

Escaping the Deliverism Trap

Recognition and Representation in Nigerian Democracy

Marco Castradori*

August 26, 2026

Abstract

Democracy depends on accountability, yet citizens find it difficult to observe and evaluate legislative effort. They therefore hold their representatives to account based on what they can see. MPs thus prioritize delivering visible, short-term, material benefits over contributing to programmatic legislation. I argue that, in contexts like Nigeria, this dynamic has settled into a stagnant “deliverist” equilibrium in which material benefits are not only desirable in themselves but also the primary evidence that citizens matter and representatives care. I test whether non-material engagement can provide alternative evidence with a cluster-randomized field experiment across three Nigerian federal constituencies, inviting citizens in some clusters to a Deliberative Town Hall (DTH) with their MP. Assignment improved citizens’ evaluations of their representative, strengthened their sense of political standing, boosted trust in politicians more broadly, satisfaction with democracy, and political hope — all without the representative delivering any immediate material benefit. It did not, however, alter the priority citizens place on immediate delivery. A vignette experiment embedded in the endline survey shows that legislative effort and respectful engagement each improve evaluations on their own. Deliberative engagement can thus broaden the foundations of trust before it changes what citizens demand. Citizens care about how they are represented, not only what they get.

*PhD candidate, Princeton University, Politics Department

“Who knows us or knows our name? In our country today, the politicians forget that it’s the people who voted for them, they don’t respect people. They won’t even come to your community again so you can talk with them. If you see them, it’s in a tinted glass car, [...] as citizens, we don’t cross their path, let alone talk to the politicians.”

– 35 y.o. female, Oyo State, Nigeria. Author interview.

1 Introduction

Effective democratic representation requires not only that elected officials act wisely on citizens’ behalf, but also that citizens can recognize and reward such action. Yet many activities central to legislative representation – negotiating, conducting oversight, and pursuing long-term policy goals – are difficult for ordinary citizens to observe. When representatives’ efforts in valuable activities remain hidden, citizens may instead rely on more visible evidence of their performance: a road, a personal transfer, or even the representative’s physical presence in the constituency – heuristics that crowd out politicians’ incentives to engage in the quieter work of programmatic policy.

How, then, can democratic institutions create political space for representatives to pursue consequential but less visible work while satisfying citizens’ understandable demand for observable evidence of effective representation? I argue that the answer lies in creating the conditions for *warranted trust* (Warren, 2018): reasons for citizens to believe that actions undertaken in their name are guided by their interests. Institutions can help create such trust by providing observable occasions on which representatives make otherwise hidden intentions and efforts legible. If citizens recognize these engagements as credible signals of care and effort, they gain grounds for evaluating representation even when policy benefits are slow, uncertain, or difficult to attribute. Democratic innovations that reliably generate such signals may therefore strengthen confidence in representative institutions while widening the political space available for less visible forms of representative work.

I apply this argument to sub-Saharan Africa, where the democratic optimism that accompanied the spread across the continent of multiparty elections in the 1990s has largely faded. Although electoral institutions have largely survived, most citizens experience democracy as stagnant, distant, and weakly responsive to their needs (Bratton et al., 2024). Nigeria exemplifies this trajectory. More than a quarter-century after the 1999 elections inaugurated the Fourth Republic, 85% of Nigerians report dissatisfaction with democracy, 92% believe the country is going in the wrong direction, and 72% describe it as either not a democracy or one with major problems.¹

In this setting, the defining condition of political representation – that representatives act for citizens in their absence (Pitkin, 1967) – gives rise to a particularly consequential problem of observability: what I call a *deliverism trap*. Because citizens cannot readily observe representatives’ intentions, effort, or legislative work, they rely on concrete and attributable evidence such as projects, transfers, and physical presence to infer the quality of representation they receive. In turn, representatives respond by prioritizing those same activities, further reinforcing tangible delivery as the sole standard of good representation. And yet, because public needs far exceed what individual legislators can ever plausibly satisfy, this apparently *responsive* equilibrium produces persistent disappointment, distrust, and deepening democratic skepticism.

Escaping this trap does not require citizens to value material outcomes less. In a context of widespread precarity, expecting them to do so would be both unreasonable and insensitive. Instead, it requires weakening the near-monopoly of continual, visible, but ultimately *small-scale* delivery as evidence of representative commitment. The aim is not to substitute for material development as an objective of good governance, but rather to create alternatives to immediate tangible delivery as

¹Afrobarometer Round 10, 2024.

evidence that effective representation is indeed taking place. When citizens receive other credible signals of representatives' care and effort, they may evaluate representation on a broader evidentiary basis and sustain warranted trust over the interval between representative effort and observable *policy* results. This can create political space for representatives to devote greater attention to legislative work whose benefits are broader, slower to emerge, and harder for individual citizens to observe.

This paper investigates that possibility through a cluster-randomized experiment involving deliberative town halls (DTHs), an institutional innovation adapted from work on deliberative democracy in the United States and other advanced democracies (Neblo et al., 2018). These settings allow citizens to observe whether a representative is willing to listen, answer criticism, and treat ordinary constituents as legitimate political interlocutors. By making these normally hidden dimensions of representation visible while withholding the material benefits through which representative commitment is ordinarily demonstrated, DTHs provide unusual inferential leverage for testing whether citizens respond to non-material evidence of representative effort and care. Constituents were randomly sampled across three Nigerian federal constituencies, and their enumeration areas were assigned to receive invitations to a DTH or to a control condition. Invited constituents gathered at local venues to address and question their member of the House of Representatives (henceforth MP), who participated live by video from the capital, Abuja. The encounter provided no direct material benefit from the MP.

I distinguish two complementary pathways through which this evidence may generate warranted trust. First, a representative-facing pathway updates citizens' evaluations of the MP's "type": her character, competence, intentions, and willingness to exert effort on their behalf. Second, a citizen-facing pathway strengthens citizens' perceived political "standing": their confidence that people like them are recognized, can make claims on political authority, and can expect a response. Together, favorable assessments of the representative and greater confidence in one's own place within the political order provide grounds for trusting that less visible political activity is conducted with citizens' interests in mind.

My results reveal a pattern consistent with the ordering proposed by the theory. Assignment to a DTH improved citizens' evaluations of the representative's character, competence, and effort. It strengthened their perceived standing within both the representative relationship and the political system more broadly. It increased trust in elected officials and produced greater satisfaction with democracy and greater hope about the political system's future performance. Most importantly, it did all this without the representative delivering immediate material benefit. The sequence stopped, however, at the final and most demanding step in the theory. Treated constituents did not reallocate the emphasis they wanted their MPs to place on immediate constituency delivery versus legislative, programmatic, and longer-term forms of representation. I explore the insights this provides about the possible limitations of a single engagement.

I embedded a vignette experiment in the endline survey to complement the field results by isolating the signals the town hall field experiment necessarily bundles together. Information about a hypothetical MP's legislative effort and his respectful engagement with constituents improved evaluations even though he delivered no project and the hypothetical bill's fate remained uncertain.

Taken together, the findings establish both the potential and the boundary of a single structured encounter. The DTH broadened the evidence citizens used to evaluate representation: participants responded to visible effort, explanation, accessibility, and recognition even in the absence of material delivery. These signals strengthened warranted trust and improved broader democratic evaluations. But the encounter did not by itself alter the equilibrium-level demands that organize representative behavior. Doing so may require repeated encounters, wider diffusion, corroborating behavior, and credible evidence that a different representative relationship will endure.

2 Theory: The Deliverism Trap and Alternative Signals

2.1 Deliverism as a Representative Equilibrium in Nigeria

The central problem developed in this paper is not simply that Nigerian legislators fail to respond to their constituents. It is that responsiveness has become organized around a narrow set of activities that are visible, local, and immediate. Citizens demand the material foundations of everyday life: roads and boreholes, jobs and scholarships, help with medical bills, school fees, wedding and burial expenses (Lindberg, 2010; Ofosu, 2024). Legislators, in turn, describe their work as dominated by unrecognized attempts to satisfy precisely these demands.² They describe the job as “thankless,”³ even as they report spending much of their time “hustling”⁴ for constituents. Still, citizens continue to regard their representatives as distant, selfish, and unresponsive. Complaints that politicians “see [citizens] as beggars [...] don’t listen to [their] voices” and “only look after their own benefits”⁵ are ubiquitous, and express a deficit of political standing as much as one of tangible goods.

This divergence emerges from what I call *deliverism*: an equilibrium representational ethos in which visible, local, and immediate material delivery serves not only as a desired outcome but as the dominant evidence that representatives are *moral* actors, committed to improving citizens’ lives. In theory, representatives can make themselves electorally valuable through credit claiming, position taking, and the communication of their activities in office (Cain et al., 1987; Grimmer et al., 2012; Mayhew, 1974). Deliverism describes an equilibrium in which the first of these strategies overwhelms the others. Because citizens cannot directly observe representatives’ intentions, effort, or much of their legislative activity, directly attributable provision is especially legible and credit-claimable, crowding out position taking and explanation as credible evidence of representation. Material delivery thus becomes an unusually persuasive sign that a representative really is trying, and that constituents really do matter. MPs thus face continual pressure to demonstrate that they are “empathetic kind of persons [that] care about the welfare of the common man,”⁶ and therefore devote most of their time to the “daily hustle of beg[ging] (ministries and agencies) for what they can take back home.”⁷ In turn, because material delivery functions as evidence of moral commitment, it shapes the public narratives through which non-delivery is understood: its absence comes to signal not merely limited competence or state capacity, but selfishness, greed, and disregard.

The deliverist logic of daily representation I describe overlaps with clientelism, but is not reducible to it. Much of the clientelism literature emphasizes quid-pro-quo exchange, vote buying, and the strategic distribution of particularistic benefits (Hicken and Nathan, 2020; Kitschelt and Wilkinson, 2007; Kramon, 2018; Stokes et al., 2013). More recent work has broadened this view by emphasizing the everyday, welfare-oriented problem-solving performed by politicians in low-capacity states (Bussell, 2019; Golden and Nazrullaeva, 2023). My emphasis on deliverism directs attention to a related but distinct problem: the pressures created when visible provision becomes the principal evidence not only of effective representation, but of whether elected officials are morally committed to those they represent.⁸

²Although elite interviews are vulnerable to strategic presentation, social desirability, and biases arising from my position as a foreign researcher, I draw on more than thirty interviews with sitting Nigerian MPs and interpret them conservatively, placing weight on themes that recur across multiple interlocutors rather than on isolated claims or anecdotes. Elite perspectives cannot simply be discarded: any plausible account of how prevailing forms of representation are reproduced – and how they might be changed – must take seriously how politicians themselves understand their roles, constraints, and obligations.

³Nigerian House of Representatives member, interview no. 01, 23 September 2025.

⁴Nigerian House of Representatives member, interview no. 07, 16 October 2025.

⁵Author interview no. K05. I worked with three bilingual research assistants – Adenike Dada, Precious Ujam, and Salisu Mahmud – to conduct interpretative interviews (Schaffer, 2014) with 120 Nigerian citizens in English, Hausa, Igbo, and Yoruba. I draw from these interviews throughout the paper to capture the complicated political views of ordinary Nigerians.

⁶Nigerian House of Representatives member, interview no. 01, 23 September 2025.

⁷Nigerian House of Representatives member, interview no. 09, 23 October 2025.

⁸Accounts of representative politics across Africa suggest that much of what I observe is relevant beyond Nigeria (Lindberg, 2010; Ofosu, 2024; Power, 2024, among others).

This evaluative logic can operate even when benefits are neither strategically targeted nor explicitly conditioned on political support, as conventional accounts of clientelist exchange would imply. Much of the distributive work Nigerian legislators describe is reactive rather than strategically initiated. It is mostly constituents themselves who seek assistance, including those who did not support the incumbent, while legislators often lack the information and administrative capacity to ration that assistance according to partisan loyalty or any clear electoral logic.⁹ As one MP explained:

“The people’s opinion is that the one [MP] that is there now, even if they didn’t participate in the electoral process or voted [for] the opposite person [...] must attend to their problem. [...] Most of them that come here were those that didn’t vote for me. [...] I just have to assist anyone that’s from the constituency.”¹⁰

To be sure, MPs remain conscious of the electoral consequences of appearing unresponsive, and some strategic targeting undoubtedly occurs. Yet deliverism is not simply an electorally efficient exchange. It is, in large part, a messy, demand-driven response to expectations that representatives perceive themselves as unable to ignore.

Those expectations are especially difficult to resist because constituent demands are both urgent and morally charged. Requests commonly arise from genuine hardship: a hospital bill, a pivotal damaged road, or a family emergency. In a context of widespread poverty, uneven state capacity, and limited access to public services, elected representatives are among the few political actors whom citizens can plausibly reach and pressure for help. MPs consequently become points of last resort, asked to resolve problems that are immediate, personal, and often only tangentially related to the formal powers of a legislator. One MP described the moral pressure embedded in these encounters:

“[They will tell you] *‘I’ve not eaten in the past two days, I’ve not eaten since morning.’* They come here, they stay from morning to evening . . . how can you leave them hungry?”¹¹

Activities such as paying school fees or medical bills may resemble textbook clientelism from a conventional academic perspective. Here, however, helping people in need is primarily understood as a moral obligation of responsive representation – socially prescribed attempts to act “in the interest of the represented, in a manner responsive to them” (Pitkin, 1967, 209). The urgency and moral force of these demands help make the equilibrium especially durable: refusing assistance risks communicating indifference, while providing it reinforces visible delivery as the standard through which representative care is recognized.

Precisely because deliverism responds to genuine needs and recognizable moral obligations, its costs cannot be understood as a simple failure of responsiveness. The problem is that this responsiveness is organized around tasks that federal legislators are poorly equipped to perform at any impactful scale. Members of Nigeria’s House of Representatives are formally national legislators, responsible for lawmaking, oversight, and budgeting. In practice, however, they are expected to act as local development agents and providers of personal assistance. Constituents often treat legislators as if they possess executive authority over infrastructure and public services, while the volume of need far exceeds the authority and resources available to any individual MP. The resulting mismatch ensures that even sustained efforts to deliver will almost always appear inadequate.

Constituency Development Funds crystallize this mismatch.¹² Introduced to give legislators some capacity to address local needs, they also reinforce the belief that MPs can and should deliver tangible projects. Yet their limited budgets cannot satisfy more than a fraction of the demands placed upon

⁹Bussell (2019) describes a similar dynamic in India.

¹⁰Nigerian House of Representatives member, interview no. 24, 11 December 2025.

¹¹Nigerian House of Representatives member, interview no. 08, 21 October 2025.

¹²Known as “zonal intervention projects” in Nigeria, constituency development funds are budgetary mechanisms that reserve public resources for local projects in legislators’ constituencies, typically giving legislators influence over project selection. They have become increasingly popular across Africa and beyond.

representatives – “just 1% of whatever it is that [constituents] need,” in the words of one MP.¹³ The institution thus expands legislators’ capacity to deliver only at the very margins, yet further validates a standard against which they are almost bound to fall short.

Crucially, those efforts also carry substantial opportunity costs. Time and resources spent responding to individualized demands cannot be devoted to legislation, oversight, coalition-building, or other activities more closely connected to legislators’ formal responsibilities and to the production of broader, longer-term gains. One MP described the resulting crowd-out of their “real work”¹⁴ directly:

“Ordinarily I should concentrate on bills, I should concentrate on ensuring that dividends of democracy gets down to the Nigerian people. But because of the constant demands back home, when I look at the law, I may not look at them holistically. Rather, I look at it more from the perspective of how it allows me to bring [funding] to my people. And you know I’m supposed to be a *federal* legislator.”¹⁵

Deliverism therefore constitutes a trap not because citizens’ needs are unreasonable¹⁶ or because legislators never respond to them, but because neither side can safely depart from the equilibrium alone. A representative who reallocates effort toward less visible legislative or programmatic work risks being judged inactive or indifferent while other politicians continue to publicize tangible benefits. Citizens who relax their demands for delivery risk surrendering the most legible form of accountability available to them without assurance that representatives, other constituencies, or the political system will change in turn. MPs therefore continue to demonstrate responsiveness through visible provision, thereby reinforcing the expectations that generate those efforts in the first place. The result is an equilibrium that is responsive in form, but persistently disappointing in practice: it cannot satisfy the scale of citizens’ material needs, yet it narrows the political space available for less visible work that might generate broader and more durable benefits.

2.2 Observability, Representative Type, and Political Standing

Deliverism persists, in part, because it addresses a real problem for citizens. They cannot directly observe their MPs’ intentions and cannot easily assess their effort or legislative work, so they have to rely on what they can observe and assess, primarily material delivery.

Citizens must therefore infer whether their representative is acting on their behalf from incomplete and imperfect evidence. In this light, visible delivery provides an especially legible signal. A borehole, road project, or even a personal visit to the constituency, are concrete and readily attributed to the representative. Such actions can be interpreted as evidence that the representative remembers the constituency, is willing to expend effort on its behalf, and possesses the competence to obtain resources. Less visible legislative and programmatic activities, by contrast, are harder to attribute and easier to interpret as inaction.

Delivery consequently carries meaning beyond its immediate material value. Citizens use visible and potentially costly actions to infer the representative’s underlying “type”: whether the representative is benevolent or selfish, hardworking or lazy, and close to or detached from ordinary people.¹⁷ In the words of one woman: “only God knows what is in [their] hearts and [their] intentions, and we [citizens] cannot differentiate between them.”¹⁸ Actions that entail personal sacrifice may therefore be

¹³Nigerian House of Representatives member, interview no. 14, 12 November 2025. The MPs I interviewed were themselves divided over whether constituency development funds should be expanded or abolished. Those favoring abolition argued that the funds blur legislative and executive responsibilities and distract MPs from lawmaking and oversight.

