Does Local Democracy Serve the Poor? Identifying the Distributive Preferences of Village Politicians in India*

Mark Schneider†

mark.allan.schneider@gmail.com

Center for the Advanced Study of India

University of Pennsylvania

Neelanjan Sircar‡

neelanjan.sircar@ashoka.edu.in

Ashoka University

*Data for this article will be made available on the corresponding author's website upon publication. Code is provided in the appendix.

†Non-Resident Fellow, Center for the Advanced Study of India at the University of Pennsylvania

‡Assistant Professor, Political Science, Ashoka University and Senior Visiting Fellow, Centre for Policy Research

Abstract

Biases in economic distribution have been widely observed in local settings, such as villages and neighborhoods, that select their leaders through elections. Existing explanations suggest that politicians target private benefits to those who are more likely to reciprocate with their votes or disproportionately responsive to local elites. We argue, in contrast, that these biases follow from the logic of "political selection," i.e., maintaining support among a minimum winning coalition. When one considers local settings, with high levels of information and dense social networks, this political selection logic is sufficient to generate elected leaders who simply have biased preferences that shape their distributive decisions. In particular, we show that while local elections select leaders who prefer to target their own supporters, elected local leaders target their poorest supporters due to a widespread preference to protect the survival of the poor. To test our theory, we develop a novel behavioral measure that isolates elected leaders' distributive preferences from electoral compulsions or elite pressures, which we implement poor villages across the Indian state of Rajasthan. According to our behavioral measure, elected leaders prefer to distribute 94% more to political supporters and 17% more to supporters one standard deviation below the mean village wealth. The results suggest local elections are consistent with significant distribution to the poor, albeit with strong political biases.

1 Introduction

In 1985, India's late Prime Minister Rajiv Gandhi declared: "Of every rupee spent by the government, only 17 paise reached the intended beneficiary."[1] The implication was clear; the centralized state lacked the capacity to efficiently target benefits to its citizens, particularly in India's many remote villages. Concerns like these sparked off a period of "democratic decentralization" in the 1990s across the developing world, which devolved authority over policy implementation to elected local governments (**??**). On one hand, decentralization was thought to empower local leaders, socially embedded in spaces, to increase efficiency in the distribution of government programs and services (**??**) while engendering responsiveness of leaders through elections. On the other hand, democratic decentralization granted local leaders sufficient discretion to favor some voters in their locality over others, which created significant biases in distribution – through partisan favoritism, ethnic favoritism, and elite capture to the detriment of the poor (**???**).

Research on developing countries broadly suggest that the state in developing countries is too weak to prevent biases in distribution that favor elites to the detriment of the poor, what is often termed "elite capture" (**??**). Moreover, research on vote buying suggests that leaders employ quid pro quo transactions of benefits for votes to induce poor voters to vote against their preferences (**??**). Fundamentally, this literature sees biases in distribution as emanating from "anti-democratic behaviors" such as the threats of elite retribution or party brokers monitoring what voter do in the polling booth.

In contrast to this pessimistic view, a large literature on local democracy suggests that simply enforcing basic "procedural democracy" with a secret ballot may be enough to ensure that local elections result in the selection of local leaders who are responsive to the preferences of the electorate (**?????**). A standard consequence of the median voter

[1] See Santosh K. Joy, "Rahul Echoes Rajiv Gandhi's Comments on Public Funds," January 17, 2008, <http://www.rediff.com/news/2008/jan/17rahul.htm>.

theorem and citizen-candidate models is that the winning candidate is likely to espouse preferences (including distributional biases) consistent with a majority (or plurality) of voters in the population. This *political selection* logic has been used to explain distributive biases in government programs (e.g., welfare spending) in democracies in the West (**???**) as well as partisan biases (**?**), but has largely been ignored in the study of democratic decentralization. In this view, biases in distribution are perfectly consistent with political selection from democratic behavior.

We argue that biases in distribution in democratic decentralization need to be understood through the lens of political selection in democracy. We claim that political selection implies that elected leaders display preferences to target benefits to their supporters, with a preference to disproportionately target benefits to the poor.

Local leaders are socially embedded in their spaces (whether it is a village or a neighborhood), so they can plausibly observe personal attributes, such as wealth and partisanship, of their constituents. Similarly, voters can observe the preferences of candidates and elected leaders due to shared history. The political selection logic suggests that voters elect leaders who have preferences to distribute benefits back to their supporters (otherwise, the voters wouldn't have voted for them). This differs from standard political economy models, where politicians are situated in large constituencies where they imperfectly observe their constituents. In such cases, the targeting of benefits towards co-partisans and supporters is due to the efficiency in targeting them through party brokers and workers (**??**).

Furthermore, unlike large constituencies, these local spaces display regular interaction and social dependence between citizens. This means that there is some understanding among citizens of who is objectively poor, and citizens are likely to show some concern for each other. The political selection logic, thus, suggests that voters elect leaders who have preferences to target the most needy as no other pattern of distribution

(e.g, targeting the rich or targeting a small ethnic group) would have such broad legitimacy. This differs from standard political economy models which see the targeting of poor voters as driven by electoral compulsions (**??**).

To understand the political selection logic, we analyze the preferences for distribution of benefits of elected local leaders from village councils in the Indian state of Rajasthan using novel behavioral measure that separates these preferences from pressures of local elites or electoral compulsions. The behavioral measure we develop uses an innovative "cross-referencing" design, which simultaneously surveys 10 randomly selected voters in a village and analyzes which of these voters leaders prefer to target with an economic benefit. We describe how to use lotteries to deduce preferences of leaders in a manner so that their decisions cannot be known to anyone else – thereby removing elite pressures or electoral compulsions from this measure of preferences.

Our results strongly confirm that the political selection logic of democracy generates targeting biases. According to our behavioral measure, elected leaders prefer to distribute 94% more to political supporters and 17% more to supporters one standard deviation below the mean village wealth. To establish the validity of our measure on the distributive behaviors of local leaders in our setting, we show that the results from our lab measure are consistent with observed targeting biases in citizen responsiveness and anti-poverty benefits in section 8 and appendix G. We also show that our results are not driven by Hawthorne effects (i.e., social desirability bias) and are robust to the ethnic and partisan characteristics of voters and leaders.

This article differs from existing research in key ways. First, much of the literature is premised upon the assumption that biases in targeted distribution result from a strategy of quid pro quo exchange and concerns about the relative efficiency of targeting different types of voters (**???**). By contrast, we argue and empirically demonstrate that observed biases in targeted distribution result from the underlying preferences of leaders in a

high-information context where efficiency concerns do not apply. Second, research on vote buying argues that local leaders target the poor at higher rates because they are more willing to sell their votes for low-value benefits (**??**). We argue that in subsistence-based settings, leaders who target the poor are more electable, and show that elecred local leaders are responsive to the poor even when they cannot feasibly gain electoral benefit from distribution. Third, we provide a counterpoint to research that highlights the "elite capture" of elected local leaders, which is understood to cause resources to be diverted from the poor in villages where poverty is pervasive (**?**). In contrast, we suggest that free and fair elections in this setting mitigate elite capture because local elections screen out leaders known to have distributive preferences that are too narrow to attract broad support among voters in the locality.

2 Distributive Preferences Under Local Democracy

In a poor village about 90 minutes outside a major city, the village's panchayat member (elected village council member), whom we shall call Mustafa, pleads his case. "I know what people like you think. I can't just keep benefits for myself and friends. I distribute them to the needy. After all, everyone in the village knows me." A young woman, recently widowed due to the tragic murder of her husband, repeats the sort of stories we hear all through the village. She tells us, "When my husband died, we had no money. Mustafa arranged for money to pay for my son's educational expenses. He is a good man." The next village over, we hear the same thing. "When my son got entry into [prestigious missionary school], we didn't have the money to pay the tuition. (Panchayat member) Farooq arranged for the money and paid our fees."

We are not used to hearing about such stories in Indian villages. The local leader in rural India is a much maligned character in the political economy literature, often

accused of engaging in egregious corruption and being beholden to elite interests (**???**). But how does this characterization explain the behavior described above? The elected leaders did not get any obvious electoral benefit from helping these beneficiaries (they were already supporters of their elected panchayat members), nor were the beneficiaries village elites. There was no obvious benefit to the elected leader in engendering turnout, as turnout in panchayat elections is already extraordinarily high, nor were there efficiency concerns in the delivery of benefits given the social proximity between villagers. In short, the elected leaders had no obvious electoral incentive to help these beneficiaries nor were they compelled to do so; they simply *wanted* to help the beneficiaries.

In contexts of weak state capacity, such as India, elected leaders are often given significant discretion in targeting selective benefits and providing general assistance to citizens. Here, their underlying preferences on whom to assist or target with selective benefits – what we refer to as *distributive preferences* – are likely to determine many observed outcomes. In this study, we are interested in the distributive preferences of local political actors selected through elections – an institutional process we refer to as *local democracy*. Since targeting behaviors are not rule-bound (i.e., programmatic) as in contexts of strong state capacity, it is important to characterize biases in the distributive preferences of elected local office holders who have such discretion.

In India and many other developing countries this is particularly pertinent to the "everyday assistance" that local leaders give to their constituents, such as filling out a form or contacting an important bureaucrat or politician to make sure a benefit is delivered, as well as in the targeting of welfare benefits, which constitutes the vast majority of their work (**??**). Unlike government schemes that only target a narrow subset of the population (e.g., pensions), everyday assistance impacts the economic well-being of all citizens. Moreover, since local representatives have broad authority over economic distribution in the village (and little policymaking power), we assume that voters' preferences

over candidates in local elections are primarily driven by concerns over the allocation of targeted benefits and personal responsiveness rather than programmatic issues.

2.1 Local Democracy and Political Selection

Local democracies have two defining features. First, because political leaders are elected by a plurality of voters, candidates must cultivate a large enough coalition to have some chance of winning the election. Second, because electoral constituencies are small, local democracy takes place in a setting of high information and dense social ties, where leaders and constituents know each other well. This differs from parliamentary or state elections, where voters have limited information on candidates and may not easily discern the demographic criteria upon which distribution is based (**?**). Since local leaders know voters personally, they can observe the demographic characteristics (e.g., economic need) of their constituents, and thus, efficiently target benefits to the voters they wish to target.[2] At the same time, voters under local democracy can observe the past behaviors of candidates, which means they can develop reasonably accurate priors on the targeting preferences of candidates and leaders prior to the election. This is especially true in local democratic settings because the most popular candidates are often those who have established reputations for effectiveness and responsiveness (See **??**).

Together, this means that voters can observe local candidates' distributive preferences and that elected local leaders possess the information required to efficiently target benefits according to voters' demographic characteristics and visible associations with political leaders in the village.[3] In contrast to models that emphasize efficiency concerns over targeting (**??**), this means that targeting biases are often a function of leaders' under-

[2]Extensive research corroborates the high-information nature of village politics (**??**).

[3]While the ballot is genuinely secret in India, research shows that local leaders can accurately identify those who belong to their local partisan networks (**?**).

lying distributive preferences in a context of discretion.[4] Moreover, as in any democratic setting selecting a single leader, a winning candidate must procure a plurality of votes which rules out extreme candidates or those who only curry favor with a few citizens (**?**).

2.2 Political Biases Under Local Democracy

Under local democracy, elected leaders in developing countries are often asked to take the place of rule-bound bureaucrats for the purposes of distribution. While this may have a positive impact on increasing accountability between citizens and local office holders, it may also generate targeting biases. These biases follow from the logic of electoral democracy, which requires candidates to develop and maintain minimum winning coalitions of voters in order to win the election. This means that voters have a strategic incentive to support a candidate that can plausibly win and who will be responsive to supporters rather than non-supporters (otherwise, there is no strategic benefit to voting for the candidate). This is how political selection generates leaders with political biases in distributive preferences.

