Do Local Leaders Serve the Poor: Identifying the Distributive Preference of Village Politicians in India

Mark Schneider[1], Neelanjan Sircar[2]

[a]Center for the Advanced Study of India, University of Pennsylvania,
[b]Centre for Policy Research, Dharam Marg, Delhi, 110021, Delhi, India

ARTICLE INFO

Keywords:
Lab Experiments
Distributive Preferences
Local Leaders
India

ABSTRACT

The prevailing literature on economic distribution primarily focuses on large constituencies in which a political actors are focused on the delivery of public goods and large cash benefits but do not personally know all of their constituents, and, thus, have low information about the attributes and political preferences of their constituents. But in local democracy, elected leaders are tasked with "everyday assistance" and have relatively high information about their constituents. In this situation, we argue that the pattern of everyday assistance is largely shaped by the distributive preferences of elected leaders. Because voters select leaders who will deliver back to them, we argue that elected leaders will display a preference to distribute to a broad group of supporters. Conditional on political support, elected leaders further conform to a social norm of targeting the most needy supporters. To measure the preferences of elected local leaders, we develop a novel behavioral measure that isolates elected leaders' distributive preferences from electoral compulsions or elite pressures, which we implement in poor villages across the Indian state of Rajasthan. We find that elected leaders prefer to distribute 94% more to political supporters and 17% more to supporters one standard deviation below the mean village wealth. This suggests that local elections are consistent with significant distribution to the poor, albeit with political biases.

1. Introduction

In a poor village about 90 minutes outside of Jaipur in the Indian state of Rajasthan, the elected village council (*panchayat*) president, whom we shall call Mustafa, receives frequent visitors asking for his help obtaining a government benefit or completing a form. He 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."

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 (Besley, Pande and Rao, 2012; Anderson, Francois and Kotwal, 2015; Srinivas, 1959). But how does this characterization explain the behavior described above? The elected leader did not get any obvious electoral benefit from helping these beneficiaries (they were already supporters of their elected village council member), nor were the beneficiaries village elites. The elected leaders simply acted on their preference to help the neediest among their supporters.

India's 73rd amendment devolved substantial authority over the targeting of government benefits and services to millions of elected village governments across the country (Crook and Manor, 1998), giving elected local leaders discretion over distribution while providing voters with access to elected local office holders who they seek for routine forms of assistance (Crook and Manor, 1998; Devarajan, Khemani and Shah, 2009; Kruks-Wisner, 2018) under a system of local democracy. At the same time, these leaders cannot help everyone in the village and must prioritize some voters over others. Proponents of democratic decentralization suggest that empowering elected local leaders, immersed in high-information local social networks ensure that leaders with such discretion are responsive to their constituents (Alderman, 2002; Alatas, Banerjee, Hanna, Olken and Tobias, 2012). On the other hand, local discretion creates a potential for targeting biases along the lines of political support, ethnicity, or wealth (Dunning and Nilekani, 2013; Bardhan and Mookherjee, 2006; Mookherjee, 2015).

In this article, we examine biases in distribution by studying the "distributive preferences" of elected local leaders among their constituents. Elected local leaders' distributive preferences are consequential because elected local leaders spend much of their time responding to citizens' requests for everyday assistance with accessing state benefits or services (Kruks-Wisner, 2018; Chauchard, 2017), which are entirely shaped by their personal decisions over who to prioritize. Local democracies are characterized by high information between leader and constituent so that local leaders can plausibly observe the personal attributes of their constituents (e.g., wealth) and voters can observe the preferences of candidates for local office by reputation or due to shared history of social interaction. With such high information, we argue that voters have an incentive to select leaders who have a predilection to redistribute to their

*Corresponding author: Neelanjan Sircar
Mark.Allan.Schneider@gmail.com (M. Schneider); neelanjan.sircar@cprindia.org (N. Sircar)
ORCID(s):

supporters, Conditional on this supporter bias, we argue that the most socially acceptable form of everyday assistance is to help those most in need – leading to significant redistribution to the poor even without institutional constraints. While this pattern of distribution mimics what has been observed in the literature on electoral targeting among core and swing voters (Calvo and Murillo, 2004), it is not driven directly by future electoral compulsions or the relative efficiency of targeting supporters over non-supporters in low information environments (Dixit and Londregan, 1996) but rather the underlying distributive preferences of local leaders.