¹⁴Nigerian House of Representatives member, interview no. 01, 23 September 2025.

¹⁵Nigerian House of Representatives member, interview no. 15, 13 November 2025. This is just one of dozens of MP testimonies to the same effect.

¹⁶Though MPs I spoke to were often resentful at the excessiveness of constituents’ requests.

¹⁷I use the term “type” as shorthand for citizens’ evaluations of a representative’s character, competence, effort, and intentions.

¹⁸Author interview no. K40, 39-year-old woman. Originally in Hausa.

especially informative. Across the sampled constituencies in this study, for example, 67% of baseline respondents agreed that an MP’s willingness to use their own money to assist constituents demonstrates that they “truly care for the people.” Similarly, more than two-thirds of baseline respondents agreed that a representative who “does not bring visible projects home, is not trying hard enough.” Delivery, accessibility, and physical presence thus serve not merely as benefits in themselves, but as observable proxies for otherwise hidden effort and intent.¹⁹ As one woman explained to me, it is “by giving us things that help life, like drinking water, electricity, roads, [that] we know we truly have representatives.”²⁰

The problem is not merely that citizens lack information about what representatives do. They also lack assurance about whose interests representatives consider when acting beyond public view. Citizens may know that an MP attends parliamentary sessions or participates in committees without trusting that this activity is undertaken with their interests at heart. Evaluating less visible representation therefore requires more than information about the quantity of legislative effort. It requires *warranted trust*: some credible reason to believe that actions undertaken in citizens’ absence will nevertheless take their interests into account (Hardin, 2002; Warren, 2018).

I suggest that such trust can rest on two complementary foundations. The first is a favorable assessment of representative type. Citizens who believe that their representative is benevolent, competent, and willing to exert effort have greater reason to interpret otherwise hidden activity as purposeful representation rather than indifference or self-enrichment. The second is a sense of political standing: the condition in which citizens can expect their interests to count in routine governance, can make themselves heard when those interests are misunderstood or neglected, and can expect political authorities to respond. A deficit of standing is not identical to a shortage of material goods, even though they may often go together (Kumar, 2022, 2025). Similarly, standing is related to, but distinct from, external political efficacy, which concerns citizens’ perception of their capacity to influence government (Craig et al., 1990). Standing refers to citizens’ relatively durable, though revisable, beliefs about their place, worth, and status in relation to political authority.²¹ Citizens infer this standing from concrete encounters with political actors and institutions, treating each interaction as evidence about how they are regarded. Standing therefore operates as a stock rather than a momentary experience: a running tally that citizens update as evidence of attention, inclusion, and respect – or disregard – accumulates over time (Hastie and Park, 1986; Lodge et al., 1989). Although no single encounter fully determines this tally, an unusually clear and credible interaction can revise it.

Political standing is important precisely because citizens cannot continually supervise their representatives. If citizens believe that their interests are ordinarily taken seriously, they have greater reason to trust that a representative’s activities, and those of the political system as a whole, will be conducted with those interests in mind, even when the activities themselves remain imperfectly observable. If citizens instead believe that they possess little standing, visible delivery becomes indispensable proof that the representative as an individual, and the system in full, have not ignored or abandoned them.

These two foundations of warranted trust correspond to distinct inferences. Beliefs about representative type concern the agent in the representative relation: *Is this a person who is willing and able to act on our behalf?* Political standing concerns the citizen’s own place within the relationship: *Are people like me among those whose interests those in power can be expected to consider?* In combination, favorable answers to these questions may reduce the burden placed on immediate material delivery as the principal evidence that representation is occurring.

¹⁹Kramon (2018) shows how valuable prospective signals of delivery capacity can be in the context of African elections.

²⁰Author interview no. E24, 46-year-old woman. Originally in Igbo.

²¹The measures of standing and external political efficacy overlap somewhat, since both capture citizens’ perceived relationship to political authority. I distinguish them analytically because efficacy concerns influence and governmental responsiveness, whereas standing concerns recognition and relational status.

2.3 Deliberative Encounters as Alternative Signals

Demands for observable evidence of representation are reasonable where information is scarce, trust is low, and poverty remains widespread. The goal, then, is not to lower citizens’ standards of evidence but to give them more and better evidence.

Deliberative town halls (DTHs) are designed to provide such evidence. Their purpose is not to substitute discussion for development, material provision, or institutional reform. Rather, they intervene in the informational and relational environment through which citizens interpret representative behavior. This form of interaction differs from ordinary constituency contact in three important ways. First, it is collective rather than individualized. Constituents participate as members of a local public, not as isolated petitioners seeking private assistance. Second, it is dialogic rather than purely performative. Citizens can ask questions, while representatives have an opportunity to explain what they have done, what they can and cannot do, and how they understand their role. Third, it makes citizens’ standing within the representative relationship observable. By positioning citizens as entitled to question their representative, and placing the representative under a public expectation of response, the encounter provides evidence that citizens’ interests and voices count.

DTHs are also theoretically distinctive in the form of participation they create. They facilitate *vertical deliberation* between citizens and political elites, rather than deliberation among citizens alone. The questions, comments, and priorities originate with citizens, while representatives must respond publicly, making their explanations visible to a broader constituency and open to collective scrutiny. Moreover, citizens are randomly assigned the opportunity to participate, rather than recruited entirely through self-selection. Actual attendance remains voluntary, of course, but the initial opportunity is not restricted to those already most politically engaged or predisposed to seek contact with their representative. DTHs are by now a well-established method, having been implemented in seven countries (Neblo et al., 2018). The Nigerian implementation, which began with an earlier pilot (Castradori et al., 2024), represents the first extension of this model into the Global South, and extends it to a setting in which citizens frequently call for the very relationship the model seeks to foster.

Although related to prior studies in citizen-representative relations in Africa, this design differs from the deliberative campaign meetings studied in Benin (Fujiwara and Wantchekon, 2013; Wantchekon and Guardado, 2024) in at least three respects: it takes place between elections rather than during a campaign, potentially reducing suspicion that the encounter is motivated solely by immediate electoral considerations; it centers the record of a sitting representative rather than the platform of a candidate; and it examines *relational* outcomes – perceived representative type and citizens’ political standing – rather than vote choice. It also differs from informational interventions that provide performance facts without direct relational contact, whose effects have been consistently weak across contexts (Dunning et al., 2019).²²

The DTH design builds on the insight that better representation, in the eyes of many Nigerians, is not only a matter of more goods. It is also a matter of access, explanation, and regular interaction. One citizen put the point in explicitly informational terms:

“The representatives need to be closer to the people because they are part of us [...] they do not need to keep away from us, they need to be part of us because we send them there. [...] [When] we interact regularly, we ask what is happening over there [and] even some of the wrong notions we have about them, about how politics is being played, will be cleared! [...] It is better for them to meet us on a regular basis to tell us what they have been doing and how they are representing us.”²³

The quotation separates interaction from material extraction. The respondent values contact

²²In much of sub-Saharan Africa, citizen-representative linkages run through traditional leaders, brokers, and other intermediaries (Baldwin, 2013, 2025; Koter, 2013). The DTH does not displace these channels, but rather adds a direct one. Recent DTHs in Malawi should help illuminate how local authority structures shape take-up of encounters like these, which is an important question for the institutionalization agenda this paper motivates.

²³Author interview no. E03, 58-year-old man.

partly because it can make representation intelligible. Citizens can ask what is happening; representatives can explain what they have done; “wrong notions” can be corrected. In a structured encounter, citizens can observe whether a representative is willing to appear before them, answer questions, explain constraints, and respond to criticism.

Other citizens I spoke to described the same need in more affective language: “sometimes people are not wanting too much,” one woman explained, “just a little attention is fine. [...] If you show people concern, *even if you are not perfect*, they will appreciate you. [...] [It shows] you are aware of them.”²⁴ Such statements do not imply that material needs are unimportant. Rather, they suggest that attention and recognition may shape how citizens interpret inevitable shortfalls in delivery. MPs too, often describe the same logic. Although they experience constituency access as burdensome, they also recognize that distance is politically dangerous. As one MP put it, “[people] trust you more when they can access you.”²⁵

Events like this provide citizens evidence about both their representative’s type and their own standing, greatly increasing the opportunity to build warranted trust. This account generates two immediate expectations. First, a deliberative encounter should improve citizens’ assessments of the representative’s character, competence, effort, accessibility, and responsiveness (**H1**). Second, it should increase citizens’ perceived standing, both vis-à-vis the representative (**H2**) and, potentially, within the political system more broadly (**H3**). A legislator is encountered not only as an individual politician but also as a representative of the broader democratic order. An encounter that affirms citizens’ standing vis-à-vis the representative may therefore also alter whether they see people like themselves as legitimate participants whose interests and voices count within the political system as a whole.

To the extent that the encounter’s effects generalize beyond the participating representative, it should also increase citizens’ trust in elected officials more broadly – especially trust grounded in perceptions of their accountability, care, and competence (**H4**). These broader changes may, in turn, improve satisfaction with democracy and hope that the political system can respond effectively in the future (**H5**).

The Sequencing Argument Summarized: In my account, representative type and political standing provide the theoretical foundations for changing representational priorities. The most downstream (and speculative) expectation of the theory is therefore that DTH assignment may increase support for legislative, programmatic, and longer-term representation relative to immediate constituency delivery (**H6**).²⁶ The argument does not, however, require all stages of this process to move simultaneously. The immediate object of a deliberative encounter like the DTH is the citizen–representative relationship. It supplies new evidence about the representative and gives citizens a direct experience of how they are regarded. Evaluations of representative type and perceived standing should therefore be the outcomes most immediately responsive to a single encounter. These changes may, in turn, extend to broader judgments about democratic responsiveness, satisfaction, and hope.

Representational priorities are more distant. Shifting them requires citizens not only to update their view of one representative, but also to become more willing to accept forms of representation whose effort and benefits remain difficult to observe, may emerge only over time, and may not provide immediate relief from pressing material needs. Citizens may reasonably wait to see whether the recognition communicated during a town hall is durable, whether the representative follows through, whether comparable opportunities recur, and whether other citizens are similarly included. Sustained changes in representational demands may therefore require repeated encounters, social diffusion, reflection, or additional evidence that the new relationship will persist.

The theory consequently implies an ordered process rather than uniform contemporaneous movement across outcomes: Deliberative engagement should first affect citizens’ evaluations of the repre-

²⁴ Author interview no. K10, 35-year-old woman.

²⁵ Nigerian House of Representatives member, interview no. 11, 4 November 2025.

²⁶ All hypotheses were pre-registered at [OSF link](#). Note, the paper presents hypotheses in the causal order the theory implies rather than in the order of the preanalysis plan; a table mapping the paper’s hypothesis labels to the registered ones is in Appendix [H](#).

sentative and their perceived standing, with the rebalancing representational priorities is the most downstream consequence. That becomes possible when citizens have accumulated sufficient reason to trust that less observable political activity is both effective and conducted with their interests at heart. The central empirical question is thus how far a single structured encounter travels along this causal sequence. Evidence of change in representative evaluations, standing, and broader democratic attitudes would indicate that the encounter affected the informational and relational foundations of the deliverism trap. An immediate change in representational priorities would constitute the furthest-reaching outcome and, if my theory of change holds, the first step to altering the incentives that sustain the deliverist equilibrium in the practice of representation.

3 Nigeria: A Demanding Context for Representative Democracy

As Nigeria approaches the 2027 general election, its democratic trajectory carries significant weight. With nearly 240 million people, it is Africa’s most populous country and one of the world’s largest electoral democracies; by mid-century, it is projected to become the world’s third most populous country. Its democratic performance therefore matters not only for its own citizens but for the wider trajectory of representative government in West Africa at a moment when democratic rule across the region is under mounting strain (Bratton et al., 2024; Opalo, 2024).

Nigeria has operated as a federal presidential republic since the return to civilian rule in 1999. The president heads the executive, while federal legislative authority rests with a bicameral National Assembly composed of a 109-member Senate and a 360-member House of Representatives. Members of the House are elected from single-member federal constituencies for four-year terms. This study focuses on these representatives, whom I refer to as MPs. Their formal responsibilities include legislation, appropriation, and executive oversight. They do not themselves administer ministries or directly implement most public services; even the constituency projects they help place in the federal budget are formally executed by government agencies. Yet they are often the most visible federal officials to whom citizens can direct demands for roads, water, employment, school fees, medical assistance, and other immediate needs. The resulting mismatch between legislators’ formal authority and the functions citizens need political actors to perform is central to the representational pressures I examine.

The sociopolitical environment makes those pressures particularly acute (LeVan, 2019). Nigeria encompasses hundreds of ethnolinguistic groups, deep religious pluralism, sharp regional inequalities, rapid urbanization, and one of the world’s largest youth populations. These features generate vibrant political competition but also intense distributional pressures. They are compounded by widespread material deprivation, uneven state capacity, and limited access to public services. Governing at this scale, across such diversity and inequality, places exceptional strain on democratic institutions and on the representatives operating within them.

To some extent, political parties are salient in this system, even though they rarely compete through differentiated programs in the conventional sense. Parties control nominations, organize electoral competition, distribute legislative positions, and structure elite careers (Demarest, 2022). At the same time, shifting coalitions, frequent party switching, and limited programmatic differentiation weaken the credibility of party labels and their relevance independent of the identities of the elites that compose them. Partisanship therefore matters, but it often operates through citizens’ evaluations of the individual politicians who carry party labels rather than through stable attachments to competing policy programs. In such a setting, a party-agnostic intervention like a DTH may be especially consequential. By engaging citizens as constituents rather than as co-partisans, it can provide credible evidence of representative commitment that cuts across, and potentially outweighs, prior judgments based on party labels.

Nigeria’s social and political features help clarify the scope conditions of my argument. The effects of deliberative encounters on perceived representative type and political standing should be greatest

where legislative work is difficult to observe, contact with political elites is rare and ordinarily individualized, and unmet demand for meaningful interaction is high. Conversely, representational priorities – what citizens value most in their representatives’ conduct – should be most resistant to change where material precarity is severe and programmatic commitments lack credibility. This divergence follows from the level at which these beliefs are formed: perceptions of representative type and standing can respond to evidence arising within a particular political relationship, whereas representational priorities reflect accumulated experience with the capacity and credibility of the political system as a whole. A direct encounter can provide new evidence about an individual representative’s intentions, effort, and regard for citizens. It can also alter citizens’ sense of whether their voices and interests count. Still, priorities are informed by broader beliefs about what the political system is capable of delivering. Where parties cannot make credible programmatic commitments and state institutions cannot reliably implement them, citizens have good reason to continue demanding benefits that an identifiable politician might plausibly secure (Keefer and Vlaicu, 2008; Kitschelt and Wilkinson, 2007). Similarly, a single DTH does little to alter broader empirical expectations about the conduct of other legislators in other constituencies: citizens have little reason to relax their demands for near-term constituency benefits unless meaningful engagement appears likely to become a durable feature of the broader political order – repeated by their own representative, adopted by other politicians, and reinforced by voters elsewhere.

Nigeria lies near the demanding end of each of these dimensions. Legislative activity undertaken in Abuja is largely hidden from ordinary constituents, while access to representatives is scarce, individualized, and widely perceived as beyond their reach. The resulting unmet demand for interaction is substantial and emotionally laden. One man expressed the stakes directly: “How can you be happy when access to people who represent you is not there, when their security details will drive you away... how can you feel good about that?”²⁷ A structured encounter like a DTH may therefore supply unusually novel and credible evidence about a representative’s type and about citizens’ standing, making movement in these relational outcomes especially plausible. At the same time, acute material need and uneven state capacity give citizens little reason to revise their broader doubts about the credibility of programmatic promises. The encounter may do little to alter the system-level conditions underlying citizens’ representational priorities, making demands for immediate delivery particularly sticky. Nigeria is therefore both fertile ground for detecting change in the theory’s proximate mechanisms and a hard test of its most downstream implication.

For these reasons, Nigeria provides both a high-stakes case and a theoretically revealing setting in which to examine how citizen-representative relationships shape democratic legitimacy and durability. If structured engagement can strengthen warranted trust under these conditions – even without immediately displacing delivery as the dominant standard of representation – it would show that citizens’ interpretations of political behavior are not fixed and identify a pathway through which repeated or institutionalized encounters might gradually reshape the everyday practice of representation.

4 Research Design and Data

To test whether structured interaction can shift the signals through which citizens evaluate representation, I conducted a cluster-randomized field experiment in three Nigerian federal constituencies. The design assigned local enumeration areas to receive invitations to a deliberative town hall with their sitting member of the House of Representatives or to a control condition. Respondents were surveyed before and after the intervention, allowing for baseline-adjusted estimates of the effect of assignment to the DTH invitation.²⁸

²⁷ Author interview no. E03, 58-year-old man.

²⁸ The study was approved by the Princeton University IRB, protocol no. 19471.

4.1 Study Sites

The intervention took place in three federal constituencies: Bende, a rural constituency in the largely Igbo-speaking Southeast; Obio-Akpor, a bustling and heterogeneous urban constituency in Port Harcourt, the largest city in the oil-rich Niger Delta; and Ikole/Oye, a predominantly agrarian, Yoruba-speaking constituency in the country’s Southwest. These sites provide variation in region, language, urbanization, and local political context, while holding constant the formal institutional role of the participating representatives: all three are members of Nigeria’s House of Representatives.²⁹

4.2 Sampling and Random Assignment

We first randomly selected a total of 80 enumeration areas (EAs) from each of our three constituencies.³⁰ The EA boundaries are drawn independently by the Nigerian Population Commission (NPopC). The EAs functioned as the unit of random treatment assignment. Randomizing at the EA rather than at the individual level served two purposes. First, the DTH invitation was implemented through local, in-person mobilization,³¹ making it natural to assign geographically defined clusters rather than isolated individuals within the same neighborhood. Second, cluster-level assignment reduced the risk that treated and control respondents living in immediate proximity would receive different invitation messages while interacting regularly with one another.