Since we are interested in the preferences of the elected leader, we may restrict our discussion to candidates with a feasible chance of winning. Such a candidate will already have a significant base of political support, perhaps through existing work in the village as a broker or intermediary or due to a family history in politics. In principle, this stable or "core" base of support may be a function of co-partisanship, co-ethnicity, or some other social tie. Because core voters are likely to have close sociopolitical ties with the leader, and given the leader's incentive to maintain their personal networks and encourage supporters with weaker ties to the leader to more closely affiliate with

[4]This is particularly true where local leaders do not face re-election incentives as is often the case in our setting of village councils in India.

the leader (**??**), it stands to reason that any candidate will have a preference for targeting their core base of voters.

In any competitive electoral system with regular alternation in power, however, a winning candidate will have to appeal beyond this core base to reach a plurality of support. Therefore, candidates who may feasibly win the election must demonstrate distributive preferences that include plausible supporters in addition to core supporters. That is, since voters strategically select candidates that are most likely to distribute back to them, elections necessarily yield leaders with preferences for targeting their supporters to the exclusion of non-supporters. Moreover, when there is no dominant ethnic group in the constituency, as in our study, one expects the minimum winning coalition to be made up a multi-ethnic coalition of core voters and plausible supporters (**??**).

Finally, as implied by the high-information context of local democracy, elected local leaders can target constituents with a variety of political characteristics efficiently, which is much harder in larger, and more "anonymous", constituencies. Thus, rather than biases in distribution toward core supporters being driven by concerns of efficiency in delivery (**???**), we argue that this may simply be a result of political selection of local democracy . That is, voters strategically select leaders who have a preference to target their core supporters.

2.3 Variation in Political Biases and Pro-Poor Targeting

It is commonly argued that proper functioning of procedural democracy can lead to a redistribution of wealth from elites to poorer classes. In the literature, two major justifications are given for this idea.

In the taxation and welfare spending literature, it is understood that the distribution of wealth in any village or polity follows a similar distribution with a small number

of "elites" who are disproportionately wealthy. This implies that the wealth of the median voter is less than the average wealth in the population, and that pivotal voters redistribute wealth to address this imbalance in wealth (**??**).

In the context of dense social ties among voters, however, another explanation emerges. Especially in areas where a significant portion of the population is living at subsistence levels, the existing social structure is often used to mitigate economic and security risks in what is referred to as a "moral economy" (**?**). In settings where a large share of the society are poor in absolute terms, the theory of the moral economy suggests that a social expectation towards protecting the *poorest* members of society is widespread. This is the case because to allow a significant portion of the community to fall below subsistence levels would have dire consequences for the entire community in terms of sustainability, health and conflict.[5]

The welfare spending literature views this redistribution in aggregate terms and there is no clear claim as to whom among the poor is most likely to be targeted. By contrast, because the moral economy entails protecting those most likely to fall below subsistence levels, the leader should target the poorest members of his coalition disproportionately. While we do not adjudicate between these two theories, we present suggestive evidence that distribution patterns are more consistent with the moral economy logic. In either case, however, the redistributive logic of democracy suggests that a coalition that excludes the poor is not politically sustainable or seen as legitimate in settings of subsistence. Candidates who are unwilling to distribute to the poor and do not have the poor in their coalition cannot plausibly win an election.

[5]Recent work that emphasizes targeted the poorest members of a locality to cultivate a pro-poor reputation is compatible with the social expectation of this theory (**?**).

2.4 Hypotheses

In sum, we expect there to be significant political biases in the distributive preferences of elected leaders in local democracy due to basic coalition theory. However, there is significant variation in the amount of political bias demonstrated towards a citizen. The imperatives of the moral economy and general pro-poor preferences imply our key comparative static: among political supporters, the poorer the voter, the more likely he or she is to be targeted.

H1. Political Bias: *The elected leader will display a preference to target her own supporters, and an even greater preference for targeting her co-partisan base of supporters.*

H2. Pro-Poor Preferences: *Among those in the leader's political coalition, poorer voters are more likely to be targeted by the leader.*

3 The Case of India

We test our theory in poor villages in Rajasthan, a rural state in Northwest India. In this section, we demonstrate that Rajasthan meets the scope conditions of local democracy and describe the institution of the village council (gram panchayat or GP), and the role that the sarpanch plays in distribution.

3.1 Local Democracy in Rural India

Although the conditions for local democracy – which requires that voters can vote according to their preferences – may not have been present in rural India in the 1950s (**?**), research suggests that this system has broken down in recent decades and that a much more democratic form of politics has taken its place. First, **?**, based on fieldwork from

rural Rajasthan, suggests that the influence of upper caste landed elites has receded with the rise of educated, often lower-caste middlemen. Second, the role of coercion in elections has become substantially weaker as the decline in the power of landlords and sharp rise in lower caste political participation attests (**?**). Along with a strengthening of the secret ballot by a vigilant, independent Election Commission (ECI) (**?**), we have seen a rise in the autonomy of the Indian voter (**?**).[6] Moreover, interviews with block-level Congress Party and BJP leaders across the state suggest that GP elections in Rajasthan are often competitive.[7] Broadly speaking, intense party competition at local and higher levels, heterogeneity in vote preferences among members of the same ethnic groups (**??**), and weak capacity to monitor votes (**?**) suggests that elections in India are free and fair.

Local democratization was concretized through the 73rd amendment of the Indian constitution, passed in 1992, which gave the Panchayat Raj (rural local government) system constitutional status, and imposed federal requirements for elections of village council (gram panchayat) members and further integration of local government and government development functions. Sarpanch in our data were elected in 2010, which was the fourth election cycle since the 73rd amendment was passed.[8] Although this varies across states, sarpanch in Rajasthan and most other North Indian states, are directly elected by a plurality of the electorate of the entire GP. According to the 2001 Census of India, sampled GPs includes approximately 1100 households on average, and sarpanch report to know 95% of sampled voters personally.

The 73rd amendment also instituted a system of rotating quotas for marginal

[6]Although local elections are managed by state election commissions rather than the federal Elections Commission of India, recent research attests to the secret ballot in GP elections and the inability of local leaders to accurately identify voters' vote intentions or past votes (**???**).

[7]GP election results were not available during fieldwork.

[8]Prior to the 73rd amendment, Rajasthan held local elections under different requirements (**?**).

groups and women for elected positions in the GP. The rotating quota system, which changes the caste and gender requirements of candidates for sarpanch, overwhelmingly precludes incumbent sarpanch from contesting for re-election.[9] This has had important consequences for village politics, although recent work and our results suggest that these quotas have not fundamentally impacted distributive outcomes (**??**).

3.2 Political Context

We conducted our study in the predominantly rural state of Rajasthan, which is a competitive state with a 2-party system that has alternated between the BJP and Congress Party in every state assembly election since 1993, usually by small margins of victory. Although GP election results were not available at the time of fieldwork, interviews with block-level BJP and Congress Party leaders across the state, suggest that GP elections in Rajasthan are often competitive. Party symbols are not permitted on the ballot in GP elections; however, parties have broadly penetrated the GP, and recent studies including this one show that partisanship is salient to local distribution.[10]

Moreover, Rajasthan's two major parties compete for the votes of the poor (**?**).[11] This differentiates Rajasthan, and India more broadly, from monopolist contexts of machine politics where the "machine" party is entrenched in power and faces little competition for the votes of the poor (**?**).

[9] To illustrate this, 91% of sarpanch in our study were serving their first term.

[10] **?** find that Rajasthani voters correctly identified the party of their sarpanch 96 percent of the time; they also find strong partisan biases in targeting.

[11] We show below that this is the case in local elections as well.

3.3 Local Leaders and Everyday Distribution

Village council presidents (sarpanch) play a central role in mediating access to the state for their constituents through everyday responsiveness to personal requests and through their formal responsibilities over the local implementation of central and state government programs including sanitation (e.g., toilets), water access (e.g., wells), the placement of local infrastructure projects (e.g., village roads), and anti-poverty programs (**??**). While the decision of sarpanch to respond to citizens' requests for mediation comes closest to our scenario of full discretion, understanding distributive preferences is also important for understanding how local leaders employ their more limited discretion over policy implementation. For example, sarpanch play a key role in the implementation of anti-poverty programs such as the Mahatma Gandhi National Rural Employment Guarantee Act (MGNREGA), which guarantees 100 days of paid labor (on government infrastructure projects) to all Indian citizens and currently constitutes the lion's share of local government budgets. Sarpanch also have more limited but non-trivial discretion over final allocation of below poverty line (BPL) cards, which are required for eligibility to benefits provided through the Public Distribution System (PDS) (**?**).

4 Design and Empirical Strategy

We developed a novel behavioral measure designed to pick up local leaders' underlying targeting preferences (i.e., biases) over real world populations.[12] We asked sarpanch to target 5 tokens any way they wished among 10 randomly selected voters in the GP to influence a non-trivial lottery prize– without asking for the reason for doing so. In what

[12]This captures local leaders proclivities to favor some of their constituents over others in the distribution of selective benefits and everyday responsiveness rather than examining the targeting of a specific policy benefit.

we refer to as a cross-referencing survey design, these voters were surveyed ahead of time and we determined leaders' targeting biases by cross-referencing data from voter and local politican surveys.[13] Crucially, unlike survey experiments or other standard surveys, the sarpanch did not have to directly admit to targeting a co-partisan, which is subject to social desirability biases. Furthermore, while survey experiments and conjoint analysis have become increasingly popular tools to measure biases, these methods are usually applied to hypothetical populations and not as effective for the measurement of targeting biases in the actual population. This makes cross-referencing likely more efficient in detecting targeting biases than these other methods, and uniquely applicable to detecting biases in the actual population–particularly in local contexts where social ties between voters and leaders are consequential.[14]

Our design continues a recent tradition of lab-in-the-field experiments (**?**) that investigate the impact of ethnicity (**?**), partisanship (**?**), and democratic selection (**?**), on targeting biases and economic distribution. Our lottery measure of distributive preferences was embedded in cross-referenced sarpanch and voter surveys conducted in 84 GPs across Rajasthan from January to February 2013. The sample frame was rural, poor contexts characterized by some degree of electoral competition and voter respondents were restricted to heads of household.[15] Specifically, we restricted sampling to subdistricts (blocks) with average margins of victory in block-level (i.e., panchayat samiti) ward elections of 15% or less that were at least 75% rural.[16] GPs with below poverty line

[13]For instance, we could ascertain a partisan bias in targeting if sarpanch targeted co-partisans at a higher rate than non-co-partisans, which we could discern by linking data on partisan preferences from voter and politician surveys.

[14]We did not conduct a direct comparison with other methods. Further work should be used to discern the relative efficiency of cross-referencing in detecting targeting biases compared to other methods.

[15]The restriction to predominantly male heads of household maximized the chance that leaders and voters interacted in the past.

[16]Block level, or panchayat samiti, elections are the second tier of local government in India and the

(BPL) rates of 20% and contested local elections were randomly sampled in blocks that met these criteria.[17] Our sample frame allows us to capture contexts of local democratic competition and subsistence societies where the implications of political selection for responsiveness to the poor is particularly important.

To identify local politicians' distributive preferences, and the targeting biases therein, we embedded a lottery with a 200 Indian Rupee ($3.64 USD) prize in a survey of sarpanch to model targeting preferences under a budget constraint.[18] Sarpanch were shown a page of names and photographs of 10 randomly sampled voters surveyed the previous day, which were obtained from publicly available voters lists. Sarpanch were given 5 tokens and asked to allocate them in any way they wished across these 10 villagers. Sarpanch were told that a lottery with a 200 rupee prize (a little more than one day of agricultural wage labor) would be held at the end of the survey, and that each token a particular voter received would make his chance of winning the prize 'much higher' and that multiple tokens could be given to the same villager. This design forced sarpanch to allocate tokens to no more than 50% of sampled villagers, which makes the measurement of targeting biases possible.

