. Coethnicity does not, however, moderate these effects. Indeed, in an interaction model using the full sample, the coefficient on the interaction term between treatment and coethnicity is statistically insignificant – in contrast to tests using vote share as an outcome. Together, this suggests that voter preferences rather than voter mobilization is the mechanism driving the ethnically motivated electoral results we discuss in the paper.

Table B.7: Voter Turnout: Treatment Effects by Ethnic Connection and Level of Performance, Official Results

VARIABLES	(1) Good Info Full Sample	(2) Good Info Coethnic (50)	(3) Good Info Non-Coethnic (50)	(4) Bad Info Full Sample	(5) Bad Info Coethnic (50)	(6) Bad Info Non-Coethnic (50)
Treatment	-0.01 (0.01)	-0.01 (0.02)	-0.00 (0.03)	0.02* (0.01)	0.01 (0.01)	-0.01 (0.03)
Constant	0.91*** (0.01)	0.89*** (0.01)	0.93*** (0.03)	0.89*** (0.01)	0.89*** (0.01)	0.91*** (0.02)
Observations	137	89	48	103	72	31
R-squared	0.52	0.56	0.32	0.67	0.71	0.81
VARIABLES	(1) Good Info Full Sample	(2) Good Info Coethnic (70)	(3) Good Info Non-Coethnic (70)	(4) Bad Info Full Sample	(5) Bad Info Coethnic (70)	(6) Bad Info Non-Coethnic (70)
Treatment	-0.01 (0.01)	-0.03 (0.03)	0.01 (0.02)	0.02* (0.01)	0.01 (0.02)	0.02 (0.02)
Constant	0.91*** (0.01)	0.92*** (0.02)	0.90*** (0.02)	0.89*** (0.01)	0.90*** (0.01)	0.89*** (0.01)
Observations	137	66	71	103	58	45
R-squared	0.52	0.47	0.72	0.67	0.78	0.66
VARIABLES	(1) Good Info Full Sample	(2) Good Info Coethnic (90)	(3) Good Info Non-Coethnic (90)	(4) Bad Info Full Sample	(5) Bad Info Coethnic (90)	(6) Bad Info Non-Coethnic (90)
Treatment	-0.01 (0.01)	-0.07 (0.05)	0.00 (0.02)	0.02* (0.01)	0.00 (0.03)	0.03 (0.02)
Constant	0.91*** (0.01)	0.94*** (0.04)	0.91*** (0.02)	0.89*** (0.01)	0.90*** (0.02)	0.88*** (0.02)
Observations	137	44	93	103	53	50
R-squared	0.52	0.64	0.58	0.67	0.79	0.59

Dependent variable is the voter turnout rate in the village/quartier. This table presents results using different cutpoints for defining a village as a coethnic village. In the top panel, villages are coded as coethnic if over 50 percent of survey respondents are coethnics of the incumbent. In the middle panel, villages are coded as coethnic if over 70 percent of survey respondents are coethnics of the incumbent. In the bottom panel, villages are coded as coethnic if over 90 percent of survey respondents are coethnics of the incumbent. Robust standard errors clustered by commune-treatment are in parentheses. Each model uses block fixed effects. *** p<0.01, ** p<0.05, * p<0.1

Table B.8: Voter Turnout: Treatment Effects by Ethnic Connection and Level of Performance, Survey Data

VARIABLES	(1) Good Info Full Sample	(2) Good Info Coethnic	(3) Good Info Non-Coethnic	(4) Bad Info Full Sample	(5) Bad Info Coethnic	(6) Bad Info Non-Coethnic
Treatment	-0.01 (0.03)	0.00 (0.04)	-0.02 (0.03)	0.04** (0.02)	0.03* (0.02)	0.03 (0.03)
Constant	0.89*** (0.03)	0.88*** (0.04)	0.90*** (0.02)	0.87*** (0.01)	0.88*** (0.02)	0.89*** (0.03)
Observations	1,906	1,176	715	1,507	1,051	453
R-squared	0.12	0.13	0.12	0.07	0.09	0.08

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

B.8 Multiple Comparisons Corrections

In our pre-analysis plan, we specified that we would first present unadjusted p values and then, for a family of hypotheses nested within one of the main theories, we would adjust p values for the multiple comparisons within that family. The key hypothesis in this paper about the moderating effect of coethnicity with the incumbent falls under our first family of hypotheses about the effect of information on voter behavior. Within this family, we registered 11 hypotheses. This subsection adjusts p values to account for these multiple comparisons.

Our pre-analysis plan specifies that we will adjust p values using a false discovery rate (FDR) correction to control the Type-1 error rate, and that we will control the FDR at the 0.05 significance level. This method requires ordering the hypothesis tests from smallest to largest nominal p value. Tables B.9 and B.10 do this for 9 of the 11 hypotheses in this family on the effect of good performance information and bad, respectively.⁴⁴

For the good performance information tests, none survive the p value adjustment exercise. The interaction with coethnicity ($p = 0.095$) as the second smallest p value, for example, would require $p < .05/2 = .025$ to be considered statistically significant. That said, the implications from the previous section are reinforced here: coethnicity remains the most statistically significant moderator of good performance information. For the poor performance information tests, the first three hypotheses survive the multiple comparison corrections. The test of coethnicity does not; as the fifth most significant test, the p value would have to be less than 0.01 for it to be considered statistically significant.

While we report these multiple comparison corrections for transparency, we also recognize they are very conservative by design, especially when accounting for the entire universe of possible tests in an experiment. We are reassured by the fact that the patterns we see in the survey data

⁴⁴The test of the hypothesis on the conditioning effect of list order of the incumbent drops out due to insufficient variation in the data. The test of the hypothesis on the conditioning effect of other media outlets also drops out because of poor data quality on accessible radio stations in each village.

with respect to the coethnicity hypothesis are almost exactly replicated in the same analyses using a different dataset. If random chance produced the relatively small p value on the coethnicity interaction term, then the probability that that same, relatively robust pattern shows up in a completely different dataset is exceedingly small.