To examine elected leaders' distributive preferences, we develop a novel behavioral measure that isolates leaders' preferences from public scrutiny. The behavioral measure we develop uses an innovative "cross-referencing" design, which simultaneously surveys 10 randomly selected voters in a village (whom leaders overwhelmingly reported to know personally) and analyzes which of these voters leaders prefer to target with an economic benefit using a unique lottery measure. Our lottery measure deduces preferences of leaders in a manner so that their decisions cannot be known to anyone else – thereby minimizing any possibility that the economic benefits will impact election outcomes.

According to our behavioral measure, elected leaders prefer to distribute 94% more to perceived political supporters (i.e., those who voted for the leader) than perceived non-supporters, and 17% more to supporters one standard deviation below the mean village wealth (than supporters at mean wealth). Moreover, non-supporters two standard deviations below mean village wealth, receive 50% less than supporters at this level of wealth, which demonstrates evidence of strong political biases toward pro-poor responsiveness. 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 sampled voters' perceptions of helpfulness of their elected village council president and biases in anti-poverty benefits in section 8 and appendix **??**. We interpret our findings to be consistent with the view that procedural democracy selects local leaders who are more responsive to their supporters, and to the poor in particular.

2. Understanding Elected Local Leaders' Distributive Preferences in Rural India

Research on distributive politics focuses on whether politicians, or their intermediaries, prioritize their co-ethnics, core or swing voters, or simply target benefits based on need (Golden and Min, 2013; Alatas, Banerjee, Hanna, Olken, Purnamasari and Wai-Poi, 2013; Hicken, 2011). This research has broadly found evidence that government benefits are targeted with partisan biases; however, scholars have rarely examined how local democracy shapes the responsiveness of local leaders in the form of everyday forms of assistance. The traditional literature on distributive politics focuses on the relative efficiency of targeting core or swing voters for future electoral gains in lower information environments where politicians cannot directly interact with all of their constituents (Dixit and Londregan, 1996; Stokes, Dunning, Nazareno and Brusco, 2013; Dunning and Nilekani, 2013), but such efficiency concerns are not pertinent to local democracy where leaders have significant interaction with and information on voters and constituents. Furthermore, the sorts of benefits described in the literature, either in the form of cash or welfare, are typically constrained by more central actors at the state or national level. Responding to everyday requests for assistance comprises a large part of elected local leaders' time in contexts where access to the state often requires responsiveness from political leaders(Kruks-Wisner, 2018). Moreover, existing research has not considered the implications of local elections on responsiveness.

2.1. Defining Local Democracy

There are two defining features for local democracies: small constituencies and free and fair elections. First, since electoral constituencies are small, local democracy takes place in a setting of high information and dense social ties, where voters know local leaders personally, have accumulated information on the reputations of candidates and their families, and have interacted with these individuals on a regular basis over many years. Moreover, since candidates for local office typically have engaged in forms of everyday assistance prior to the election as brokers or members of elite families (Kruks-Wisner, 2018), voters in this setting have the ability to observe leaders' reputations for responsiveness prior to the election. Similarly, local leaders know their constituents well in this setting. Since local leaders overwhelmingly know voters personally, they can observe the demographic characteristics (e.g., economic need, ethnicity) of their constituents, and target benefits to the voters they wish to target (Alderman, 2002; Alatas et al., 2012). This differs from the large constituencies of parliamentary or state elections, where voters have limited information on candidates or their personal preferences; are unlikely to know their representatives personally, and may not easily discern the demographic criteria upon which distribution is based (Chandra, 2004). Indeed, these are the conditions that oblige politicians to reach voters through intermediaries (Dunning and Nilekani, 2013).

The second feature of local democracy is that local leaders are selected by voters through free and fair elections. Under local democracy, voters have the ability to select leaders whose distributive preferences align with their own free from coercion. This means that winning candidates must cultivate a coalition of electoral support large enough to win the election. Since voters from a particular group have an incentive to select leaders whose distributive preferences include themselves (Ansolabehere and Snyder, Jr., 2006), free and fair local elections make it possible for the local electorate to use the ballot box to select leaders they expect to be responsive to them.

2.2. Political Bias

For local candidates to hold office, they must win a plurality of votes in the election. Consistent with standard political economy models (Ansolabehere and Snyder, Jr., 2006; Downs, 1957), voters have an incentive to select leaders who will be responsive to them (and not exclude them). Similarly, leaders have an incentive to cultivate reputations for responsiveness through past behavior to attract supporters (i.e., those who voted for them) prior to the election. Since elections provide a mechanism for voters to select leaders who display inclusive preferences, we should expect elected leaders to be more responsive to their supporters than their non-supporters irrespective of other demographic traits.