Practically, we included each voter survey respondent's name on slips of paper once and added one additional slip per token given to the respondent. Thus, if a sarpanch gave all of his five tokens to one person, the probability that this individual's name was picked was approximately six times that of all other sampled respondents from his GP. If he gave one token to each person, villagers who received tokens were seven percentage points more likely to win the prize than those who received no tokens.

---

lowest level where party symbols are allowed on the ballot; it is also the lowest level where election data was available.

[17]Further details on the sampling procedure are provided in appendix A.

[18]Note that although our lottery prize is relatively modest, a large literature in economics on lab games shows that increasing the size of payoffs has no effect on distributive behavior (**??**).

To ensure that our measure was not influenced by political incentives or other constraints, token allocations was kept secret from voters (sarpanch were told that this would be the case). Moreover, since every voter had some chance of winning the lottery (even if they received no tokens), villagers could not infer how the sarpanch allocated tokens from observing the winner of the lottery. We dispersed the prize as an unannounced electronic payment in the form of mobile phone credit after the conclusion of the survey; sarpanch were not told who won the lottery and lottery winners themselves may not have realized a lottery took place. In sum, the design of our behavioral measure is not plausibly shaped by electoral incentives by design.[19] While we removed electoral incentives, we argue that our measure captures routine distributive decisions and demonstrate that this is the case in section 8.

4.1 Predictors and Cross-Referenced Measures

The dependent variable in the analysis is the number of tokens given to an individual. Our analyses rely on a number of predictors discussed below. To test for whether sarpanch prioritize their supporters, we asked the sarpanch whether each of the voters in his GP voted for him. If the sarpanch answered in the affirmative, the individual was coded as a perceived electoral supporter.[20] To capture partisan ties, we asked voters and sarpanch whether or not they feel close to any particular party, and then asked them to name the party to which they feel close. When the voter reported that he or she feels close to the same party reported by the sarpanch, the voter was coded as a "co-partisan." The ethnicity measure categorizes the sarpanch into politically salient caste categories

[19]Sarpanch were largely ineligible for re-election scheduled for 2015: two years after the conclusion of the survey.

[20]This measure broadly captures voters who the sarpanch does not consider to be certain non-supporters.

and Muslim religion based on voters' self-reported identities.[21] We defined a co-ethnic as any voter who fell into the same category as the sarpanch. Finally, to understand distributive preferences vis-à-vis the wealth of the voter, we constructed a scale based on an item response model of observable assets of the voter. The cross-referenced measures of co-partisanship, co-ethnicity, asset wealth were specifically designed to minimize social desirability concerns in the measurement of targeting biases.

4.2 Statistical Model

The "comparative static" of interest, of asset wealth conditional political affiliation, is measured at the GP level. This is a non-causal exercise since the attractiveness of allocating to a voter is dependent upon his/her relative attributes as compared to others in the same GP. The key observation that allows for identification of the empirical model is that mean allocation in a GP is always identical, the number of tokens divided by the number of potential receivers, or $5/10 = 0.5$. If all the predictors are centered around their means in the GP, the constant term in a regression is fixed. In particular, let y_{iv} denote the allocation given to potential receiver i in GP $v \in \{1, \ldots, V\}$. Consider predictors $x_1, \ldots, x_J$. Let us denote the mean of predictor x_j in GP v as $\overline{x}_{jv}$. Since the number of tokens is in the form of count data, a Poisson regression (accounting for overdispersion) is appropriate. A quasipoisson regression model provides the same mean function as poisson regression, λ_i, for observation i, but allows for overdispersion by estimating variance $\sigma^2\lambda_i$ at observation i.[22] Because the relative impact of each variable is likely to be different in each GP, we fit a hierarchical model which varies coefficients by GP. The

[21] We used two different definitions of co-ethnicity, jati and varna, which yield substantively similar results and present results on the former. Reflecting ethnic politics in India, Muslims were coded as a separate category in both measures.

[22] In the standard poisson distribution, the variance is fixed at λ_i, the same as the mean.

model can be written as below:

$$y_i \sim Poisson(\lambda_i, \sigma^2) \text{ where } \sigma^2 \text{ denotes an overdispersion parameter} \tag{4.1}$$

$$\lambda_i = \exp(\beta_0 + \beta_1(x_{1iv} - \overline{x}_{iv}) + \cdots + \beta_J(x_{Jiv} - \overline{x}_{Jv}))$$

$$y_i = \lambda_i + \varepsilon_i \text{ where } \varepsilon_i \sim N(0, \sigma^2\lambda_i)$$

$$\beta_{jv} = \beta_j + b_{jv}; \quad b_{jv} \sim N(0, \sigma_V^2) \quad \text{s.t.} \quad v \in \{1, \ldots, V\}$$

where β, σ and σ_V denote parameters in the regression model, and x_{iv} denotes a predictor for individual i in GP v.

5 Scope Conditions

As argued in section 2, we are particularly interested in understanding how local democracy functions in subsistence-based societies. We begin the section by demonstrating that our theoretical scope conditions are satisfied in the sample, namely: 1) politics is reasonably competitive at the local level; 2) a sizeable subset of sarpanch have preferences that are likely known to constituents; 3) a significant proportion of voters can be characterized as poor; and 4) the relative wealth of citizens in the GP is known to the sarpanch.

5.1 Characterizing the Sample

Our theory of local democracy is built upon the assumption of free and fair elections in a largely subsistence-based population in the context of reasonably high information about citizens of the GP from the sarpanch and vice versa. We assess whether our sample meets these scope conditions.

In order to construct an asset wealth measure, we relied on readily verifiable information, i.e., those things that could be confirmed by the enumerator. The measure is constructed upon whether the respondent owns: 1) a "pucca"/"semi-pucca" dwelling or permanent dwelling structure; 2) a scooter/motorcycle; 3) a bicycle; 4) a television; 5) proper toilet facilities; 6) a refrigerator; 7) a fan; 8) mobile phone; and 9) electric pump set. Table 1 displays the average for each of these (binary) items in the population and compares them against census (or national sample) estimates. The average levels observed in sampled villages (in 2013) are broadly lower than those reported at an all-India level two years before with the exception of scooters and the rapidly growing mobile phone. This suggests that our village sample is quite poor even by average Indian standards (and certainly by most absolute standards).

Item	Mean in Sample	Census/NSS 2011
Pucca House	0.73	0.82
Scooter	0.26	0.21
Bicycle	0.26	0.45
Television	0.33	0.47
Proper Toilet	0.15	0.47
Refrigerator	0.10	0.17*
Electric Fan	0.63	0.66*
Mobile Phone	0.82	0.63
Electric Pump	0.19	–

Table 1: Mean Levels of Assets

* Data are adapted from the 66^{th} round of the National Sample Survey (NSS) because they are not included in the 2011 Indian Census. Data on electric pumps are not available in either dataset.

Each of the items above is a binary variable, and a 2-parameter Rasch model (**?**) was fit using Markov Chain Monte-Carlo (MCMC) using the program `JAGS` to construct a raw asset score.[23] The raw asset score gives approximately ten different "scores," suggesting reasonably high levels of correlation between owning these assets.

We look at the relationship between these 10 values on our asset index and the percentage of the sample at that asset value owning a refrigerator or a proper toilet (two natural markers of economic development). The results are shown in table 2. In both cases, even the 80^{th} percentile of wealth does not meet the all-India averages for those amenities. Taken together, this implies that a substantial proportion of these villages are very poor, and, at least in terms of asset ownership, our sample displays a significant

[23] Let $y_{ik} \in \{0,1\}$ denote a binary outcome variable for person i and object k, $1 \leq k \leq K$. A two parameter Rasch model fits:

$$P(y_{ik} = 1) = logit^{-1}(\alpha_i - \beta_k)$$

where β_k is a parameter placing the object on a wealth scale and α_i is the value of the asset index for individual i.

level of inequality.

Asset Score	Sample Size	Percentile	% With Toilet	% With Refrigerator
-2.2	34	4	0	0
-1.5	101	16	0	0
-0.9	155	35	0	1
-0.3	161	54	1	4
0.4	150	72	3	13
0.9	92	83	9	22
1.4	75	92	32	36
1.9	39	96	74	67
2.4	25	99	100	100
3.0	7	100	100	100

Table 2: Toilet/Refrigerator Ownership by Asset Score

For the analysis, we generated an asset index by standardizing the raw asset score to have mean 0 and standard deviation 1 within each GP. The value of the asset index for an individual can be interpreted as the number of standard deviations the individual's asset score differs from the mean asset score in the GP. Since the asset index is a function of the average asset wealth in the GP, the index has no meaning in terms of *aggregate* wealth, only in terms of *relative* wealth.

To establish that sarpanch observe the relative wealth of their constituents, we asked the sarpanch to rank individuals from wealthiest (1) to least wealthy (10). Despite some small non-linearities in the middle, figure 1 shows that our constructed asset measure is broadly consistent with the ranking provided by the sarpanch. The person rated the poorest is on average 0.56 standard deviations poorer than the mean individual in the GP according to our asset measure, and the person rated the wealthiest is on average 0.47 standard deviations wealthier than the average person according to our asset wealth measure.

Figure 1: Sarpanch Assessments of Wealth vs. Asset Measure

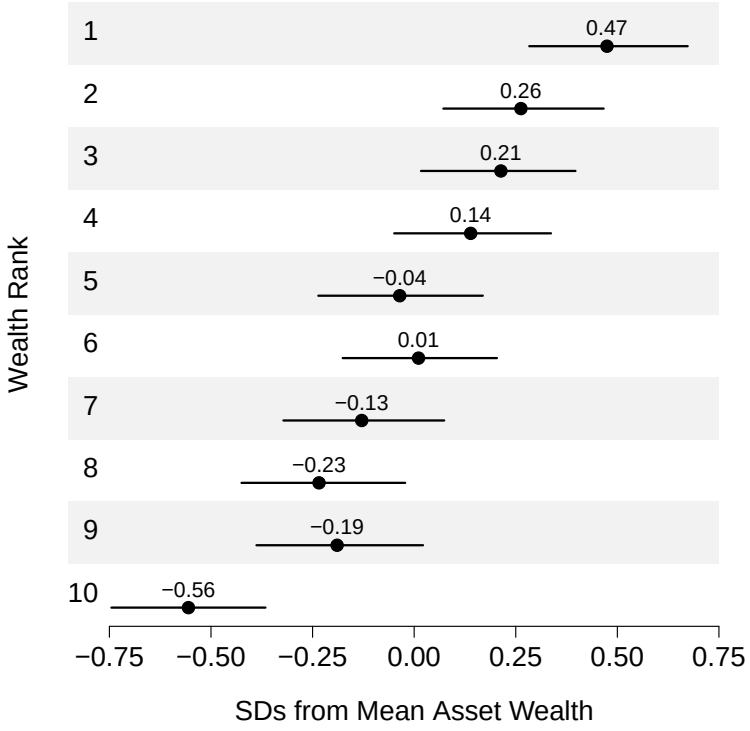

In a context of local democracy, leaders are directly able to assess the wealth of their constituents, and this strongly related to objective measures of observable wealth, as shown in figure 1. Rather than relying on local proxies, or brokers, for information about wealth as in much of the literature (**?**), both voters and leaders understand that distribution can be based on commonly observed levels of wealth. This implies that leaders can target the poor without much risk of misallocation, and that voters can reliably assess how well the leader is targeting the poorest citizens.

