



WIDER Working Paper 2021/125

Clientelistic politics and pro-poor targeting

Rules versus discretionary budgets

Dilip Mookherjee¹ and Anusha Nath²

July 2021

Abstract: Past research has provided evidence of clientelistic politics in delivery of programme benefits by local governments, or *gram panchayats* (GPs), and manipulation of GP programme budgets by legislators and elected officials at upper tiers in West Bengal, India. Using household panel survey data spanning 1998–2008, we examine the consequences of clientelism for distributive equity. We find that targeting of anti-poverty programmes was progressive both within and across GPs and is explained by greater ‘vote responsiveness’ of poor households to receipt of welfare benefits. Across-GP allocations were more progressive than those of a rule-based formula recommended by the Third State Finance Commission based on GP demographic characteristics. Moreover, alternative formulae for across-GP budgets obtained by varying weights on GP characteristics used in the State Finance Commission formula would have only marginally improved pro-poor targeting. Hence, there is not much scope for improving pro-poor targeting of private benefits by transitioning to formula-based budgeting.

Key words: clientelism, governance, targeting, budgeting

JEL classification: H40, H75, H76, P48

Acknowledgements: This is the revised version of a paper presented to the UNU-WIDER conference on Clientelistic Politics and Development, the IEG-DSE seminar and Calcutta University workshop on political economy. Participants provided many useful comments and questions, especially Steve Wilkinson and Rachel Gisselquist. We are grateful to Pranab Bardhan, Sandip Mitra, and Abhirup Sarkar for past collaborations and numerous discussions on West Bengal panchayats. For financial support, we thank the EDI network and UNU-WIDER, in addition to the MacArthur Foundation, the US National Science Foundation, and the International Growth Center for previous grants funding various rounds of the household surveys. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

An earlier version of this paper was published as [Federal Reserve Bank of Minneapolis Staff Report 624](#) and an EDI working paper.

¹ Department of Economics, Boston University, USA, corresponding author: dilipm@bu.edu; ² Federal Reserve Bank of Minneapolis, USA

This study has been prepared within the UNU-WIDER project [Clientelist politics and economic development – theories, perspectives and new directions](#).

Copyright © UNU-WIDER 2021

UNU-WIDER employs a fair use policy for reasonable reproduction of UNU-WIDER copyrighted content—such as the reproduction of a table or a figure, and/or text not exceeding 400 words—with due acknowledgement of the original source, without requiring explicit permission from the copyright holder.

Information and requests: publications@wider.unu.edu

ISSN 1798-7237 ISBN 978-92-9267-065-8

<https://doi.org/10.35188/UNU-WIDER/2021/065-8>

Typescript prepared by Mary Lukkonen and Siméon Rapin.

United Nations University World Institute for Development Economics Research provides economic analysis and policy advice with the aim of promoting sustainable and equitable development. The Institute began operations in 1985 in Helsinki, Finland, as the first research and training centre of the United Nations University. Today it is a unique blend of think tank, research institute, and UN agency—providing a range of services from policy advice to governments as well as freely available original research.

The Institute is funded through income from an endowment fund with additional contributions to its work programme from Finland, Sweden, and the United Kingdom as well as earmarked contributions for specific projects from a variety of donors.

Katajanokanlaituri 6 B, 00160 Helsinki, Finland

The views expressed in this paper are those of the author(s), and do not necessarily reflect the views of the Institute or the United Nations University, nor the programme/project donors.

1 Introduction

A hallmark of good governance is the successful delivery of welfare benefits to those most in need. This requires suitable institutions and the devolution of decision-making authority to those with suitable information regarding deservingness of different regions and household units within those regions and the incentive to prioritize the needy. An important argument in favour of decentralized governance has been the superiority of local information. On the other hand, there are concerns about lack of accountability or local government officials' perverse incentives (World Bank 2004, Mookherjee 2015). Accountability concerns arise from evidence of political distortions such as elite capture or political clientelism (Mansuri and Rao 2013, Bardhan and Mookherjee 2012). These raise questions regarding the suitable design of delivery mechanisms and the extent to which authority should be delegated to local governments.

We address this question in the context of rural West Bengal, a state in eastern India. We examine whether moving from discretionary allocation of benefits across local government to formula-based allocations would improve the targeting of anti-poverty programmes. Recent research has found increasing evidence of political clientelism in the delivery of benefits by West Bengal local governments.¹ Using household data covering 2004–11, Bardhan et al. (2020) showed votes of household heads responded to receipt of excludable private benefits disbursed by local governments, or *gram panchayats* (GPs), at the bottom-most tier, but not to provision of non-excludable local public goods. Mirroring this, middle tiers of government at the district and block level responded to increased political competition by manipulating lower tier GP's programme budgets for private benefits but not for infrastructure programmes.² In particular, GPs controlled by the same party at both tiers received higher budgets, while those controlled by rival parties experienced severe cuts. Dey and Sen (2016) and Shenoy and Zimmerman (2020) provide evidence of a similar phenomenon during the post-2011 period, during which there was a different ruling party in most areas: winners of close election races raised employment programme scales only in aligned GPs, presumably rewarding GP areas and leaders that helped deliver votes for their party.

Hence, there is clear evidence that discretionary control over benefit distribution is exercised opportunistically in West Bengal, both within and across GPs. We examine the resulting consequences for pro-poor targeting of welfare benefits for which the poorest households are the intended beneficiaries. Using a panel household survey spanning 1998–2008, we evaluate the distribution of benefits in relation to proxy measures of the deservingness of households. We then estimate possible impacts on pro-poor targeting from switching to a formula-bound programmatic system of transfers that would remove scope for local officials' discretion.

Conceptually, the extent of likely improvement from a centralized formula would depend on the informational advantage of local officials relative to information contained in budgeting formulae, in conjunction with the targeting incentives of the former. At one extreme, a centralized formula-based programme could achieve perfect targeting if the state had perfect information about the distribution of socio-economic status (SES) across individual households and could costlessly deliver benefits directly to them. In practice, upper level governments (ULGs) at the national or state level in India have neither such information nor the capacity to transfer benefits directly to households. The level of disaggregation of governments' information regarding economic backwardness is low, being limited to village census records supplemented by household sample surveys that are representative at best at the district level.

¹ See Bardhan et al. (2010, 2015, 2020); Bardhan and Mookherjee (2012); Dey and Sen (2016); Shenoy and Zimmerman (2020).

² The causal effect of changing political competition was identified by comparing changes in the budgets of GPs redistricted in 2007 to more contested state assembly constituencies with changes in the budgets of others not redistricted or those redistricted to less contested constituencies.

Moreover, a large fraction of the rural poor do not have functioning bank accounts. Even the biometric citizen identification Aadhar cards, which have been rolled out nationwide over the past decade, have yet to achieve universal coverage, cannot be integrated with bank accounts, and contain many errors.³

Hence, GPs have traditionally been delegated the task of identifying the SES of households within their jurisdiction, selecting beneficiaries, and delivering various benefit (mostly in-kind) programmes. In such a system the information and incentives of government officials would determine how well benefits are targeted. Middle level governments (MLGs hereafter) at block and district levels are responsible for allocating programme budgets across GPs within their jurisdiction, based on their knowledge of the distribution of poverty and need across GP areas. Owing to weaknesses in the informational and delivery capacity of ULGs, a formula-bound programme would perforce have to devolve within-GP allocation powers to GPs. Hence, the scope of programmatic policy reforms would be restricted to determining GP programme budgets, thereby affecting resource allocations *across* rather than *within* GPs. A recent World Bank programme for strengthening local governance involving 1000 GPs in West Bengal was based on direct grants to GPs determined by transparent formulae; this programme constitutes an example of such an approach.⁴

Imperfections in the information on which formula-bound GP budgets would be based would inevitably cause targeting errors. There would be errors both of inclusion (prosperous villages with few poor households that are misclassified as poor villages would end up receiving large budgets) and of exclusion (poor villages misclassified as prosperous would fail to qualify for programme grants). It is a priori unclear whether the formula-bound programme would generate better pro-poor targeting compared with that of the existing discretionary system. The net result would depend on (a) the superiority of ‘local soft’ information available to MLGs relative to the ‘hard’ information available to ULGs, and (b) incentives generated by political clientelism for MLGs to target benefits towards truly poor areas.