Within each constituency, sampled EAs were randomly assigned to treatment or control with equal probability.³² Respondents in treatment EAs were eligible to receive invitations to attend the DTH with their sitting MP. Respondents in control EAs did not receive an invitation. The target baseline sample included approximately 600 respondents across the 80 enumeration areas per constituency, with 7 or 8³³ respondents interviewed in each EA.³⁴

Because assignment occurred at the EA level, the primary estimand is the intent-to-treat effect of residing in a community assigned to receive an invitation. This estimand captures the effect of the invitation strategy as implemented, including imperfect attendance among invited respondents and any within-community diffusion generated by the mobilization process.

4.3 The Deliberative Town Hall Intervention

The treatment consisted of an invitation to attend a deliberative town hall with the respondent’s sitting House member.³⁵ Because reliable individual internet access is rare, constituents gathered in five local venues, equipped and staffed by the research team, while the MP joined live by video from

²⁹The three MPs and their constituencies were selected during earlier fieldwork in Nigeria, during which I interviewed more than thirty sitting MPs about their experiences in office. Although they vary substantially in party affiliation, constituency characteristics, and political ambitions, they clearly constitute a purposive rather than representative or randomly selected sample.

³⁰240 total enumeration areas were therefore included in the initial sample. For reasons ranging from enumerator security to en masse out-migration out of sampled villages due to flooding, a total of nine EAs across the three constituencies could not be sampled. Unless otherwise stated, all analyses in this paper thus use a sample size of 231 total EAs. See Appendix B for details on missing EAs.

³¹We collaborated with an experienced Nigerian pollster to recruit local enumerators to conduct all baseline and endline surveys in person. All invitations for DTHs were delivered in-person at the end of the baseline survey by the enumerators.

³²In the constituencies of Bende and Obio-Akpor, there was no further stratification. Because Ikole/Oye federal constituency is composed of two local governments (LGAs) we stratified treatment assignment by LGA such that each had 20 EAs for each condition.

³³Which EAs had 7 or 8 was randomized by flipping a coin.

³⁴Note: The final sample exceeds the pre-registered target. The pre-analysis plan registered approximately 480 respondents per constituency (1,440 in total). Due to a misunderstanding with the polling company, enumerators interviewed 7–8 respondents per EA, yielding approximately 600 per constituency. None of the sampling procedures, treatment assignment, or EA selection differed from the pre-registered protocol, and I was not aware of the discrepancy until after receiving the data. Because the sampling process followed the pre-registered steps and the sample exceeds the registered target, the discrepancy increases precision without threatening the validity of the preregistered analyses.

³⁵See example images in Appendix A.

Abuja. Throughout the event, the research team collected and organized questions, and ensured that the discussion included both citizen concerns and representative responses. The format was intended to create a public but structured encounter: more accessible to the average citizen than ordinary elite visits; more structured than individualized petitioning that characterizes how citizens encounter political representatives (when they do). The questions gathered from each of the five venues were passed on one by one to the MP who then answered them live on video for participants to hear. Each event lasted between 62 and 76 minutes.³⁶ At the conclusion of the session with the MP, attendees participated in focus group discussions with the other constituents gathered in their respective venues. Below, I draw from the transcripts of these discussions to gain a richer understanding of participants' interpretation of the day's proceedings.

The DTH format had three central components. First, the encounter was collective. Constituents participated alongside others from their area, which shifted the interaction away from individualized petitioning and toward a more public discussion of representation. Second, the encounter was dialogic. Citizens were able to raise questions and concerns, while the representative had the opportunity to explain his activities, constraints, and understanding of his role. Third, the encounter was structured. The discussion was organized such that the MP was exposed to citizen demands, but the event did not become simply a queue for personal requests.

This structure was central to the theory of the intervention. Ordinary constituency access often reinforces deliverism because it occurs through individualized claims for urgent assistance.³⁷ By contrast, the DTH was designed to preserve the relational value of access while changing the form of interaction. Citizens could see and question their representative, but in a setting that also made explanation, listening, and collective accountability more salient. The absence of physical co-presence also made clear that the encounter would not involve the immediate distribution of cash or other material benefits. The goal, therefore, was to spotlight non-material dimensions of representation.³⁸

4.4 Survey Waves and Compliance

The study included a baseline survey before the DTH and an endline survey after the intervention.³⁹ The baseline survey measured pre-treatment attitudes toward representation, evaluations of the MP, perceived political standing, and demographic and political covariates. The endline survey repeated the main outcome measures, allowing treatment effects to be estimated while adjusting for baseline values of the same outcomes.

Among respondents selected in each EA assigned to treatment, not all attended the DTH.⁴⁰ This noncompliance is substantively important. The main analysis therefore reports intent-to-treat effects of assignment to the DTH invitation, preserving the original randomization. In secondary analyses, I estimate the effect of attendance among compliers using assignment as an instrument for attendance (Appendix I.4). These estimates should be interpreted with the usual exclusion restriction in mind: among some possible threats to the assumption, the invitation may itself signal effort or accessibility even for respondents who do not attend. For this reason, the ITT estimates remain the primary

³⁶See Appendix A for some further details about each DTH implementation.

³⁷These are most often conducted through WhatsApp. Most MPs I spoke to carried several phones (one carried four) which they use to communicate with and keep track of constituent needs.

³⁸Participants were given a stipend equivalent to approximately USD 5 to cover their travel expenses to reach their assigned venues.

³⁹The baseline survey was conducted between 10 and 5 days from the scheduled town hall in the relevant constituency. The endline 5-10 days after.

⁴⁰Even so, compliance (measured as invited respondents individually verified as attending, via check-in against the sample list) across the three constituencies reached 62.5% – nearly two-thirds of all those assigned to treatment. This is an exceptionally high rate for an intervention that required participants to travel physically to a venue and devote several hours to the event. Comparable deliberative town hall experiments in advanced democracies have generated substantially lower participation rates – approximately 37% in the United States (Neblo et al., 2018). Compliance in Nigeria was therefore more than 25 percentage points higher, despite the considerable practical costs of attending. That so many randomly invited citizens were willing to bear those costs is itself substantively revealing, suggesting an unusually strong unmet demand for more direct and engaging forms of democratic politics.

evidence.

Spillovers are also possible. Respondents in control areas may hear about the DTH from acquaintances, social media, or local political networks; respondents in treatment areas who do not attend may hear about the event from neighbors. Such spillovers would tend to blur the contrast between treatment and control conditions. The estimates should therefore be interpreted as the effect of assignment to an invitation strategy under realistic conditions of political communication, not as the effect of perfectly isolated exposure.⁴¹

4.5 Outcome Measures

The main mechanism-related outcomes measure two theoretically distinct forms of updating. The first concerns evaluations of the representative’s type: whether citizens view the MP as caring, hard-working, competent, accessible, and oriented toward the public interest. The second concerns political standing: whether citizens feel that people like them are listened to, taken seriously, and able to make claims on their own representative, and on government as a whole.

Additional outcomes examine satisfaction with democracy, democratic optimism and a range of other political and social attitudes.

The final outcome battery measures citizens’ representational priorities. These items ask whether citizens place greater weight on immediate constituency delivery or on broader forms of representation, including legislative work, programmatic performance, and long-term development. The related hypothesis is that assignment to the DTH reduces the exclusive priority placed on immediate, visible, constituency-targeted delivery and increases the weight citizens place on other dimensions of democratic representation.⁴²

The next section describes the construction of these measures and the estimating equations used to evaluate treatment effects.

4.6 Model Specifications

Three of the principal item batteries used to operationalize the outcomes in my theory lack established measurement properties in existing public-opinion research.⁴³ It was therefore not appropriate to impose a unidimensional measurement structure a priori or to assume that all items reflected a single latent construct. Accordingly, I preregistered an item-stacked specification as the primary analytical approach, rather than relying exclusively on a composite index that aggregates the full battery into a single outcome.

The stacked specification restructures each battery into respondent-item observations, such that each respondent contributes one observation for every item answered. Post-treatment item responses are then jointly regressed on treatment assignment and their corresponding baseline values, with item fixed effects accounting for systematic differences in response levels across questions and constituency fixed effects accounting for differences between study sites. The treatment coefficient therefore estimates the average intention-to-treat effect across the outcomes included in the battery, with standard errors clustered at the level of treatment assignment (the EA). This approach does not require the items to be interchangeable indicators of a single latent construct or require the imposition of index weights derived from their covariance structure. It also uses the available item-level information, including responses from participants who may be missing other items in the battery, while generally providing greater precision than a series of separate item-by-item tests (which I separately estimate and present in Appendix F).

⁴¹For what it is worth, should the theory I propose hold water, any spillover across treatment conditions should dampen the magnitude of the estimated treatment effects. The possibility of spillover, if at all relevant, should make identifying statistically significant treatment effects more difficult.

⁴²The exact wording of all the items analysed in this draft are in Appendix C.

⁴³**Standing_MP** measures respondents’ perceived political standing vis-à-vis their own representative; **Standing_Gov** captures their perceived standing within the political system more broadly; and **Rep_Priorities** assesses the relative importance citizens assign to different dimensions of legislative representation.

For analyses on stacked items, I therefore estimate the following linear model:

$$Y_{iej1} = \alpha + \beta T_e + \rho Y_{iej0} + \lambda_j + \delta_c + \varepsilon_{iej}, \quad (1)$$

where Y_{iej1} is the endline response of respondent i in enumeration area e to item j , Y_{iej0} is the respondent's baseline response to the same item, T_e is an indicator for assignment of enumeration area e to treatment, λ_j denotes item fixed effects, and δ_c denotes constituency fixed effects. The item fixed effects account for average differences in response levels across items, while the same-item baseline covariate adjusts for each respondent's pre-treatment response to that item. The coefficient β estimates the average ITT effect across the stacked item family. Standard errors are clustered by enumeration area, the level of treatment assignment.

PCA Indices. For batteries whose items exhibit sufficient internal consistency to support their aggregation into a common scale, I follow the preregistered procedure and construct an index using principal component analysis (PCA). This procedure applies to two outcomes: the **Rep_Type** battery, which uses conventional items to assess the representative's character, competence, and intentions (Fiske and Durante, 2014; Laustsen and Bor, 2017), and the **Trust** battery, which measures generalized trust in elected officials across dimensions of accountability, care, and competence. For both batteries, the item distributions, pairwise correlations, PCA loadings, and coefficient theta indicate that the responses are sufficiently dominated by a common first dimension to justify summarizing them as a single scale.⁴⁴ I therefore use the resulting PCA indices as the primary outcomes, rather than estimating the effects using stacked item-level specifications.

After recoding all items so that higher values indicate more favorable evaluations of the representative, I standardize the items and estimate the principal-component loadings using baseline data. The same standardization parameters and loadings are then applied to the endline responses, ensuring that changes in the index are not mechanically induced by treatment-related changes in the covariance structure. Missing item responses are handled using the preregistered mean-imputation procedure before the index scores are calculated.

For each index, I estimate the treatment effect using the following baseline-adjusted specification:

$$Y_{ie1} = \alpha + \beta T_e + \rho Y_{ie0} + \delta_c + \varepsilon_{ie}, \quad (2)$$

where Y_{ie1} and Y_{ie0} are the endline and baseline index scores, respectively. As above, β estimates the ITT effect, constituency fixed effects are included in the pooled analysis, and standard errors are clustered by enumeration area.

5 DTH Results

The experimental results reveal a clear distinction between the two stages of the proposed causal process: despite its explicitly non-material character, assignment to a DTH improved citizens' evaluations of their representative, strengthened their perceived political standing, and their trust in elected officials. It also increased democratic satisfaction and political hope. It did not, however, alter representational priorities, suggesting that a single structured encounter can change how citizens perceive political authorities and their relationship to them without changing what they demand from their representatives.

⁴⁴Coefficient theta exceeded 0.6 for both batteries. Full reliability diagnostics are reported in the appendix.

Representative Type

I begin by examining whether assignment to a deliberative town hall changed citizens' evaluations of their representative (H1).⁴⁵

I argue that citizens and MPs get stuck in the deliverism trap because of a dearth of evidence for accountability. So before asking whether the treatment improved evaluations of the MP, it is worth asking whether it made evaluation possible at all. Citizens had a "Don't know enough to tell" (DK) option on the survey questions about the MPs' character and competence.

Consistent with the preregistered expectation, assignment to a DTH reduced the probability of DKs across the eight representative-type items by 6.55 percentage points ($p < 0.01$). The effect is concentrated in Obio-Akpor, where treated respondents went from being unable to answer an average of four of the eight items at baseline, to under two by the endline. Many citizens changed their judgments of the MP after the DTH, but many others formed judgments for the first time. Though a happy normative result for democracy, this also creates a measurement challenge in my study. I coded respondents who answered fewer than five of the items as missing for the index. Thus, the treatment created a larger pool of respondents with a substantive Type index value. The constituency-level `Rep.Type` estimates should be read with that caution in mind. In the appendix, I report analyses (including Lee bounds (Lee, 2009)) probing the consequences for inference.⁴⁶

Figure 1 presents the intent-to-treat estimates for each constituency separately and for the pooled sample, with constituency fixed effects. Across the pooled sample, assignment to a DTH improved citizens' evaluations of their elected representative. The aggregate effect appears to be concentrated in Obio-Akpor and Ikole/Oye, where the estimated effects are positive, while the estimate in Bende is close to zero. These results provide initial evidence that structured political encounters can improve citizens' assessments of their representatives even in the absence of direct material provision of the type that so often characterizes elite-citizen relations (in expectation, and in practice). By giving MPs an opportunity to explain their legislative activities and plans, respond to constituents' questions, and demonstrate their willingness to engage, the DTH made otherwise difficult-to-observe representative qualities more visible.

Political Standing

Hypotheses 2 and 3 predict that exposure to a deliberative town hall increases citizens' sense of political standing. I define political standing as citizens' relatively durable, though revisable, belief that people like them are recognized as legitimate political claimants whose views merit consideration and who are entitled to address political authorities and expect a response.

Political standing may operate differently within citizens' direct relationship with their MP than within their broader relationship with democratic government. I therefore measure the concept using two parallel batteries. The batteries share the same underlying structure but differ in the object of evaluation: one refers to "our representative," while the other refers to "elected officials" more generally. Table 1 presents the estimated treatment effects for MP-specific and generalized political standing, using the stacked model specification described in section 4.6.

At the aggregate level, both measures move in the hypothesized direction. Following the DTH, respondents assigned to treatment reported greater political standing than respondents assigned to the control condition. The intervention thus appears to have strengthened citizens' sense that their

⁴⁵The index is constructed from the `Rep.Type` battery reported in the Appendix. The battery asks respondents to evaluate their MP on eight qualities: dishonest, trustworthy, selfish, understands people like them, hardworking, inaccessible, unqualified, and intelligent. All items are coded so that higher values indicate a more favorable evaluation of the MP. A figure in the Appendix shows treatment effects by individual item.

⁴⁶The pooled treatment difference in the endline index's *availability* is small and insignificant (1.5 points, $p = 0.60$). The effect is specific to Obio-Akpor (+19.1 points, $p < 0.001$; Bende -0.6, Ikole/Oye -7.8, neither distinguishable from zero). Under the monotonicity assumption that treatment can induce but not suppress availability, the complete-case Obio-Akpor estimate of 0.157 SD has a Lee-identified set of [-0.102, 0.503] SD. Of the 706 respondents excluded from the both-wave index sample, 534 were already below the availability threshold at baseline; the post-treatment selection concern is confined to the remaining 172. Appendix D reports the full analysis.




Figure 1: Intent-to-treat effects on evaluations of representative type

interests would be considered, that they could make themselves heard, and that political authorities would respond – both in relation to their individual MP and to elected officials more broadly.

The constituency-specific estimates are also consistently positive, although their precision varies. In Obio-Akpor, the estimated effects on both MP-specific and generalized standing are statistically distinguishable from zero at conventional levels. The corresponding estimates in Bende are positive but less precisely estimated and do not reach conventional levels of statistical significance. Overall, the consistency in the direction of the estimates across constituencies and measures supports the expectation that structured engagement can strengthen citizens’ perceived standing in their relationships with political authorities.

Generalized Trust in Elected Officials

My theory treats warranted trust as the hinge connecting perceptions of representatives’ qualities and citizens’ political standing to eventual shifts in representational priorities (H4). Although related, the three constructs capture distinct stages in this process: *type* concerns citizens’ assessments of the particular MP; *standing*, their perceptions of their own position within the political system; and *generalized trust*, the broader downstream attitude these two foundations jointly support. Citizens evince warranted trust when they believe that actions undertaken in their name – especially those they cannot directly observe – are guided by their interests.

To assess the effect of treatment on trust, the endline survey included a battery measuring generalized trust in elected officials. Following the preregistered procedure, I combined these items into a PCA index. Table 2 presents baseline-adjusted intent-to-treat estimates. Assignment to a DTH increased generalized trust by approximately 0.3 standard deviations in the pooled sample. The constituency-level pattern mirrors that of the other relational outcomes: effects are largest in Ikole/Oye, positive and statistically significant in Obio/Akpor, and approximately zero in Bende, where the MP’s late arrival likely weakened the encounter’s signal of respect, as discussed further below.