Table B.9: Comparing Outcomes within the Same Family of Hypotheses, Good Performance Information

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7) Turnout	(8)	(9)
Treatment	0.067* (0.056)	0.007 (0.894)	0.114** (0.014)	0.105 (0.118)	0.055 (0.208)	0.011 (0.868)	-0.011 (0.703)	0.083* (0.090)	0.062* (0.069)
Treatment x Incumbent Coethnicity		0.110* (0.095)							
Coethnic with incumbent			-0.124** (0.041)						
Treatment x Female				-0.095 (0.118)					
Female					0.077 (0.166)				
Treatment x Incumbent Preference					-0.006 (0.587)				
Preference for Incumbent Party						-0.002 (0.877)			
Treatment x Education						0.003 (0.601)			
Years of education							-0.003 (0.595)		
Treatment x Positive Prior							0.031 (0.694)		
Positive Prior								-0.059 (0.385)	
Treatment x Electoral Competition								-0.014 (0.845)	
Electoral Competitiveness									-0.065 (0.440)
Treatment x Personal Help									-0.008 (0.957)
Received personal assistance from incumbent									0.031 (0.826)
Constant	0.448*** (0.000)	0.517*** (0.000)	0.409*** (0.000)	0.399*** (0.000)	0.459*** (0.000)	0.547*** (0.000)	0.890*** (0.000)	0.476*** (0.000)	0.451*** (0.000)
Observations	1,672	1,657	1,671	886	1,667	810	1,906	1,660	1,657
R-squared	0.181	0.184	0.182	0.209	0.181	0.239	0.115	0.181	0.180

Robust standard errors clustered by village; uncorrected p values in parentheses. Models include block fixed effects.

*** p<0.01, ** p<0.05, * p<0.1. Incumbent vote share is the DV in all but model 7.

Table B.10: Comparing Outcomes within the Same Family of Hypotheses, Poor Performance Information

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
			Turnout						
Treatment	-0.095 (0.105)	-0.008 (0.904)	0.039** (0.011)	-0.086 (0.121)	-0.169* (0.059)	-0.116** (0.050)	-0.140* (0.094)	-0.007 (0.897)	-0.090 (0.124)
Treatment x Positive Prior		-0.249*** (0.003)							
Positive Prior			0.135* (0.066)						
Treatment x Education				-0.020*** (0.005)					
Years of education					0.013** (0.038)				
Treatment x Incumbent Coethnicty						0.124 (0.187)			
Coethnict with incumbent							-0.107 (0.236)		
Treatment x Female								0.067 (0.288)	
Female									0.004 (0.937)
Treatment x Electoral Competition									0.121 (0.297)
Electoral Competitiveness									-0.100 (0.343)
Treatment x Incumbent Preference									-0.006 (0.519)
Preference for Incumbent Party									0.012 (0.155)
Treatment x Personal Help									0.023 (0.858)
Received personal assistance from incumbent									-0.068 (0.524)
Constant	0.543*** (0.000)	0.514*** (0.000)	0.875*** (0.000)	0.572*** (0.000)	0.642*** (0.000)	0.570*** (0.000)	0.617*** (0.000)	0.463*** (0.000)	0.578*** (0.000)
Observations	631	1,324	1,507	1,358	1,356	1,358	1,330	632	1,357
R-squared	0.238	0.244	0.074	0.243	0.246	0.247	0.251	0.267	0.245

Robust standard errors clustered by village; uncorrected p values in parentheses. Models include block fixed effects.

*** p<0.01, ** p<0.05, * p<0.1. Incumbent vote share is the DV in all but model 3.

B.9 Addressing differential attrition and response bias in survey

Using data from the endline survey to estimate treatment effects is subject to two potential problems of bias. First, there was substantial attrition between the baseline and endline surveys and differential attrition patterns across treatment and control could lead us to make biased inference. Second, the question of whether and for whom the respondent voted for is subject to social desirability bias. Depending on the respondent's perception of the enumerator and research project's partisan leanings, they may have been tempted to dissemble when answering the vote choice question. Our use of administrative data to evaluate impacts of treatment – and the similarity in findings across data sources mitigates the potential problems introduced by these sources of bias. However, we explore here the extent of the problems in the survey data and the direction of the bias potentially introduced.

B.9.1 Differential attrition

It is unsurprising that, with about half of the participants surveyed in person at baseline being unresponsive or unavailable for the telephone endline survey, the group of individuals who did participate would be different than those who attrited. Indeed, as shown in columns 1-2 of Table B.11, the endline participants are more likely to be male, better educated, and wealthier.⁴⁵ While these differences limit the generalizability of our findings to a more specific subset of Beninese, they do not necessarily imply problems of bias for making causal inference. That said, there is differential attrition across treatment and control groups that could induce such bias.

Examining key covariates collected on participants at baseline, we find that women are significantly more likely to attrit in treatment than in control and coethnics are significantly less likely to attrit in treatment than in control (see columns 3-6 of Table B.11). To the extent there is gender balance across ethnic groups, we are less concerned about our main findings being threatened by

⁴⁵Here, priors are measured on a 4-point scale where higher numbers indicate beliefs about better legislative performance. *Good News* is a binary indicator of whether the commune's deputy scored better than the local average on our performance index used in the treatment.

the first pattern; they could, however, be subject to bias as a result of the second.

Table B.11: Differential Rates of Attrition by Pre-treatment Covariates and Treatment

VARIABLES	(1) Endline	(2) Vote choice	(3) Endline	(4) Vote choice	(5) Endline	(6) Vote choice	(7) Endline	(8) Vote choice
Treatment	-0.039 (0.035)	-0.025 (0.032)	-0.077*** (0.029)	-0.062** (0.028)	0.037 (0.035)	0.035 (0.035)	-0.068 (0.051)	-0.060 (0.041)
Female	0.043*** (0.017)	0.060*** (0.017)	-0.008 (0.027)	0.026 (0.029)				
Coethnic with incumbent	-0.008 (0.026)	0.008 (0.026)			0.074** (0.035)	0.077** (0.034)		
Years of education	-0.006*** (0.002)	-0.004* (0.002)						
Urban	-0.069 (0.100)	-0.061 (0.091)						
Poverty Level	0.042** (0.016)	0.051*** (0.016)						
Positive Prior	0.016 (0.018)	0.001 (0.017)					0.002 (0.020)	-0.014 (0.020)
Positive Prior x Good News	-0.015 (0.024)	-0.015 (0.024)					0.015 (0.044)	-0.003 (0.044)
Treatment x Female		0.094*** (0.030)	0.067** (0.032)					
Treatment x Coethnicity				-0.110*** (0.040)	-0.103*** (0.038)			
Treatment x Positive Prior							0.032 (0.038)	0.033 (0.038)
Treatment x Good News							0.108 (0.139)	0.056 (0.137)
Treatment x Positive Prior x Good News							-0.037 (0.045)	-0.013 (0.046)
Constant	0.451*** (0.059)	0.517*** (0.057)	0.471*** (0.025)	0.515*** (0.024)	0.420*** (0.028)	0.479*** (0.029)	0.410*** (0.074)	0.533*** (0.072)
Observations	2,698	2,698	6,128	6,128	6,072	6,072	2,721	2,721
R-squared	0.129	0.109	0.121	0.101	0.115	0.094	0.121	0.099

In parentheses, robust standard errors clustered by village.
 Models include block fixed effects. *** p<0.01, ** p<0.05, * p<0.1.