As the previous literature indicates, the latter is likely to depend in turn on whether elections in poorer regions are less contested or feature different patterns of political alignment between MLGs and ULGs. For instance, improvements in pro-poor targeting would result from a transition to formula-based budgets if elections in poorer areas were less contested, or resulted in a lack of vertical alignment of political control. Also relevant is the relative responsiveness of the votes of the poor and non-poor to benefit delivery. Some have argued that clientelism creates a bias in favour of distributing benefits towards the poor, since their votes are cheaper to ‘buy’. Others have argued that the votes of the poor are determined more by ‘identity’ considerations and less by actual governance performance, while non-poor and better educated voters are more prone to swing based on benefits received. It is therefore hard to predict a priori whether political opportunism for MLGs in a clientelistic setting would translate into a pro- or anti-poor bias.

Hence, the effect of moving to formula-based GP budgets is an empirical question, which we address in this paper. It is based on actual targeting patterns estimated on the basis of household panel surveys in a sample of 59 GPs covering 2,400 households over a 10- year period from 1998 to 2008. Besides declarations of benefits received by household heads, the surveys include household demographic, asset, and income information which allow us to classify households into categories of ultra-poor, moderately poor, and marginally poor. Our definition of these categories is based on whether three, two, or one of the following criteria are satisfied by any given household: if it is landless (owns no agricultural land), if the head has no education (zero years of schooling), and if the household belongs to a scheduled caste or tribe (SC/ST). Apart from capturing the multidimensionality of poverty, this method accurately measures the depth of poverty: the distribution of annual reported income, the value of land owned, or

³ For a recent discussion of these problems, see Dreze et al. (2020).

⁴ See <https://projects.worldbank.org/en/projects-operations/project-detail/P159427>.

of the reported value of the dwelling of successive classes are ordered by first order stochastic dominance.

The within-GP targeting pattern (which conditions on the budget the GP receives from MLGs) for anti-poverty programmes in our data reveals a clear bias in favour of poor households. Poorer households were more likely to receive either an employment benefit or any of the other anti-poverty benefits (low income housing and sanitation, below-poverty-line (BPL) cards entitling holders to subsidized grains and fuel, subsidized loans). On the other hand, the allocation of subsidized farm inputs, an agricultural development programme rather than a welfare programme, was biased in favour of the non-poor, who owned more agricultural land. Hence, the targeting of within-GP allocations appears to be in the ‘right’ direction, varying with the extent to which the corresponding benefit would be likely to benefit the recipient.

For all programmes, increased GP programme budgets (proxied by per household benefits distributed in the GP) resulted in near-uniform increases in allocations to all households irrespective of poverty status. The targeting patterns are robust to varying specifications, including functional form (linear versus Poisson), controls for village characteristics or inclusion of year, and GP or district fixed effects. The results for the linear specification are also unchanged in an instrumental variable (IV) regression in which we instrument for the per household GP benefit by the corresponding per household GP benefit in all other GPs in the same district in that year (à la Levitt-Snyder (1997)), while controlling for district fixed effects. The fact that, conditional on GP budgets, the targeting patterns are unaffected by replacing GP fixed effects with district fixed effects is consistent with the hypothesis that GP budgets represent the primary channel by which MLGs’ actions affect targeting. And the robustness of targeting patterns with respect to the potential endogeneity of GP budgets indicates that the estimated impact of GP budgets can be interpreted causally. One can then use the estimates to predict the targeting impacts of changing the way GP budgets are determined.

Next, we examine how observed GP budgets varied across GPs. The budgets were also progressive: GPs with a higher proportion of ultra or moderately poor households were allocated higher budgets. This indicates that the political incentives of elected officials were aligned in favour of delivering welfare benefits to the poor. To explain this result, we rely on the model of clientelistic allocation in Bardhan et al. (2020). Within GPs, officials of both the incumbent and the challenger party are motivated to deliver benefits to those who are most likely to respond with their votes in the subsequent election. Using data on political support expressed by household heads and extending the method used in Bardhan et al. (2020), we find that the political support of poorer households was more responsive to benefits than that of non-poor households. This is consistent with the common wisdom regarding clientelism (Stokes 2005; Stokes et al. 2013), as well as with the observed intra-GP targeting patterns. Regarding across-GP allocation decisions made by MLGs, the model predicts that the progressivity of these allocations depends on how electoral competition and vertical alignment (of political control between GPs and upper tiers) vary across regions with different poverty rates. We do not find evidence of a significant correlation between either competitiveness or alignment and the poverty rates across GP areas. Hence, we infer that the progressivity of a cross-GP budget allocations was driven primarily by the higher vote responsiveness of poor households.

The observed across-GP allocations turn out to be more progressive than those of the formula for the allocation of fiscal grants to GPs recommended by the Third State Finance Commission (SFC) of West Bengal (State Finance Commission 2008). The SFC formula incorporates seven village characteristics from the census and some household surveys: population size, SC/ST proportion, proportion of female illiterates, a food insecurity index, proportion of agricultural workers, village infrastructure, and population density. Across GPs, SFC-recommended grants turned out to be less positively correlated (compared with actual allocations) with the village proportion of (at least moderately) poor households.

This suggests that transitioning to GP budgets based on the SFC formula would have resulted in less pro-poor targeting. To verify this, we use the estimated within-GP targeting pattern to predict how the expected number of benefits would have changed for any given household in the sample. We aggregate this to estimate the state-wide share of benefits accruing to different poverty groups. The exact results depend on some details regarding the specific method of budget reallocation and the estimation procedure. Budgets could be reallocated across GPs within each district, or across all GPs in the state. Budget balancing within the GP could be achieved by proportionally scaling predicted changes in within-GP allocations (*proportional scaling*). Alternatively, the allocations for poor groups could be predicted on the basis of the estimated within-GP targeting patterns, with the non-poor picking up the slack treated as residual claimants (*residual scaling*). The results turn out to be qualitatively similar across these different approaches. With proportional scaling, the resulting impacts on targeting are negligible, while in the case of residual scaling, poor groups end up with fewer expected welfare benefits under a system based on the SFC-formula.

Finally, we examine whether variations on the weights used in the SFC formula could have improved targeting beyond the observed allocations. For employment benefits and proportional scaling, we estimate that the share of the ultra-poor could at best have been increased from 18.4% to 19.2%, and that of the moderately poor from 35.9% to 36.3%. The changes in shares of non-employment anti-poverty benefits are of a similar order of magnitude.

In summary, the scope for improving pro-poor targeting by switching to formula-based GP budgets is limited at best, as long as the formula is based on indicators used by the West Bengal SFC. This owes partly to a degree of pro-poor accountability in West Bengal's local government and partly to local official's superior information about the distribution of need compared with measures utilized by the SFC. For formula-based budgeting to achieve further improvements, it would have to rely on better information regarding ownership of key assets of land and education at the household level.