Voters and sarpanch in our sample have reasonably close ties. As stated, there is an average of only 1100 households per GP in sampled GPs, and sarpanch reported to know 95 percent of sampled voters personally. Moreover, candidates for sarpanch often served as unelected fixers or elected GP ward members prior to contesting elections for sarpanch (**??**), with 31 percent of sarpanch in our data serving as GP ward

representatives previously.[24] An additional 32 percent had a family member currently or previously in elective office, which provides voters with information on candidates' families' distributive preferences. This provides strong evidence that voters in our sample can feasibly surmise the distributive preferences of candidates for sarpanch prior to election day.

Finally, local democracy requires some degree of competition for the screening mechanism of elections to take effect. At the outset, the sample frame includes GPs that were considered moderately or very competitive by block-level party leaders and non-competitive GPs were excluded. Second, we coded partisan competition at the polling booth level for each polling booth in our sample for the 2014 parliamentary election.[25] The median (and average) effective number of parties/candidates (ENP) at the polling booth level is 2.1. If two parties each received exactly 50% vote share, ENP would take the value of 2; as such, ENP values greater than 2 are typically seen as a reasonable measure of a competitive electoral scenario. Third, while GP election data is unavailable, 90% of sampled sarpanch were serving their first term and interviews suggest that these elections are often hotly contested.[26] Given that our electoral setting displays high levels of alternation and competition, we can be reasonably certain that voters are making genuine choices and that their preferences and strategic incentives are reflected in their elected leaders.

[24]Ward representatives are elected council members of the GP, which is led by the sarpanch.

[25]While the BJP was a decisive overall winner in this election – the only one that could be accurately linked to our data at the time of writing – we show substantial multi-party competition at the local level even in this case.

[26]This level of first-term sarpanch is plausibly due to a system of rotating caste and gender quotas, which often makes the incumbent ineligible to run for reelection; nonetheless, this means that there is substantial alternation in power at the GP level.

6 Characterizing Political Biases in Allocation

In characterizing targeting biases, we remind the reader that an "unbiased" allocation would put the expected number of tokens at 0.5. Anything above this value can be viewed as evidence for a *premium* in allocation for the voter. At first blush, there seem to be a strong premium for perceived political support. The average perceived non-supporter received 0.26 tokens, while the average perceived supporter received 0.61 tokens. We regard the set of voters that report being co-partisans of the sarpanch in addition to supporters as a more stable base of supporters due to partisan affinity.[27] When we further subdivide political support by co-partisanship, we see quite a bit of variation. Co-partisan supporters receive 0.81 tokens on average, while non-copartisan supporters receive 0.51 tokens on average. Non-supporters do not receive many tokens on average, whether co-partisan (0.32) or not (0.22). Similarly, we show in appendix D that non-supporters also receive substantially fewer tokens when co-ethnicity is taken into account. At the same time, these aggregates may be correlated to relative asset wealth, so we must measure these effects within our modeling context. Appendix C reports 8 different regression models, that adhere to the empirical strategy above, controlling for relative asset wealth in a GP, as well as electoral support, co-partisanship, and co-ethnicity between voter and sarpanch measured in various ways.

[27]In Indian village politics, parties have an incentive to make sure their supporters vote for a leader from the same party, so that this leader can hook into the larger party organization across the state.

Figure 2: Expected Tokens and Electoral Support

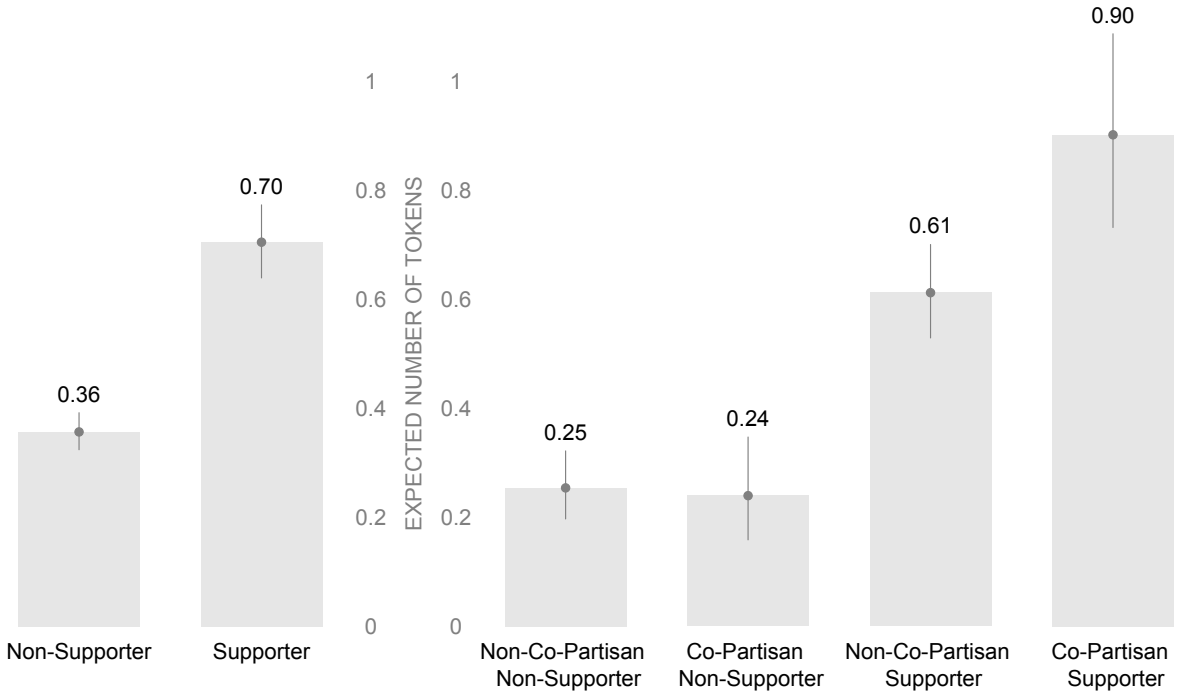

Figure 2 reports the estimated expected number of tokens for perceived electoral supporters and non-supporters, further subdividing the effects by co-partisanship (with 95% posterior/credible intervals). The expected number of tokens for supporters and non-supporters are derived from our core model, assuming that the individual has mean wealth in the GP, and that the mean number of supporters in the GP is held at the sample mean of supporters (68%). The expected number of tokens for the interaction between co-partisanship and political support is derived from a more complicated model that controls for the two-way interactions between support and co-partisanship, as well as interactions with relative asset wealth, as shown in column 4 of appendix C, calculating predicted values at mean GP wealth and the sample mean for each of the categories.

At the mean level of GP wealth, a supporter is predicted to receive nearly twice as much on average (94%), as compared to a non-supporter. To test whether the strength of

the sociopolitical tie affects the level of allocation, we test whether there is a discernible increase in allocation to co-partisan supporters (whom we view as more stable supporters). When further subdivided by co-partisanship, we see that co-partisanship has little effect on allocation to non-supporters. However, co-partisan supporters are predicted to receive 48% more allocation than non-copartisan supporters at the mean level of GP wealth.[28] Taken together, our results strongly confirm the expectation of targeting biases towards political supporters in hypothesis 1; this effect is particularly pronounced for more stable co-partisan political supporters.[29]

7 Targeting of the Poor in Allocation

In this section, we investigate our second hypothesis (and our key comparative static) that among political supporters, the poorer the voter, the more likely he or she is targeted. Figure 3 displays the the expected number of tokens at the mean level of electoral support (68%) estimated from our core model (with 95% posterior/credible bands). The graph shows that targeting of the poorest citizens among both supporters and non-supporters, with far more pronounced targeting of the poorest among electoral supporters. The gap between supporters and non-supporters in allocation is much greater, with even the wealthiest supporters predicted to receive more than the poorest non-supporters in each model. Among political supporters, a one standard deviation decrease in wealth from the mean is associated with a 17% increase in the expected number of tokens. This supports our expectation that elected leaders will bias distribution

[28]These differences are highly significant with 99% or more of the difference in the posteriors being bounded away from zero.

[29]We note that we do not find statistically significant effects for co-ethnicity. A full discussion of the results is included in Appendix E. This provides some support for the idea that minimum winning coalitions in multi-ethnic societies tend be built around political and partisan identities (**?**).

toward their poor supporters while excluding poor non-supporters even when they are extremely poor.

Figure 3: Electoral Support vs. Asset Wealth

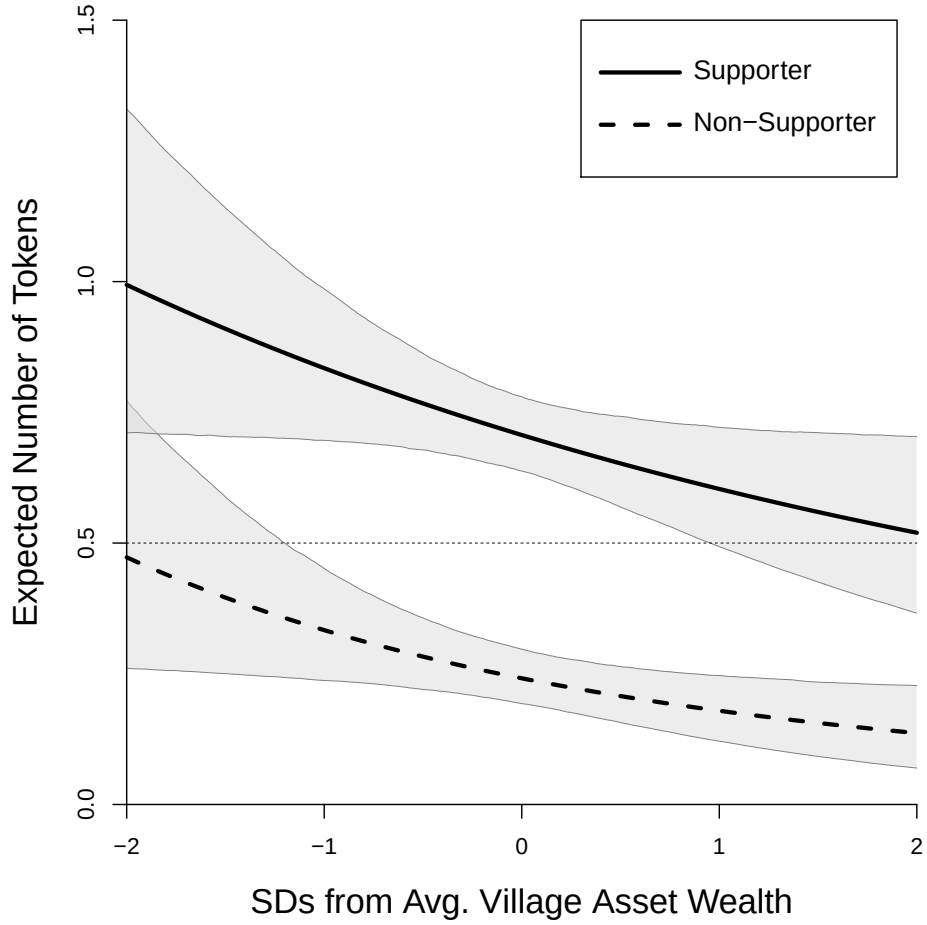

7.1 Welfare Implications for the Poor

We have shown that the expected allocation is substantially greater among poorer supporters of the elected leader, while poor non-supporters are often excluded. But what are the aggregate welfare implications of this pattern of allocation, i.e., does it result in overall targeting towards the poorest in the GP?