Table 1: Treatment Effects on Political Standing

Sample Model:	Treatment Effects on Political Standing			
	Pooled sample (1)	Bende (2)	Obio/Akpor (3)	Ikole/Oye (4)
<i>Panel A: Political standing vis-a-vis the individual MP</i>				
Assigned to Treatment	0.246*** (0.078)	0.112 (0.116)	0.302*** (0.099)	0.325* (0.177)
Baseline adjustment	Yes	Yes	Yes	Yes
Item fixed effects	Yes	Yes	Yes	Yes
Constituency fixed effects	Yes	No	No	No
Observations	11,454	3,731	3,947	3,776
EA clusters	231	78	78	75
<i>Panel B: Generalized political standing</i>				
Assigned to Treatment	0.281*** (0.073)	0.123 (0.101)	0.313*** (0.093)	0.415** (0.168)
Baseline adjustment	Yes	Yes	Yes	Yes
Item fixed effects	Yes	Yes	Yes	Yes
Constituency fixed effects	Yes	No	No	No
Observations	11,647	3,754	4,091	3,802
EA clusters	231	78	78	75

Each panel reports a stacked respondent-item OLS model across the six items in the corresponding battery. All models adjust for the respondent's baseline response to the same item and include item fixed effects. The pooled models include constituency fixed effects.

All items are 5-point Likert scales. Standard errors are clustered at the EA level.

*Observations are respondent-item observations. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.*

Table 2: Treatment Effects on Generalized Trust in Elected Officials

Dependent Variable: Sample: Model:	Generalized Trust			
	Overall (1)	Bende (2)	Obio-Akpor (3)	Ikole/Oye (4)
Assigned to Treatment	0.3031*** (0.0985)	0.0287 (0.1402)	0.2830** (0.1202)	0.6138*** (0.2190)
Baseline adjustment	Yes	Yes	Yes	Yes
Constituency fixed effects	Yes			
Observations	1,948	627	685	636
EA clusters	231	78	78	75

*Standard errors are clustered at the EA level. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

The outcome is the first principal component of six generalized-trust items, oriented so that higher values indicate greater trust in elected officials.

Democratic Evaluations

I next examine three secondary outcomes: satisfaction with democracy, hope that democracy will improve over the next five years, and hope that government can solve the country’s problems. These outcomes are not central to the paper’s main theoretical predictions, but nevertheless capture broader orientations that may respond to the same relational changes. If the DTH improves citizens’ evaluations of their representative and strengthens their sense that they are recognized and taken seriously by political authorities, it may also increase satisfaction with democracy and confidence in the political system’s future capacity to govern effectively.

Figure 2 presents the intent-to-treat effects on all three outcomes. In the pooled sample, assignment to the DTH produced positive and statistically significant effects on satisfaction with democracy, hope that democracy will improve over the next five years, and hope that government can solve problems. The consistency of these effects across the three measures suggests that the intervention influenced not only respondents’ immediate relationship with political authorities but also their broader assessments of democratic performance and political possibility.

The constituency-specific estimates are uniformly positive, although their precision varies. The effects in Ikole/Oye are consistently statistically distinguishable from zero across all three outcomes. Estimates in the other constituencies point in the same direction but do not reach conventional levels of statistical significance. Taken together, these results complement the findings on representative type and political standing: the more favorable evaluations generated by the DTH appear to extend beyond citizens’ immediate relationship with their representative to their broader satisfaction with democracy and hope in government’s future performance.

Priorities for Representation

The next step in the proposed causal chain is that improvements in citizens’ evaluations of their representative and in their sense of political standing should ultimately reshape what they demand from representation. In particular, Hypothesis 6 predicts that these changes will reduce the relative weight citizens place on short-term, local, and “deliverist” behavior and increase their support for legislative, programmatic, and longer-term forms of representation. Such a shift matters because the weight placed on immediate and visible constituency delivery can weaken incentives for effective legislative action and impede cooperation around less tangible or longer-term goals.

The results do not support this hypothesis. Figure 3 presents the intent-to-treat estimates from the stacked model for the **Rep.Priorities** battery. The pooled estimate is not statistically distinguishable from zero, nor are any of the constituency-specific estimates. The estimates are precise enough to be informative: the pooled 95% confidence interval of $[-0.165, +0.055]$ on five-point items excludes even modest shifts toward less deliverist priorities.

The item-level estimates presented in Appendix F yield the same substantive conclusion. Assignment to the DTH produced little discernible movement on any individual item in the battery. The absence of an aggregate effect therefore does not appear to conceal offsetting changes across different dimensions of representational priorities.

How should we interpret this null finding? The DTH did not move the theory’s final link. One might reasonably interpret this pattern as revealing something important about the nature of citizens’ representational preferences. ‘Type’ is entirely dyadic between a given constituent and her representative. And while standing and trust are broader in scope, citizens can still update their baseline attitudes from the ambient cultural baseline through their direct experience with a particular MP as an exponent of the broader political system.

Representational priorities are different in kind. Demanding less immediate delivery from one’s own MP is sensible only if citizens expect an alternative mode of representation to persist and spread: their own MP must continue engaging, other MPs must follow suit, and other constituents must begin holding their representatives to the same standard. Without these mutually reinforcing expectations, a single encounter, however well received, is unlikely to create the common knowledge on which a new




Figure 2: Intent-to-treat effects on satisfaction with democracy and political hope. All measured on a 5-point scale.




Figure 3: Intent-to-treat effects on representational priorities. 5-point scale.

representational equilibrium depends (Chwe, 2001; Pinker, 2025). Moreover, evidence from African democracies suggests that citizens condition their political preferences on perceived state capacity, implying that even widespread deliberative engagement may fail to alter representational priorities if citizens doubt the state’s capacity for programmatic delivery (Gottlieb, 2024).

In the pre-analysis plan, I registered changes in representational priorities as directional hypotheses, so the null results fail to support them. At the same time, I described the underlying process in explicitly gradualist terms, anticipating that citizens “may become gradually less reliant on visible delivery” as alternative evidence about the representative’s type and their own standing accumulates. Changes in priorities were thus expected to flow through these upstream outcomes. The observed pattern is consistent with, but does not establish, that ordering: the upstream outcomes moved; the equilibrium-level expectation did not.

Participants themselves underscored the possible importance of repetition and diffusion, often conditioning their optimism on the practice continuing and spreading to other politicians. As one participant reflected:

“I am happy with it because I have not seen this type before, how I wish other politicians or other group will imitate this pattern I’m telling you the truth, things will be more better [sic].”⁴⁷

The null finding therefore leaves open an important question for future research: would repetition, publicity, and diffusion help transform changes in political relationships into changes in the representational demands that sustain the deliverist equilibrium? If so, how much diffusion would be necessary?

Constituency-Level Variation

Although the pooled estimates reveal a broadly consistent pattern, the magnitude and precision of the effects vary across constituencies. Bende, in particular, stands out most clearly on the main outcomes discussed in this paper. The estimated effect on representative type is close to zero, and the effects on both measures of political standing are positive but smaller and less precisely estimated than those observed elsewhere. Consistently, the ITT estimates for generalized trust follow the same pattern: positive change in Obio-Akpor and Ikole/Oye, not so in Bende.

Although the available evidence cannot fully explain this constituency-level pattern, I tentatively advance a consideration that may add context to this pattern and could itself be theoretically informative. In Bende, the MP joined the DTH 93 minutes later than scheduled; in both Obio-Akpor and Ikole/Oye, the MPs were perfectly on time. The delay is especially relevant to outcomes such as representative type and political standing. An encounter intended to communicate attentiveness, respect, and willingness to engage may be far less effective at conveying those qualities when constituents are required to wait for an extended period. Indeed, lateness may itself be interpreted as a signal of inaccessibility or insufficient regard, partially offsetting the positive signals conveyed once the discussion begins. The design cannot isolate the independent effect of the delay from other differences between constituencies, but participants in Bende themselves identified it as relevant:

“What I want to say is that, yes what he did is good [...] if he can do it regularly, it will make the people happy and [they will] know ‘yes we are accepted in this now.’ [...] but let it be started on time.”⁴⁸

The remark is revealing because its language closely echoes the concept of political standing, while its final qualification underscores that the meaning citizens draw from engagement depends not only on whether it occurs, but also on how it is enacted – including whether their time is treated with respect.

⁴⁷DTH participant. Umueze Hall, Amamba Uzuakoli, Bende.

⁴⁸DTH participant. Bende LGA Headquarters Hall, Bende.

Of course, this interpretation is necessarily speculative. The timing of the MP’s arrival was not experimentally assigned, and constituencies differ in other respects that could also account for variation in the estimates. Nor do the constituency-specific estimates establish statistically distinguishable treatment effects across sites without a direct test of those differences. Nevertheless, the pattern in Bende is consistent with the broader theoretical claim that the effects of deliberative encounters depend not only on the information exchanged but also on whether the encounter communicates recognition and respect. In this sense, the findings point to a potential limitation of canonical deliberative democratic theory, which has tended to privilege the rational and informational qualities of deliberation while paying less attention to its relational dimensions.

One might further worry that identity dynamics are driving the results. Constituents of the same ethnic group as the MP may have had that identity activated by the event. I tested for this by analyzing differences in treatment effects between constituents who were co-ethnic with the MP and those who were not. The MP in Obio-Akpor is from a very small ethnic group that was not among the eleven survey choices. So, to be maximally conservative, I coded all constituents who answered anything other than one of the eleven as co-ethnic. The data do not support the identity interpretation. Co-ethnics and non-co-ethnics had indistinguishable conditional treatment effects on all five outcomes.⁴⁹ The one marginally significant difference (Obio-Akpor) runs opposite to the identity account. These comparisons are underpowered, so I do not claim that they constitute evidence of precise equality. The pattern, however, is consistent with my theory’s relational claim. The recognition conveyed by the encounter is general and may actually be stronger for citizens who do not have *ex ante* identity reasons to trust the MP.

How Participants Understood the Encounter

Immediately after each event, we conducted focus groups at all fifteen venues to understand how participants experienced and interpreted the encounter. While their personal accounts cannot independently identify causal effects they can help triangulate the experimental findings and clarify the processes participants themselves believed were at work. Three recurring patterns closely track the mechanisms proposed by the theory: participants updated their beliefs about the representative’s character and effort, described a heightened sense of recognition and political standing, and distinguished material delivery from the evidence of care and commitment that delivery ordinarily provides. The discussion below illustrates these patterns with selected quotations.⁵⁰

Type updating. Participants’ accounts suggest that the encounter did more than provide new information about what their representative had done. It created an opportunity to observe him directly and, from that behavior, draw inferences about his effort and underlying commitment to the constituency. A clear expression of this updating came from a participant in Ikole/Oye: “before, I thought he was not working, but with what I saw today, I can see that he is really working.”⁵¹ These revised judgments extended beyond assessments of performance. Participants also treated the representative’s willingness to appear, listen, and respond as evidence of his motivations. As one participant explained, “For him to have the zeal to communicate with his people, I believe that he has his people at heart and their development.”⁵² A participant in Bende drew a similar connection, saying that the meeting had “built our confidence and our hope more and more” because it showed that the representative was “very, very serious with [the] people of Bende.”⁵³

The interactive and unusually spontaneous character of the encounter appears to have made these inferences especially credible. Participants were not simply hearing another prepared account of the

⁴⁹All models adjust for the baseline outcome and include constituency fixed effects. Standard errors are clustered at the EA level. See Appendix I.3 for the full results.

⁵⁰I classified the excerpts by their dominant theme; Appendix J reports the full classification and selection procedure.

⁵¹DTH participant, Believer Church of Christ Hall, Ikole/Oye.

⁵²DTH participant, Centre for Nuclear Energy Study Hall, Obio-Akpor.

⁵³DTH participant, Ndiaokpor Hall, Bende.

representative's record; they could watch him respond publicly to questions posed by constituents and judge how he handled demands he could not fully anticipate. One participant contrasted this experience with the representative's previous televised appearances: "Before now, I used to think that the questions he answered [on TV were] pre-arranged questions. But seeing him answer questions in real time gives me more confidence."⁵⁴ Because answering unrehearsed questions in public is difficult to stage, the representative's conduct during the encounter carried informational weight that a prepared speech or mediated appearance might not. The format therefore allowed participants to treat not only what the representative said, but how he behaved under conditions of limited control, as credible evidence of his effort, competence, and concern.

Standing. Participants also described the encounter in a language that closely aligns with the concept of political standing. For some, the representative's willingness to enter into direct exchange with ordinary constituents was itself an act of recognition and respect. "I think for him to be able to participate in this kind of thing, that means he values us," one participant explained; "He values the masses. He values our opinions."⁵⁵ Other participants situated their experience against a political status quo in which meaningful access is ordinarily reserved for people with status and connections. A participant in Ikole/Oye observed, "All the governments we've passed through [...] only children of big people get [the] opportunity to speak. But in a gathering like this, we're all here. We all have voting rights, we all deliberated on how the town can be better, and we saw you (the MP) reacting with us. You gave less privileged people a chance to open up."⁵⁶ Her account links recognition not simply to the opportunity to speak, but to an experience of political equality: people from different ethnic communities and social positions entered the same forum, exercised the same voice, and received a response from political authority.

The same idea appeared in a phrase that recurred throughout the discussions: being "carried along." In Nigerian English, the expression is ubiquitous in everyday political complaints and functions as an emic description of something close to political standing. To be carried along is not merely to be informed about what government is doing. It is to be treated as someone whose interests warrant consideration, whose voice counts, and who is entitled to make claims of those in power. As one participant put it, "I feel that, as a Rivers person, I am being carried along. So if these kinds of events happen, after now, [and] it still continues to happen, our voices can be heard, which is one of our constitutional rights. So we have rights to keep asking. If things are not being done, we still have the right to keep asking."⁵⁷ Recognition here did not produce passive gratitude. It strengthened the participant's sense of entitlement to continue questioning the representative and demanding an account of his performance.

This understanding was especially clear when the moderator asked what participants would say to friends or neighbors who doubted that the event mattered. One participant answered by identifying the opportunity to address the representative as a concrete political gain: "What have I achieved? What have I gained? I have gotten the opportunity to [voice] out my mind, my opinion to Mr. Chinda. I have had an opportunity to speak with him – even not physical or private – but I have every opportunity to speak with him, to iron out my views, to encourage him on what he has done and what to do more."⁵⁸ What mattered, on this account, was not privileged access or an individualized benefit, but being placed in a relationship in which an ordinary citizen could speak, be acknowledged, and offer both criticism and encouragement. The encounter thus altered how participants understood their position relative to the representative: not simply as distant, passive, and powerless recipients of government action, but as citizens with standing to be heard.

⁵⁴DTH participant, Nuton Sofia Hall, Obio-Akpor.

⁵⁵DTH participant, R and B Hall, Obio-Akpor.

⁵⁶DTH participant, Yeye Oba Hall, Ikole/Oye.

⁵⁷Rivers is the state in which Obio-Akpor constituency is located. DTH participant, Nuton Sofia Hall, Obio-Akpor.

⁵⁸DTH participant, R and B Hall, Obio-Akpor.

Hints of the Signal Substitution Logic. Although participation in a single DTH did not alter our measures of representational expectations, the focus-group transcripts reveal a plausible pathway through which repeated encounters might do so over time. Several participants spontaneously treated engagement, explanation, and visible effort as meaningful alternatives to receiving an immediate material benefit. A participant in Bende captured the distinction plainly: “Today, we are not carrying [a] bag of rice, and he didn’t give us money, but we are satisfied with what he told us.”⁵⁹ The absence of a transfer did not make the encounter empty. What the representative said, and his willingness to answer for his conduct, supplied a different kind of evidence about the relationship. A participant in Obio-Akpor developed this logic further: “[through this event] you are very close to them, not because you can give out money ... You create an avenue whereby you reach out to the people that you are leading ... He [hears] us. We talk to him. He knows what is happening in his area, not just giving out money [sic].”⁶⁰ On this account, closeness between citizens and their representative need not be produced through personal transfers. It can also be enacted through reciprocal communication: citizens make their circumstances known, the representative listens and responds, and both sides demonstrate their continued connection to one another. Engaging constituents with respect can send the comforting signal that “At least he has tried. We have seen efforts, so he’s still with us.”⁶¹ Effort and engagement thus became observable evidence that the representative had not withdrawn from the relationship, even when the encounter delivered no material reward.

Undoubtedly, participants continued to value projects for their practical benefits and to treat them as powerful evidence of a representative’s intentions. Indeed, the moral logic at the heart of deliverism remained clearly visible. As one Bende participant explained, “Our people say seeing is believing. At least in Ugwueke, we have a hall that he built for us to indicate that he loves us.”⁶² The hall mattered not only as useful infrastructure but as visible proof of love and commitment, the core moral logic of deliverism. The encounter did not displace delivery as either an objective or a signal. It demonstrated, perhaps more modestly, that delivery need not remain the only credible evidence available to citizens when judging whether a representative cares about them. Participants distinguished, without prompting, between delivery as evidence of care and delivery as an end – accepting substitutes for the former while retaining the latter. This helps reconcile the favorable changes in representative evaluations and political standing with the null estimates for representational priorities: citizens could revise what they believed about the representative without immediately revising what they ultimately wanted him to provide.

For such a substitution to reshape representational demands, however, participants suggested that engagement would need to recur. A single encounter could still be interpreted as a temporary performance, particularly as an election approached. One participant explained that town halls “will make the people you are representing know your heart,” but immediately conditioned that judgment on their continuation: “What he did now is what is expected of him all the time, so that people will not feel he is doing it to win the people’s hearts because of the upcoming elections ... They should try and do this time to time so people will not think that way.”⁶³ Repetition would make engagement harder to dismiss as electoral theater and allow evidence of responsiveness to accumulate across encounters. Citizens themselves therefore described standing as something closer to a stock than a momentary reaction: it is built through repeated recognition and cannot be durably transformed by a single event. While the qualitative evidence does not overturn the null causal estimates, it does locate them within a longer process that suggests repeated encounters could gradually broaden the evidentiary basis on which citizens judge representation.

⁵⁹DTH participant, Amafor Civic Hall, Bende.

⁶⁰DTH participant, Nirvana Hotel Hall, Obio-Akpor.