Perceived supporters includes both core voters – who voted for the leader and share closer socio-political ties to the leader – and swing voters – who voted for the leader but do not share close socio-political ties with the leader. In our context, where re-election rates are low,[1] there is reason to expect leaders to place a premium on voters who share their partisan affiliation. Research on core targeting suggests that politicians will target members of their co-partisan networks to mobilize support for partisan elections or maintain their coalitions (Calvo and Murillo, 2013; Diaz-Cayeros, Estevez and Magaloni, 2016). In the context of local democracy, local leaders are likely to favor voters in their co-partisan networks to reward voters who are more closely affiliated with them either because they can be mobilized for higher-level partisan elections or because they simply have an underlying preference for voters who share their partisan preferences. Alternatively, research on swing targeting suggests that parties target pivotal voters because targeted benefits are most likely to impact the vote choice of these voters (Dixit and Londregan, 1996). Despite the concern that decentralization encourages using state resources for vote buying purposes (Khemani, 2010), we do not expect to see biases in preferences toward swing voters (non-co-partisan supporters) over core voters (co-partisan supporters) because the benefit in our study was designed to be private.

H1. Supporter Bias: *The elected leader will display a preference to target her own supporters over non-supporters*

H2. Co-Partisan Supporter(Core Targeting) Bias: *The elected leader will display a preference to target co-partisan supporters over non-co-partisan supporters*

H3. Non-Co-Partisan Supporter (Swing Targeting) Bias: *The elected leader will display a preference to target non-co-partisan supporters over co-partisan supporters*

Research on ethnic politics suggests that politicians will disproportionately target their co-ethnics in low-information environments where it is difficult for voters to discern other attributes that determine targeting outcomes (Chandra, 2004).This is not the case in local democracies where voters can observe the socio-economic status of members of their localities (Alatas et al., 2012). However, leaders may still have a preference to use their discretion to favor supporters who are also members of their ethnic group because co-ethnics share close social ties with the leader in this setting. This may be viewed as another way to characterize core supporters who share socio-political ties with the leader.

H4. Co-Ethnic Bias: *The elected leader will display a preference for targeting co-ethnics irrespective of political support*

H5. Co-Ethnic Supporter Bias: *The elected leader will display a preference for targeting co-ethnic supporters (over non-coethnic supporters)*

2.3. Bias Based on Wealth

In the taxation and welfare spending literature, it is understood that pivotal voters –who select the winner of the election – redistribute wealth from the rich to the less affluent in the middle of the income distribution (Meltzer and Richard, 1981; Alesina and Rodrik, 1994). In the context of local democracy, this means that local elections should eliminate leaders who display narrow preferences for targeting a small number of relatively wealthy voters.[2]

On the other hand, existing research proposes two sets of expectations for pro-poor targeting. First, research on moral economy suggests that a social expectation towards protecting the *poorest* members of society is widespread where a significant portion of the population is living at subsistence levels, as in our context of subsistence villages (Scott, 1976). 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. According to this theory, we should expect pro-poor targeting of the poorest constituents irrespective of whether a voter supported the leader.

The second expectation for pro-poor targeting is that local leaders will display targeting biases toward their poorest supporters. This expectation is consistent with several explanations. First,pro-poor targeting preferences have been linked to incentives among local leaders in poor contexts to cultivate reputations for responsiveness to the needy and avoid negative perceptions of elite favoritism (Auerbach and Thachil, 2022). According to this work, local leaders have a preference to signal that they are helpful to the poor, but limit this to those in their political networks. Second, pro-poor targeting among supporters is also consistent with research on core targeting, which suggests that poor supporters are targeted because their support is the least expensive to maintain (Diaz-Cayeros et al., 2016). Third, this expectation is consistent with work on democratic responsiveness,

[1] Only 19% of village council presidents in our data were serving a second term

[2] This is consistent with research on the targeting of anti-poverty benefits, which shows that siphoning of anti-poverty benefits to the local elite is relatively minimal.(Alatas et al., 2013).

which suggests that elected leaders are responsive to their constituents' preferences, which are more likely to be pro-poor when it comes to low-value subsistence benefits (Arora, George, Rao and Sharan, 2023).

H6. Pro-Poor Targeting (Moral Economy): *The elected leader will display a preference to target the poorest voters in the locality irrespective of the voters' (perceived) political behavior*

H7. Poor Supporter Targeting: *The elected leader will display a preference to target the poorest voters in the locality conditional on (perceived) political supporters*

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 criteria of local democracy and describe the institution of the village council (gram panchayat or GP), and the role that the village council president (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 (Srinivas, 1959), 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, Krishna (2003), 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 (Yadav, 1999). The autonomy of the Indian voter is well-documented (Sircar, 2015). Moreover, interviews with block-level Congress Party and BJP leaders across the state suggest that GP elections in Rajasthan are often competitive. Broadly speaking, intense party competition at local and higher levels, heterogeneity in vote preferences among members of the same ethnic groups (Dunning and Nilekani, 2013; Thachil, 2014), and weak capacity to monitor votes (Schneider, 2019) 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.[3] 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 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 (Parthasarathy, 2017). This has had important consequences for village politics, although recent work and our results suggest that these quotas have not fundamentally impacted distributive outcomes (Dunning and Nilekani, 2013).