At the outset, it is important to note that in 81% of GPs, the sarpanch allocated a

token to an individual with a raw asset score less than zero, i.e., an individual likely living at subsistence levels who is poorer than the median citizen in our sample. In appendix E, the regression coefficient on the relative asset wealth of the voter remains remarkably consistent over each of the eight models, with the various models predicting a 21-23% increase in allocation to a voter with asset wealth one standard deviation below the GP mean, holding all else constant. This implies, that even controlling for the most relevant voter characteristics, substantial targeting towards the poorest citizens of the GP is observed.

Figure 4: Expected Number of Tokens vs. Asset Wealth

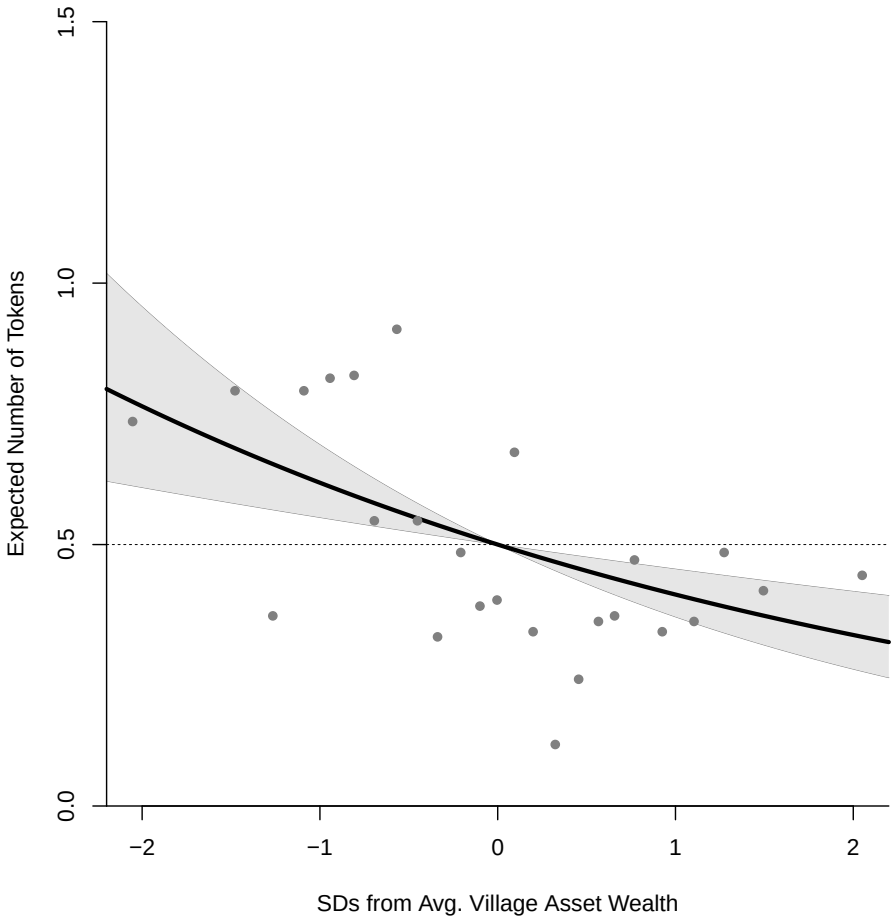

Figure 4 displays the predictions generated from our core model, which controls for relative asset wealth, perceived political support, and the interaction between the two (column 2 in Appendix C). In order to generate the figure, we assumed the level of support was at the GP mean and generated curves from the fixed coefficients in the models. The gray region around the curve is the 95% posterior interval (generated from the model) at each level of wealth.

The gray points in figure 4 are the binned averages of tokens across 25 bins (approximately 34 observations per bin), with cutpoints spaced every 4 percentile points, over the distribution of relative asset wealth. That is, the points display the average number of tokens given to individuals included in a particular bin of relative asset wealth. The effects are quite strong; in the bottom 40 percentile of relative asset wealth, only one binned average is below 0.5, and in the top 40 percentile of relative asset wealth, no binned average is above 0.5. We also repeat the laboratory measure of discretionary distribution for anti-poverty benefits in appendix G, yielding substantially similar biases (and aggregate welfare impacts) to those described above.

8 Robustness and Connection to Actual Distribution

It is plausible that our results are driven by characteristics of sampled sarpanch that may be associated with pro-poor targeting–such as whether the sarpanch is female, a member of the lower castes, or affiliated with the Congress Party–rather than the selection effects of the institution of local democracy.[30] In appendix D, we show these patterns

[30]The Congress party generally has a poorer, lower caste voter base than the BJP and lower caste leaders may be more pro-poor than upper caste leaders; moreover, research on experimental games shows that women are more likely to have preferences of inequality aversion than men (**?**).

do not hold in our data.[31] To test more complicated hypotheses about the impact of caste/religious identity or wealth (landedness) of the sarpanch, we consider a large set of possible voter and sarpanch characteristics as confounders to our core model, and the magnitudes/significance of the variables of interest remain very similar to previous models.

Another concern is that the distribution we measure is driven by "Hawthorne effects," that is, sarpanch behave in a way that would satisfy the researcher. In order to understand whether this occurred, we coded whether the sarpanch self-identified whether each of the voters placed in front of him was a supporter after the distribution had taken place. If the sarpanch were attempting to display distribution that is socially desirable, we would not expect to see allocation towards such supporters (since it obviously is a deviation from any programmatic ideal of distribution). In order to test whether our results are driven by Hawthorne effects, we calculate the percentage of GPs in which, according to our data, sarpanch target their supporters more heavily. Supporters were targeted more heavily than non-supporters in 87% of GPs.

Moreover, if Hawthorne effects are strong, we should see that sarpanch provide tokens to the those whom they personally identify as one of the two poorest individuals in the village, even when they are non-supporters. This was often not the case; only 40% of non-supporters ranked as among the two poorest individuals received any tokens as compared to 75% of supporters. This demonstrates that our behavioral method is effective in picking up social biases that differ from any programmatic ideal, and that the observed results aren't purely driven by Hawthorne effects.

A final concern is that our lab method is too disconnected from, and thus has little

[31]The "comparative static" part of our claim is important here. Congress supporters tend to be much poorer, so in aggregate Congress sarpanch target poorer voters. But *conditional on political support* both Congress and non-Congress sarpanch display similar levels of pro-poor behavior.

relevance for, the actual distribution of everyday benefits and help. Since this form of discretionary assistance is not characterized by a single large benefit but rather general brokerage or help, we looked to understand the relationship of our measure to general notions of "helpfulness" rather than a single government-regulated benefit (in which the sarpanch would have limited discretion). In particular, we compared our observed lab behavior to voters' perceptions of sarpanch behavior. We asked voters: "Do you believe the sarpanch would help you if you approached him/her for help?" We find a very strong relationship between our token-based measure and voter perceptions of helpfulness. If the voter did not believe the sarpanch to be helpful, she received an average of 0.39 tokens, whereas a voter that believed the sarpanch to be helpful received an average of 0.57 tokens. This constitutes a significant difference, and a 48% increase in allocation associated with those who found the sarpanch to be helpful. This suggests that our measure has a natural real-world analogue. We also note that when we model sarpanch discretion for anti-poverty benefits (yielding similar political and pro-poor biases), the resulting behavior tracks the true distribution of anti-poverty benefits well (see appendix G).

9 Discussion

This article shows that local democracy selects leaders with distributive preferences that display political biases but are sufficiently broad to include a plurality of voters in the locality, and in the context of subsistence-based societies, select those who prioritize the poorest members of that coalition. We demonstrate that this is the case with evidence from a behavioral measure embedded in a cross-referenced survey and analyze this data with a statistical method that appropriately considers the complexity of this data. Our method has several advantages over existing methods, which are susceptible to flawed measurement due to social desirability bias and cannot easily discern motivations for

distribution over the actual population.

This study advances research on distributive politics in several important ways. First, we consider the consequences of local democracy – a high-information context where efficiency concerns do not apply – on distribution. Existing models suggest that pivotal swing voters should be targeted where this can be done efficiently (**???**). However, we find a baseline preference for core targeting due to a political selection logic, even where efficiency concerns don't apply. Moreover, while existing theory suggests that core targeting can be an effective electoral strategy for coalition maintenance (**?**) or vote mobilization (**?**), our results cannot plausibly be shaped by such strategic concerns given the private nature of our measure of distributive preferences.

Second, we provide an alternative logic for the targeting of poor citizens that is rooted in representation in subsistence-based societies, where a moral economy to target the extreme poor is likely to be salient and broadly viewed as the most the legitimate distribution. Our logic differs from research on vote buying, which suggests that the poor are most likely to be targeted because they are most responsive to low-value handouts and other targeted benefits (**???**). In contrast to "perverse accountability" posited in the vote buying literature, our theory suggests that the poor are targeted in subsistence-based societies precisely because of a properly functioning democracy, namely because pivotal voters with pro-poor preferences select pro-poor leaders. This reflects broad evidence that the ballot is secret and that vote monitoring and other aspects of quid pro quo exchange are unfeasible in India (**??**), while our behavioral measure removes electoral incentives by design. Although we do not question that the construction and maintenance of local networks is strategic, we argue that pro-poor leaders are likely to be responsive to the poor irrespective of future electoral benefit.

Third, this article contributes to research on local distribution in contexts of decentralization. Much of this work is focused on the allocation of benefits from a small

number of welfare programs by elected local leaders (see e.g., **?????**). Although this work often infers local leaders' preferences and strategies from these outcomes, local leaders such as sarpanch have limited authority over such outcomes (**?**). By developing a measure that uniquely captures the personal proclivities of sarpanch, this article captures biases in targeting relevant to a wider range of responsiveness over which the sarpanch has near-complete discretion. This is important because while research shows that anti-poverty policy benefits–which involve a variety of political actors and bureaucratic constraints beyond the village– have often failed to reach the poor, our results suggest that elite capture at the level of elected local leaders is quite exaggerated (see also **?**).

Local democracy crucially matters for our results. Wealthier households do not receive greater priority by the elected representative because the vote of each voter receives equal weight, irrespective of the personal characteristics of the voter. Without elections, households are differentiated, and the leader prefers to hold sway over the wealthiest, highest status households in the area since this maximizes the extent of his influence. Recent work that employs our method in the Indian state of Bihar (**?**) corroborates this intuition, finding that unelected leaders are systematically biased towards the wealthiest citizens, the *opposite* of the finding in this paper.

Above all, our findings suggests that in settings of subsistence, procedural democracy at the local level leads to the political selection of leaders who prefer substantial targeting to the poor. This is important because where state capacity is weak, as is the case in rural India and many other contexts in the developing world, the screening mechanism that local elections provide may be the best assurance of post-election distribution and everyday responsiveness to the poor. At the same time, contexts of discretion are characterized by serious political biases in targeting, which leads to the exclusion of poor non-supporters. While this is consistent with democratic responsiveness in a non-

programmatic setting, this means that without strengthening bureaucratic oversight and bottom-up social pressure, the poorest citizens who lack political ties to elected leaders are likely to be excluded in local democracies.

A Sampling Strategy

The voter and sarpanch surveys sampled 96 gram panchayats in seven districts, twelve blocks and six of Rajasthan's seven administrative divisions.[32] As mentioned in the article, one GP President, or sarpanch, could not be interviewed, which yielded a sample of 95 sarpanch. The sample in this article was further reduced to 84 sarpanch on account of coding mistakes on the tokens measure made by our survey team.