Related to this point, it is important to note that we are not addressing the broader question of the overall anti-poverty effects of clientelism. Our analysis concerns only discretionary budgeting's effects on the pro-poor targeting of private benefits within a clientelistic regime. By focusing on pro-poor targeting, or vertical equity, we are ignoring horizontal equity considerations, that is, the allocation of benefits between different poor groups, either between or within villages. Indeed, by showing how this allocation seems to have been manipulated for political purposes, the existing literature has already demonstrated patterns of unfairness. Another important dimension we have ignored in this paper is insurance with respect to uncertain shocks to household or village needs. Moreover, as often alleged, clientelism could cause under supply of local public goods essential for long-term reduction of poverty and undermine political competition, transparency, state legitimacy, and the rule of law.

Our work relates to some recent literature studying the implications of moving from discretionary to formula-based programme grants in Brazil (Azulai 2017; Finan and Mazzocco 2020) and in drought relief declarations in south Indian states (Tarquinio 2020). The results of these papers indicate more significant targeting benefits than we find in West Bengal, thus suggesting that the expected results of transitioning to formula-based budgets are context-specific. On the other hand, our main result concerning pro-poor targeting of political clientelism echoes broader arguments made by Holland (2017) concerning redistributive benefits of 'forbearance,' or the lack of enforcement of property laws against specific citizens for political reasons that occurs in many Latin American countries. In similar vein, Alatas et al. (2012) show that the benefits of targeting that could be achieved by formulae based on household based proxies of poverty in Indonesia would be only marginally superior to those achieved by local community groups. Their focus, however, is on within-village targeting, whereas our paper deals with the implications of alternative ways of deciding across-village allocations.

Section 2 provides details of the setting and describes the data. Section 3 then presents evidence on within-GP targeting patterns, and Section 4 on across-GP targeting and how it would be impacted by switching to formula-based GP budgets. Finally Section 5 concludes with some qualifications and directions for future research.

2 Context, data and descriptive statistics

Each Indian state has a hierarchy of local governments, or panchayats, in rural areas. The panchayats that deliver diverse in-kind benefits to households living in villages. Most of these programmes are financed by central and state governments. District-level governments, called *zilla parishads* (ZPs), allocate funds to middle-tier governments at the ‘block’ level, which comprises an elected body, *panchayat samiti* (PS), and appointed bureaucrats in the Block Development Offices. The middle tier then allocates funds to bottom-tier gram panchayats within their block, which in turn distribute benefits across and within villages in their jurisdiction. Each GP oversees 10–15 villages, and each village in turn includes an average of 300 households. GPs also administer rural infrastructure projects, in which they employ the local population. Despite being subject to oversight both below (from village assembly meetings) and above (from middle level governments that approve projects and expenditures and audit accounts), GPs exercise considerable discretion in their allocation and project decisions. MLG officials face considerably less scrutiny, as there are no stated criteria for horizontal allocation of funds or project approvals across GPs reporting to them. The near-complete absence of any transparency in across-GP allocations confers substantial discretionary authority to MLG officials.

Our data on programme benefits received by households come from two rounds of longitudinal household surveys carried out in 2004 and 2011. The survey includes 89 villages in 57 GPs spread through all 18 agricultural districts of West Bengal and has been used in previous papers (Bardhan et al. 2020). There are over 2,400 households in the sample, amounting to approximately 25 households per village. Households within a village were selected by sampling randomly in different land strata. Table 1 provides a summary of the demographic characteristics of these households. Over half own no agricultural land, nearly one in three belong to a Scheduled Caste (SC) or Scheduled Tribe (ST), and one-third of household heads have no education. Agricultural cultivation is the primary occupation among the landed, while the landless are primarily workers relying on labour earnings.

Table 1: Summary statistics—demographics

Agri land owned (acres)	No. of households	Characteristics of head of households				
		Avg. age	% males	Years of schooling	% SC/ST	% in agriculture
Landless	1214	45	88	6.6	37.4	26
0-1.5	658	48	88	7.8	38.9	65
1.5-2.5	95	56	92	10.8	22.4	82
2.5-5	258	58	93	11.1	27.1	72
5-10	148	60	89	12.5	26.1	66
> 10	29	59	100	13.9	30.9	72
All	2402	49	89	8.0	35.4	47

Note: this table provides demographic characteristics of the head of households (who were the main respondents to the survey) in 2004. % in agriculture refers to percentage of household heads whose primary occupation is agriculture.

Source: authors’ calculations based on survey data.

The period of our study is 1998–2008, spanning two consecutive elected local governments. Since our focus is on political clientelism, we focus attention on excludable private benefit programmes distributed by the GP. The most important of these are programmes offering *employment* in local infrastructure

construction, such as Jawahar Rozgar Yojana (JRY), the National Rural Employment Guarantee Act (NREGA), and the Members of Parliament Local Area Development Scheme (MPLADS). Mostly carried out in the lean agricultural season between March and July, they provide employed households the opportunity to earn a wage set statutorily above the average market wage rate. In years of low rainfall, when private employment opportunities and wages are low, they constitute an important source of income protection for poor households. Other anti-poverty programmes earmarked exclusively for low SES households include subsidized loans, housing/toilet construction subsidies, and Below Poverty Line (BPL) cards entitling holders to subsidized food grains and other household items. GPs also help distribute agricultural minikits that contain subsidized seeds, fertilizers, and pesticides, but their circulation is an agricultural development programme rather than an anti-poverty programme. We will see that the targeting patterns for these farm subsidies differ substantially from all the other programmes. Table 2 shows the percentage of households receiving at least one benefit in the two panchayat terms.

Table 2: Percentage of households receiving at least one benefit

	1998–2003	2004–08
Employment	6.77	24.22
Non-employment anti-poverty	35.12	22.33
Farm subsidy	0.97	7.21

Source: authors' calculations based on survey data.

Our data include different dimensions of low socio-economic status (SES): whether a household belongs to an SC or ST, whether it is landless, and whether the head of household has no education. Depending on whether all, two, or none of these conditions apply, we classify each household as belonging to one of four groups: ultra-poor, moderately poor, marginally poor, and non-poor. These categories measure the number of dimensions in which a household is poor. They also correspond to more standard measures used to measure the depth of poverty. Table 3 shows regressions of annual reported income, acres of agricultural land owned, and the value of the principal dwelling of the household on dummies for these different poverty classes, after controlling for village fixed effects. Compared with the non-poor, households in any of the poverty groups earn significantly lower incomes, own less land, and own less valuable homes on average.

Table 3: Income/wealth variations across poverty groups

	Reported income (rupees lakhs) (1)	Agricultural land (acres) (2)	Value of house (rupees lakhs) (3)
Ultra poor	-0.477*** (0.080)	-2.897*** (0.246)	-1.263*** (0.152)
Moderately poor	-0.397*** (0.052)	-2.519*** (0.201)	-0.989*** (0.129)
Marginally poor	-0.263*** (0.051)	-1.775*** (0.197)	-0.565*** (0.111)
Observations	2256	2256	1691
Adjusted R^2	0.097	0.302	0.238
Mean dependent variable	0.371	1.241	0.848
SD dependent variable	0.759	2.388	1.214
Village fixed effects	YES	YES	YES

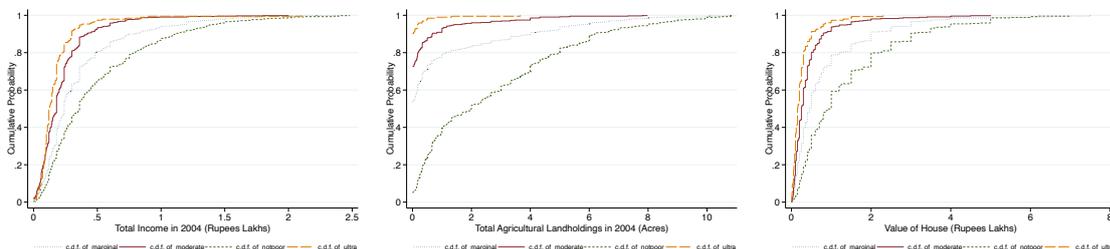
Note: this table examines the relationship between our poverty measures and reported income/wealth in the 2004 household survey. The precise reported measure used is indicated at the top of each column. All specifications include village fixed effects. Robust standard errors are in parentheses, clustered at GP level.