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.[4]

Moreover, Rajasthan's two major parties compete for the votes of the poor (Thachil, 2014). 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 (Calvo and Murillo, 2004).

3.3. Local Leaders and Everyday Distribution

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 (Bohlken, 2016; Pattenden, 2011). 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 have 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) (Niehaus and Atanassova, 2013).

[3] Prior to the 73rd amendment, Rajasthan held local elections under different requirements (Narain, 1964).

[4] Dunning and Nilekani (2013) find that Rajasthani voters correctly identified the party of their sarpanch 96 percent of the time; they also find strong partisan biases in targeting.

4. Design and Empirical Strategy

We developed a novel behavioral measure designed to pick up local leaders' targeting preferences (i.e., biases) over real world populations. 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 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. Crucially, unlike survey experiments or other standard surveys, the sarpanch did not have to explicitly reveal preferences 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. Cross-referencing, in contrast, is uniquely applicable to detecting biases in the actual population–particularly in local contexts where social ties between voters and leaders are consequential.

Our design continues a recent tradition of lab-in-the-field experiments (Grossman, 2011) that investigate the impact of ethnicity (Habyarimana, Humphreys, Posner and Weinstein, 2009), partisanship (Fowler and Kam, 2007), and democratic selection (Baldassarri and Grossman, 2011) 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.[5] Specifically, we restricted sampling to sub-districts (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.[6] GPs with below poverty line (BPL) rates of 20% and contested local elections were randomly sampled in blocks that met these criteria.[7] 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.[8] 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.

To ensure that our measure isolated the distributive preferences of the sarpanch from public pressures or other external 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 and neither lottery winners or sarpanch were notified of the winner. 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 were not informed that they won the lottery and may not have realized a lottery took place. In sum, our behavioral measure is designed to minimize if not eliminate short-term electoral incentives.[9] To establish the connection between preferences and distributive outcomes, we demonstrate that voters' perceptions of helpfulness from the leader are correlated with distributive preferences 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) "supporter". 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 and Muslim religion based on voters'

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

[6] Block level, or panchayat samiti, elections are the second tier of local government in India and the lowest level where party symbols are allowed on the ballot; it is also the lowest level where election data was available.

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

[8] 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 (Alatas et al., 2012, 2013).

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

self-reported identities.[10] 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.[11] 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 model can be written as below:

$$y_i \sim Poisson(\lambda_i, \sigma^2) \tag{1}$$

where σ^2 denotes an overdispersion parameter

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

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

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

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 elected local leaders' distributive preferences in the context of local democracies in subsistence-based societies. We begin the section by demonstrating that features of our definition of local democracy 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; and 3) the relative wealth of citizens in the GP is known to the sarpanch. We also establish that our sample include subsistence villages by demonstrating that a significant proportion of voters can be characterized as poor.

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.

5.1. Characterizing the Sample

Our definition of local democracy is built upon the assumption of free and fair elections 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 criteria. 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. Table 1 displays the average for each of these 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).

Each of the items above is a binary variable, and a 2-parameter Rasch model (Gelman and Hill, 2007) was fit using Markov Chain Monte-Carlo (MCMC) using the program JAGS to construct a raw asset score.[12] The raw asset score gives approximately ten different "scores," suggesting reasonably high levels of correlation between owning these assets.

[10] 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.

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

[12] 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.

Table 2
Toilet/Refrigerator (%) by Asset Score

Score	N	Percentile	Toilet(%)	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