The sample generalizes to voters and local politicians in rural contexts with a moderately high share of households below the poverty line and moderate inter-party competition. To build the sample frame for this population, we used 2001 census data on the rural composition of blocks,[33] data from the Government of Rajasthan on the share of below poverty line (BPL) households across blocks in 2001, and Election Commission data on political competition in panchayat samiti election– the tier of the panchayat raj system above the gram panchayat, which aligns with administrative blocks.[34]

We restricted the sample to blocks with a 75 percent rural population according to the 2001 census to reduce the chance of sampling GPs that function as suburbs, and excluded blocks with less than 20 percent of households in the BPL category in 2001 to ensure that the chance of sampling voters eligible for anti-poverty programs at random was non-trivial. This ensures that our sample is one of pervasive poverty and that the lottery benefit is salient in this population. We also excluded blocks where the median margin of victory across elections to all ward representative elections to the Panchayat Samiti– a sub-district, or block, level electoral body one tier above the GP– was greater

[32]Rajasthan has 33 districts, 249 blocks, 7 administrative divisions, and 9177 gram panchayats in all.

[33]Government data on the share of BPL households across gram panchayats was from 2001. More recent data was not available at the time of fieldwork in 2013.

[34]This is the lowest level of aggregation at which election commission data is available from a central source and the lowest level that permits party symbols on the ballot.

than 15 percent to increase the chance that we selected competitive GPs.[35]

After this restriction was applied, approximately 60 of 249 blocks were eligible for sampling in the state. Logistical concerns required that we sample two blocks in each district to the extent possible. This reduced the list to approximately 50 blocks. I randomly sampled one district in 5 of Rajasthan?s seven divisions from a pool of districts in which three or more blocks were eligible for sampling according to these criteria. Two blocks were randomly selected from the pool of eligible blocks in each district. In Udaipur, the sixth division selected, three eligible blocks did not exist in any one district; As a practical alternative, we randomly selected one block each from two neighboring districts in the division: Udaipur and Rajsamand.

Once 12 blocks were sampled, one of us collected data on political competition across gram panchayats through interviews.[36] Members of the research team interviewed block party presidents– party organizers immersed in the politics of gram panchayats in their block? who were asked to characterize the level of competition between Congress and the BJP as non-competitive, somewhat competitive, or very competitive. Of the 452 GPs in 12 sampled blocks, 180 were described as non-competitive, 133 as somewhat competitive, and 139 as very competitive. To increase the chance that the target population would be sampled, given resource constraints, non-competitive GPs were dropped from the pool for sampling. In each block, 4 GPs were randomly sampled among those coded as somewhat competitive and among those coded very competitive respectively.

Subsequently, one ward in each sampled GP (with an average of 100 households

[35]Each member of this block-level legislative body is elected from one single ward and elected according to a first past the post electoral rule. We use the median margin of victory across ward elections to the Panchayat Samiti as gram panchayat electoral data could not be obtained during fieldwork.

[36]This was necessary because electoral commission data on gram panchayat elections is not available from a centralized source.

per ward) were randomly sampled.[37] We randomly sampled household in sampled wards using the gram panchayat voters? list, which is public information provided by the Election Commission. We sampled (predominantly male) heads of household in randomly sampled households because they are generally the household member most engaged in village politics and citizen-state relations.[38] The elite survey was fielded the day after the vote survey was completed in a given GP.

[37]This was done according to the design of another article from this survey project which required that all sampled voters lived in one GP member's ward.

[38]To identify heads of household, interviewers were instructed to request to speak to the head of household upon approaching each sampled household. If heads of household were not at home, interviewers were instructors to either interview them in the fields in which many of them worked or to return to the household later in the day. If they did not return, supervisors provided alternative respondents who were also randomly selected from a voters list.

B Descriptive Statistics

Table 3: Voter Characteristics

Statistic	N	Mean	St. Dev.	Min	Max
Upper Caste	839	0.094	0.292	0	1
Rajput	839	0.105	0.307	0	1
Jat	839	0.105	0.307	0	1
Other Backward Caste	839	0.316	0.465	0	1
Scheduled Caste	839	0.167	0.373	0	1
Scheduled Tribe	839	0.068	0.252	0	1
Muslim	839	0.086	0.280	0	1
Illiterate	839	0.327	0.469	0	1
Some Primary School	839	0.230	0.421	0	1
Class 5 Pass	839	0.194	0.396	0	1
Class 8 Pass	839	0.138	0.345	0	1
Class 10 Pass	839	0.050	0.218	0	1
College Degree	839	0.089	0.285	0	1
Supporter	839	0.682	0.466	0	1
Co-Partisan	839	0.352	0.478	0	1

Table 4: Sarpanch Characteristics

Statistic	N	Mean	St. Dev.	Min	Max
Upper Caste	84	0.107	0.311	0	1
Rajput	84	0.155	0.364	0	1
Jat	84	0.083	0.278	0	1
Other Backward Caste	84	0.238	0.428	0	1
Scheduled Caste	84	0.202	0.404	0	1
Scheduled Tribe	84	0.048	0.214	0	1
Muslim	84	0.048	0.214	0	1
Illiterate	84	0.167	0.375	0	1
Some Primary School	84	0.226	0.421	0	1
Class 5 Pass	84	0.226	0.421	0	1
Class 8 Pass	84	0.143	0.352	0	1
Class 10 Pass	84	0.036	0.187	0	1
College Degree	84	0.202	0.404	0	1
Congress Member	84	0.619	0.489	0	1
BJP Member	84	0.333	0.474	0	1
Landless	84	0.167	0.375	0	1

C Regression Results

Table 5: Regression Results

	Dependent variable:			
	Expected Number of Tokens			
	(1)	(2)	(3)	(4)
Assets	-0.239^{*}	-0.212^{***}	-0.192^{***}	-0.192^{***}
	(0.140)	(0.068)	(0.065)	(0.068)
Supporter		1.091^{***}		
		(0.193)		
Supporter x Assets		0.165		
		(0.216)		
Non-Co-Partisan Supporter			0.928^{***}	0.904^{***}
			(0.195)	(0.194)
Co-Partisan Non-Supporter			-0.045	-0.057
			(0.287)	(0.286)
Co-Partisan Supporter			1.352^{***}	1.298^{***}
			(0.206)	(0.209)
Non-Co-Partisan Supporter x Assets				0.203
				(0.229)
Co-Partisan Non-Supporter x Assets				-0.408
				(0.339)
Co-Partisan Supporter x Assets				-0.075
				(0.237)
σ^2	3.353	0.674	0.620	0.589
Observations	839	839	839	839
Number of GP	84	84	84	84
pD	835.6	497.1	507.1	496.7
DIC	1947.5	1857.5	1872.0	1861.0
Note:				$^{*}\underline{\pi} < 0.1$; $^{**}\underline{\pi} < 0.05$; $^{***}\underline{\pi} < 0.01$

The regressions described above follow the protocol described in section 4.2. Results report estimates from a 3750 posterior simulations from a regression model estimated in a Bayesian framework through Markov Chain Monte Carlo (MCMC) with 3 chains and diffuse priors on all parameters, using the program `JAGS`. Standard deviations of the posteriors on the respective parameters are given in parentheses. Statistical significance in the model is given with respect to the posterior distribution. In particular, let $\widehat{\pi}$ be a vector of values drawn from the posterior distribution of a parameter of interest. Then, we define $\underline{\pi} = 2 * P(\widehat{\pi} < 0)$. The deviance information criterion (DIC) is a measure of fit that is defined as the sum of one-half of the estimated variance of deviance (pD) and the expected value of the deviance. The lower value of DIC is taken to be a better fit, with pD entering as a penalty for overfitting the data.

Table 6: Regression Results (continued)

	Dependent variable:			
	Expected Number of Tokens			
	(5)	(6)	(7)	(8)
Assets	-0.188^{***}	-0.201^{***}	-0.200^{***}	-0.193^{***}
	(0.067)	(0.065)	(0.064)	(0.069)
Non-Co-Ethnic Supporter	1.142^{***}	1.133^{***}		
	(0.206)	(0.212)		
Co-Ethnic Non-Supporter	0.412	0.458		
	(0.348)	(0.369)		
Co-Ethnic Supporter	1.307^{***}	1.331^{***}		
	(0.272)	(0.298)		
Non-Co-Ethnic Supporter x Assets		0.085		
		(0.222)		
Co-Ethnic Non-Supporter x Assets		0.107		
		(0.327)		
Co-Ethnic Supporter x Assets		0.291		
		(0.324)		
Co-Partisan Non-Co-Ethnic Non-Supporter			−0.029	−0.199
			(0.332)	(0.356)
Non-Co-Partisan Co-Ethnic Non-Supporter			0.503	0.390
			(0.353)	(0.375)
Co-Partisan Co-Ethnic Non-Supporter			0.641	0.436
			(0.526)	(0.554)
Non-Co-Partisan Non-Co-Ethnic Supporter			0.975^{***}	0.924^{***}
			(0.220)	(0.218)
Co-Partisan Non-Co-Ethnic Supporter			1.414^{***}	1.303^{***}
			(0.218)	(0.222)
Non-Co-Partisan Co-Ethnic Supporter			1.176^{***}	1.141^{***}
			(0.312)	(0.315)
Co-Partisan Co-Ethnic Supporter			1.670^{***}	1.604^{***}
			(0.353)	(0.350)
Co-Partisan Non-Co-Ethnic Non-Supporter x Assets				−0.452
				(0.394)
Non-Co-Partisan Co-Ethnic Non-Supporter x Assets				−0.069
				(0.392)
Co-Partisan Co-Ethnic Non-Supporter x Assets				−0.646
				(0.587)
Non-Co-Partisan Non-Co-Ethnic Supporter x Assets				0.152
				(0.235)
Co-Partisan Non-Co-Ethnic Supporter x Assets				−0.190
				(0.241)
Non-Co-Partisan Co-Ethnic Supporter x Assets				0.169
				(0.331)
Co-Partisan Co-Ethnic Supporter x Assets				−0.073
				(0.432)
σ^2	0.685	0.677	0.616	0.578
Observations	839	839	839	839
Number of GP	84	84	84	84
pD	520.9	514.7	456.8	535.0
DIC	1877.4	1871.5	1816.5	1901.2
Note:				$^{*}\underline{\pi} < 0.1$; $^{**}\underline{\pi} < 0.05$; $^{***}\underline{\pi} < 0.01$

The regressions described above follow the protocol described in section 4.2. Results report estimates from a 3750 posterior simulations from a regression model estimated in a Bayesian framework through Markov Chain Monte Carlo (MCMC) with 3 chains and diffuse priors on all parameters, using the program `JAGS`. Standard deviations of the posteriors on the respective parameters are given in parentheses. Statistical significance in the model is given with respect to the posterior distribution. In particular, let $\hat{\pi}$ be a vector of values drawn from the posterior distribution of a parameter of interest. Then, we define $\underline{\pi} = 2 * P(\hat{\pi} < 0)$. The deviance information criterion (DIC) is a measure of fit that is defined as the sum of one-half of the estimated variance of deviance (pD) and the expected value of the deviance. The lower value of DIC is taken to be a better fit, with pD entering as a penalty for overfitting the data.

D Robustness

D.1 Partisan and Gender Identity of the Sarpanch

While a large literature finds that local female leaders may display different political and distributional preferences – likely more pro-poor than male leaders (Duflo and Chattopadhyay, 2003), we find little evidence in our data. Recipients of tokens from male sarpanch were about 0.26 standard deviations below the mean wealth according to our asset measure, while recipients from female sarpanch were about 0.14 standard deviations below the mean. The difference was not statistically significant ($p = 0.36$). At first blush, it seems that the sarpanch with Congress Party affiliations are more pro-poor. Recipients of tokens from a Congress sarpanch were approximately 0.28 standard deviations below the mean, while recipients from non-Congress sarpanch were only 0.07 standard deviations below the mean (although the difference is not significant with $p = 0.12$). But this ignores, the "comparative" aspect of our claim. It turns out that Congress supporters are typically much poorer than non-Congress supporters, and conditioning on the relative wealth of co-partisans, Congress and non-Congress sarpanch both demonstrate pro-poor behavior. The average co-partisan of a Congress sarpanch is 0.17 standard deviations below mean GP wealth, and average co-partisan of a non-Congress sarpanch is 0.22 standard deviations *above* the mean GP wealth.