⁶¹DTH participant, Nuton Sofia Hall, Obio-Akpor.

⁶²DTH participant, Alayi Hall, Bende.

⁶³DTH participant, Bende LGA Headquarters Hall, Bende.

6 Vignette Survey Experiment

Separately from the DTH field experiment, I test the theory’s key mechanisms through a vignette experiment embedded in the endline survey. The experiment isolates the signals that the DTH field encounter necessarily bundles together. Respondents evaluated how a hypothetical Member of the House of Representatives from another constituency conducted himself during legislative deliberations over a national infrastructure bill addressing soil erosion, flooding, and related disasters.⁶⁴ All versions held constant the policy problem, institutional constraints, and the MP’s formal support for the bill while varying the additional information provided about his conduct.

Each respondent was randomly assigned to one of four conditions: a common low-information baseline or one of three theoretically motivated departures from it. The first described the MP undertaking substantively valuable legislative work, including engaging relevant committees, following up with legislative leaders, and pressing for progress on the bill, that remained largely outside public view and difficult for citizens to verify. The second described comparable legislative effort as publicly visible and independently verifiable through reporting and discussion within the constituency. Comparing these conditions tests whether information about non-deliverist effort improves evaluations of representative intent and commitment, and whether such effort becomes a more credible substitute for material delivery when it is publicly legible and externally verifiable. The third departure instead described repeated and respectful engagement with constituents, including listening to their concerns, discussing trade-offs, and explaining institutional constraints. It makes no mention of additional legislative ‘effort.’ This condition captures the citizen-facing mechanism through which recognition and respectful treatment communicate dignity, inclusion, and equal political standing; relational dimensions of representation that may shape citizens’ evaluations even in the absence of material delivery.

The design is intentionally asymmetric rather than fully factorial. Its purpose is not to estimate every possible interaction among representational attributes, but to examine how a small number of theoretically salient signals alter citizens’ judgments relative to a shared low-information baseline. This approach reflects the limited observability of representation in practice: citizens seldom encounter multiple clear and independently verifiable signals of legislative conduct simultaneously. Random assignment to the vignette conditions was conducted at the individual level and thus independently of the EA-level DTH assignment. This design permits causal identification of how specific cues about legislative effort, its observability, and recognition shape evaluations while the broader political scenario remains fixed.

After reading the vignette, respondents assessed whether the MP was “trying to help his constituents,” how much blame he deserved if the bill failed to pass, and their overall approval of him. For ease of interpretation, the blame measure below is reverse-coded so that higher values indicate less blame and, consequently, more favorable evaluations. Together, these outcomes test whether non-deliverist signals – particularly legislative effort and respectful engagement – can substitute for visible material delivery in shaping citizens’ evaluations of political representation, at least in a hypothetical setting.

6.1 Vignette Results

Figure 4 presents the estimated effects of the three informational conditions relative to the common low-information baseline. All three conditions produced significantly more favorable evaluations of the hypothetical MP across all three outcomes. Respondents who received information about the MP’s legislative effort expressed greater approval, were more likely to believe that he was trying to help his constituents, and assigned him less blame for the bill’s possible failure. Information about respectful engagement with constituents produced the same general pattern.

These results support a central premise of the theory: citizens’ evaluations of representation are not determined exclusively by the delivery of visible material benefits. Even though none of the

⁶⁴See Appendix G for the full vignette text.




Figure 4: Effects of vignette conditions on evaluations of the hypothetical MP. Estimates relative to low-information baseline and include constituency FEs. ‘Approval’ is on a 5 point scale. The others: 1-4.

experimental conditions described the MP successfully delivering a project or securing passage of the bill, information about his legislative effort and respectful engagement improved evaluations across every outcome. Citizens therefore appear willing to draw favorable inferences from non-material and non-deliverist forms of representation. Legislative work can signal that a representative is acting on constituents’ behalf, while respectful engagement can communicate recognition and political standing. Both types of signals reduce the tendency to interpret the absence of a successful policy outcome as evidence of indifference or insufficient effort.

The positive effect of the low-observability effort condition is particularly notable. Once respondents were informed that the MP was undertaking meaningful legislative work, they evaluated him more favorably even when that work remained difficult for citizens to verify independently. This indicates that information about otherwise hidden legislative action can alter judgments. My theory further predicts, however, that such effort should provide a more credible and effective signal when it is publicly observable. Public visibility allows citizens to corroborate the representative’s conduct, makes the signal less dependent on the representative’s own claims, and facilitates its transmission through the constituency. The high-observability condition should therefore generate more favorable evaluations than the low-observability condition, even when the underlying legislative effort is comparable.

Table 3 tests this implication directly. For each outcome, I use a Wald test of linear restrictions to assess whether the estimated effects of the high- and low-observability conditions differ statistically from one another. The table reports the effects of the low- and high-observability effort conditions relative to the baseline and the contrast between them.

The results are directionally consistent across the three outcomes. Relative to low-observability effort, high-observability effort increased approval by 0.103 points, increased the perception that the MP was trying to help by 0.128 points, and reduced blame by an additional 0.096 points. The high-observability coefficient is therefore larger in every comparison, as predicted. However, only the

Table 3: Comparison of high- and low-observability legislative-effort conditions.

Outcome	Low	High	Difference (H-L)	SE	p	N
Approval	0.632	0.735	0.103	0.077	0.1790	1938
Politician is trying	0.482	0.610	0.128	0.054	0.0187	1898
Reduced blame	0.471	0.567	0.096	0.081	0.2337	1374

The blame attribution outcome had a significantly higher number of "Don't Know" responses which explains the lower N.

difference for perceived effort is statistically distinguishable from zero at conventional levels ($p = 0.0187$). The differences for approval and reduced blame are positive but estimated too imprecisely to reject equality between the two effort conditions.

The evidence for the additional importance of observability is therefore suggestive rather than conclusive. Its consistent direction across all three outcomes aligns with my theoretical expectation that publicly legible legislative work provides a stronger signal of representative commitment. Still, the failure to detect statistically significant differences for two of the three outcomes cautions against claiming that observability uniformly amplifies the effect of legislative effort.⁶⁵ The stronger conclusion from the experiment is that information about non-deliverist effort and recognition substantially improves citizens' evaluations relative to a low-information baseline. The more qualified conclusion is that making legislative effort publicly verifiable may further strengthen this effect, especially by increasing citizens' confidence that the representative is genuinely trying to act on their behalf.

7 Discussion and Conclusion

Where legislative activity is difficult to observe, citizens may rely on tangible material delivery as evidence that a representative is at all working on their behalf. Visible material delivery, on this account, matters not only because citizens want and need the things being delivered. It also communicates effort, concern, and a willingness to recognize the constituency, and its people, as politically consequential – in one sense, the very core promise of democratic governance. This dual role can sustain a deliverist equilibrium in which citizens demand immediate and visible benefits, representatives organize their behavior around providing them, and less observable forms of representation receive comparatively little weight. This dynamic neither yields sufficient material delivery to reliably make constituents' lives better, nor encourages the kinds of forward-looking democratic cooperation that might contribute to development over the medium to long term. The result is a particularly discouraging form of democratic stasis: life remains difficult for most Nigerians, while their faith in democracy progressively weakens. For anyone concerned with democratic consolidation on the continent, this account is troubling.

This paper asks whether democratic innovations can partially loosen that connection by making representative effort and benevolence visible even when they do not immediately produce tangible rewards. The claim is not that engagement can replace material development, but that it can broaden the evidence citizens use to interpret political conduct and strengthen a representative relationship otherwise judged exclusively through constant, tangible, but inevitably small scale material delivery.

The results support the earlier stages of this argument, but not the completion of the proposed sequence. Assignment to a Deliberative Town Hall improved citizens' evaluations of their representa-

⁶⁵The vignette only imperfectly reproduces differences in observability as they arise outside this survey-experimental, hypothetical setting. In both effort conditions, respondents are explicitly informed that the representative undertook the same legislative work; the conditions differ only in whether ordinary citizens could independently verify it. This stipulation may attenuate the contrast, since in everyday politics unobservable effort is not merely difficult to verify but often uncertain to have occurred at all. The consistently positive, though unevenly estimated, effect of high-observability may therefore represent a conservative test of the importance of public verifiability.

tive, increased their perceived standing in relation to both the representative and government more generally, raised their level of trust towards politicians broadly, and strengthened satisfaction with democracy and political hope – all without delivering a project or private benefit. It did not, however, reduce the weight citizens placed on immediate constituency delivery or increase their relative support for programmatic, nationally oriented, or longer-term representation. This pattern of results identifies the main constraint with some precision. The deliverism trap’s informational and relational conditions (beliefs about type, standing, and trustworthiness) responded to a single, relatively inexpensive, non-material event. The equilibrium expectation did not change materially, and there is perhaps little reason to believe it should. Norms built on equilibria mostly change from repeated, public, coordinated changes in practice. The question becomes how much such engagement with what kind of backing would change them. That is more of a question about institutional design than political culture.

This null result admits of at least three distinct, though potentially complementary, interpretations. They operate, respectively, at the level of the political system, social equilibrium, and individual belief. First, a *credibility account* suggests that citizens continue to prioritize constituency delivery because thin programmatic commitments and weak state capacity make material benefits more reliable. On this view, even routinized engagement may not alter priorities absent credible signals of broader institutional change. Second, a *coordination account* would suggest that one event, no matter how well received, is not enough to escape from deliverism’s sub-optimal equilibrium. On this view, if citizens and MPs came to expect engagement to become widespread and routine, priorities might eventually shift. Third, an *additive account* suggests that process and recognition need not displace material priorities at all. Citizens may incorporate listening, giving account, respect, and legislative effort into their evaluations while continuing to value material delivery just as strongly. If so, institutionalizing such practices would broaden the criteria of good representation without necessarily reordering them. The denominator would grow but the material component of the numerator would stay the same.

My experimental estimates cannot adjudicate among these accounts, which likely operate simultaneously. My qualitative evidence, however, bears on two of them. Participants spontaneously conditioned their optimism on the practice recurring and spreading, as the coordination account predicts, and they distinguished delivery as evidence of care from delivery as an end, as the additive account implies. The credibility account rests instead on the structural features of the setting described in Section 3. The current study contains little evidence that bears on it. The distinction matters because the accounts’ institutional implications differ. Under credibility, change requires increased state capacity and programmatic offerings. Under coordination, repetition, spread, and publicity might suffice, identifying a potentially novel and inexpensive path to democratic gains. Under additivity, broadening the evidence basis for representation is itself an achievement, and the only question is how much impact it might have at scale. Future research could distinguish them by varying the recurrence and public visibility of engagement, comparing settings with different levels of programmatic credibility and state capacity, and tracing whether sustained engagement ultimately displaces delivery priorities or simply adds new dimensions to citizens’ expectations.

Perhaps, though, mere contact with the MP, rather than the content of the encounter, accounts for these effects. Two features of the evidence suggest that these effects reflect the content of the encounter rather than contact per se. The recognition-only vignette condition improved evaluations without conveying any information about legislative effort, and in Bende, where the encounter’s conduct undercut the respect signal it was designed to send, contact alone failed to produce the effects observed elsewhere. While observational, the null effects in Bende correspond with an encounter where the MP arrived 93 minutes late, and participants explicitly said it undercut the respect the event was designed to convey. That said, I cannot rule out unobserved constituency-level confounders.

One way to interpret the overall pattern of results is to distinguish more carefully between delivery *as evidence* – a signal of something – and delivery *as an objective*. The results suggest that non-material forms of engagement can partially substitute for delivery as evidence that a representative cares, is making an effort, and takes citizens seriously. They do not indicate that engagement

can substitute for the short-term, locally oriented material outcomes that citizens want their representatives to pursue. Citizens may become more willing to believe that an MP is hardworking or acting in good faith after hearing him explain his activities, while still believing that the ultimate test of representation is whether roads are repaired, employment is created, or assistance reaches the constituency, as regularly and visibly as possible. In a setting of considerable material need, this combination of judgments is not especially surprising.

Its political consequences could nevertheless be important. If citizens can recognize engagement as credible evidence of representative effort even in the absence of immediate delivery, relations between citizens and representatives need not depend entirely on the continuous provision of visible and targeted goods. This may create some space for representatives to retain citizens' confidence while undertaking legislative, programmatic, or longer-term work whose benefits are less immediate and more difficult to observe, if potentially greater.

For debates about the resilience of African democracies, this implication is both encouraging and sobering. It is encouraging because relatively modest practices of listening, explanation, and recognition may strengthen citizens' sense that they count within the political system, sustaining democratic satisfaction and hope even where governments cannot immediately satisfy extensive material demands. It is sobering because procedural inclusion cannot – and should not – indefinitely compensate for failures of material performance. Making democracy more resilient therefore requires neither engagement alone nor delivery alone, but institutions that combine regular and meaningful citizen-representative interaction with a greater capacity to produce the outcomes that citizens reasonably expect. Meaningful interaction may be especially valuable where the record of material performance is mixed and open to interpretation. Under such conditions, whether disappointing outcomes erode democratic support may depend partly on whether citizens attribute them to a hardworking and well-intentioned government constrained by limited capacity, or to one that possesses the capacity to act but is dishonest, selfish, or indifferent to their welfare. The contribution of practices such as Deliberative Town Halls may therefore be to broaden the foundations on which democratic legitimacy rests, making it less exclusively dependent on immediate material provision even if they cannot, and should not be expected to, eliminate that dependence altogether.

The generalized effects on standing, trust, and satisfaction with democracy further suggest a possible micro-foundation for democratic legitimacy. Citizens encounter democratic government partly through their interactions with particular representative agents of the system. Being heard by an MP may provide information about that individual, but it may also provide information about whether the political system is accessible and responsive more generally. The fact that the effects extended beyond MP-specific evaluations is therefore consistent with an account in which experiences of political recognition accumulate into broader beliefs about democracy. If that is true, then satisfaction with democracy may depend not only on institutional performance or policy outputs in the aggregate, but also on whether citizens experience political authority as something to which they can speak and from which they can reasonably anticipate a response.⁶⁶

These findings also contain a normative ambiguity. Greater trust, political hope, and satisfaction with democracy are not unconditionally desirable if they are produced by symbolic political performances unconnected to subsequent behavior. Deliberative encounters could generate warranted trust by making genuine effort visible and allowing citizens to scrutinize representatives' reasoning. But they could also be used to legitimate weak performance, particularly if representatives learn that the appearance of listening can substitute for more difficult forms of responsiveness (Ding, 2020). The argument of this paper is therefore not that DTHs should replace material development, legislative effectiveness, or electoral accountability. In contrast to more ambitious theories of democratic deliberation (Landemore, 2020, *inter alia*), the goal is a more modest one: not to displace electoral institutions of representation, but to augment them by improving the informational and relational conditions un-

⁶⁶Soss (1999) shows citizens translate their individual experiences with state bureaucracies into inferences about how government works more generally, which become the basis for broader political orientations. My conjecture loosely mirrors this logic.

der which they operate. Meaningful engagement may *complement* other forms of representation by giving citizens better evidence with which to interpret political conduct and by recognizing them as participants whose judgments count. Whether the trust generated by such encounters remains warranted will depend on what representatives do afterward. Such engagement makes a positive contribution insofar as it equips citizens with a broader set of criteria by which to judge – and ultimately reward or sanction – the quality of representation they receive.

There are several reasons for caution in interpreting my results. For example, the endline survey followed relatively soon after the intervention (5-10 days). The observed effects could represent an enduring revision of citizens’ beliefs, but they could also capture temporary enthusiasm produced by the event. The present results cannot confirm my theoretical prediction that political standing is a durable stock that can be built through repeated encounters or a more temporary judgment that fades once ordinary political relations resume. Future work along these lines should seek longer term answers to these questions of persistence. The argument, however, does not imply that a single DTH *should* durably transform citizens’ evaluations of government or their own political standing. Rather, the claim is that a practice of representation structured around more regular DTH-like engagements could reshape how representative democracy is experienced and understood. This experiment is therefore intended as a proof of concept about the potential role of deliberative innovations in young democracies, not as a claim to a one-shot solution.

A second concern is that respondents may have inferred the aims of the study and reported more favorable attitudes in response to perceived researcher expectations, rather than genuinely revising their attitudes and beliefs. I cannot rule out this possibility entirely. Yet, I would argue that the pattern of results is difficult to reconcile with a simple account of experimenter demand. Treatment effects were concentrated on evaluations of representative type, political standing, trust and democratic performance, while the battery measuring priorities for effective representation remained consistently unaffected. The effects were also weakest in Bende, where the MP’s lengthy delay plausibly undermined the signals of attention and respect that the encounter was intended to convey. If respondents were merely seeking to please the researchers, it is unclear why that tendency should have stopped at the priorities battery or varied with the quality of the encounter itself. One might worry, however, that the modest stipend participants received (approximately USD 5, to cover travel to their assigned venues) produced the favorable responses rather than the encounter itself. Three features of the design and results cut against this interpretation: (i) participants were explicitly told the payment came from the research budget, not from the MP; (ii) Bende participants received the same stipend and showed null effects on type and trust – if payment drove the results, it should have worked there too; and (iii) the vignette experiment, which involved no payment, reproduces the signal effects.⁶⁷

A third limitation concerns the supply side of institutionalization. The three participating MPs agreed to stand for unscripted public questioning. In that respect, they may constitute a selected sample of the representative type the encounter is designed to reveal. The institutionalization question is partly a question about supply-side uptake. Which representatives would adopt such a practice, and under what incentives, remains an open question.