Figure 1: Sarpanch Assessments of Wealth vs. Asset Measure

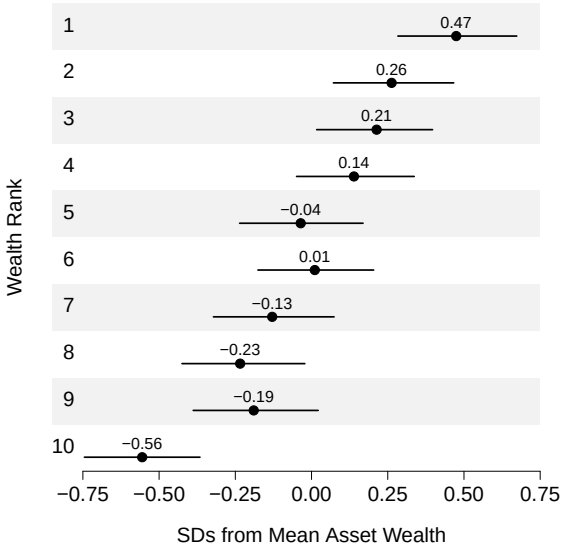

We look at the relationship between these 10 values on our asset index and the percentage of the sample at that asset value owning a refrigerator or a proper toilet (two natural markers of economic development). The results are shown in table 2. In both cases, even the 80^{th} percentile of wealth does 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 level of inequality.

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.

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 (Stokes et al., 2013), 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 (Kruks-Wisner, 2018; Pattenden, 2011), with 31 percent of sarpanch in our data serving as GP ward representatives previously.[13] 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. 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. Given that our electoral setting displays high levels of alternation and competition, we can

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

be reasonably certain that voters are making genuine choices and that their preferences and strategic incentives are reflected in their elected leaders.

6. Characterizing Political and Ethnic 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 closer socio-political tie among supporters due to partisan affinity.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 **??** 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.

Figure 2: Expected Tokens and Electoral Support

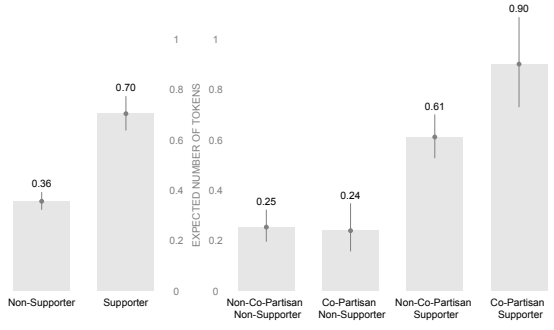

Figure 2 reports the estimated expected number of tokens for perceived electoral supporters and non-supporters, further subdividing the effects by co-partisanship (with 95% posterior/credible intervals). The expected number of tokens for supporters and non-supporters are derived from our core model, 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 **??**, 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 core 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.[14] Taken together, our results strongly confirm the expectation of targeting biases towards political supporters (H1); this effect is particularly pronounced for more closely aligned core co-partisan political supporters (H2). We do not find evidence that elected local leaders favor swing voters (non-copartisan supporters) over core supporters (H3).

Figure 3: Electoral Support and Co-Ethnicity

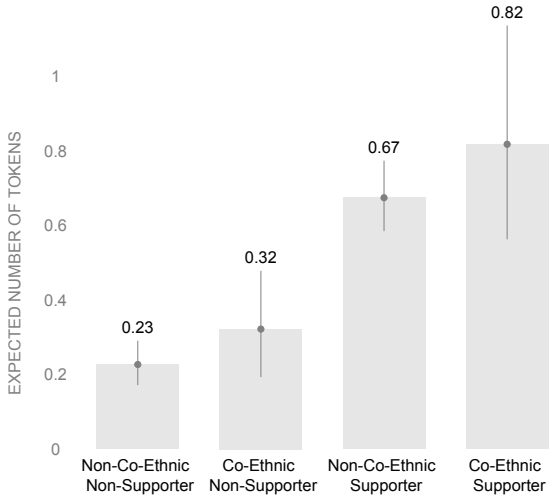

In figure 3, we test the effect of co-ethnic targeting, further subdivided by supporters, but we do not find statistically significant effects (and thus no support for H4 and H5). A full discussion of the results is included in Appendix **??**. This provides some support for the idea that minimum winning coalitions in multi-ethnic societies tend be built around political and partisan identities (Dunning and Nilekani, 2013).

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

7. Targeting of the Poor in Allocation

In this section, we investigate hypotheses on biases in targeting toward the poor among supporters and irrespective of perceived political support. Figure 4 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 results for targeting the poorest citizens among both supporters and non-supporters, with far more pronounced targeting of the poorest among 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 toward their poor supporters (H7), but exclude poor non-supporters even when they are extremely poor– contrary to the expectation of the moral economy hypothesis (H6).