D.2 Ethnic Identity and Class Effects of the Sarpanch

Another competing hypothesis is that certain ethnic identities (measured as caste and religious identities of voters and sarpanch) and class identities (measured by education and land ownership) yield affinities to explain our results. Using a large set of predictors (and noting that variation in sarpanch effects enter in as interactions in the regression), we show that the magnitudes of our key variables of interest (the relative asset wealth of the individual's household and co-partisanship) have similar magnitudes to our core models.

Table 7: Coefficients for Robustness Regression

Dependent Variable: Expected Number of Tokens

Assets	-0.199^{*}	Assets x OBC Sarpanch	0.009
	(0.121)		(0.201)
Supporter	1.285^{***}	Assets x SC Sarpanch	0.005
	(0.271)		(0.245)
Assets x Supporter	0.100	Assets x ST Sarpanch	-0.418
	(0.214)		(0.414)
Rajput Voter	0.033	Assets x Muslim Sarpanch	0.254
	(0.380)		(0.466)
Jat Voter	-0.340	Assets x Meena Sarpanch	0.214
	(0.378)		(0.286)
OBC Voter	-0.006	Assets x Illiterate Sarpanch	0.151
	(0.295)		(0.267)
SC Voter	-0.065	Assets x Landless Sarpanch	-0.076
	(0.328)		(0.253)
ST Voter	0.257	Supporter x OBC Sarpanch	-0.008
	(0.440)		(0.499)
Muslim Voter	0.093	Supporter x SC Sarpanch	-0.066
	(0.525)		(0.746)
Meena Voter	-0.006	Supporter x ST Sarpanch	-1.387
	(0.495)		(1.651)
Literate Voter	-0.572	Supporter x Muslim Sarpanch	-2.422^{***}
	(0.377)		(1.186)
Some Primary School Voter	0.594	Supporter x Meena Sarpanch	-0.277
	(0.381)		(0.651)
Class 5 Pass Voter	0.613	Supporter x Illiterate Sarpanch	0.326
	(0.382)		(0.832)
Class 8 Pass Voter	0.722^{*}	Supporter x Landless Sarpanch	-0.247
	(0.386)		(0.725)
Class 10 Pass Voter	0.582		
	(0.434)		
		Observations	839
		Number of GP	84
		pD	518.5
σ^2	0.583	DIC	1892.1

Note: $^{*}\underline{\pi} < 0.1$; $^{**}\underline{\pi} < 0.05$; $^{***}\underline{\pi} < 0.01$

The regression described above follow the protocol described in section 4.2. The table on the left reports coefficients from voter-side variables, and the table on the right reports (interacted) coefficients by sarpanch characteristics. Results report estimates from a 3750 posterior simulations from a regression model estimated in a Bayesian framework through Markov Chain Monte Carlo (MCMC) with 3 chains and diffuse priors on all parameters, using the program `JAGS`. Standard deviations of the posteriors on the respective parameters are given in parentheses. Statistical significance in the model is given with respect to the posterior distribution. In particular, let $\widehat{\pi}$ be a vector of values drawn from the posterior distribution of a parameter of interest. Then, we define $\underline{\pi} = 2 * P(\widehat{\pi} < 0)$. The deviance information criterion (DIC) is a measure of fit that is defined as the sum of one-half of the estimated variance of deviance (pD) and the expected value of the deviance. The lower value of DIC is taken to be a better fit, with pD entering as a penalty for overfitting the data.

E The Role of Ethnicity

The effect of co-ethnicity on allocation is less pronounced than that of co-partisan supporters. Among non-supporters, a non-co-ethnic receives 0.20 tokens on average, while a co-ethnic receives 0.41 tokens on average. Among supporters, a non-co-ethnic receives 0.60 tokens on average, while a co-ethnic receives 0.68 tokens on average. Once again, in order to disentangle these effects from relative asset wealth, we calculate the impact of co-ethnicity on allocation through our modeling framework.

[FIGURE ABOUT HERE]

Figure 5: Electoral Support and Co-Ethnicity

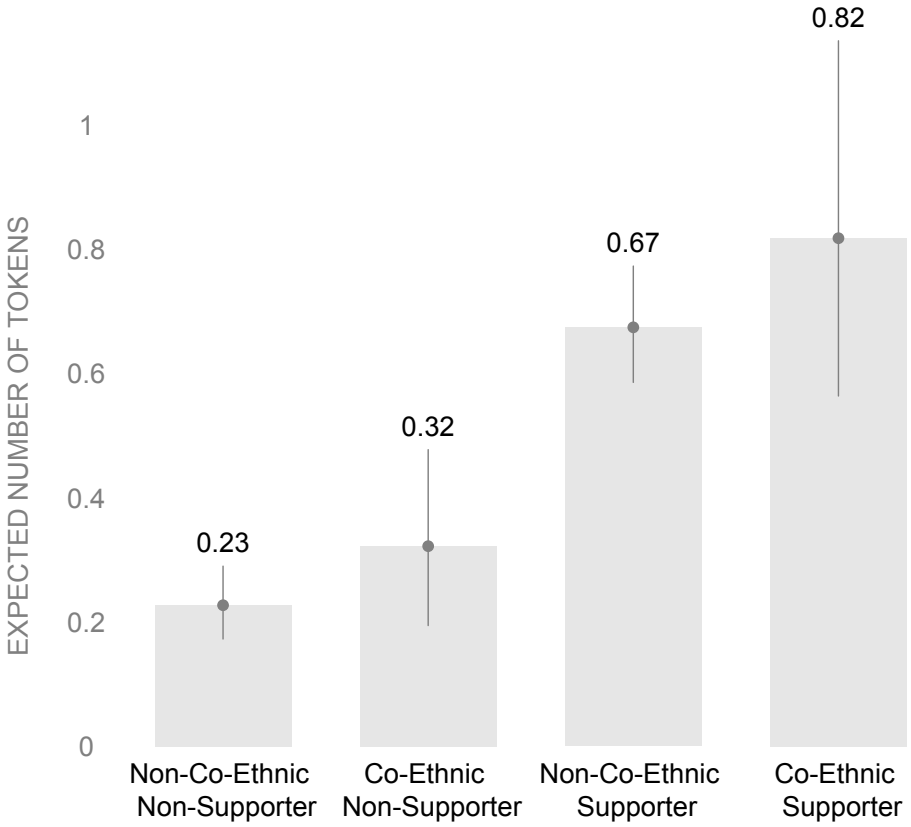

Figure 5 displays predicted average token allocation from a model that controls for relative asset wealth, political support, co-ethnicity, and the interactions between these variables, as displayed in column 6 of appendix C. The impact of co-ethnicity is statistically insignificant for both non-supporters and supporters, providing evidence for the assertion that elected leaders do not typically have preferences to narrowly focus on one ethnic group. This provides further support for the idea that minimum winning coalitions in multi-ethnic societies tend be built around political and partisan identities (**?**).[39]

F Relevant Code

F.1 Code for Item Response Model (Asset Index)

R Code:

```
N <- length(dat2$gpnumber[valid])

gp <- as.numeric(as.factor(dat2$gpnumber[valid]))
n.gp <- max(gp)
y <- cbind(pacca, scooter, bicycle, tv, toilet, fridge, fan, mobile, pump)[valid,]
K <- ncol(y)
item <- NULL; for (i in 1:K) item <- c(item, rep(i, N))
person <- rep(1:N, K)
y <- as.vector(y)
n <- length(y)


itr.data <- list("y", "n", "person", "item", "N", "K")
itr.inits <- function(){
list(a.raw=rnorm(N), b.raw=rnorm(K), sigma.person=runif(1,0,3),
sigma.item=runif(1,0,3), mu.a.raw=rnorm(1), mu.b.raw=rnorm(1))}
itr.par <- c("a", "b", "sigma.person", "sigma.item", "mu.b.raw")
itr.model2p <- jags(data=itr.data, inits=itr.inits, parameters.to.save=itr.par,
model.file="itemresponse2p.txt", n.iter=5000)

assets <- itr.model2p$BUGS$mean$a
```

[39]It is also worth noting that a very complicated model that interacts across political support, co-partisanship, co-ethnicity, and asset wealth, reported in column 8 of appendix C, finds a significant co-partisan effect among non-coethnics but not among co-ethnics. In these models, too, co-ethnicity is not a significant predictor.

JAGS Code – itemresponse2p.txt

```
model{
  for (i in 1:n){
    y[i] ~ dbern(p[i])
    logit(p[i]) <- mu[i]
    mu[i]  <- a[person[i]] - b[item[i]]
}
  for (i in 1:N){
    a.raw[i] ~ dnorm(0, tau.person)
    a[i] <- a.raw[i]
}
  for (i in 1:K){
    b.raw[i] ~ dnorm(mu.b.raw, tau.item)
    b[i] <- b.raw[i]
}
    mu.a.raw ~ dnorm(0,.0001)
    mu.b.raw ~ dnorm(0,.0001)
    tau.item <- pow(sigma.item, -1)
    tau.person <- pow(sigma.person, -1)
    sigma.person ~ dunif(0,100)
    sigma.item ~ dunif(0,100)
}
```

F.2 Code for Regression Model (>1 predictor)

R Code

```
X <- as.matrix(Xadjmat[[i]])  ## GP-mean-adjusted matrix

y <- dat2$tokens_s[valid]

gp <- as.numeric(as.factor(dat2$gpnumber[valid]))
n.gp <- max(gp)
K <- ncol(X)
W <- diag(K)
n <- length(y)

cons <- rep(NA, length(gp))
for (i in 1:length(gp)){
con --s[i] <- 5/sum(gp == gp[i]) }

token.data <- list("y", "X", "W", "n", "gp", "n.gp", "K", "cons")
token.inits <- function(){
list(Tau.B=diag(K), mu.beta=rnorm(K), sigma.epsilon=runif(1,0,100))}
```

```
token.par <- c( "mu.beta", "B", "Sigma.B", "sigma.epsilon")
token.model <- jags(data=token.data, inits=token.inits,
parameters.to.save=token.par, model.file="qpoismultilevel.txt", n.iter=20000)
```

JAGS Code – qpoismultilevel.txt

```
model{
 for (i in 1:n){
     y[i] ~ dpois(lambda[i])
     log(lambda[i]) <- log(cons[i]) + X[i,] %*% B[gp[i],1:K] + epsilon[i]
     epsilon[i] ~ dnorm(0,tau.epsilon)
}
for (j in 1:n.gp){
B[j,1:K] ~ dmnorm(B.hat[j,], Tau.B[,])
B.hat[j,1:K] <- mu.beta[]
}
for (j in 1:K){
mu.beta[j] ~ dnorm(0,.0001)
}
Sigma.B[1:K,1:K] <- inverse(Tau.B[,])
Tau.B[1:K,1:K] ~ dwish(W[,], df)
df <- K+1
tau.epsilon <- pow(sigma.epsilon, -2)
sigma.epsilon ~ dunif(0,100)
}
```

G Anti-Poverty Benefits

Our protocol was primarily geared towards understanding the underlying distributive preferences of sarpanch, which captures the targeting biases that local leaders will apply to the extent that they have discretion over everyday allocation. In weaker state capacity contexts, like in this study, these preferences likely have a relationship to behavior related to benefits with serious institutional constraints. In order to understand the role of personal preferences in distributive outcomes, we designed a "pro-poor cue." In this exercise, we asked the sarpanch to repeat the exercise above, but in a manner as if they were newly allocating below-poverty-line (BPL) benefits, i.e., welfare benefits in the Indian system. We also stipulated that no economic benefits would accrue to recipients of tokens in this exercise. This was done to remove discernible economic incentives for biased targeting. The pro-poor cue, thus, was designed to maximally remove biases from personal preferences in distribution in a weak state capacity scenario, but, as we will see below, such biases still persist in the data. While this may seem like a weak constraint, our results below demonstrate that this "pro-poor" cue has discernible effects on behavior, and observed behavior in this pro-poor cue exercise is quite related to actual distribution of benefits.