Taken together, the results point to a narrower theory of political change than the strongest version of my argument might have anticipated. Alternative signals can change how citizens interpret representatives and how they understand their own position in relation to government. They can make less-visible, well-intentioned political action appear as credible effort rather than straightforward disregard, and they can produce greater trust and optimism about the responsiveness of democratic institutions. But changing the objectives citizens assign to representation may require those signals to be repeated, corroborated through subsequent behavior, and accompanied by material or institutional changes that make reliance on immediate delivery less necessary. Democratic innovation may therefore strengthen the foundations of warranted trust before it alters the demands citizens place

⁶⁷Though somewhat probative, Bende is not a true placebo test for the stipend. One can always argue that a positive effect of the payment was offset by the MP’s lateness, differences in baseline popularity, or other site characteristics. Because the delay was unanticipated rather than designed, however, the comparison still provides useful, if imperfect, triangulation.

upon representatives. Whether it can eventually do the latter remains a central question raised, but not yet answered, by these findings.

References

- Baldwin, K. (2013). Why Vote with the Chief? Political Connections and Public Goods Provision in Zambia. *American Journal of Political Science*, 57(4):794–809.
- Baldwin, K. (2025). Traditional Authorities Around the World: Toward a Better Understanding of Their De Facto Power, Downward Accountability, and State Recognition. *Annual Review of Political Science*, 28(Volume 28, 2025):95–114.
- Bratton, M., Mattes, R., Gyimah-Boadi, E., Lewis, P., and Asunka, J. (2024). Democracy at Risk - the people’s perspective.
- Bussell, J. (2019). Clients or Constituents? A Theory of Assistance in Patronage Democracies. In Bussell, J., editor, *Clients and Constituents: Political Responsiveness in Patronage Democracies*, page 0. Oxford University Press.
- Cain, B., Ferejohn, J., and Fiorina, M. (1987). *The Personal Vote: Constituency Service and Electoral Independence*. Harvard University Press.
- Castradori, M., Minozzi, W., and Neblo, M. A. (2024). Beyond the Clientelist Horizon.
- Chwe, M. S.-Y. (2001). *Rational Ritual: Culture, Coordination, and Common Knowledge*. Princeton University Press.
- Craig, S. C., Niemi, R. G., and Silver, G. E. (1990). Political efficacy and trust: A report on the NES pilot study items. *Political Behavior*, 12(3):289–314.
- Demarest, L. (2022). Elite clientelism in Nigeria: The role of parties in weakening legislator-voter ties. *Party Politics*, 28(5):939–953.
- Ding, I. (2020). Performative Governance. *World Politics*, 72(4):525–556.
- Dunning, T., Grossman, G., Humphreys, M., Hyde, S. D., McIntosh, C., Nellis, G., Adida, C. L., Arias, E., Bicalho, C., Boas, T. C., et al. (2019). Voter information campaigns and political accountability: Cumulative findings from a preregistered meta-analysis of coordinated trials. *Science advances*, 5(7):eaaw2612.
- Fiske, S. T. and Durante, F. (2014). Never trust a politician? Collective distrust, relational accountability, and voter response. In Prooijen, J.-W. v. and Lange, P. A. M. v., editors, *Power, Politics, and Paranoia: Why People are Suspicious of their Leaders*, pages 91–105. Cambridge University Press, Cambridge.
- Fujiwara, T. and Wantchekon, L. (2013). Can informed public deliberation overcome clientelism? Experimental evidence from Benin. *American Economic Journal: Applied Economics*, 5(4):241–255.
- Golden, M. A. and Nazrullaeva, E. (2023). The Puzzle of Clientelism. *Cambridge Core*.
- Gottlieb, J. (2024). How Economic Informality Constrains Demand for Programmatic Policy. *American Journal of Political Science*, 68(1):271–288. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajps.12746>.
- Grimmer, J., Messing, S., and Westwood, S. J. (2012). How Words and Money Cultivate a Personal Vote: The Effect of Legislator Credit Claiming on Constituent Credit Allocation. *American Political Science Review*, 106(4):703–719.
- Hardin, R. (2002). *Trust and trustworthiness*. Russell Sage Foundation.

- Hastie, R. and Park, B. (1986). The relationship between memory and judgment depends on whether the judgment task is memory-based or on-line. *Psychological Review*, 93(3):258–268.
- Hicken, A. and Nathan, N. L. (2020). Clientelism’s Red Herrings: Dead Ends and New Directions in the Study of Nonprogrammatic Politics. *Annual Review of Political Science*, 23(Volume 23, 2020):277–294.
- Keefer, P. and Vlaicu, R. (2008). Democracy, Credibility, and Clientelism. *The Journal of Law, Economics, and Organization*, 24(2):371–406.
- Kitschelt, H. and Wilkinson, S. I. (2007). *Patrons, clients and policies: Patterns of democratic accountability and political competition*. Cambridge University Press.
- Koter, D. (2013). King Makers: Local Leaders and Ethnic Politics in Africa. *World Politics*, 65(2):187–232.
- Kramon, E. (2018). *Money for votes: The causes and consequences of electoral clientelism in Africa*. Cambridge University Press.
- Kumar, T. (2022). Home Price Subsidies Increase Local-Level Political Participation in Urban India. *The Journal of Politics*, 84(2):831–845.
- Kumar, T. (2025). *Building Social Mobility: How Subsidized Homeownership Creates Wealth, Dignity, and Voice in India*. Cambridge Studies in Comparative Politics. Cambridge University Press, Cambridge.
- Landemore, H. (2020). *Open Democracy: Reinventing Popular Rule for the Twenty-First Century*. Princeton University Press.
- Laustsen, L. and Bor, A. (2017). The relative weight of character traits in political candidate evaluations: Warmth is more important than competence, leadership and integrity. *Electoral Studies*, 49:96–107.
- Lee, D. S. (2009). Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects. *The Review of Economic Studies*, 76(3):1071–1102.
- LeVan, A. C. (2019). *Contemporary Nigerian Politics: Competition in a Time of Transition and Terror*. Cambridge University Press, Cambridge.
- Lindberg, S. I. (2010). What accountability pressures do MPs in Africa face and how do they respond? Evidence from Ghana. *The Journal of Modern African Studies*, 48(1):117–142.
- Lodge, M., McGraw, K. M., and Stroh, P. (1989). An Impression-Driven Model of Candidate Evaluation. *American Political Science Review*, 83(2):399–419.
- Mayhew, D. R. (1974). *Congress: The Electoral Connection*. Yale University Press, second edition.
- Neblo, M. A., Esterling, K. M., and Lazer, D. M. (2018). *Politics with the people: Building a directly representative democracy*. Cambridge University Press.
- Ofosu, G. K. (2024). What do voters want from their legislators? evidence from ghana. *African Affairs*, 123(491):165–192.
- Opalo, K. O. (2024). The Truth About Africa’s Coups. *Journal of Democracy*, 35(2):93–107.
- Pinker, S. (2025). *When Everyone Knows That Everyone Knows...: Common Knowledge and the Mysteries of Money, Power, and Everyday Life*. Scribner, New York, NY.

- Pitkin, H. F. (1967). *The Concept of Representation*. University of California Press.
- Power, G. (2024). Finding Ways Out of the Gyratory Circus: The Human Side of Political Development. In Power, G., editor, *Inside the Political Mind: The Human Side of Politics and How It Shapes Development*, page 0. Oxford University Press.
- Schaffer, F. C. (2014). Ordinary Language Interviewing. In *Interpretation and Method*. Routledge, 2 edition. Num Pages: 11.
- Soss, J. (1999). Lessons of Welfare: Policy Design, Political Learning, and Political Action. *American Political Science Review*, 93(2):363–380.
- Stokes, S. C., Dunning, T., Nazareno, M., and Brusco, V. (2013). *Brokers, Voters, and Clientelism*. Cambridge Studies in Comparative Politics. Cambridge University Press, Cambridge.
- Wantchekon, L. and Guardado, J. (2024). Deliberation and ethical voting behavior: Evidence from a campaign experiment in benin. *American Political Science Review*, 118(3):1261–77.
- Warren, M. (2018). Trust and Democracy. In Uslander, E. M., editor, *The Oxford Handbook of Social and Political Trust*. Oxford University Press.

Appendix

A DTH Implementation

This section provides supplementary information about the implementation of the deliberative town halls.

A.1 Implementation Across Constituencies

Table 4: Basic details of the Deliberative Town Halls

	Bende	Obio-Akpor	Ikole/Oye
Date	June 20	July 11	July 18
Participants	205	235	238
Duration	64 minutes	62 minutes	76 minutes
MP delay from scheduled start	93 minutes	0 minutes	0 minutes
Attendance rate among invited	62.5%	61.7%	63.6%

A.2 Photographs from the DTHs



Figure 5: Example venue in Bende.



Figure 6: Example venue in Ikole/Oye.

B Missing EAs Details

Table 5: Enumeration Areas Not Covered During Data Collection

Constituency	EA Code	Treatment Assignment	Reason EA Was Not Covered
Ikole/Oye	229	Treatment	EA could not be accessed due to an ongoing conflict in the border area between Ekiti and Kwara States.
Ikole/Oye	179	Control	EA was identified as non-residential, with no eligible households available for interview.
Ikole/Oye	185	Treatment	EA was inaccessible due to an ongoing community dispute.
Ikole/Oye	163	Control	EA could not be accessed due to security concerns related to reported kidnapping and banditry in the area.
Ikole/Oye	176	Treatment	EA was identified as non-residential, with no eligible households available for interview.
Obio/Akpor	87	Treatment	EA is located within a restricted military zone and could not be accessed by the field team.
Obio/Akpor	108	Control	EA is located within a gated estate and the field team was not granted access to conduct interviews.
Bende	9	Treatment	EA was identified as non-residential, with no eligible households available for interview.
Bende	55	Treatment	EA could not be accessed due to safety concerns arising from an ongoing land dispute in the area.

C Survey Measures and Coding

This section reports the exact wording, response options, and coding of the principal outcome measures analyzed in the paper. Unless otherwise indicated, all measures were recoded such that higher values represent more favorable evaluations or greater support for the relevant construct.

C.1 Representative-Type Battery

Respondents were asked the following question:

Thinking about Hon. [NAME], in your opinion, how well does each of the following words or phrases describe him?

The order of items A–H was randomized across respondents.

Table 6: Representative-Type Battery

Item	Word or phrase	Very well	Quite well	Not very well	Not well at all	Do not know	Refused
A	Dishonest	1	2	3	4	888	999
B	Trustworthy	1	2	3	4	888	999
C	Selfish	1	2	3	4	888	999
D	Hardworking	1	2	3	4	888	999
E	Understands people like you	1	2	3	4	888	999
F	Inaccessible	1	2	3	4	888	999
G	Not qualified	1	2	3	4	888	999
H	Intelligent	1	2	3	4	888	999

Notes: “Do not know” and “Refused” were not read aloud by enumerators and are treated as missing in the main analysis, as pre-registered. For analysis, all items are coded so that higher values indicate a more favorable evaluation of the representative. Given the original response scale, the positively worded items—Trustworthy, Hardworking, Understands people like you, and Intelligent—are reverse-coded.

C.1.1 Construction of the Representative-Type Index

Responses outside the valid range of 1-4, including “Do not know” (888) and “Refused” (999), were treated as missing. Because the original response scale assigns lower values to stronger endorsement, the positively worded items—Trustworthy, Hardworking, Understands people like you, and Intelligent—were reverse-coded. After recoding, higher values on every item indicate a more favorable evaluation of the representative.

I constructed the index using principal component analysis (PCA). Following the pre-analysis plan, I first examined item distributions, pairwise correlations, PCA loadings, and coefficient theta to assess whether the battery was sufficiently dominated by a common first dimension to justify summarizing it as a single scale. The measurement model was estimated using an untreated measurement sample comprising all baseline responses and endline responses from control-group respondents. Item means, standardization parameters, and PCA weights were estimated exclusively from this sample, preventing treatment-induced responses from shaping the index. The items were centered and standardized to unit variance before estimating the PCA, and the first principal component was retained as the representative-type index. The resulting scoring rule was then applied to all respondents at both baseline and endline.

At each wave, an index score was constructed only for respondents with valid responses to more than half of the eight items, that is, at least five items. Respondents missing four or more items were assigned a missing index value. For eligible respondents missing one or more items, each missing response was replaced with the corresponding item mean estimated from the untreated measurement

sample before the PCA weights were applied. As a robustness check, I also constructed the index using complete cases only, excluding respondents missing any component item.

The sign of a principal component is arbitrary. I therefore oriented the first component so that the sum of its item loadings was positive. Given the prior coding of the items, higher index values consequently represent more favorable evaluations of the representative’s character, competence, and effort.

The same baseline scoring rule was used to construct the baseline and endline indices. Specifically, the baseline item means, item standard deviations, and first-component loadings were held fixed when scoring both waves. Missing endline responses were imputed using the corresponding baseline item means rather than means calculated from the endline data. Thus, any change in the index across waves reflects changes in respondents’ answers rather than changes in the construction or weighting of the index. Respondents who answered at least one item in a given wave received an index score; respondents missing the entire battery in that wave were assigned a missing index value.

Finally, the raw first-component scores from both waves were standardized using the mean and standard deviation of the baseline score distribution among respondents who answered at least one baseline item. The baseline index therefore has a mean of (approximately) zero and a standard deviation of (approximately) one. Endline scores remain expressed in baseline standard-deviation units, making the magnitude of treatment effects directly interpretable relative to the baseline distribution.

As a measure of internal consistency, I calculated coefficient theta from the baseline PCA. The resulting coefficient was 0.802, indicating a high degree of common variation among the eight items.

C.2 Political Standing vis-à-vis the MP

Respondents were asked the following question:

For each of the following statements, please state the extent to which you agree or disagree.

Enumerators were instructed to probe respondents for the strength of their agreement or disagreement. All items used the following response scale: Strongly disagree (1), Somewhat disagree (2), Neither agree nor disagree (3), Somewhat agree (4), and Strongly agree (5).

Table 7: Political Standing vis-à-vis the MP Battery

Item	Statement	Coding
A	Our representative takes the opinions of ordinary people seriously.	As stated
B	Ordinary people are ignored when our representative decides what is important.	Reversed
C	There is nothing ordinary people can do to stop our representative from doing what he wants.	Reversed
D	Regular people have ways of making their voices heard by our representative.	As stated
E	If people from my area went to our representative with a problem, he would know he has to take us seriously.	As stated
F	Our representative does not act like he depends on the support of ordinary people.	Reversed

Notes: Items B, C, and F were reverse-coded for the analysis. All items are therefore coded such that higher values indicate greater perceived political standing vis-à-vis the representative. All items ranged from 1-5.

C.3 Generalized Political Standing

Respondents were asked the following question:

For each of the following statements, please state the extent to which you agree or disagree.

Enumerators were instructed to probe respondents for the strength of their agreement or disagreement. All items used the following response scale: Strongly disagree (1), Somewhat disagree (2), Neither agree nor disagree (3), Somewhat agree (4), and Strongly agree (5).

Table 8: Generalized Political Standing Battery

Item	Statement	Coding
A	Nigerian officials care about what the masses think.	As stated
B	Our leaders treat ordinary people like beggars, who can be ignored.	Reversed
C	People like me do not have a real chance to speak to those in government.	Reversed
D	If the government ignores the masses, people like us can make them pay attention.	As stated
E	Government leaders know they will lose support if they are not responsive when the masses speak up.	As stated
F	Even when people like us try to speak to our leaders, it usually does not lead to real change.	Reversed

Notes: Items B, C, and F were reverse-coded for the analysis. All items are therefore coded such that higher values indicate greater perceived political standing vis-à-vis the government. All items ranged from 1-5.

C.4 Generalized-Trust Battery

Respondents were asked the following question:

To what extent do you agree with the following statements?

Table 9: Generalized-Trust Battery

Item	Statement	Strongly disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Strongly agree
A	I trust elected officials because they need to do what the people want to get re-elected.	1	2	3	4	5
B	Elected officials cannot be trusted because they do not fear the people.	1	2	3	4	5
C	I trust elected officials because they genuinely care about people like me.	1	2	3	4	5
D	Elected officials cannot be trusted because they always take care of themselves first.	1	2	3	4	5
E	I trust that elected officials know how to solve Nigeria’s problems.	1	2	3	4	5
F	I don’t trust elected officials because they have no idea how to fix what’s wrong with our politics.	1	2	3	4	5

Notes: Responses outside the valid range of 1–5 were treated as missing. For analysis, all items are coded so that higher values indicate greater generalized trust in elected officials. The negatively worded items—B, D, and F—are reverse-coded.

C.4.1 Construction of the Generalized-Trust Index

Responses outside the valid range of 1-5 were treated as missing. Because higher values on the original response scale indicate stronger agreement, the negatively worded items were reverse-coded. After recoding, higher values on every item indicate greater generalized trust in elected officials.

I constructed the index using principal component analysis (PCA). Following the pre-analysis plan, I first examined item distributions, pairwise correlations, PCA loadings, and coefficient theta to assess

whether the battery was sufficiently dominated by a common first dimension to justify summarizing it as a single scale. The measurement model was estimated using an untreated measurement sample comprising all baseline responses and endline responses from control-group respondents. Item means, standardization parameters, coefficient theta, and PCA weights were estimated exclusively from this sample, preventing treatment-induced responses from shaping the measurement model. The items were centered and standardized to unit variance before estimating the PCA, and the first principal component was retained as the generalized-trust index.

The sign of a principal component is arbitrary. I therefore oriented the first component so that the sum of its item loadings was positive. Given the prior coding of the items, higher index values consequently represent greater generalized trust in elected officials.

At each wave, an index score was constructed only for respondents with valid responses to more than half of the six items – that is, at least four items. Respondents missing three or more items were assigned a missing index value. For eligible respondents missing one or two items, each missing response was replaced with the corresponding item mean estimated from the untreated measurement sample. The items were then centered and standardized using the means and standard deviations estimated from that same sample before the first-component weights were applied. These parameters and weights were held fixed when scoring all respondents at both baseline and endline. Thus, differences between waves reflect changes in respondents’ answers rather than changes in the construction or weighting of the index. As a robustness check, I also constructed a complete-case version of the index that excluded respondents missing any component item.