Source: authors' calculations based on survey data.

Figure 1 depicts the distribution of income and wealth by poverty groups. For each of the measures of socio-economic status, the distributions across poverty groups are ordered by first order stochastic dom-

inance. This supports our interpretation of the poverty groups: ultra and moderately poor households have a higher depth of poverty compared with marginally poor groups. Hence, we will use these as definitions of poverty for the remainder of the paper.

Figure 1: Distribution of income and wealth by poverty groups



Source: authors' calculations based on survey data.

Table 4: Poverty groups—demographic share and share of reported benefits

Group	Demographic share	Share of reported benefits		
		Employment	Anti-poverty	Farm subsidy
Ultra poor	8.53	18.38	12.37	1.59
Moderately poor	27.56	35.91	31.51	12.70
Marginally poor	38.33	30.64	33.71	42.33
Non-poor	25.58	15.07	22.41	43.39

Source: authors' calculations based on survey data.

Table 4 provides the demographic shares and the share of benefits for each group. In our sample, the proportions of households that were ultra-poor, moderately poor, or marginally poor were 8.5%, 27.6%, and 38.3%, respectively. The shares of employment and non-employment anti-poverty benefits for ultra and moderately poor households were higher than their demographic shares. However, the opposite is the case for farm subsidies.

3 Within-GP targeting

In this section we examine targeting patterns within GPs. We start with the following Poisson count regression specification for each type of benefit k :

$$b_{ikpgt} = \exp(\beta_k * B_{kgt} + \sum_p \delta_{pk} d_{ip} + \sum_l \gamma_{kl} * X_{v(i)l} + \eta_{kg} + \alpha_{kt}),$$

where

- b_{ikpgt} : number of benefits of type k received by household i belonging to group p in GP g in year t ;
- B_{kgt} : GP g budget estimate (per HH number of benefits of type k in g sample) in year t ;
- d_{ip} : dummy for poverty group p of i ;
- $X_{v(i)l}$: i 's village $v(i)$ characteristic l (population, distribution);
- η_{kg} and α_{kt} : GP/district and year dummies, respectively.

Table 5 presents the results for each type of programme, along with a corresponding linear (OLS) specification. The coefficients of the Poisson regression (expected increase in log benefits associated with a

unit increase in the regressor) have a different interpretation from that in the OLS regression (expected increase in benefits associated with a unit change in regressor); thus, the two are not directly comparable. The regressors include the household's poverty status (with the non-poor serving as the default group); the GP budget (proxied by the number of benefits per household in the GP sample for that year); and a number of characteristics of the village in which the household resides, includes size (number of households in the village) and the proportion of households in each poverty group in the village. 'Villages' are defined by the census; they correspond to sub-units within the GP jurisdiction. Each GP jurisdiction includes between 8 and 15 villages. Controls include either district or GP fixed effects and year dummies. Standard errors are clustered at the GP level. We show results for three programmes: employment programmes, benefits aggregated across all other anti-poverty programmes, and subsidized farm inputs.

Note first that the estimated coefficients of household poverty status change little across the GP and district fixed effect versions of the Poisson regression (first two columns for each programme). Moreover, the Poisson and OLS linear regression versions with district fixed effects (second and third columns in each set) yield qualitatively similar results. Time-varying across-GP targeting differences are driven by corresponding temporal variations in their respective programme budgets, whereas the other non-time-varying regressors capture within-GP targeting patterns. In the specification used in this table, the underlying assumption is that the within- and across-GP targeting patterns are orthogonal; we relax this assumption later. Table 5 shows that the within-GP targeting of anti-poverty programme benefits is progressive: poorer households receive more benefits. The pattern is exactly the opposite for subsidized farm inputs. The distribution patterns therefore tend to allocate each type of programme by prioritizing those who would benefit the most from them.

Table 5: Within-GP targeting poisson regression—GP vs district fixed effects

	Dependent variable: number of benefits received								
	Employment benefit			Non-employment anti-poverty programmes			Subsidized farm inputs		
	Poisson (1)	Poisson (2)	OLS (3)	Poisson (4)	Poisson (5)	OLS (6)	Poisson (7)	Poisson (8)	OLS (9)
GP benefits k	0.162*** (0.028)	0.142*** (0.019)	0.011*** (0.002)	0.124*** (0.021)	0.109*** (0.014)	0.010*** (0.002)	0.137** (0.055)	0.112*** (0.034)	0.009*** (0.002)
Ultra poor	1.484*** (0.197)	1.492*** (0.199)	0.057*** (0.009)	0.655*** (0.121)	0.658*** (0.121)	0.046*** (0.010)	-2.119*** (0.718)	-2.141*** (0.717)	-0.011*** (0.004)
Moderately poor	1.053*** (0.170)	1.071*** (0.174)	0.033*** (0.007)	0.532*** (0.096)	0.536*** (0.096)	0.034*** (0.007)	-1.245*** (0.417)	-1.258*** (0.417)	-0.009** (0.004)
Marginally poor	0.520*** (0.142)	0.531*** (0.144)	0.014*** (0.004)	0.219*** (0.071)	0.221*** (0.071)	0.014*** (0.004)	-0.406** (0.177)	-0.413** (0.176)	-0.004* (0.003)
Number HH in village	0.002*** (0.000)	-0.000 (0.001)	-0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.003*** (0.001)	-0.001 (0.001)	-0.000 (0.000)
Proportion of ultra poor	-1.210 (1.307)	- (0.972)	-0.087*** (0.033)	0.534 (1.117)	-1.150 (1.223)	-0.086 (0.060)	2.522 (1.970)	-3.215 (2.328)	-0.022 (0.013)
Proportion of moderately poor	-0.444 (0.754)	-0.745 (0.540)	-0.022 (0.018)	-0.139 (0.739)	-0.613 (0.644)	-0.044 (0.036)	1.422 (1.117)	1.042 (1.121)	0.006 (0.009)
Proportion of marginally poor	-0.963* (0.502)	-0.568 (0.453)	-0.023 (0.016)	-0.032 (0.410)	-0.436 (0.429)	-0.022 (0.025)	-0.995 (1.270)	-1.268 (1.033)	-0.002 (0.007)
Observations	25025	25025	25025	25025	25025	25025	25025	25025	25025
Mean dependent variable	0.033	0.033	0.033	0.064	0.064	0.064	0.008	0.008	0.008
SD dependent variable	0.194	0.194	0.194	0.262	0.262	0.262	0.087	0.087	0.087
Year FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
GP FE	YES	NO	NO	YES	NO	NO	YES	NO	NO
District FE	NO	YES	YES	NO	YES	YES	NO	YES	YES

Note: observations are at the household-year level, 1998–2008. Dependent variable in columns (1)–(3) is the number of employment benefits received by the household in year t . For columns (4)–(6), the dependent variable is the number of non-employment anti-poverty benefits and for columns (7)–(9), it is the number of subsidized farm inputs. For each type of benefit, the first two columns report the results from a poisson regression while the third column reports estimates from an OLS regression. Regression coefficients in Poisson regressions can be interpreted as the change in log of expected number of benefits associated with a unit change in each regressor. Each specification includes year fixed effects. Whether the specification includes GP fixed effects or district fixed effects is indicated at the bottom of the table. Robust standard errors are in parentheses, clustered at GP level.