We consider how the differential patterns of attrition with respect to ethnicity might bias the direction of our results. Among coethnics, treated individuals are less likely to attrit than people in the control group so these treated respondents are more representative of the population. The more narrow set of control group participants, being more urban, wealthy and more educated on average, may be more likely to vote for the incumbent either because they already know the incumbent is a good performer or because economic voting would lead to greater pro-incumbency among the wealthy. Both possibilities would bias against finding a treatment effect among good news communes; toward finding an effect in bad news communes. On the other hand, the narrower set of control participants may also be more likely to be critical of the incumbent – more educated citizens are often found to be more distrusting of government producing greater anti-incumbency. This possibility would bias us toward finding a treatment effect in good news communes; against finding an effect in bad news communes.

It is not obvious which direction we should expect the bias go in, so we control for predictors of attrition in the aggregate to mitigate the imbalance across the samples – at least on observable characteristics. Table B.12 shows that our main findings from Table 3 are robust to adding these controls.

Table B.12: Replicating Main Results Controlling for Predictors of Attrition

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	Good Info Full Sample	Good Info Coethnic	Good Info Non-Coethnic	Bad Info Full Sample	Bad Info Coethnic	Bad Info Non-Coethnic
Treatment	0.07* (0.04)	0.12*** (0.04)	0.00 (0.06)	-0.08 (0.06)	-0.03 (0.05)	-0.19** (0.09)
Female	-0.00 (0.02)	-0.01 (0.03)	0.02 (0.04)	0.06** (0.03)	0.08*** (0.03)	0.01 (0.05)
Years of education	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.01)
Poverty Level	-0.02 (0.02)	-0.03 (0.02)	0.02 (0.04)	-0.04 (0.03)	-0.06* (0.03)	0.02 (0.05)
Constant	0.45*** (0.03)	0.41*** (0.04)	0.49*** (0.05)	0.54*** (0.06)	0.50*** (0.05)	0.61*** (0.10)
Observations	1,667	1,028	625	1,324	914	408
R-squared	0.18	0.22	0.19	0.24	0.28	0.23

In parentheses, robust standard errors clustered by village.

Models include block fixed effects. *** p<0.01, ** p<0.05, * p<0.1.

B.9.2 Social desirability bias

We do find some evidence of potential response bias in reporting vote share for the incumbent in our survey measure. In treated areas that received bad news about the incumbent, the reported vote share is 8 percentage points lower than the official vote share – a pattern consistent with under-reporting of votes for the incumbent due to having received bad performance information about the incumbent. Similarly in good news areas, there is evidence consistent with over-reporting of votes for the incumbent.

These patterns are consequential for making inference about the overall effect of treatment from the survey data. However, if response bias affects ethnic groups equally, then the main findings of the paper that compare outcomes by ethnicity are less subject to concerns about biased interpretation. To test whether ethnic groups differentially generate response bias in the reporting of vote share for the incumbent, we test whether the coethnicity of the enumerator conditions outcomes. The idea here is that the main reason coethnics might dissemble at different rates is because they perceive the normatively correct answer differently depending on who is asking the question. For instance, a respondent might be more likely to falsely report voting for a high-performing coethnic incumbent if the enumerator is also a coethnic because the norm of rewarding a coethnic politician is more likely to be reinforced when speaking with another coethnic.

Table B.13 demonstrates that coethnicity mostly does not condition reporting of vote share for the incumbent. In tests of the full sample of good news or bad news areas, we interact enumerator coethnicity with incumbent coethnicity (which is what we think should matter most). The coefficient on enumerator coethnicity and on its interaction terms with incumbent coethnicity are not close to statistical significance in these models (columns 1 and 4). This alleviates the concern that differential response bias by ethnicity on the survey could be driving the main results in the paper. That said, when we divide the sample into whether respondents are coethnics of the enumerator, we see some evidence that enumerator coethnicity matters in column 5. Here, irrespective of treatment group, coethnics with the enumerator are more likely to report voting for the incumbent even after

receiving bad news. While this is some indication of the existence of response bias, the difference in bias across coethnics and non-coethnics of the incumbent is not statistically significant as we see in column 4 which mitigates the concern.

Table B.13: Is Reported Vote Choice Conditioned by Enumerator Coethnicity?