Figure 4: Electoral Support vs. Asset Wealth

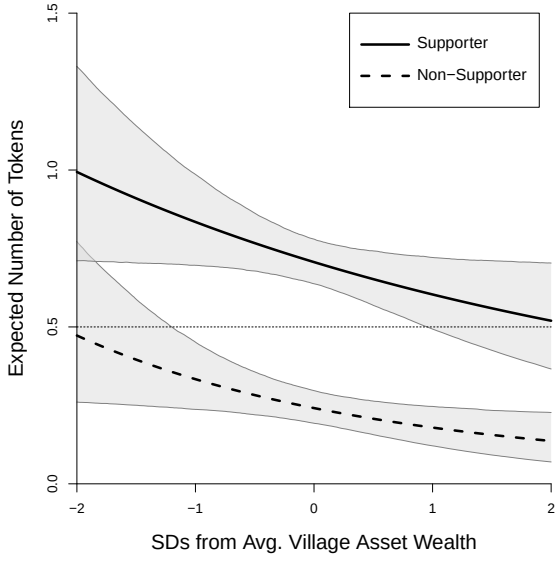

8. Connection to Actual Distribution

We establish the external validity of our lab measure 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.

9. Robustness

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. In appendix **??**, we show these patterns do not hold in our data.[15] 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.

10. 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

[15] Cconditional on political support both Congress and non-Congress sarpanch display similar levels of pro-poor behavior.

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 5: Expected Number of Tokens vs. Asset Wealth

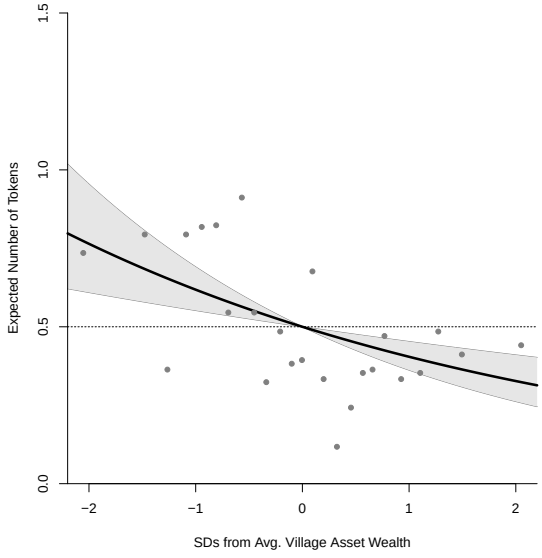

Figure 6 displays the predictions generated from our core model, which controls for relative asset wealth, perceived political support, and the interaction between the two (column 2 in Appendix **??**). In order to generate the figure, we assumed the level of support was at the GP mean and generated 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 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. 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.

To fit closer to the programmatic demands of welfare schemes, we designed a "pro-poor cue." In this exercise, we asked the sarpanch to repeat the tokens exercise, but in a manner as if they were newly allocating below-poverty-line (BPL) benefits, i.e., welfare benefits in the Indian system. Figure **??** plots the estimated impact of the asset measure on expected number of tokens for the voter.

Figure 6: Pro-Poor Cue: Expected Number of Tokens vs. Asset Wealth

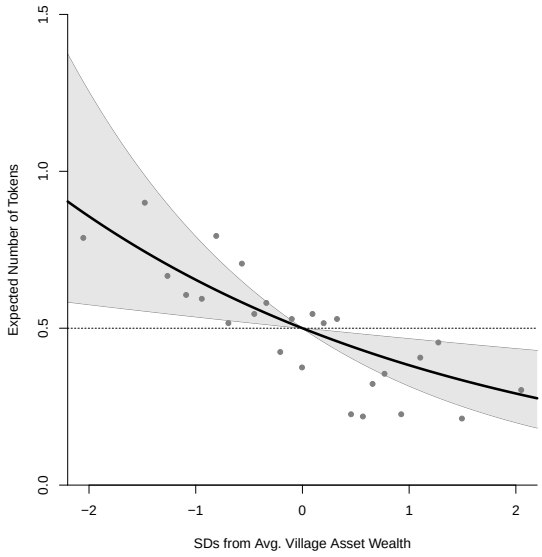

The model predicts an 19% increase in allocation without the pro-poor cue and a 23% increase in allocation for those one standard deviation below mean wealth in the GP. Consistent with our expectations, the coefficient on asset wealth is significant in both regressions, with the magnitude greater when there is an explicit pro-poor (BPL) cue. This demonstrates the noticeable targeting of the 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 surrounding the allocation of anti-poverty benefits.

11. Discussion

This article shows that elected local leaders have distributive preferences that display political biases toward their supporters, and even more so those supporters who share their partisan preferences. We also show that in the context of subsistence-based societies, elected local leaders prioritize the poorest members their coalition of supporters. We do not find evidence that elected local leaders favor their co-ethnics or the relatively poor irrespective of their perceived political support. 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 results are consistent with the view that local politicians hold strong biases toward those they perceived to be their supporters. This results in exclusion of the neediest voters who the leader perceived to be non-supporters. At the same time, local leaders are responsive to a large share of the poorest villagers, which suggests that elected leaders are broadly responsive to the poor in their localities by virtue of their preference for responsiveness to those they view as plausible supporters.