Source: authors' calculations based on survey data.

A higher proportion of poor households residing in the village generally tends to lower benefits received by a representative household, though these estimates tend to lack statistical significance. These negative effects are more pronounced in the version with district rather than GP fixed effects. Since the regression conditions on the GP programme budget, it is likely to arise mechanically from the GP budget constraint, combined with the progressive pattern of targeting within the GP. Since poorer households are more likely to receive benefits than non-poor ones, a GP with a larger fraction of poor households and with a given budget will have fewer resources available to distribute to non-poor households. It should not necessarily be interpreted as a form of regressivity in the across-GP targeting pattern, which will be manifested in the allocation of budgets across GPs (which will be examined in the next Section).

In order to simulate the within-GP effects of changes in GP budgets, it is important to obtain an unbiased estimate of the causal impact of changing these budgets. The preceding regression estimate of the GP budget effect is subject to various possible biases. First, the GP budget is not directly observed and is measured with error by its proxy, the per household benefit in the sample. The resulting measurement error could result in a downward (attenuation) bias. Second, the per capita benefit measure in the GP includes each household in the sample, thereby mechanically inducing a positive bias. Third, GP budget allocations may not be exogenous, as they could be driven by the political considerations of officials in upper level governments. These unobserved political considerations (competitive stakes, political alignment, responsiveness of votes to programme benefits) could possibly vary across GPs and may be systematically correlated with the regressors, thereby biasing the coefficient estimates in Table 5.

To deal with these concerns, Table 6 presents an instrumental variable (IV) regression for the linear specification, in which we instrument for the GP budget by average per household programme scale in all **other** GPs in the district. This approach is similar to the instrument used in Levitt and Snyder (1997) and Bardhan et al. (2020). This reflects factors less likely to be correlated with GP-specific unobserved political attributes, such as the scale of the programme budget at the district level (determined by financing constraints at the district level) and political attributes of other GPs in the district with which the GP in question is competing for funds. As explained in some detail in Levitt and Snyder (1997) and Bardhan et al. (2020), under plausible assumptions, the resulting IV estimate will exhibit smaller bias, which tends to vanish as the number of GPs per district becomes large.⁵

⁵ See Bardhan et al. (2020) for details of the first stage regressions and the strength of the instrument in predicting variation in GP budgets.

Table 6: Within-GP targeting regressions with district fixed effects—IV version

	Dependent variable: number of benefits received					
	Employment benefit		Non-employment anti-poverty programmes		Subsidized farm inputs	
	OLS (1)	IV (2)	OLS (3)	IV (4)	OLS (5)	IV (6)
GP benefits k	0.011*** (0.002)	0.014*** (0.003)	0.010*** (0.002)	0.018*** (0.007)	0.009*** (0.002)	0.012*** (0.003)
Ultra poor	0.057*** (0.009)	0.057*** (0.009)	0.046*** (0.010)	0.046*** (0.010)	-0.011*** (0.004)	-0.011*** (0.004)
Moderately poor	0.033*** (0.007)	0.033*** (0.007)	0.034*** (0.007)	0.034*** (0.007)	-0.009** (0.004)	-0.009** (0.004)
Marginally poor	0.014*** (0.004)	0.014*** (0.004)	0.014*** (0.004)	0.014*** (0.004)	-0.004* (0.003)	-0.004* (0.003)
Number HH in village	-0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
Proportion of ultra poor	-0.087*** (0.033)	-0.114*** (0.040)	-0.086 (0.060)	-0.199 (0.126)	-0.022 (0.013)	-0.029* (0.015)
Proportion of moderately poor	-0.022 (0.018)	-0.031 (0.019)	-0.044 (0.036)	-0.068 (0.046)	0.006 (0.009)	0.003 (0.009)
Proportion of marginally poor	-0.023 (0.016)	-0.028 (0.018)	-0.022 (0.025)	-0.033 (0.033)	-0.002 (0.007)	-0.004 (0.007)
Observations	25025	25025	25025	25025	25025	25025
Adjusted R ²	0.085	0.079	0.054	0.037	0.092	0.085
Mean dependent variable	0.033	0.033	0.064	0.064	0.008	0.008
SD dependent variable	0.194	0.194	0.262	0.262	0.087	0.087
F-test of excluded instruments (p-value)		15.18 (0.00)		4.08 (0.05)		10.29 (0.00)
Rank test (p-value)		5.86 (0.02)		2.87 (0.09)		4.03 (0.04)
Weak-instrument-robust AR test [†] (p-value)		12.37 (0.00)		6.85 (0.01)		6.92 (0.01)

Note: * p<0.10, ** p<0.05, *** p<0.01. † Ho: $\beta=0$ and Ho: instruments valid i.e. $E(Zu)=0$. Observations are at the household-year level, 1998–2008. Dependent variable in columns (1)–(2) is number of employment benefits received by the household in year t . For columns (3)–(4), the dependent variable is non-employment anti-poverty benefits and for columns (5)–(6), it is number of subsidized farm inputs. For each type of benefit, the first column reports the results from an OLS regression while the second column reports estimates from an IV regression. Each specification includes year and district fixed effects. Robust standard errors are in parentheses, clustered at GP level.

Source: authors' calculations based on survey data.

The IV regression includes both year and district fixed effects. For each programme, the first column is the OLS regression (reproduced from the corresponding third column in Table 5), and the second column is the corresponding IV regression. It is evident that the OLS and IV estimates are very close to each other. Hence, the bias in the OLS regression does not appear sizeable. In what follows, we shall assume that there is no endogeneity bias in the estimated marginal impact of increasing the GP budget.

Our preferred model for predicting the number of benefits received by households when GP budgets are reallocated according to the formula is the Poisson regression model. This method is appropriate because the log transformation in the Poisson model guarantees that the predicted number of benefits is non-negative. We enrich the specification in Table 5 to allow for interactions between GP budget and household poverty status. Table 7 shows that these interaction coefficients are negative, implying that while poor households continue to receive priority, this priority diminishes as the GP budget expands—increases in the budget are directed more towards non-poor households. These coefficients, however, are quantitatively negligible compared with the corresponding coefficients of the poverty status dummies themselves. Even though there is relatively little heterogeneity in the effect of varying GP budgets across

different poverty groups, we will use this extended version of the model in order to improve the accuracy of the predictions.