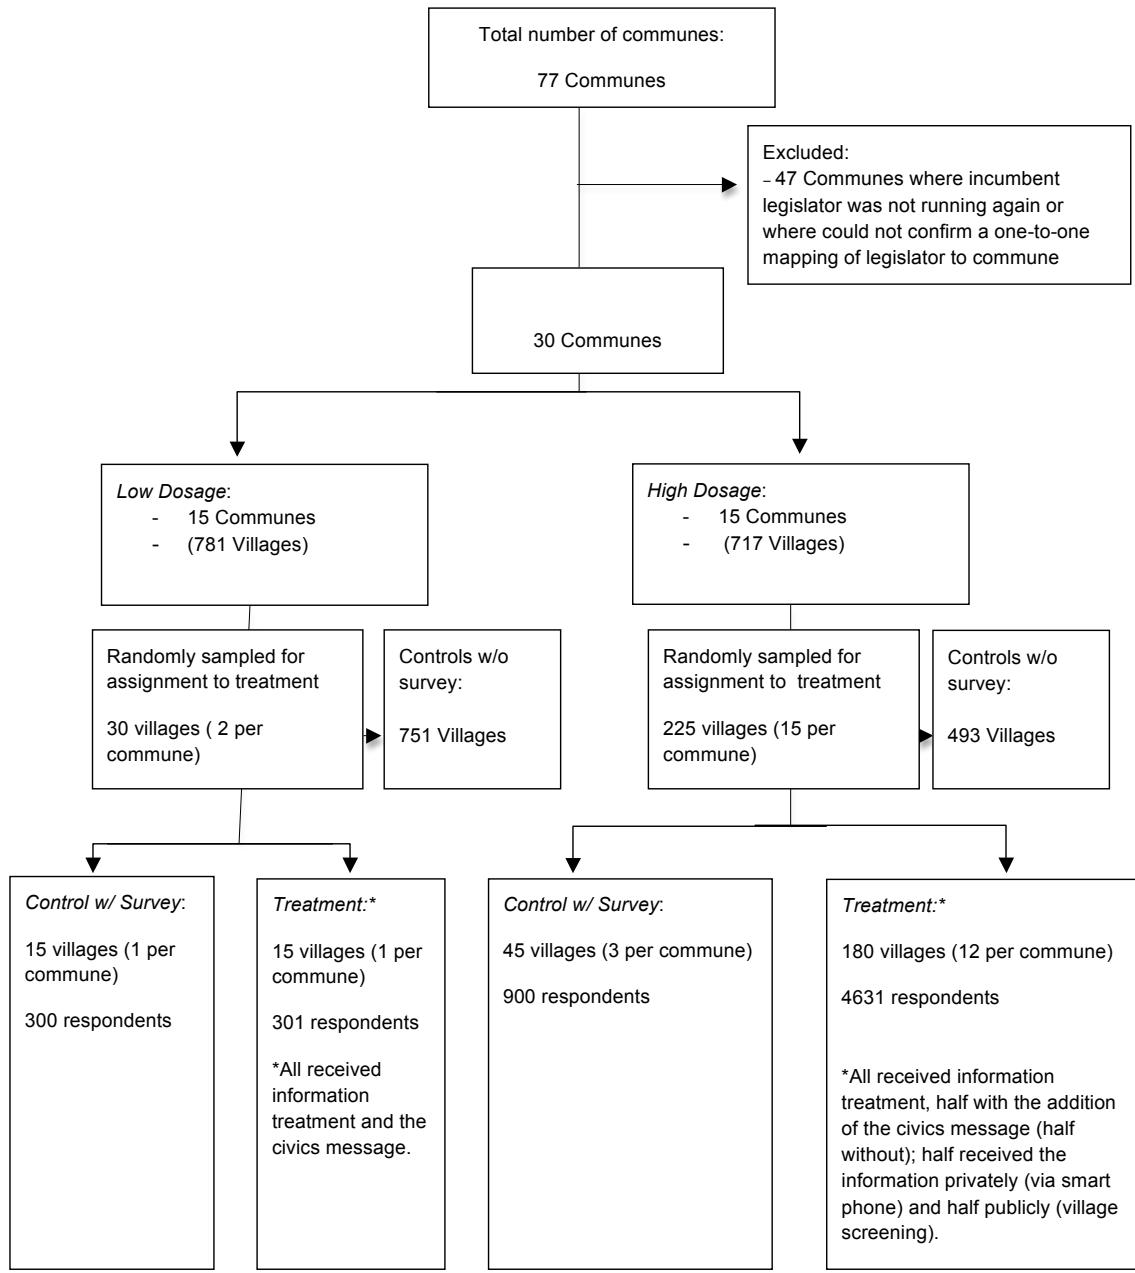
VARIABLES	(1) Good Info Full Sample	(2) Good Info Coethnic	(3) Good Info Non-Coethnic	(4) Bad Info Full Sample	(5) Bad Info Coethnic	(6) Bad Info Non-Coethnic
Treatment	0.02 (0.06)	0.12** (0.06)	0.03 (0.06)	-0.16 (0.11)	-0.01 (0.08)	-0.16 (0.10)
Coethnic with Enumerator	0.13 (0.12)	-0.02 (0.10)	0.16 (0.13)	0.02 (0.14)	0.21* (0.12)	0.00 (0.15)
Treatment x Enumerator Coethnicity	-0.13 (0.13)	-0.01 (0.10)	-0.13 (0.13)	-0.04 (0.13)	-0.07 (0.10)	-0.07 (0.14)
Coethnic with incumbent	-0.10 (0.08)			-0.16 (0.12)		
Treatment x Incumbent Coethnicity	0.08 (0.09)			0.14 (0.13)		
Enumerator Coethnicity x Incumbent Coethnicity	-0.15 (0.16)			0.11 (0.16)		
Treatment x Enumerator Coethnicity x Incumbent Coethnicity	0.15 (0.17)			-0.04 (0.17)		
Constant	0.50*** (0.05)	0.41*** (0.05)	0.47*** (0.06)	0.64*** (0.11)	0.46*** (0.08)	0.60*** (0.10)
Observations	1,657	1,030	627	1,356	948	408
R-squared	0.19	0.22	0.19	0.25	0.28	0.23

In parentheses, robust standard errors clustered by village.

Models include block fixed effects. *** p<0.01, ** p<0.05, * p<0.1.

C Further Details on Experimental Design

C.1 Geographic Sampling and Randomization Procedures



Note: Randomized 30 communes to high/low dosage, blocking on North/South, high/low incumbent performance. After randomly sampling villages within each commune, randomly assigned villages to treatment/control within three blocks (nested within communes): urban, rural/electorally competitive, rural/electorally uncompetitive.

C.2 Survey Sampling Procedures

The sampling procedure for the baseline survey occurred as follows: enumerators used a random walk procedure to select compounds in which to administer the baseline survey and/or intervention. Within compounds, individual respondents were randomly selected from the list of adult members of the compound while alternating on gender. As a condition of participating in the baseline survey, respondents had to have access to a cell phone.⁴⁶ Respondents were then recontacted by phone during the endline survey. A total of 3,419 individuals participated in the baseline *and* endline surveys (6,132 in the baseline), with an additional 6,174 receiving the intervention (or an invitation to the public screening) but no survey. To maximally harmonize public and private treatments, we endeavored to treat the same numbers of individuals per village across conditions. Thus, we provided the private treatment to 40 individuals in each private village even though we surveyed only a random half of those.

In each Private village, 20 people were randomly selected both to take the survey and to receive the intervention, 20 people were randomly selected only to receive the intervention, and 10 people were randomly selected to serve as control individuals and thus took the survey but received no intervention. In Public villages, 20 people were randomly selected to take the survey and be invited to receive the intervention at a public workshop, while an additional 40 people were randomly selected only to be invited to receive the intervention. We sampled in this way so as to ensure that roughly the same number of people would be treated in both Private and Public villages.⁴⁷

⁴⁶ According to the most recent round of Afrobarometer surveys in Benin (2014), approximately 77% of Beninois own a cell phone (<http://www.afrobarometer.org>). Still, respondents were not required to own a cell phone to participate in our intervention. The cell phone to which they had access could belong to a friend or relative.

⁴⁷ On average, 55 individuals in Public villages attended the video screenings (range from 20 to 70), indicating a reasonable balance of treated individuals across Public and Private villages.

C.3 Additional Notes on Ethical Considerations

In designing the experiment, we followed the principles agreed to by the entire Metaketa Initiative, as noted in the paper. We obtained permission from the traditional authority in each village and also individually from all subjects who participated in the survey (in both treatment and control villages) and received the information intervention. Important, we also sought and obtained permission from the President of the National Assembly, an actor who could evaluate risks to the candidates running the communes in the experimental sample. Following the guidelines for the Metaketa initiative, we did not seek consent from every individual legislator. Since the purpose of the study was to evaluate whether information interventions can help voters hold their representatives more accountable, we viewed it as a disservice to voters and to Beninois society more generally to allow poorly performing politicians, for instance, to act as veto players over the study. To ensure that the information we distributed was objective, context-appropriate and reliable, we utilized data collected by the legislature itself long before any knowledge of its use in an intervention like ours. We note that this data is meant to be publicly available.