Finally, the raw first-component scores from both waves were standardized using the mean and standard deviation of the baseline score distribution among respondents who satisfied the index’s missing-data requirement. The baseline index therefore has a mean of approximately zero and a standard deviation of approximately one. Endline scores remain expressed in baseline standard-deviation units, making treatment effects directly interpretable relative to the baseline distribution.

The coefficient theta estimated from the untreated measurement sample was 0.66, exceeding the preregistered threshold for treating the six items as indicators of a common underlying construct.

C.5 Democratic Evaluations

Table 10: Satisfaction with Democracy and Political Hope

Measure	Exact wording	Response scale
Satisfaction with democracy	All in all, how satisfied are you with how Nigerian democracy works today?	1–5
Hope in democracy over the next five years	In general, how hopeful are you about the next 5 years of Nigerian democracy?	1–5
Hope that government can solve problems	To what extent do you agree with this statement: If government officials work together with people, they are capable of improving Nigeria’s problems.	1–5

Representational Priorities Battery

Please, for each of the following pairs of statements, indicate which one you most agree with. Responses were recorded using the following scale:

Table 11: Response Scale for the Representational Priorities Battery

Code	Response category
1	Strongly agree with Statement 1
2	Agree with Statement 1
3	Neither agree nor disagree with either statement
4	Agree with Statement 2
5	Strongly agree with Statement 2

Item	Statement 1	Statement 2
a	It is more important for my representative in the National Assembly to fight to bring development projects and money to this constituency.	It is more important for my representative in the National Assembly to work together with other representatives to pass policies that benefit the whole country according to need.
b	My representative should prioritize helping individual constituents with personal problems when they need assistance. <i>Note to enumerator: Examples can include paying school fees for children of poor families or helping with hospital bills.</i>	My representative should focus on making laws to improve the constituency and Nigeria as a whole, even if that means this limits his or her attention to the individual needs and requests of constituents.
c	A good member of the National Assembly spends more time in the constituency working with local people to solve problems.	A good member of the National Assembly spends more time in the capital working with other members to solve problems.
d	I would prefer my representative to focus on delivering visible projects immediately, even if they are not very big.	I would prefer my representative to focus on long-term policies that will create bigger impacts later on, even if the effects are not visible immediately.
e	It is better for my representative to make sure people get what they need now, even if the same problems keep coming back.	It is better for my representative to deal with the causes of problems, even if people get less help in the short term.
f	Representatives in the National Assembly are elected to make important decisions on the people's behalf. They should listen to ordinary people but, in the end, they should be the ones making the important decisions.	Representatives in the National Assembly can take care of the day-to-day political decisions, but on important decisions, they should almost always simply follow what the public wants.

Representational Priorities Battery (continued)

Item	Statement 1	Statement 2
g	In a democracy, representatives should guide citizens as parents guide their children.	In a democracy, representatives should be guided by citizens as public servants accountable to the people.

D Treatment-Induced Selection into the Representative-Type Index

D.1 Measurement and Sample Construction

The representative-type battery asked respondents to evaluate their MP using eight characteristics: dishonest, trustworthy, selfish, understands people like them, hardworking, inaccessible, unqualified, and intelligent. Each item had a four-category substantive response scale ranging from 1 to 4. Items were oriented so that higher values consistently indicated a more favorable evaluation of the representative. In addition to the substantive response categories, respondents could answer “Don’t know enough to tell.” For the preregistered analysis of uncertainty, I coded this response as one and a valid substantive response as zero. Other invalid or missing responses were left missing rather than classified as “Don’t know enough to tell.”

For construction of the **Rep_Type** index, “Don’t know enough to tell” responses were treated as missing. Following the preregistered rule, an index score was calculated only when a respondent provided substantive answers to at least five of the eight items. Respondents answering fewer than five items received a missing index value at that wave. For eligible respondents with between one and three missing items, each missing response was replaced with the corresponding item mean estimated from the untreated measurement sample. The items were then centered and standardized using parameters estimated from that measurement sample, and the first principal component was used as the representative-type index. The same item means, standardization parameters, and PCA weights were used to score responses at baseline and endline. This ensures that any change in the index reflects changes in respondents’ answers rather than changes in the measurement rule across waves.

The primary treatment-effect specification uses a model in which the endline index is regressed on treatment assignment and the baseline value of the same index. A respondent must therefore have a valid index at both waves to enter this analysis. Table 13 describes the resulting sample. Of the 2,014 respondents in the panel, 1,308 had an available **Rep_Type** index at both baseline and endline. The remaining 706 respondents were excluded from the both-wave analysis. Most of these respondents (534 of the 706) were already below the five-item availability threshold at baseline. The other 172 had an available baseline index but fell below the threshold at endline.

Table 13: Construction of the Both-Wave Representative-Type Sample

Sample stage	Respondents
Full panel	2,014
Index available at both waves	1,308
Excluded from both-wave analysis	706
Already ineligible at baseline	534
Eligible at baseline but unavailable at endline	172

Notes: The representative-type index is available at a given wave when a respondent provides substantive answers to at least five of the eight component items. “Don’t know enough to tell” responses are treated as missing for purposes of index construction. Respondents must have an available index at both baseline and endline to enter the baseline-adjusted analysis.

The exclusions were unevenly distributed across the three constituencies: 55 respondents were excluded in Bende, 401 in Obio-Akpor, and 250 in Ikole/Oye. This distribution reflects substantial cross-constituency differences in respondents’ initial ability to evaluate their MP, with the uncertainty particularly pronounced in Obio-Akpor.

Note that the 534 respondents whose index was already unavailable at baseline could not enter the sample regardless of their endline responses. Because their exclusion is determined by a pre-treatment measurement, it does not itself constitute treatment-induced selection into the both-wave sample.

The direct post-treatment selection concern is instead confined to the 172 respondents who were eligible at baseline but whose endline answers determined whether they remained in the analysis. The subsequent analyses therefore examine whether assignment to a DTH affected endline index availability among the 1,480 respondents who crossed the availability threshold at baseline.

D.2 Treatment Effects on “Don’t Know Enough to Tell” Responses

Table 14: Treatment Effects on “Don’t Know Enough to Tell” Responses

Sample	Control mean (%)	Treatment effect (pp)	SE (pp)	95% CI (pp)	<i>p</i>	Respondents	Item responses
Pooled sample	19.7	-6.55	2.01	[-10.50, -2.60]	0.001	1,972	15,354
Bende	3.3	1.94	1.33	[-0.72, 4.59]	0.150	627	4,966
Obio-Akpor	44.5	-20.78	3.95	[-28.65, -12.91]	< 0.001	686	5,311
Ikole/Oye	11.3	-0.29	3.53	[-7.33, 6.75]	0.934	659	5,077

Notes: The outcome equals one when a respondent selected “Don’t know enough to tell” and zero when the respondent provided a substantive response from 1 to 4. Each row reports the treatment coefficient from a baseline-adjusted stacked linear probability model with item fixed effects. The pooled model additionally includes constituency fixed effects. Standard errors are clustered by enumeration area. The control mean is the unadjusted endline percentage in the corresponding estimation sample. Negative treatment effects indicate a reduction in insufficient-information responses.

D.3 Treatment Effects on Index Availability

Table 15: Treatment Effects on Endline Representative-Type Index Availability

Sample	Control availability (%)	Treatment availability (%)	Treatment effect (pp)	SE (pp)	95% CI (pp)	<i>p</i>	Respondents	EA clusters
Pooled	88.0	88.7	1.50	2.81	[-4.04, 7.04]	0.595	1,480	226
Bende	96.7	96.1	-0.61	1.91	[-4.42, 3.19]	0.749	604	78
Obio-Akpor	72.3	91.3	19.07	5.04	[9.02, 29.12]	< 0.001	351	74
Ikole/Oye	87.3	79.5	-7.82	6.63	[-21.04, 5.40]	0.242	525	74

Notes: The sample is restricted to respondents whose representative-type index was available at baseline. The outcome equals one when a respondent provided substantive answers to at least five of the eight representative-type items at endline and zero otherwise. Control and treatment availability rates are unadjusted percentages. Treatment effects are estimated using linear probability models. The pooled model includes constituency fixed effects; constituency-specific models include no additional fixed effects. Standard errors are clustered by enumeration area. Positive estimates indicate that assignment increased the probability that the endline index was available.

Distribution of Endline Representative-Type Responses

The dashed line marks the five-item threshold required to construct the index




Figure 7: Distribution of Endline Rep_Type Available Items

D.4 Selection into the Representative-Type Analysis Sample

Table 16: Composition of the Representative-Type Analysis Sample

Sample	Assignment	Baseline eligible	Retained	Excluded after baseline	Retained (%)	Baseline mean: retained	Baseline mean: excluded
Pooled sample	Control	683	601	82	88.0	-0.23	-0.45
Pooled sample	Treatment	797	707	90	88.7	0.11	-0.07
Bende	Control	300	290	10	96.7	-0.23	-0.08
Bende	Treatment	304	292	12	96.1	0.27	0.12
Obio-Akpor	Control	155	112	43	72.3	-0.44	-0.65
Obio-Akpor	Treatment	196	179	17	91.3	-0.12	-0.09
Ikole/Oye	Control	228	199	29	87.3	-0.11	-0.28
Ikole/Oye	Treatment	297	236	61	79.5	0.08	-0.11

Notes: The sample is restricted to respondents who provided substantive answers to at least five of the eight representative-type items at baseline. A respondent is retained when the index is also available at endline.

Respondents excluded after baseline provided fewer than five substantive answers at endline and therefore cannot enter the baseline-adjusted analysis. The final two columns report mean baseline representative-type index scores, in standard-deviation units, among retained and subsequently excluded respondents. Pooled counts repeat the constituency-specific observations and should not be added to the constituency rows.

Table 17: Complete-Case Treatment Effects on Representative Type

Sample	Control mean	Treatment effect	SE	95% CI	<i>p</i>	Respondents	EA clusters
Pooled sample	0.172	0.155	0.079	[-0.001, 0.312]	0.052	1,308	221
Bende	0.200	0.024	0.123	[-0.221, 0.270]	0.843	582	78
Obio-Akpor	-0.057	0.186	0.146	[-0.105, 0.478]	0.207	291	72
Ikole/Oye	0.261	0.277	0.124	[0.030, 0.525]	0.029	435	71

Notes: The dependent variable is the endline representative-type PCA index, standardized in the untreated measurement sample. The analysis is restricted to respondents whose index is available at both baseline and endline.

All models adjust for the baseline value of the index. The pooled model additionally includes constituency fixed effects. Standard errors are clustered by enumeration area. Treatment effects are reported in standard-deviation units.

These complete-case estimates condition on post-treatment index availability and are therefore the estimates examined in the subsequent Lee-bounds analysis.

D.5 Lee Bounds for Obio-Akpor

Table 18: Lee Bounds for the Representative-Type Effect in Obio-Akpor

Quantity	Value	95% bootstrap interval
<i>Panel A: Selection and trimming diagnostics</i>		
Baseline-eligible control respondents	155	—
Baseline-eligible treatment respondents	196	—
Observed control outcomes at endline	112 (72.3%)	—
Observed treatment outcomes at endline	179 (91.3%)	—
Treated outcome distribution trimmed	20.9%	—
<i>Panel B: Baseline-adjusted estimates</i>		
Baseline-adjusted complete-case difference	0.157	[-0.121, 0.410]
Baseline-adjusted Lee lower bound	-0.102	[-0.432, 0.192]
Baseline-adjusted Lee upper bound	0.503	[0.204, 0.753]
Lee-identified set	[-0.102, 0.503]	[-0.432, 0.753]

Notes: The sample contains Obio-Akpor respondents whose representative-type index was available at baseline. The bounds impose positive monotonicity: assignment to a deliberative town hall may make the endline index available but may not make an otherwise available index unavailable. Because endline availability was higher under treatment, the treated outcome distribution is trimmed until its effective availability rate equals the control rate. For the lower bound, the highest treated outcomes are removed; for the upper bound, the lowest treated outcomes are removed. Baseline-adjusted estimates residualize the endline index using the relationship between baseline and endline outcomes estimated among observed control respondents. Estimates are reported in standard-deviation units. Bootstrap intervals are percentile intervals from 2,000 valid EA-cluster bootstrap replications out of 2,000, with EAs resampled separately within treatment arms. Replications in which the treatment availability rate fell below the control rate did not satisfy the imposed positive-monotonicity direction and were omitted.

D.6 Sensitivity to Index-Construction Decisions

Table 19: Sensitivity to Representative-Type Index Construction

Sample	Index specification	Availability effect (pp) [95% CI]	Index effect (SD) [95% CI]	Baseline eligible	Index sample	EA clusters
Pooled sample	PCA, at least 4 items	4.33 [1.16, 7.50]	0.152 [0.004, 0.299]	1,531	1,440	225
Pooled sample	PCA, at least 5 items (preregistered)	1.50 [-4.04, 7.04]	0.155 [-0.001, 0.312]	1,480	1,308	221
Pooled sample	PCA, at least 6 items	2.84 [-3.32, 9.00]	0.167 [-0.000, 0.335]	1,375	1,172	214
Pooled sample	PCA, all 8 items	1.66 [-6.07, 9.39]	0.185 [0.020, 0.351]	1,060	826	196
Pooled sample	Equal-weighted, mean imputation, at least 5 items	1.50 [-4.04, 7.04]	0.165 [0.003, 0.326]	1,480	1,308	221
Pooled sample	Equal-weighted, available-item mean, at least 5	1.50 [-4.04, 7.04]	0.164 [0.007, 0.322]	1,480	1,308	221
Bende	PCA, at least 4 items	-0.16 [-4.21, 3.90]	0.009 [-0.228, 0.247]	619	598	78
Bende	PCA, at least 5 items (preregistered)	-0.61 [-4.42, 3.19]	0.024 [-0.221, 0.270]	604	582	78
Bende	PCA, at least 6 items	-2.15 [-6.81, 2.52]	0.025 [-0.227, 0.277]	575	543	78
Bende	PCA, all 8 items	-3.07 [-12.81, 6.68]	0.106 [-0.123, 0.335]	478	406	76
Bende	Equal-weighted, mean imputation, at least 5 items	-0.61 [-4.42, 3.19]	0.037 [-0.220, 0.293]	604	582	78
Bende	Equal-weighted, available-item mean, at least 5	-0.61 [-4.42, 3.19]	0.043 [-0.204, 0.291]	604	582	78
Obio-Akpor	PCA, at least 4 items	15.73 [6.21, 25.25]	0.210 [-0.061, 0.480]	374	320	73
Obio-Akpor	PCA, at least 5 items (preregistered)	19.07 [9.02, 29.12]	0.186 [-0.105, 0.478]	351	291	72
Obio-Akpor	PCA, at least 6 items	24.99 [14.80, 35.18]	0.182 [-0.148, 0.512]	307	242	69
Obio-Akpor	PCA, all 8 items	22.46 [5.35, 39.56]	0.164 [-0.304, 0.632]	200	127	57
Obio-Akpor	Equal-weighted, mean imputation, at least 5 items	19.07 [9.02, 29.12]	0.183 [-0.132, 0.498]	351	291	72
Obio-Akpor	Equal-weighted, available-item mean, at least 5	19.07 [9.02, 29.12]	0.195 [-0.123, 0.512]	351	291	72
Ikole/Oye	PCA, at least 4 items	1.64 [-1.96, 5.24]	0.249 [0.023, 0.475]	538	522	74
Ikole/Oye	PCA, at least 5 items (preregistered)	-7.82 [-21.04, 5.40]	0.277 [0.030, 0.525]	525	435	71
Ikole/Oye	PCA, at least 6 items	-5.00 [-19.84, 9.84]	0.320 [0.053, 0.587]	493	387	67
Ikole/Oye	PCA, all 8 items	-3.19 [-18.24, 11.87]	0.249 [0.008, 0.491]	382	293	63
Ikole/Oye	Equal-weighted, mean imputation, at least 5 items	-7.82 [-21.04, 5.40]	0.293 [0.046, 0.541]	525	435	71
Ikole/Oye	Equal-weighted, available-item mean, at least 5	-7.82 [-21.04, 5.40]	0.278 [0.038, 0.517]	525	435	71

Notes: The availability outcome equals one when the number of substantive responses meets the threshold specified in the second column. Availability models are estimated among respondents meeting the same threshold at baseline. Index models are baseline-adjusted models restricted to respondents meeting the relevant threshold at both waves. The preregistered specification uses the first principal component, requires at least five substantive answers, and replaces remaining missing items with untreated-measurement-sample item means. The available-item specification averages only the standardized items answered by each respondent. All index scores are standardized in the untreated measurement sample. Pooled models include constituency fixed effects. Standard errors are clustered by enumeration area.

E Regression Table for Coefficient Plots in Main Body

Table 20: Intent-to-Treat Effects on Evaluations of Representative Type

	Representative-Type PCA Index			
	Pooled	Bende	Obio-Akpor	Ikole/Oye
	(1)	(2)	(3)	(4)
Assigned to Treatment	0.157** (0.072)	-0.008 (0.114)	0.259** (0.119)	0.248** (0.114)
Baseline index adjustment	Yes	Yes	Yes	Yes
Constituency fixed effects	Yes			
Observations	1,548	623	383	542
EA clusters	227	78	75	74
Adjusted R ²	0.282	0.364	0.214	0.195

Standard errors clustered at the EA level.