Table 7: Within-GP targeting—Poisson prediction model

	Dependent variable: number of benefits received		
	Employment benefit	Non-employment anti-poverty programmes	Subsidized farm inputs
	(2)	(3)	(4)
GP budget (per HH benefit)	0.183*** (0.027)	0.147*** (0.022)	0.154*** (0.059)
Ultra poor	1.867*** (0.203)	0.870*** (0.116)	-1.164* (0.608)
Moderately poor	1.258*** (0.198)	0.742*** (0.081)	-0.755* (0.431)
Marginally poor	0.554*** (0.165)	0.411*** (0.073)	-0.225 (0.200)
GP Benefits * Ultra poor	-0.045*** (0.009)	-0.029*** (0.009)	-0.255*** (0.083)
GP Benefits * Moderately poor	-0.025*** (0.007)	-0.028*** (0.009)	-0.053** (0.024)
GP Benefits * Marginally poor	-0.009 (0.010)	-0.027*** (0.006)	-0.017* (0.009)
Number HH in village	0.002*** (0.000)	0.000 (0.000)	-0.003*** (0.001)
Proportion of ultra poor	-1.375 (1.333)	0.465 (1.111)	2.859 (1.936)
Proportion of moderately poor	-0.449 (0.741)	-0.205 (0.736)	1.190 (1.116)
Proportion of marginally poor	-0.903* (0.492)	-0.109 (0.410)	-1.152 (1.245)
Observations	25025	25025	25025
Mean dependent variable	0.033	0.064	0.008
SD dependent variable	0.194	0.262	0.087
Year fixed effects	YES	YES	YES
District fixed effects	YES	YES	YES

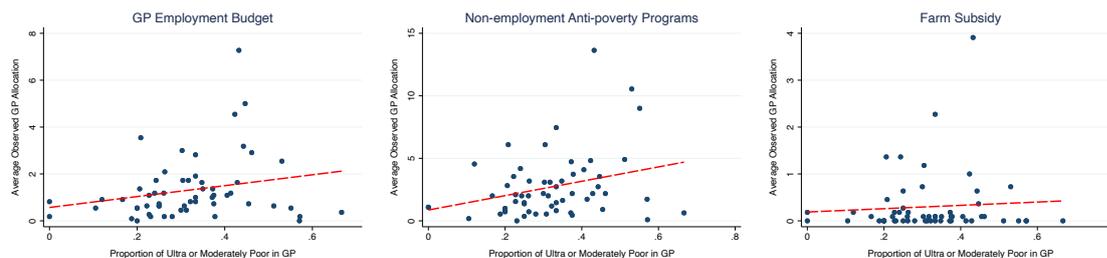
Note: observations are at the household-year level, 1998–2008. Dependent variable in column (1) is the number of employment benefits received by the household in year t , column (2) is the number of non-employment anti-poverty benefits, and column (3) is the number of subsidized farm inputs. Each specification is estimated using a Poisson regression model and the coefficients can be interpreted as the change in log of expected number of benefits associated with a unit change in each regressor. Each specification includes year and GP fixed effects. Robust standard errors are in parentheses, clustered at GP level.

Source: authors' calculations based on survey data.

4 Across-GP targeting

In this section, we examine the targeting patterns in across-GP observed allocations. Figure 2 plots estimated GP budgets against the proportion of households in the village that are ultra or moderately poor, with the red dashed line showing the corresponding OLS linear regression. These regressions all show a positive slope, indicating that the across-GP allocation was progressive.

Figure 2: Across-GP budget variations with GP poverty



Source: authors' calculations based on survey data.

4.1 Explaining the progressivity of targeting patterns

To shed light on the role of clientelism in driving the progressive allocation of programme benefits, we refer back to the theoretical model of two-party electoral competition in a two-tier (middle and lower) government hierarchy in Bardhan et al. (2020). Elections are held at both tiers, based on a first-past-the-post contest. The middle tier allocates programme budgets across different GPs at the lower tier, while elected GP officials allocate their assigned budgets across households within the GP. Officials at both tiers use their discretionary allocation powers to maximize the likelihood of their respective party's re-election. Voters assign credit for benefits received to the party controlling the GP, a plausible consequence of the budgeting process's lack of transparency. With a standard model of probabilistic voting, GP officials of either party allocate their assigned budgets to households most likely to respond with their votes to benefits they receive. Hence, within-GP targeting is biased in favour of households with stronger 'vote responsiveness' or 'swing propensity'. Within-GP targeting would therefore tend to be pro-poor if poorer households were more responsive.

We construct political support data from ballots cast by heads of household in the 2011 survey. The process simulated the official 'secret ballot' voting process. The households were provided sample ballots marked with symbols of principal political parties participating in local elections. The names of the respondents did not appear on the ballots and were instead replaced by a number assigned by a security code available only to the PIs. The respondents were given the ballot and a locked box. They were allowed to go into a separate room, cast their vote by putting their ballots in the locked box and then return the box to the interviewer. The survey was conducted shortly after the state assembly elections in 2011.

Table 8 reports the results for voting responsiveness to receipt of private benefits (aggregating all three categories of private programme) for 2009–11 for two groups: poor (combining ultra and moderately poor groups) and less poor (combining marginally poor and non-poor) households. The OLS results in column (1) show that a one standard deviation increase in private benefits received by poor households resulted in a 3.6% higher likelihood for the head of the household to vote for the GP incumbent. Consistent with the results in Bardhan et al. (2020), our findings show no voting responsiveness for public good benefits received, as predicted by the clientelist theory (since public good benefits being non-exclusionary cannot be used as a clientelist instrument to generate votes). Column 3 shows the corresponding OLS estimates for the less poor. While the coefficient of public benefits fails to be positive and significant, the coefficient of private benefits is one-third of the magnitude of the corresponding coefficient for poor households and fails to be statistically significant.

The second and fourth columns show the corresponding IV estimates when benefit distribution within the GP is instrumented by per household supply in the district excluding the GP in question, again in line with the IV strategy in Levitt and Snyder (1997) and Bardhan et al. (2020). The IV estimates are substantially larger in magnitude than the OLS estimates, but the qualitative pattern remains the same: only private benefits matter for votes, and they matter much more for poor households. Hence, the

greater vote responsiveness of the poor is robust to endogeneity concerns for the supply of benefits and helps explain why within-GP targeting tends to be pro-poor.

Table 8: Effect of benefits on votes for incumbent in 2011 straw polls

	Dependent variable: whether respondent voted for the incumbent party in majority at the GP			
	Poor households		Marginally poor and non-poor households	
	OLS (1)	IV (2)	OLS (3)	IV (4)
Private benefits	0.036** (0.014)	0.221** (0.095)	0.011 (0.013)	0.141 (0.104)
Public benefits	0.011 (0.023)	-0.146 (0.134)	-0.024 (0.018)	-0.072 (0.113)
Observations	891	891	1492	1492
Adjusted R^2	0.170	0.019	0.192	0.144
Mean votes for Left	0.511	0.511	0.521	0.521
SD votes for Left	0.500	0.500	0.500	0.500
F-test of excluded instruments (p-value)		7.83, 3.44 (0.00, 0.00)		9.31, 5.35 (0.00, 0.00)
Rank test (p-value)		7.65 (0.10)		6.18 (0.18)
Weak-instrument-robust AR test [†] (p-value)		11.15 (0.05)		7.06 (0.22)

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. [†] Ho: $\beta_{private} = 0$ and $\beta_{public} = 0$ and Ho: $E(Zu) = 0$. The dependent variable is whether the respondent voted for the incumbent party in majority at the GP in our 2011 straw polls. Private and public benefits are standardized and aggregated over period 2009-2011. All specifications include household (HH) characteristics, GP characteristics, and district fixed effects. HH Characteristics include SC/ST, religion, landlessness, occupation, and level of education of household head. GP characteristics include dummy for left GP, dummy for left panchayat samiti (PS) and dummy for alignment between GP and PS. Robust standard errors are in parentheses, clustered at village level in (1) and (3).

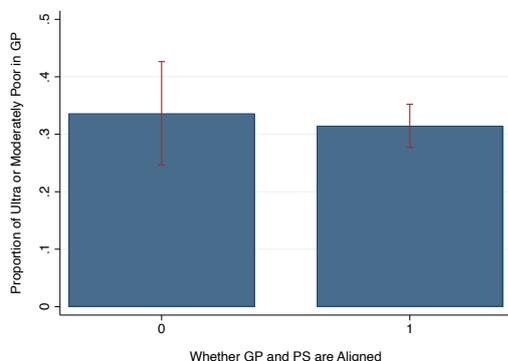
Source: authors' calculations based on survey data.