This study advances research on distributive politics in several important ways. First, we establish the importance of understanding the personal preferences of leaders in a way that transcends quid pro quo logic, and demonstrate a strong pattern of democratic responsiveness among elected local leaders. Although existing research suggests that elected local leaders in India are likely to be unresponsive to their constituents, our results suggest reason for optimism when it comes to private benefits and everyday responsiveness.

Second, existing models suggest that politicians should prioritize pivotal swing voters where this can be done efficiently (Dixit and Londregan, 1996; Cox and McCubbins, 1986; Stokes, 2005). However, we find a strong preference for targeting core (co-partisan) supporters due to an in-group targeting logic. Moreover, while existing theory suggests that core targeting can be an effective electoral strategy for coalition maintenance (Diaz-Cayeros et al., 2016) or vote mobilization (Nichter, 2008; Jensenius and Chhibber, 2023), our results are not driven by direct attempts at electoral mobilization given the secrecy of token allocation and weak re-election incentives in our context.

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., Besley, Pande, Rahman and Rao, 2004; Bardhan and Mookherjee, 2006; Galasso and Ravallion, 2005; Olken, 2006; Dunning and Nilekani, 2013). 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 full discretion. This is important because while research shows that anti-poverty policy benefits have often failed to reach the poor, our results suggest that targeting decisions without these constraints can still result in substantial pro-poor targeting.

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 (Sircar and Chauchard, 2018) 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, elected leaders may prefer substantial targeting to the poor conditional on plausible political support. 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.

References

Alatas, V., Banerjee, A., Hanna, R., Olken, B.A., Purnamasari, R., Wai-Poi, M., 2013. Does Elite Capture Matter? Local elites and targeted welfare programs in Indonesia. NBER Working paper .

Alatas, V., Banerjee, A., Hanna, R., Olken, B.A., Tobias, J., 2012. Targeting the Poor: Evidence from a field experiment in Indonesia. American Economic Review 102, 1206–1240.

Alderman, H., 2002. Do local officials know something we don't? decentralization of targeted transfers in albania. Journal of Public Economics 83, 375–404.

Alesina, A., Rodrik, D., 1994. Distributive politics and economic growth. Quarterly Journal of Economics 109, 465–490.

Anderson, S., Francois, P., Kotwal, A., 2015. Clientelism in indian villages. American Economic Review 105, 1780–1816.

Ansolabehere, S., Snyder, Jr., J., 2006. Party control of state government and the distribution of public expenditures. The Scandinavian Journal of Economics 108, 547–569.

Arora, A., George, S., Rao, V., Sharan, M., 2023. The added value of local democracy: Evidence from a natural experiment in india .

Auerbach, A.M., Thachil, T., 2022. Migrants and Machine Politics: How India's Urban Poor Seek Representation and Responsiveness. Princeton University Press.

Baldassarri, D., Grossman, G., 2011. Centralized sanctioning and legitimate authority promote cooperation in humans. Proceedings of the National Academy of Sciences 108, 11023–11027.

Bardhan, P., Mookherjee, D., 2006. Pro-poor targeting and accountability of local governments in west bengal. Journal of Development Economics 79, 303–327.

Besley, T., Pande, R., Rahman, L., Rao, V., 2004. The politics of public good provision: Evidence from indian local governments. Journal of the European Economic Association 2, 416–426.

Besley, T., Pande, R., Rao, V., 2012. Just rewards? local politics and public resource allocation in south india. The World Bank Economic Review 26, 191–216.

Bohlken, A.T., 2016. Democratization from Above: The logic of local democracy in the developing world. Cambridge University Press.

Calvo, E., Murillo, M.V., 2004. Who delivers? partisan clients in the argentine electoral market. American Journal of Political Science 48, 742–757.

Calvo, E., Murillo, M.V., 2013. When parties meet voters: Assessing political linkages through partisan networks and distributive expectations in argentina and chile. Comparative Political Studies 46, 851–882.

Chandra, K., 2004. Why Ethnic Parties Succeed. Cambridge University Press, New York City.

Chauchard, S., 2017. Why Representation Matters: The meaning of ethnic quotas in rural India. Cambridge University Press.

Cox, G.W., McCubbins, M., 1986. Electoral politics as a redistributive game. The Journal of Politics 48, 370–389.