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 21: Intent-to-Treat Effects on Representational Priorities

Sample Model:	Representational Priorities			
	Pooled	Bende	Obio-Akpor	Ikole/Oye
	(1)	(2)	(3)	(4)
Assigned to Treatment	-0.062 (0.075)	-0.084 (0.111)	-0.086 (0.113)	-0.008 (0.158)
Baseline adjustment	Yes	Yes	Yes	Yes
<i>Fixed-effects</i>				
Item	Yes	Yes	Yes	Yes
Constituency	Yes			
<i>Fit statistics</i>				
Observations	9,894	3,140	3,454	3,300
EA clusters	231	78	78	75
Adjusted R ²	0.110	0.126	0.066	0.081

Standard errors clustered at the EA level.

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 22: Intent-to-Treat Effects on Satisfaction with Democracy and Hope

Satisfaction with democracy and hope				
Sample Model:	Pooled (1)	Bende (2)	Obio-Akpor (3)	Ikole/Oye (4)
<i>Panel A: Satisfaction with democracy</i>				
Assigned to Treatment	0.302*** (0.106)	0.245 (0.163)	0.139 (0.169)	0.522** (0.214)
Baseline adjustment	Yes	Yes	Yes	Yes
<i>Fixed-effects</i>				
Constituency	Yes			
<i>Fit statistics</i>				
Observations	1,957	622	687	648
EA clusters	231	78	78	75
Adjusted R ²	0.135	0.123	0.077	0.082
<i>Panel B: Hope in democracy (next 5 years)</i>				
Assigned to Treatment	0.359*** (0.095)	0.266 (0.187)	0.238* (0.142)	0.546*** (0.156)
Baseline adjustment	Yes	Yes	Yes	Yes
<i>Fixed-effects</i>				
Constituency	Yes			
<i>Fit statistics</i>				
Observations	1,949	620	684	645
EA clusters	231	78	78	75
Adjusted R ²	0.223	0.169	0.127	0.108
<i>Panel C: Hope government can solve problems</i>				
Assigned to Treatment	0.257*** (0.083)	0.189 (0.141)	0.071 (0.105)	0.507*** (0.171)
Baseline adjustment	Yes	Yes	Yes	Yes
<i>Fixed-effects</i>				
Constituency	Yes			
<i>Fit statistics</i>				
Observations	1,963	622	687	654
EA clusters	231	78	78	75
Adjusted R ²	0.138	0.178	0.112	0.134

Standard errors clustered at the EA level.

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

F Item-Level Field-Experiment Results

The figures in this section report item-level intent-to-treat estimates. All models adjust for the corresponding baseline response and include constituency fixed effects. Standard errors are clustered at the enumeration-area level. Items are coded so that positive estimates indicate movement in the theoretically favorable direction.

F.1 Representative Type




Figure 8: Item-Level Effects on Evaluations of Representative Type

F.2 Political Standing vis-à-vis the MP




Figure 9: Item-Level Effects on Political Standing vis-à-vis the MP

F.3 Generalized Political Standing




Figure 10: Item-Level Effects on Generalized Political Standing

F.4 Representational Priorities




Figure 11: Item-Level Effects on Representational Priorities

G Vignette Experiment

This section provides the complete text of the vignette experiment, including the common introductory scenario, the four experimental conditions, and the outcome measures. Endline respondents were randomly assigned to one of the four conditions with equal probability.

G.1 Common Introductory Text

All respondents first received the following information:

I now want to tell you about a Member of the House of Representatives representing a constituency in another part of Nigeria. For privacy, we will call him Honourable X.

As you know, in many parts of Nigeria, communities face serious problems from soil erosion, flooding, and other climate-related disasters. Addressing these problems often requires large-scale infrastructure such as drainage systems, erosion control, and repairs to roads and bridges. Local governments alone usually cannot cover the costs, so funding decisions are often made at the national level.

Recently, the National Assembly has been considering a bill to approve significant funding for a national program to address soil erosion, flooding, and related disasters. The bill would direct resources to affected areas across the country, but funds are limited, and not every community can receive projects immediately.

Respondents were then randomly assigned to receive one of the following four versions of the vignette.

G.2 Experimental Conditions

G.2.1 Condition 1: Baseline (Low Information)

Honourable X did not participate in the debate and negotiations but voted in favour of the bill.

G.2.2 Condition 2: Legislative Effort with Low Observability

Honourable X voted in favour of the bill in the National Assembly. In addition, on his social media channels, he has repeatedly stated that he was very active in pushing the bill forward inside the legislature. For example, he posted updates saying that he raised the issue in internal committee discussions and pressed legislative leaders to treat the bill as a priority.

G.2.3 Condition 3: Legislative Effort with High Observability

Honourable X voted in favour of the bill in the National Assembly and was very active in advancing it within the legislature. A viral video showed him sharply confronting the Minister of Works to secure a fair distribution of the funds across Nigeria. These actions by Honourable X were widely discussed in independent media and are well known among his constituents.

G.2.4 Condition 4: Recognition and Respectful Engagement

Honourable X voted in favour of the bill in the National Assembly. In addition, he repeatedly met with ordinary citizens—not just wealthy ones—listened to their concerns and suggestions, and answered questions about flooding, erosion, and infrastructure needs. He treated all participants respectfully and made sure that everyone was included.

G.3 Vignette Outcome Measures

After receiving the assigned vignette, respondents were told:

Based on this information, please answer the following questions about Honourable X.

G.3.1 Approval

If you lived in the constituency represented by Honourable X, how much would you approve or disapprove of the way Honourable X is representing his constituency?

Table 23: Response Options for Vignette Approval

Response option	Code
Strongly disapprove	1
Somewhat disapprove	2
Neither approve nor disapprove	3
Somewhat approve	4
Strongly approve	5
Do not know/Refused	999

Notes: Responses coded 999 were treated as missing.

G.3.2 Perceived Effort

To what extent do you think Honourable X is genuinely trying to do what is best for people in his constituency?

Table 24: Response Options for Perceived Effort

Response option	Code
Not trying at all	1
Not trying much	2
Trying somewhat	3
Trying a lot	4
Do not know/Refused	999

Notes: Responses coded 999 were treated as missing.

G.3.3 Blame for Legislative Failure

Now imagine that the bill did not pass in the National Assembly, so funds were not made available in different parts of Nigeria to address problems like erosion and flooding. In this scenario, how much would you blame Honourable X, given what you have heard?

Table 25: Response Options for Blame

Response option	Code
No blame	1
A little blame	2
Some blame	3
A lot of blame	4
Do not know/Refused	999

Notes: Responses coded 999 were treated as missing. In the analysis, the blame outcome was reverse-coded so that higher values indicate less blame and therefore a more favorable evaluation of the hypothetical representative.

H Pre-Analysis Plan Crosswalk

Table 26: Crosswalk Between Paper and Pre-Analysis Plan Hypothesis Labels

Paper Label	PAP Label	Construct
H1	H2	Representative Type
H2	H3	Political standing vis-à-vis the MP
H3	H4	Generalized political standing
H4	S4	Generalized trust in elected officials
H5	S2	Satisfaction with democracy and political hope
H6	H1	Representational priorities
—	H2a	Warmth sub-index (Appendix I.1)
—	H2b	Competence sub-index (Appendix I.1)
—	H2c	“Do not know” responses to the type battery (Appendix D)

I Additional Pre-Registered Analyses

I.1 Warmth and Competence Sub-Indices (PAP H2a, H2b)

Sample Model:	Representative Warmth PCA Index			
	Pooled (1)	Bende (2)	Obio-Akpor (3)	Ikole/Oye (4)
Assigned to DTH	0.153** (0.068)	0.000 (0.105)	0.213 (0.137)	0.267** (0.104)
Baseline adjustment	Yes	Yes	Yes	Yes
<i>Fixed-effects</i>				
Constituency	Yes			
<i>Fit statistics</i>				
Observations	1,502	617	352	533
EA clusters	226	78	74	74

Standard errors clustered at the EA level.
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Sample Model:	Representative Competence PCA Index			
	Pooled (1)	Bende (2)	Obio-Akpor (3)	Ikole/Oye (4)
Assigned to DTH	0.178** (0.075)	0.082 (0.130)	0.285** (0.119)	0.197* (0.116)
Baseline adjustment	Yes	Yes	Yes	Yes
<i>Fixed-effects</i>				
Constituency	Yes			
<i>Fit statistics</i>				
Observations	1,533	617	377	539
EA clusters	227	78	75	74
Adjusted R ²	0.235	0.295	0.195	0.159

Standard errors clustered at the EA level.
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

I.2 Multiple-Testing Adjustment Across the Registered Primary Tests

Table 27: Treatment Effects with Benjamini–Hochberg False Discovery Rate Correction

Outcome	Estimate	Std. Error	Raw p -value	BH-adjusted p -value	FDR < .05	FDR < .10	Observations
Representational priorities	-0.062	0.075	0.412	0.494	No	No	9,894
Overall representative type	0.155	0.079	0.052	0.078	No	Yes	1,308
Standing vis-a-vis the MP	0.246	0.078	0.002	0.005	Yes	Yes	11,454
Generalized political standing	0.281	0.073	0.000	0.001	Yes	Yes	11,647

The table reports intention-to-treat estimates with EA-clustered standard errors.

I.3 Registered Moderators: Gender and Coethnicity

Table 28: Treatment Effects by Respondent Gender

Outcome	Treatment effect Men	Treatment effect Women	Female-male difference	SE	95% CI	p-value	N
Overall representative type	0.156	0.156	0.001	0.072	[-0.141, 0.142]	0.994	1,308
Standing vis-a-vis the MP	0.217	0.277	0.060	0.070	[-0.079, 0.199]	0.397	11,454
Generalized political standing	0.238	0.327	0.089	0.064	[-0.037, 0.215]	0.166	11,647
Representational priorities (A-E)	-0.027	-0.096	-0.070	0.082	[-0.231, 0.091]	0.394	9,894
Generalized trust in elected officials	0.238	0.372	0.134	0.088	[-0.040, 0.307]	0.131	1,948

Notes: Gender is coded as an indicator equal to one for women and zero for men. The female-male difference is the coefficient on the interaction between treatment assignment and the female indicator. The treatment-effect columns report the corresponding conditional treatment effects. All models adjust for the baseline outcome and include constituency fixed effects. Models for standing and representational priorities stack the component items and additionally include item fixed effects. Representational priorities include the preregistered items A-E. Standard errors are clustered at the EA level. PCA-index estimates are expressed in untreated-measurement-sample standard-deviation units; stacked estimates are expressed in the original response scale. For stacked models, N counts respondent-item observations. Significance codes for the female-male difference are *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

Table 29: Treatment Effects by Co-Ethnicity with the Representative

Outcome	Treatment effect Non-coethnic	Treatment effect Coethnic	Coethnic-non-coethnic difference	SE	95% CI	p-value	N
Overall representative type	0.079	0.150	0.071	0.228	[-0.378, 0.521]	0.754	1,017
Standing vis-a-vis the MP	0.758	0.171	-0.586	0.402	[-1.382, 0.209]	0.147	7,507
Generalized political standing	0.611	0.232	-0.379	0.360	[-1.091, 0.333]	0.294	7,556
Representational priorities (A-E)	-0.142	-0.045	0.097	0.252	[-0.401, 0.596]	0.700	6,440
Generalized trust in elected officials	1.192	0.263	-0.929	0.571	[-2.057, 0.199]	0.106	1,263

Notes: The sample is restricted to Bende and Ikole/Oye. Respondents are classified as coethnic with their representative when Ethnicity-pre equals 'Igbo' in Bende or 'Yoruba' in Ikole/Oye. Respondents with other nonmissing ethnicity codes are classified as non-coethnic. The coethnic-non-coethnic difference is the coefficient on the interaction between treatment assignment and co-ethnicity. The first two numeric columns report the corresponding conditional treatment effects. All models adjust for the baseline outcome and include constituency fixed effects. Models for standing and representational priorities stack the component items and additionally include item fixed effects. Representational priorities include preregistered items A-E. Standard errors are clustered at the EA level. PCA-index estimates are expressed in untreated-measurement-sample standard-deviation units; stacked estimates are expressed in the original response scale. For stacked models, N counts respondent-item observations. Significance codes for the interaction are *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

Table 30: Treatment Effects by Co-Ethnicity with the Representative

Outcome	Treatment effect Non-coethnic	Treatment effect Coethnic	Coethnic–non-coethnic difference	SE	95% CI	p-value	N
Overall representative type	0.132	0.161	0.028	0.147	[−0.261, 0.318]	0.847	1,308
Standing vis-a-vis the MP	0.419	0.171	−0.248*	0.132	[−0.509, 0.012]	0.062	11,454
Generalized political standing	0.387	0.232	−0.155	0.120	[−0.391, 0.081]	0.196	11,647
Representational priorities	−0.087	−0.049	0.038	0.127	[−0.212, 0.288]	0.764	9,894
Generalized trust in elected officials	0.388	0.265	−0.124	0.165	[−0.448, 0.201]	0.453	1,948

Notes: The analysis includes All three constituencies.

Respondents are classified as coethnic with their representative when Ethnicity-pre equals Igbo in Bende, or Yoruba in Iko/Oye. The MP in Obio-Akpor comes from an ethnic group so small as to not be included in the more than 10 options on the Afrobarometer ethnicity list for Nigeria. For this reason, we do not have exact co-ethnicity coding for this constituency. This table shows results coding all those who selected "Other" for ethnicity as co-ethnics of the MP. Note this is a conservative estimate of the difference between co-ethnics and non-. This analysis is necessarily imperfect and should be interpreted carefully. Respondents with other nonmissing ethnicity codes are classified as non-coethnic. The coethnic–non-coethnic difference is the coefficient on the interaction between treatment assignment and co-ethnicity. The first two numeric columns report the corresponding conditional treatment effects. All models adjust for the baseline outcome and include constituency fixed effects. Models for standing and representational priorities stack the component items and additionally include item fixed effects. Representational priorities include preregistered items A–E. Standard errors are clustered at the EA level. PCA-index estimates are expressed in untreated-measurement-sample standard-deviation units; stacked estimates are expressed in the original response scale. For stacked models, N counts respondent-item observations. Significance codes for the interaction are *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

I.4 Complier Average Causal Effects

Table 31: Complier Average Causal Effects of Town-Hall Attendance

Outcome	CACE	SE	<i>p</i> -value	First-stage effect	<i>N</i>
Overall representative type	0.236*	0.120	0.052	0.659***	1,308
Standing vis-a-vis the MP	0.389***	0.124	0.002	0.634***	11,454
Generalized political standing	0.443***	0.116	< 0.001	0.634***	11,647
Representational priorities	-0.097	0.119	0.412	0.632***	9,894
Generalized trust in elected officials	0.480***	0.158	0.003	0.631***	1,948

Notes: The table reports two-stage least-squares estimates of the effect of town-hall attendance among compliers. Attendance is the endogenous treatment and randomized treatment assignment is the instrument. The first-stage column reports the estimated effect of treatment assignment on attendance. All models adjust for the baseline outcome and include constituency fixed effects. Models for political standing and representational priorities stack the component items and additionally include item fixed effects. Representational priorities include preregistered items A–E. Standard errors are clustered at the EA level. PCA-index estimates are expressed in untreated-measurement-sample standard-deviation units; stacked estimates are expressed in the original response scale. For stacked models, *N* counts respondent-item observations. First-stage estimates may differ slightly because each model uses its outcome-specific estimation sample. Significance codes: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

J Focus Group Quotation Classification

The qualitative material comprises 143 excerpts selected in a preliminary reading from discussions conducted immediately after the DTHs in 15 venues across Bende, Obio-Akpor, and Ikole/Oye.⁶⁸ Each separately recorded excerpt was treated as the unit of analysis. Information identifying the Hall and constituency was retained for every excerpt. The classification scheme was developed iteratively and was theory-informed. Recurring ideas in the excerpts were consolidated into ten broad themes, several of which correspond to distinctions developed in the paper, including perceived representative character, political standing, warranted trust, evaluation of effort independently of outcomes, and the relationship between engagement and material delivery. When an excerpt addressed multiple ideas, it was assigned according to its dominant evaluative or substantive emphasis. The categories are therefore mutually exclusive for organizational purposes, but not substantively exclusive: an excerpt classified under one theme may also contain language relevant to another.

Table 32: Classification Scheme for Post-DTH Focus-Group Excerpts

Theme	Definition	Excerpts
Direct access, listening, and two-way dialogue	Opportunities to see the representative directly, speak freely, ask questions, and be heard at the grass-roots.	22
Care, empathy, humility, and genuine intent	Inferences about the representative’s character and motives, including concern, humility, seriousness, responsibility, and willingness to listen.	22
Transparency, accountability, and information	Engagement as a means of explaining actions, reporting back, obtaining feedback, making performance visible, and enabling accountability.	11
Trust, confidence, and revised evaluations	Expressions of trust or confidence and statements that the encounter changed, clarified, or strengthened evaluations of the representative.	15
Recognition, voice, and democratic inclusion	Feeling valued, respected, represented, or “carried along,” including the inclusion of ordinary, rural, nonpartisan, or less privileged citizens.	13
Hope, satisfaction, and optimism	Satisfaction with the encounter, renewed hope, optimism about future performance, and faith in political development.	13
Effort, legislative work, and evaluation beyond outcomes	Recognition of effort, advocacy, lawmaking, and representative work even when a desired project or policy outcome was not secured.	13
Material delivery and distributive priorities	Material projects and transfers as evidence of performance, alongside reflections on allocation, priority-setting, and engagement without money or gifts.	8
Critiques of ordinary political practice	Accounts of exclusion, disrespect, unresponsiveness, election-only contact, broken promises, elite favoritism, and democratic failure.	11
Calls to repeat, expand, and institutionalize engagement	Requests for town halls to occur regularly, reach additional communities, and be adopted by other representatives rather than appearing only near elections.	15
Total		143

⁶⁸I received the transcripts a few days ago, so this analysis has been very preliminary thus far.