We now turn to the across-GP targeting pattern, resulting from GP budgetary allocations made by officials at the upper tier. The Bardhan et al. (2020) model shows that the optimal allocation to ensure their re-election is one in which the allocation for a given programme k to GP g is increasing in $[C_{a(g)} * A_{a(g),g} * v_{kg}]$, where $C_{a(g)}$ denotes **competitiveness** of assembly constituency $a(g)$ in which g is located, $A_{a(g),g} \in \{-1, 1\}$ is **alignment** of party controlling $a(g)$ with party controlling GP g , and v_{kg} is the **marginal responsiveness of votes** in GP g to programme k budget. A GP with positive (resp. negative) alignment is controlled by the same (rival) party; hence, allocating a larger budget to such a GP ensures an increase in votes for one's own (resp. the rival) party in the electoral contest at the upper tier. Therefore, the targeting is biased in favour of (resp. against) positively (resp. negatively) aligned GPs. The extent of such bias increases as the electoral contest becomes tighter, and marginal vote swings have a larger role in affecting which party wins. As poorer voters are more responsive, this factor by itself induces a pro-poor bias. Hence, a sufficient condition for across-GP targeting patterns to be progressive is that electoral competitiveness and alignment exhibit either zero or positive correlation with GP poverty rates.

Figure 3 examines how GP poverty varied with alignment (between control of GP and the next upper tier, the *panchayat samity* (PS)), taking two possible values: zero (not aligned) and one (aligned). It shows the average proportion of poor households is very similar for aligned and non-aligned GPs. Figure 4 plots the victory margin in 2011 assembly elections on the vertical axis and number of ultra or moderately poor households on the horizontal axis. The plots show that there is no relationship between GP poverty and electoral competition. Moreover, this lack of correlation does not differ significantly between aligned and non-aligned GPs.

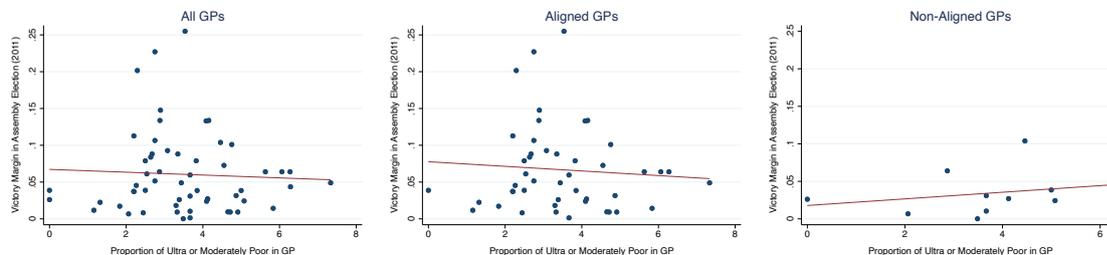
In summary, electoral competition and alignment exhibit negligible correlation with GP poverty rates. Hence, the progressivity of across-GP budget allocations appear to have been driven primarily by a higher voting responsiveness of poor households to receipt of private benefits.

Figure 3: GP poverty and alignment



Source: authors' calculations based on survey data.

Figure 4: GP poverty, electoral competition, and alignment



Source: authors' calculations based on survey data.

4.2 Targeting implications of formula-based budgets

We now address the question whether pro-poor targeting would have improved if the allocation of programme budgets to GPs had been determined by the formula recommended by the Third State Finance Commission (SFC, State Finance Commission 2008). The SFC's recommendations were based on the following GP variables, drawn from the village census and other household surveys:

GP_{1g} : weighted population share of g , the sum of undifferentiated population (which receives a weight of 0.500), and SC/ST population (a weight of 0.098);

GP_{2g} : female non-literates' share of g ;

GP_{3g} : food insecurity share of g , calculated from 12 proxy indicators collected in the Rural Household Survey of 2005, based on survey responses to questions such as "do you get less than one square meal per day for major part of the year?" ;

GP_{4g} : population share of marginal workers, those employed for less than 183 days of work in any of the four categories: cultivation, agricultural labour, household-based economic activities, and others;

GP_{5g} : total population without drinking water or paved approach or power supply share of g ;

GP_{6g} : sparseness of population (inverse of population density) share of g .

Table 9 shows how well these characteristics predict the proportion of households in different poverty groups in any given GP. The ultra-poor ratio is rising in the SC/ST proportion and population sparseness, but it does not significantly vary with the other SFC characteristics; the overall R-squared of this regression is 45%. So most of the variation in ultra-poor incidence is not explained. A larger fraction of variation (about two-thirds) in the moderately poor proportion is explained; most of this predictive

power comes from a sharp positive slope with respect to village population size. The size of the other two groups is less precisely predicted (R-squared below 40%) by the SFC characteristics; none of the individual characteristics are individually significant. These facts highlight the paucity of information available to construct formulae for programmatic GP budgets.

The specific formula recommended by the SFC for budget b_g to be allocated to GP g is

$$b_g = 0.598 * GP_{1g} + \sum_{i=2}^4 0.100 * GP_{ig} + \sum_{j=5}^6 0.051 * GP_{jg}. \quad (1)$$

Table 9: Demographic share of poverty groups and SCF GP characteristics

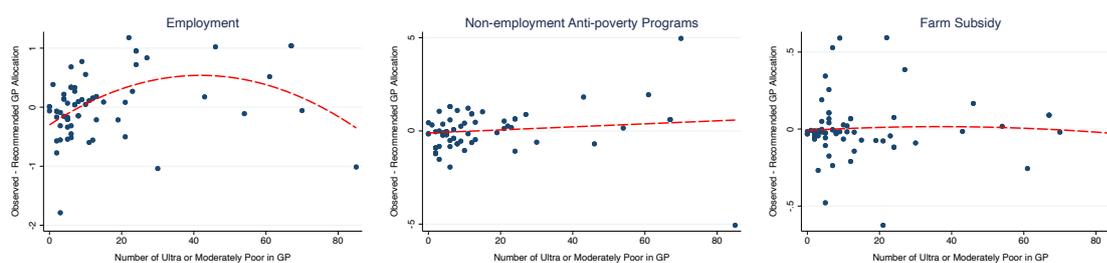
	Ultra poor (1)	Moderately poor (2)	Marginally poor (3)	Non-poor (4)
Population	0.013 (0.111)	0.472** (0.178)	0.042 (0.790)	0.172 (0.836)
SC/ST	0.141** (0.063)	0.021 (0.143)	-1.896 (1.450)	-2.086 (1.489)
Female illiteracy	-0.106 (0.212)	0.335 (0.276)	1.453 (1.216)	1.455 (1.051)
Food insecurity	-0.030 (0.042)	-0.054 (0.090)	-0.491 (0.315)	-0.109 (0.331)
Lack of infrastructure	-0.032 (0.239)	-0.230 (0.344)	0.881 (1.533)	0.469 (1.406)
Marginal workers	-0.029 (0.085)	-0.040 (0.147)	1.100 (0.805)	0.889 (0.844)
Sparseness of population	0.435** (0.180)	0.266 (0.229)	0.409 (0.706)	0.707 (0.885)
Observations	56	56	56	56
Adjusted R^2	0.449	0.649	0.387	0.333

Note: this table examines the relationship between our poverty measures and the components of the State Finance Commission formula. Observations are at GP level. Robust standard errors are in parentheses.

Source: authors' calculations based on survey data.