C.4 Treatment Conditions

In addition to randomizing whether voters received legislative performance information or not, our study design manipulated the content of the treatment message to vary the signal of the importance of the legislative performance dimension. Treated participants were shown a video with either *only* the information about relative legislator performance (*Info Only*), or that same information *plus* an additional message highlighting the importance of legislative performance to voter welfare (*Civics*).⁴⁸

⁴⁸For those in the Civics experimental condition, the video contained a civics message that first described the main responsibilities of legislative deputies, namely, their responsibility for legislation, executive oversight and representation. It then provided three concrete examples of how legislative performance (or lack thereof) can impact voter welfare. A positive example of good legislation was the passage of an anti-graft law requiring public servants to disclose assets. A negative example of a missed opportunity was the failure of the legislature to vote on and pass a health

We also varied the method by which the information was disseminated. Treated participants received the intervention either privately by watching a video on a smartphone in the respondent's household (*Private*) or publicly through the screening of the same video via a projector in a public location in the village or quarter (*Public*).

In addition, we varied dosage — the density of treated villages in the commune. Participants were told during the intervention how many other villages in their commune were receiving legislative performance information. In high dosage communes, we randomly assigned 3 villages to each of the four combinations of content plus dissemination method (Info-Only/Private, Info-Only/Public, Civics/Private, Civics/Public), for a total of 12 villages treated with legislative performance information. In low-dosage communes, we assigned only one village to treatment (Civics/Public).

In high dosage communes, we implement a factorial design with four treatment conditions (Info-Only/Private, Info-Only/Public, Civics/Private, and Civics/Public) and a control group. Because our pre-specified hypotheses make predictions about how ethnicity conditions the effect of any of these treatments, we collapse these categories into one *Treatment* group in the main analyses of this paper. We introduce these distinctions here in order to be clear about our experimental design and because, as we detail below, these different treatments will become useful in that they allow us to test competing explanations for the main results we present.

C.4.1 Information Only Condition Script (in English)

Especially if you live outside the capital, it is not always a simple matter to know how well your national deputy is performing in Porto Novo. This information, however, can be useful during election time when you decide on who you want to represent you in the national assembly.

insurance scheme that was proposed in 2008. Finally, a positive example of executive oversight detailed how the legislature opposed changes to the Constitution proposed by the president that would expand his power. The *Civics* treatment was designed to encourage voters to switch to voting along a legislative performance dimension. The text of the video script for the *Civics* condition can be found in the appendix.

We have collected data from the national assembly here in Benin about how all the deputies perform their responsibilities. In light of the upcoming elections on March 31, we want to share with you some information about your own deputy, and how he performs relative to other deputies in Benin. There are many ways we could evaluate the performance of a deputy in the national assembly, but we choose to focus on two key aspects that directly correspond to a deputy's formal responsibilities: his performance in plenary sessions and his performance in committee. Working with researchers in the United States, the Centre d'Études et de Promotion de la Démocratie à Cotonou worked for several months to gather information about the legislative performance of Benin's 83 deputies during the previous 4-year mandate. We obtained reports from the President of the National Assembly that detail all the activities undertaken by the assembly and its deputies. We use the information in these reports to evaluate the performance of each deputy so that we can present to you a clear and concise report of how your deputy is doing. As I mentioned, we evaluate deputies on two themes: their performance in plenary sessions and their performance in committees. Plenary sessions are when deputies meet in the national assembly to vote on laws and oversee activities by the president and his government. The assembly holds an average of X sessions per year. We evaluate a deputy's performance in plenary sessions first by his rate of attendance at these meetings. Rates of attendance vary from X% to Y%. Second, we assess plenary performance by whether the deputy poses questions, oral or written, about the laws being discussed or the president's activities being evaluated as a sign of their active participation. The majority of deputies, X%, never ask any questions. However, some deputies are very involved asking up to 70 questions.

Much of the legislative work of deputies gets done outside of plenary sessions and in committees. These committees are organized by theme such as committees on finance or on education where deputies meet to discuss how to make laws pertaining to that them. We measure a deputy's performance in committee by whether or not the deputy is a member of the committee, by how productive the committee is or how many laws it considered, by how many meetings the committee held, and by the deputy's rate of attendance at those meetings.

I've just given you a lot of information about how we measure and evaluate the performance of deputies. I know that it is a lot to keep in your head which is why we try to condense the information as much as possible. So, we have created two scores: one for plenary performance and one for committee performance—that summarize how your deputy is doing on each aspect of his job relative to other deputies in Benin. These scores are just a combination of all the information I mentioned. Later, I will present the total score for your deputy on each of the two themes, and the combined score. I will also tell you all the ingredients that went into creating these scores for your deputy. If you want to know more, you can always contact CEPRODE at the number provided to you.

Benin is comprised of 12 départements. In each département, there are two constituencies (circonscriptions). All the towns and villages in a constituency vote together to elect 2 to 5 deputies to represent them. The number of deputies each constituency gets is based on population size. For example, your constituency is here and has this many deputies. The other constituency in your département has this many deputies. Today, we will provide you with information about the performance of [NAME OF DEPUTY]. Though he is one of [NUMBER] deputies in your constituency, our sources tell us that he is the most important deputy in this commune, [NAME OF COMMUNE].

First, I will tell you about how [NAME] performed in plenary sessions. He attended [RAW NUMBER] sessions of the X total plenary sessions. In other words, he attended [NUMBER OUT OF TEN] plenary sessions. He asked [RAW NUMBER] questions during these sessions. Remember, while most deputies ask no questions, some ask up to 70. Combining these two measures, we give [NAME] a score of [NUMBER] out of 100 on the plenary performance index. As you can see, your deputy did [BETTER/WORSE] on this measure than other deputies in your département. And he did [BETTER/WORSE] than the national average for all the deputies in Benin.

Second, I will tell you about how [NAME] performed in committee sessions. [NAME] [IS/NOT] a member of a legislative committee. [HE IS A MEMBER OF X COMMITTEE]. This committee is one of the [MORE/LESS] productive committees and treated [RAW NUMBER] of

laws during its tenure. This committee held [RAW NUMBER] meetings. Your deputy, [NAME], attended [NUMBER OUT OF TEN] of these meetings.

Combining these measures, we give [NAME] a score of [NUMBER] out of 100 on the committee performance index. As you can see, your deputy did [BETTER/WORSE] on this measure than other deputies in your département. And he did [BETTER/WORSE] than the national average for all the deputies in Benin. If we combine the scores for our two indices together, we see that your deputy, [NAME] performed [BETTER/WORSE] in total than other deputies in your département, and [BETTER/WORSE] than the national average for all the deputies in Benin.

So, to summarize all the information I have just told you: Your deputy, [NAME], is [MUCH/A LITTLE] [BETTER/WORSE] than other deputies in Benin when it comes to performing his legislative responsibilities. This is mostly because he: a. Does(n't) attend plenary sessions, b. Does(n't) participate actively in plenary sessions, c. Is(n't) a member on a committee, d. Does(n't) attend committee meetings.