G.1 Asset Effects of the "Pro-Poor" Cue

Figure 6 plots the estimated impact of the asset measure on expected number of tokens for the voter, comparing models without (regular tokens) and with (BPL tokens) an explicit cue for targeting the poor . As described above, the asset measure is normalized to have mean 0 and standard deviation 1 inside each GP. The curves and coefficients are to be understood with respect to standard deviations from the mean asset wealth among sampled voters in the GP. For instance, a value of -1 for the asset measure means that

the voter is one standard deviation below the mean asset wealth in the GP. The model predicts an 19% increase in allocation without the pro-poor cue and a 23% increase in allocation.

Figure 6: Expected Number of Tokens vs. Asset Wealth Comparison

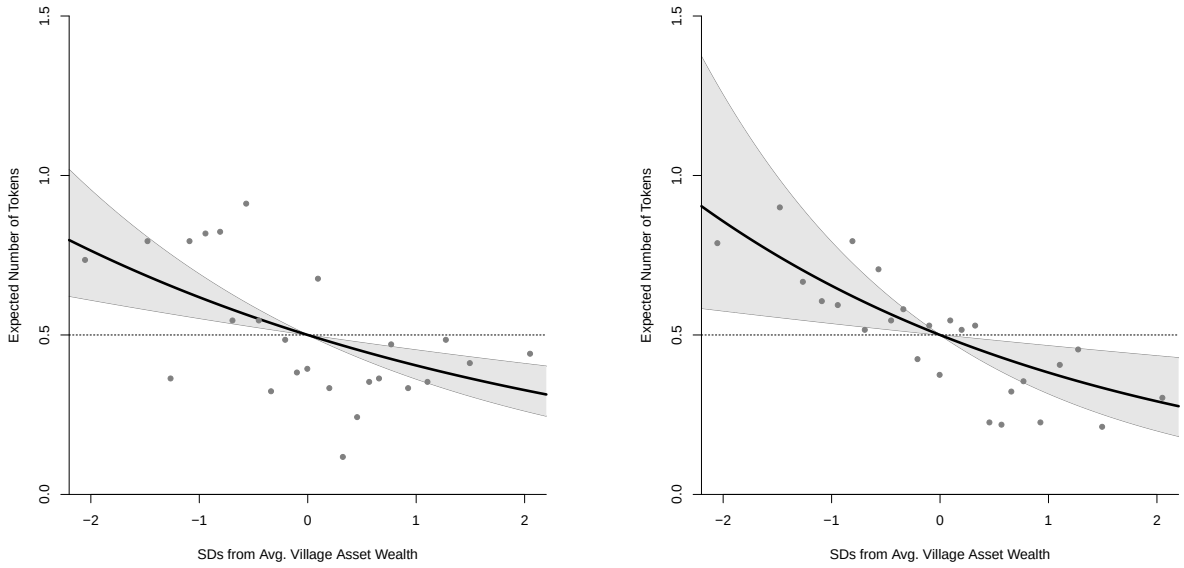

(a) Regular Tokens

(b) BPL Tokens

The gray points in figure 6 are the binned averages of tokens across 25 bins (approximately 34 observations per bin), with cutpoints spaced every 4 percentile points, over the distribution of relative asset wealth. That is, the points display the average number of tokens given to individuals included in a particular bin of relative asset wealth. Consistent with our expectations, the coefficient on asset wealth is significant in both regressions, with the magnitude greater when there is an explicit pro-poor cue. This demonstrates the noticeable targeting of poorest voters in the data regardless of cue, and provides some evidence that sarpanch are further responsive to explicit pro-poor cues, perhaps due to institutional prerogatives.

G.2 Political Biases under the "Pro-Poor" Cue

Figure 7 reports the estimated expected number of tokens for perceived electoral supporters and non-supporters without and with the pro-poor cue and further subdivides the effects by co-partisanship. As in the main text, under regular tokens the sarpanch believes the voter supported him in the last election, then he is willing to give significantly more tokens to that voter as compared to a non-supporter. When we further subdivide the results by whether the voter is a co-partisan of the sarpanch, we see more nuanced results. When there is no pro-poor cue, the sarpanch allocates more towards co-partisans; however, when we introduce a pro-poor cue, this co-partisan effect disappears, suggesting that the impact of sociopolitical ties are impacted by institutional constraints. Even in the case of the supporter effect, while the magnitude is large, the difference is not significant under a pro-poor cue.

G.3 Comparison to Actual Distribution

A natural concern is that our pro-poor cue is too disconnected from, and thus has little relevance for, the actual distribution of anti-poverty benefits. In order to understand the applicability of our measured preferences for actual distribution, we compared our lab behavior to the actual distribution of benefits. In particular, we focus our comparison on whether voters received two benefits, below poverty line (BPL) status and Indira Awas Yojana (IAY) benefits. The first benefit entitles a household to purchase foodstuffs at a reduced price, and the second benefit entitles households to build a home using a government grant. There are only a small number of households that receive IAY benefits, and they must have BPL status to quality for these benefits. As such, the intended recipients of IAY benefits are particularly needy households that should be targeted more heavily. We verified receipt of a BPL card by asking respondents to show interviewers their ration cards. Although IAY benefits were self-reported, new homes

Figure 7: Political Biases Comparison

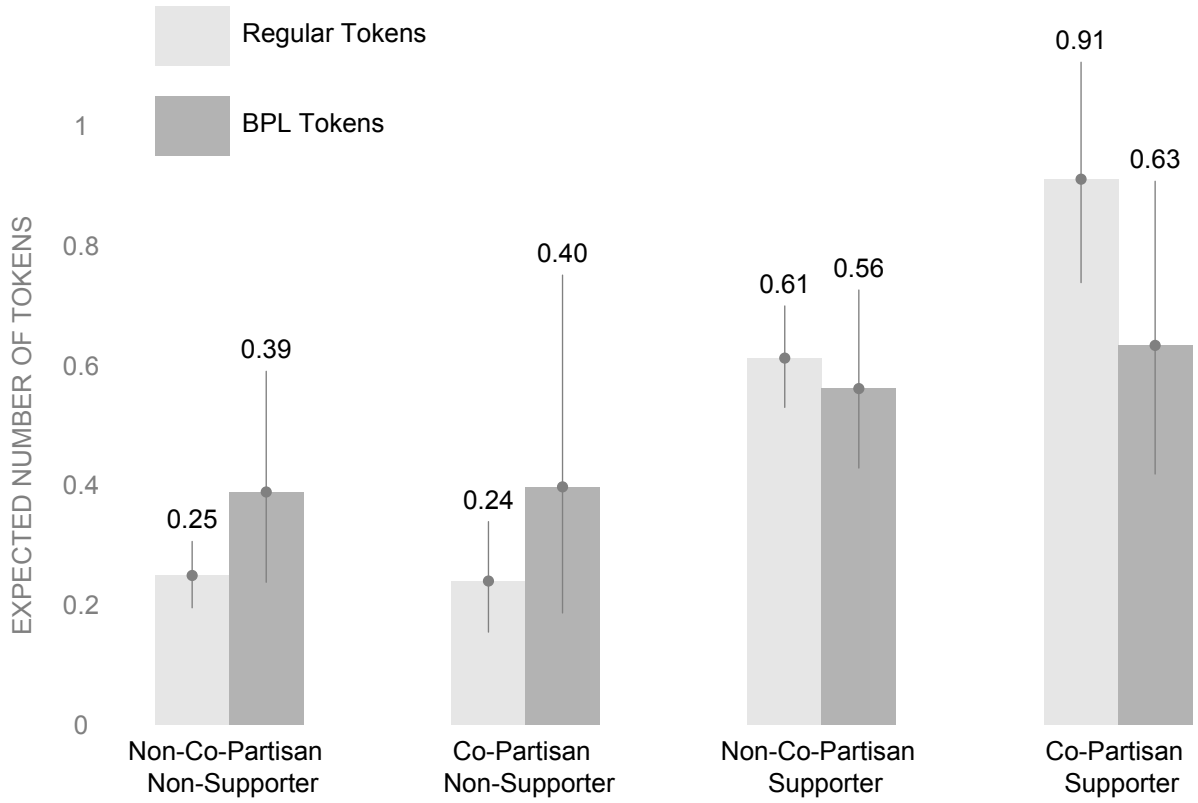

built through this program can be visibly identified as beneficiaries.

Figure 8: Relation Between Lab Measures and Actual Distribution

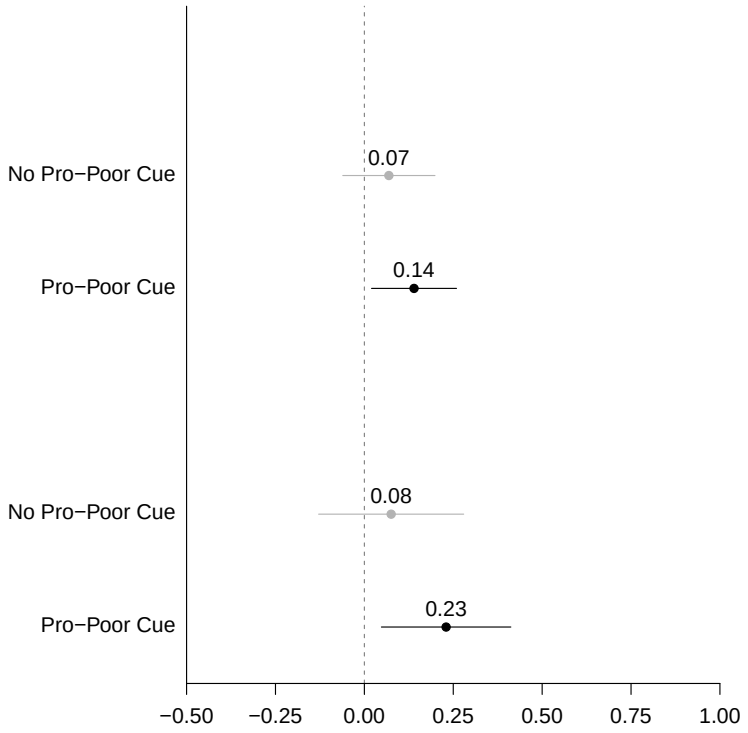

Figure 8 displays the coefficients of an overdispersed Poisson model, relating whether a voter has a benefit (BPL and/or IAY) and whether he or she received a token, using the regression formulation described above. While the coefficients are positive, when there is no pro-poor cue, voters do not receive significantly more tokens if they have a benefit. On the other hand, when there is a pro-poor cue, we find that voters who have benefits are also much more likely to receive a token, and the effects are significant. Consistent with the discussion above, the estimated coefficients are much larger for the IAY benefits than for BPL status. Having BPL status raises the expected number of tokens to a voter by 15% under the pro-poor cue, and receipt of IAY benefits raises the expected number of tokens to a voter by 26% under the pro-poor cue. This provides very strong evidence that our lab setup, when removing disincentives to allocate to the poor (i.e., institutional constraints), can be reasonably associated with actual distribution. Furthermore, we believe our basic setup, without a pro-poor cue, reasonably approximates underlying

distributive preferences where the leaders are not constrained by the pressures of future electoral motivations and have low social or institutional pressures to distribute benefits in a particular manner.