Crook, R.C., Manor, J., 1998. Democracy and Decentralisation in South Asia and West Africa: Participation, accountability and performance. Cambridge University Press.

Devarajan, S., Khemani, S., Shah, S., 2009. The politics of partial decentralization , 79–101.

Diaz-Cayeros, A., Estevez, F., Magaloni, B., 2016. The Political Logic of Poverty Relief: Electoral strategies and social policy in Mexico. Cambridge University Press.

Dixit, A., Londregan, J., 1996. The determinants of success of special interests in redistributive politics. the Journal of Politics 58, 1132–1155.

Downs, A., 1957. An Economic Theory of Democracy. Harper.

Dunning, T., Nilekani, J., 2013. Ethnic quotas and political mobilization: Caste, parties, and distribution in indian village councils. American Political Science Review 107, 35–56.

Fowler, J.H., Kam, C.D., 2007. Beyond the Self: Social Identity, Altruism, and Political Participation. The Journal of Politics 69, 813–827.

Galasso, E., Ravallion, M., 2005. Decentralized targeting of an antipoverty program. Journal of Public economics 89, 705–727.

Gelman, A., Hill, J., 2007. Data Analysis Using Regression and Multilevel Hierarchical Models. Cambridge University Press, New York City.

Golden, M., Min, B., 2013. Distributive politics around the world. Annual Review of Political Science 16.

Grossman, G., 2011. Lab-in-the-field Experiments. Newsletter of the APSA Experimental Section 2, 1–26.

Habyarimana, J., Humphreys, M., Posner, D.N., Weinstein, J.M., 2009. Coethnicity: Diversity and the Dillemmas of Collective Action. Russell Sage Foundation.

Hicken, A., 2011. Clientelism. Annual Review of Political Science 14, 289–310.

Jensenius, F.R., Chhibber, P., 2023. Privileging one's own? voting patterns and politicized spending in india. Comparative Political Studies 56, 503–529.

Khemani, S., 2010. Political capture of decentralization: Vote-buying through grants-financed local jurisdictions. World Bank Policy Research Working Paper .

Krishna, A., 2003. What is happening to caste? a view from some north indian villages. The Journal of Asian Studies 62, 1171–1193.

Kruks-Wisner, G., 2018. Claiming the State: Active citizenship and social welfare in rural India. Cambridge University Press.

Meltzer, A.H., Richard, S.F., 1981. A rational theory of the size of government. Journal of Political Economy 89, 914–927.

Mookherjee, D., 2015. Political decentralization. economics 7, 231–249.

Narain, I., 1964. Democratic decentralization and rural leadership in india: The rajasthan experiment. Asian Survey 4, 1013–1022.

Nichter, S., 2008. Vote buying or turnout buying? machine politics and the secret ballot. American Political Science Review 102, 19–31.

Niehaus, P., Atanassova, A., 2013. Targeting with agents. American Economic Journal: Economic Policy 5, 206–238.

Olken, B.A., 2006. Corruption and the costs of redistribution: Micro-evidence from indonesia. Journal of Public Economics 90, 853–870.

Parthasarathy, R., 2017. Ethnic quotas as term-limits: Caste and distributive politics in south india. Comparative Political Studies 50, 1735–1767.

Pattenden, J., 2011. Gatekeeping as accumulation and domination: Decentralization and class relations in rural south india. Journal of Agrarian Change 11, 164–194.

Schneider, M., 2019. Do local leaders know their voters? a test of guessability in india. Electoral Studies 61, 102049– 102060.

Scott, J., 1976. The Moral Economy of the Peasant: Rebellion and subsistence in Southeast Asia. Yale University Press.

Sircar, N., 2015. Tale of two villages: Kinship networks and preference change in rural india. Center for the Advanced Study of India Working Paper No. 15-02 .

Sircar, N., Chauchard, S., 2018. Courting votes without party workers: The effect of political competition on partisan networks in rural india. Presented at the 2017 Indian Political Economy Workshop, Carnegie Endowment for International Peace .

Srinivas, M., 1959. The dominant caste in rampura. American Anthropologist 61, 1–16.

Stokes, S., 2005. Perverse accountability: A formal model of machine politics with evidence from argentina. American Political Science Review 99, 315–325.

Stokes, S., Dunning, T., Nazareno, M., Brusco, V., 2013. Brokers, Voters, and Clientelism: The puzzle of distributive politics. Cambridge University Press.

Thachil, T., 2014. Elite Parties, Poor Voters: How social services win votes in India. Cambridge University Press.

Yadav, Y., 1999. Electoral politics in the time of change: India's third electoral system, 1989-99. Economic and Political Weekly , 2393–2399.