C.4.2 Civics Condition Script (in English)

NOTE: The second half of this script is identical to the information only script.

I would like to talk to you about the National Assembly: specifically about the roles and responsibilities of deputies elected to the National Assembly and about how their activities in the National Assembly in Porto Novo can affect you and your family.

There are 83 deputies elected across the country, including the deputies from this constituency. Deputies are charged with three main roles. First, they are responsible for legislation, which means making laws that can have an impact on your daily life. Second, deputies are responsible for oversight—that is, for holding the President accountable, for making sure that he respects the laws and people of Benin. Third, deputies are responsible for representation—that is, for conveying your needs to the government and for explaining the actions of the government to you.

Let me discuss each of these responsibilities in turn. Some of this you may know already and

some may be new information.

As I mentioned, the first main responsibility of deputies is Legislation. Either the President (and his ministers) or individual deputies can have an idea for a new law. They write that idea down as a proposed law, called a "bill." The President or Deputy submits the bill to the head of the National Assembly. After the head of the National Assembly declares it admissible, the bill is sent to a committee made up of deputies who have expertise in the matters raised by the bill. For instance, if the bill concerns education, it will be sent to the Committee on Education, Culture, Employment and Social Affairs for study. That committee then meets in order to study and review the bill carefully and issues a report about the bill that is then circulated and presented to all of the members of the National Assembly. Members of the National Assembly then debate the committee's report and each article of the bill in a full session in the capital in Porto Novo. During this time, individual deputies can make public statements about their positions as to whether the bill is good or bad for their constituents and for Benin as a whole. They can try to persuade other deputies to vote a certain way. After the debate, the deputies then each vote to pass or not to pass the bill. A bill passes if a majority of deputies present vote "yes" to the bill. The National Assembly passes approximately 25 laws each year. It is important to note that only deputies who attend their assigned committee meetings and who attend and participate in the full sessions of parliament can influence which laws pass and which do not.

The second main responsibility of deputies is Oversight. As I mentioned, oversight means holding the President accountable and making sure that he respects the laws and people of Benin. One very important way in which deputies are authorized by law to engage in oversight is by intervening in the process by which the national budget is crafted each year. In fact, by law, the National Assembly is the institution that can oversee the President's budget and make sure that it reflects the needs of the people of Benin. Each year the President proposes a budget—that is, he proposes the total amount of money that will be spent on executing national policies and projects in that coming year, and he proposes how that money will be divided across projects and across different parts of the country. The most important committees in the National Assembly

for overseeing this proposal for spending money are the Finance and the Planning committees. Deputies on these two committees are supposed to meet regularly in order to analyze and study the proposed budget. These two committees review more bills than any other committees in the National Assembly. They can make recommendations to the President about ways to amend the budget before it is presented to the National Assembly. All deputies can also vote to approve the President's budget once it is presented to the full Assembly. When the budget is implemented, the National Assembly can make recommendations to the President if they observe that the budget is not being spent properly.

Another important way in which deputies can engage in oversight is by making sure that any proposed legislation or ordinance put forward by the President is in compliance with the Constitution and with all electoral laws. Deputies on the Legal Committee of the National Assembly are charged with studying any bills that would change rules about elections or the powers of the President and with making reports on their legality to the full National Assembly. This committee reviews the third largest set of bills each year, after the sets reviewed by the Finance and Planning committees. If any change is proposed to the Constitution of Benin, at least three-fourths of the members of the National Assembly have to vote to approve the change before it can move forward.

The third main responsibility of deputies is Representation. As you know, deputies are elected to serve particular constituencies, including the constituency in which this village is located. As citizens, you are very busy with meeting your daily needs and those of your family. You cannot travel to the capital to tell the President what your needs are. Instead, that is part of the deputies' job. They are charged with communicating your needs and the needs of other voters in this constituency to the National Government. The deputies can do this by raising questions and concerns about national legislation in their assigned committee meetings and when bills are debated in full sessions of the National Assembly. During those times, deputies can make clear to other politicians whether or not the law is in your best interests—that is, whether or not it is in the interests of voters in the deputies' home constituencies. Deputies can also come up with new ideas for legislation, based on their understanding of your needs. If deemed admissible for review by the head of the

National Assembly, these new ideas—written down as bills—will then be reviewed by committees and debated by all deputies who attend the full meetings of the National Assembly. Again, it is important to note that only deputies who attend committee meetings and full National Assembly meetings, and who participate by asking questions and voicing your concerns, can fulfill their responsibilities of legislating, engaging in oversight and representing your needs in the capital.

Now, you may still be thinking that none of these activities has much to do with your welfare. But let me give you some examples of ways in which what deputies do in Porto Novo does matter for the quality of your life and that of your family.

One example is the anti-graft law that the National Assembly passed in August 2011. This is a law that requires Benin's top leaders, civil servants, central directors of the administration, project managers and accountants of any public body to disclose their assets when they enter and leave office. The law is intended to help prevent corruption so that the money in the national budget is spent on you, the citizen, and not on lining the pockets of powerful people. The deputies in the National Assembly are the ones who had to review and approve this law. Their work in the National Assembly in Porto Novo is thus relevant to ensuring that resources get to you.

Here is another example. Le Régime d'assurance maladie universelle (RAMU) is a proposed national program that would help the people of Benin access healthcare. It would help you if you have trouble paying for medical treatment. The consequences of getting sick can be financially disastrous for you and your family if you do not have the money to pay for healthcare. If it becomes law, RAMU would help you. It would help the poor; it would help farmers; it would help students; it would help taxi and moto drivers; it would help people who are informally employed. It would cover visits to the doctor, costs of staying in or being treated in the hospital, costs of medication, transportation to the hospital or doctor and tests to know if you are sick.

The idea for the program was conceived by the Council of Ministers in 2008, and since 2011, there have been small versions of the program operating in some villages, known as “zones sanitaires.” In 2014, the President established a National Steering Committee. But in order to become a program that operates across the whole country, RAMU has to be approved by a vote in Par-

liament. In other words, the national deputies have to do the work of evaluating and voting on the proposed law before it can become an implemented national program that can help you pay if you get sick. The President has said that RAMU is a national priority. But the performance of the national deputies is crucial if the proposal is actually to become law. Whether your national deputy shows up and participates in Parliament has an impact on whether RAMU becomes law and thus on whether you and your family get help if you are sick.

Third, let me give you an example of Parliament's important role in presidential oversight. In 2009, President Boni Yayi sent a proposal (known as a "projet de loi") to the National Assembly that sought to revise Benin's constitution, which has not been changed since it was enacted in 1990. The proposal was again sent to the National Assembly in 2013. The proposal went to the Law Committee of the National Assembly but did not make it to the Assembly for a vote. Members of the public began opposing the proposal fearing it was a way for President Yayi to extend his presidential mandate. Benin citizens came together to voice their opinion and created movements against the changes such as the "Red Wednesday" movement ("Mercredi rouge"). Leaders of several political parties came forward to oppose the proposals, stating that the public was not ready for a change to the constitution, and citing more important issues for the president to concentrate on. Even pro-presidential members of the Law Committee were against the changes, and large majority of the committee's members voted to reject the proposal on September 24th, 2013. In this sense, legislative representatives not only exercised their right as a check to the executive but also represented the public interest which voiced its disapproval of any constitutional changes. These are just a few examples of how the performance of your national deputy—his participation in committees and in plenary sessions of the National Assembly, his willingness to ask questions and voice positions on legislation and to exercise presidential oversight—are important for your daily lives.

But, especially if you live outside the capital, it is not always a simple matter to know how well your national deputy is performing in Porto Novo. This information, however, can be useful during election time when you decide on who you want to represent you in the national assembly.

We have collected data from the national assembly here in Benin about how all the deputies perform their responsibilities. In light of the upcoming elections on March 31, we want to share with you some information about your own deputy, and how he performs relative to other deputies in Benin. There are many ways we could evaluate the performance of a deputy in the national assembly, but we choose to focus on two key aspects that directly correspond to a deputy's formal responsibilities: his performance in plenary sessions and his performance in committee. Working with researchers in the United States, the Centre d'Études et de Promotion de la Démocratie à Cotonou worked for several months to gather information about the legislative performance of Benin's 83 deputies during the previous 4-year mandate. We obtained reports from the President of the National Assembly that detail all the activities undertaken by the assembly and its deputies. We use the information in these reports to evaluate the performance of each deputy so that we can present to you a clear and concise report of how your deputy is doing. As I mentioned, we evaluate deputies on two themes: their performance in plenary sessions and their performance in committees. Plenary sessions are when deputies meet in the national assembly to vote on laws and oversee activities by the president and his government. The assembly holds an average of X sessions per year. We evaluate a deputy's performance in plenary sessions first by his rate of attendance at these meetings. Rates of attendance vary from X% to Y%. Second, we assess plenary performance by whether the deputy poses questions, oral or written, about the laws being discussed or the president's activities being evaluated as a sign of their active participation. The majority of deputies, X%, never ask any questions. However, some deputies are very involved asking up to 70 questions.

Much of the legislative work of deputies gets done outside of plenary sessions and in committees. These committees are organized by theme such as committees on finance or on education where deputies meet to discuss how to make laws pertaining to that them. We measure a deputy's performance in committee by whether or not the deputy is a member of the committee, by how productive the committee is or how many laws it considered, by how many meetings the committee held, and by the deputy's rate of attendance at those meetings.

I've just given you a lot of information about how we measure and evaluate the performance of deputies. I know that it is a lot to keep in your head which is why we try to condense the information as much as possible. So, we have created two scores: one for plenary performance and one for committee performance—that summarize how your deputy is doing on each aspect of his job relative to other deputies in Benin. These scores are just a combination of all the information I mentioned. Later, I will present the total score for your deputy on each of the two themes, and the combined score. I will also tell you all the ingredients that went into creating these scores for your deputy. If you want to know more, you can always contact CEPRODE at the number provided to you.

Benin is comprised of 12 départements. In each département, there are two constituencies (circonscriptions). All the towns and villages in a constituency vote together to elect 2 to 5 deputies to represent them. The number of deputies each constituency gets is based on population size. For example, your constituency is here and has this many deputies. The other constituency in your département has this many deputies. Today, we will provide you with information about the performance of [NAME OF DEPUTY]. Though he is one of [NUMBER] deputies in your constituency, our sources tell us that he is the most important deputy in this commune, [NAME OF COMMUNE].

First, I will tell you about how [NAME] performed in plenary sessions. He attended [RAW NUMBER] sessions of the X total plenary sessions. In other words, he attended [NUMBER OUT OF TEN] plenary sessions. He asked [RAW NUMBER] questions during these sessions. Remember, while most deputies ask no questions, some ask up to 70. Combining these two measures, we give [NAME] a score of [NUMBER] out of 100 on the plenary performance index. As you can see, your deputy did [BETTER/WORSE] on this measure than other deputies in your département. And he did [BETTER/WORSE] than the national average for all the deputies in Benin.

Second, I will tell you about how [NAME] performed in committee sessions. [NAME] [IS/NOT] a member of a legislative committee. [HE IS A MEMBER OF X COMMITTEE]. This committee is one of the [MORE/LESS] productive committees and treated [RAW NUMBER] of

laws during its tenure. This committee held [RAW NUMBER] meetings. Your deputy, [NAME], attended [NUMBER OUT OF TEN] of these meetings.

Combining these measures, we give [NAME] a score of [NUMBER] out of 100 on the committee performance index. As you can see, your deputy did [BETTER/WORSE] on this measure than other deputies in your département. And he did [BETTER/WORSE] than the national average for all the deputies in Benin. If we combine the scores for our two indices together, we see that your deputy, [NAME] performed [BETTER/WORSE] in total than other deputies in your département, and [BETTER/WORSE] than the national average for all the deputies in Benin.

So, to summarize all the information I have just told you: Your deputy, [NAME], is [MUCH/A LITTLE] [BETTER/WORSE] than other deputies in Benin when it comes to performing his legislative responsibilities. This is mostly because he: a. Does(n't) attend plenary sessions, b. Does(n't) participate actively in plenary sessions, c. Is(n't) a member on a committee, d. Does(n't) attend committee meetings.

C.4.3 Dosage Treatment

Our sample included the 30 communes for which we could strongly confirm a one-to-one mapping of commune to incumbent candidate. Of those 30, we randomly assigned 15 to each dosage category (high or low).⁴⁹ Within communes, the unit of randomization was the rural village or urban quarter. These units are the lowest level of social and territorial organization.

Within each low dosage commune, we randomly select two villages and then randomly assign one to Civics/Public and the other to control. Because in each high-dosage commune 12 villages were assigned to treatment and three to control, and in each low-dosage commune one village was assigned to treatment and one to control, the sample comprises 225 villages in high-dosage communes and 30 villages in low-dosage communes, for a total of 255.

⁴⁹When assigning dosage, we blocked on incumbent legislative performance, which is observed at the commune level and on region (north/south). Within 4 blocks (high and low performance in the north and south) of communes, we assigned half to high-dosage and half to low-dosage treatment.

Before the treatment videos were shown, participants in high and low dosage communes were told the following:

1. High Dosage: You have been selected through a random process to participate in a research study about the performance of your deputies in the National Assembly. Your community is one of 12 villages or quartiers in your commune receiving this exact same video. Many other communes in Benin are also part of the study.
2. Low Dosage: You have been selected through a random process to participate in a research study about the performance of your deputies in the National Assembly. Your community is the only one in the commune receiving this information.

After the video treatment videos were shown, participants were told the following:

1. High Dosage: Remember, your village/quartier is one of 12 villages or quartiers in your commune receiving this video.
2. Low Dosage: Remember, your village/quartier is the only one in the commune receiving this video.

C.4.4 Linking the Results to the Companion Paper

Appendix C.4 describes each of the treatment arms in our experimental design. In a companion paper , we present the full results from this design. The analysis produces two main results. First, we find that in low dosage communes, positive performance information actually had a *negative* impact on the vote share of the incumbent. Second, we find that in high dosage communes, positive performance information has the expected positive effect in the civics treatment condition (the *Civics* treatment was designed to encourage voters to value the legislative performance dimension). The first finding was unanticipated; the second, consistent with our expectations.

The interpretation of the results centers on the notion that politicians engage in different types of activities and that these dimensions of performance may be perceived as substitutes rather than complements in the minds of voters. In our context, legislators generally engage in legislative activities and constituency service that involves the provision of local collective and private goods to voters. We leverage qualitative evidence collected in Benin that supports the interpretation that voters in the experiment perceived “good” legislative performance information as “bad” news about constituency service (which they ultimately care most about), unless the information is distributed widely across the commune and in combination with the civics treatment.

In this Appendix section, we discuss how the results in this paper relate to those presented in the companion paper. First, we discuss the results in low dosage communes (communes with only one treated village), where positive performance information reduced the vote share of incumbents in treated relative to control villages. In this treatment condition, positive information generated a negative response to the incumbent. In the current paper, we show that, aggregating across treatment conditions, negative information has a negative impact on vote share when provided to non-coethnics of the incumbent, but no impact on incumbent vote share when the information was provided to the incumbent’s coethnics. If the interpretation of legislative and constituency service dimensions as substitutes is true, we should expect that *the negative effect of positive information in low dosage communes will be larger (in absolute value) among non-coethnics and close to zero*

Table C.1: Treatment Effects by Ethnic Connection and Level of Performance, Official Results Data

VARIABLES	(1) Good Info Full Sample	(2) Good Info Coethnic (70)	(3) Good Info Non-Coethnic (70)	(4) Good Info Coethnic (90)	(5) Good Info Non-Coethnic (90)
Treatment	-0.06 (0.13)	0.05 (0.17)	-0.36* (0.16)	0.13 (0.26)	-0.17 (0.23)
Constant	0.53*** (0.09)	0.47*** (0.12)	0.63*** (0.09)	0.53** (0.19)	0.49** (0.14)
Observations	12	7	5	5	7
R-squared	0.84	0.91	0.97	0.90	0.84

The sample only includes low dosage, good news communes.

among coethnics (because only non-coethnics punish negative information). Table C.1 shows that this is indeed the case. In low dosage, positive information communes, the average treatment effect is negative.⁵⁰ Columns 2 and 3 show that this negative effect is driven by a large negative effect among non-coethnic villages.⁵¹

The second set of main results shows that information has the expected effect when it is distributed widely across the commune and voters are told as much (high dosage) *and* the civics treatment is included. Tables C.2 and C.3 thus present the coethnicity analyses using only data from the control group and the high dosage/civics treatment arm. The results are consistent with, and in some instances stronger than (in terms of the magnitude of the effects),⁵² the results presented in the main body of this paper.

⁵⁰In the companion paper, we use all villages/quarters in a commune as the control group (even those where we did not conduct surveys), which increases our statistical power. The negative coefficient presented in that paper is statistically significant. In this paper, we can only analyze the sample where baseline surveys were conducted because we only have ethnicity information for these villages/quarters.

⁵¹Due to the small sample sizes, the p-values should be interpreted with caution. We emphasize the direction and the magnitude of the effects rather than their statistical significance.

⁵²The results are generally less statistically strong because of the much smaller sample sizes.

Table C.2: Effect of *Civics* Treatment by Ethnic Connection and Level of Performance in High-Dosage Communes, Official Results

VARIABLES	(1) Good Performance Coethnic (50)	(2) Good Performance Non-Coethnic (50)	(3) Bad Performance Coethnic (50)	(4) Bad Performance Non-Coethnic (50)
Treatment	0.090 (0.058)	-0.044 (0.097)	0.008 (0.092)	-0.276 (0.206)
Constant	0.467*** (0.045)	0.474*** (0.060)	0.488*** (0.082)	0.579*** (0.160)
Observations	46	29	38	15
R-squared	0.778	0.594	0.769	0.750
VARIABLES	(1) Good Performance Coethnic (70)	(2) Good Performance Non-Coethnic (70)	(3) Bad Performance Coethnic (70)	(4) Bad Performance Non-Coethnic (70)
Treatment	0.061 (0.077)	-0.001 (0.101)	0.054 (0.074)	-0.193 (0.155)
Constant	0.446*** (0.059)	0.505*** (0.067)	0.499*** (0.065)	0.514*** (0.126)
Observations	37	38	31	22
R-squared	0.832	0.629	0.823	0.684
VARIABLES	(1) Good Performance Coethnic (90)	(2) Good Performance Non-Coethnic (90)	(3) Bad Performance Coethnic (90)	(4) Bad Performance Non-Coethnic (90)
Treatment	0.107 (0.090)	0.039 (0.080)	0.059 (0.074)	-0.192 (0.120)
Constant	0.399*** (0.069)	0.483*** (0.056)	0.500*** (0.065)	0.516*** (0.096)
Observations	25	50	28	25
R-squared	0.849	0.605	0.864	0.691

The sample only includes control and civics condition units in high dosage. This table presents results using different cutpoints for defining a village as a coethnic village. In the top panel, villages are coded as coethnics if over 50 percent of survey respondents are coethnics of the incumbent. In the middle panel, villages are coded as coethnics if over 70 percent of survey respondents are coethnics of the incumbent. In the bottom panel, villages are coded as coethnics if over 90 percent of survey respondents are coethnics of the incumbent. Robust standard errors clustered by village are in parentheses. Each model uses block fixed effects. *** p<0.01, ** p<0.05, * p<0.1

Table C.3: Effect of *Civics* Treatment by Ethnic Connection and Level of Performance in High-Dosage Communes, Survey Data

VARIABLES	(1)	(2)	(3)	(4)
	Good Performance Coethnic	Good Performance Non-Coethnic	Bad Performance Coethnic	Bad Performance Non-Coethnic
Treatment	0.132** (0.051)	0.003 (0.065)	-0.020 (0.051)	-0.243** (0.112)
Constant	0.381*** (0.045)	0.516*** (0.050)	0.546*** (0.044)	0.555*** (0.095)
Observations	543	336	473	227
R-squared	0.201	0.185	0.323	0.227

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The sample only includes control and civics condition units in high dosage.