We apply this formula to calculate recommended budgets, upon assigning weights to GPs based on their scores using (1) and reallocating district programme scales across these GPs in the same ratio as their respective weights. The deviation of the observed from the recommended GP budgets are plotted in Figure 5 against the proportion of (ultra or moderately) poor households within the GP. For non-linear relationships, we fit a quadratic regression whose predicted values are depicted by the red dashed line. Over the relevant range of GPs in which less than 50% of households are poor, we see that the regression for employment benefits is upward sloping. For other anti-poverty benefits, it is upward sloping over the entire range. Hence, the SFC-recommended budgets for anti-poor programmes were less progressive than the observed allocations. The political discretion of ULGs therefore induced a more pro-poor across-GP allocation than would have resulted from the formula recommended by the SFC.

Figure 5: Deviation of observed from SFC-recommended GP budgets



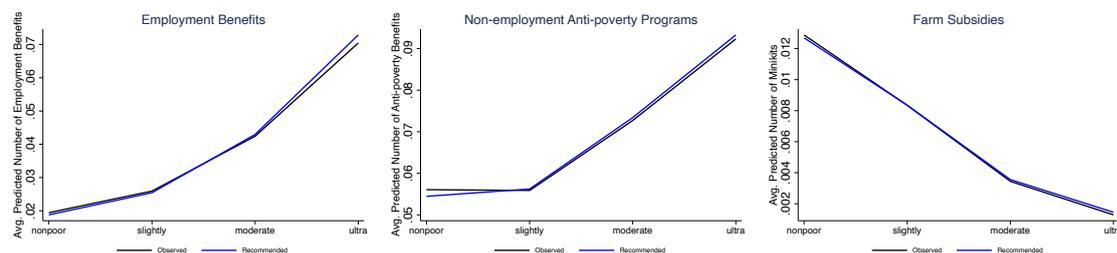
Source: authors' calculations based on survey data.

Next, we examine the consequences for targeting at the more disaggregated household level. Using the within-GP targeting pattern estimates shown in Table 7, we predict the number of benefits each household would have received had the observed GP budget been replaced by the SFC-recommended budget. The within-GP targeting pattern is described by the estimates in Table 7. There is no guarantee that the corresponding estimates of benefits received by each group generated independently for these groups will add up exactly to the incremental budget allocated. To ensure the GP budget remains balanced, we need to adjust the predicted benefits suitably. In one approach, which we call *proportional scaling*, we scale the predicted benefits for all four groups by the same proportion in such a way as to ensure budget balance. In the other method, called *residual scaling*, we generate the estimates for the three poor groups independently from the within-GP targeting regression, then adjust the benefits for the non-poor to ensure budget balance.

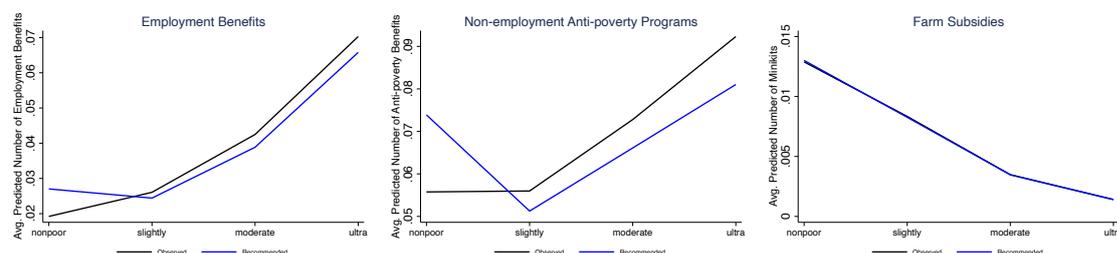
We subsequently aggregate the observed and predicted benefits from formula-based grants across the entire sample, and compare them for the average household in a given group. These results are shown in Figure 6. They confirm what one might expect from the greater progressivity of the observed GP budgets compared with the recommended ones: that the use of the SFC formula would not have improved pro-poor targeting. Under proportional scaling, average targeting patterns are practically unchanged, while under residual scaling, the poor would have been worse off with formula-based budgets.

Figure 6: Comparing observed targeting with predicted targeting under SFC formula-based within-district reallocation of GP budgets

[a] Proportional scaling



[b] Residual scaling



Source: authors' calculations based on survey data.

The corresponding implications for a related but different measure of targeting—the aggregate share of benefits delivered to poor groups—are shown in Table 10. Under proportional scaling, the SFC formula would marginally increase the aggregate share of ultra poor and moderately poor households for all three types of programmes. With residual scaling, on the other hand, targeting to all the poor groups would deteriorate for all welfare benefits.

Table 10: Group shares under observed and recommended allocations with within-district formula-based reallocation

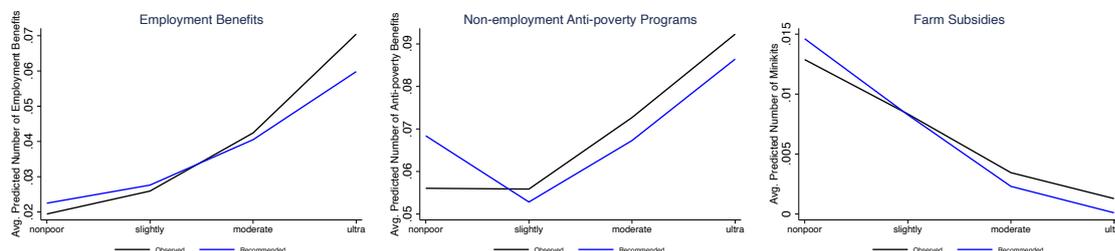
Group	Demographic share	Employment		Non-emp anti-pov.		Farm subsidy ec.	
		Observed	Rec.	Observed	Rec.	Observed	ec.
[a] Proportional scaling							
Ultra poor	8.53	18.42	19.06	12.37	12.49	01.45	01.64
Moderately poor	27.56	35.86	36.30	31.47	31.77	12.58	12.98
Marginally poor	38.33	30.48	29.90	33.64	33.85	42.35	42.39
Non-poor	25.58	15.24	14.74	22.53	21.88	43.62	42.99
[b] Residual scaling							
Ultra poor	8.53	18.42	17.19	12.37	10.85	1.45	1.58
Moderately poor	27.56	35.86	32.84	31.47	28.61	12.58	12.62
Marginally poor	38.33	30.48	28.71	33.64	30.86	42.35	41.85
Non-poor	25.58	15.24	21.26	22.53	29.68	43.62	43.95

Source: authors' calculations based on survey data.

The preceding exercise concerned the impacts of reallocating GP budgets within each district, but did not incorporate reallocations across districts. We now examine the consequences of reallocating across GPs across the entire state, using the SFC formula. The predicted impacts (under the proportional scaling method) on per household benefits for each group are shown in Figure 7 and on the average group shares in Table 11. The effects turn out to be similar to and somewhat larger than the corresponding

impacts of within-district reallocations. For this reason, in the rest of the paper, we focus on the effects of within-district reallocations.

Figure 7: Comparing observed targeting with predicted targeting under SFC formula-based state-wide reallocation of GP budgets, proportional scaling



Source: authors' calculations based on survey data.

Table 11: Group shares under observed allocation vs. recommended formula-based state-wide reallocation of GP budgets, proportional scaling

Group	Demog. share	Employment		Non-emp anti-pov.		Farm subsidy	
		Observed	Rec.	Observed	Rec.	Observed	Rec.
Ultra poor	8.53	18.42	15.64	12.37	11.58	01.45	00.12
Moderately poor	27.56	35.86	34.24	31.47	29.13	12.58	08.41
Marginally poor	38.33	30.48	32.48	33.64	31.81	42.35	41.97
Non-poor	25.58	15.24	17.64	22.53	27.49	43.62	49.50

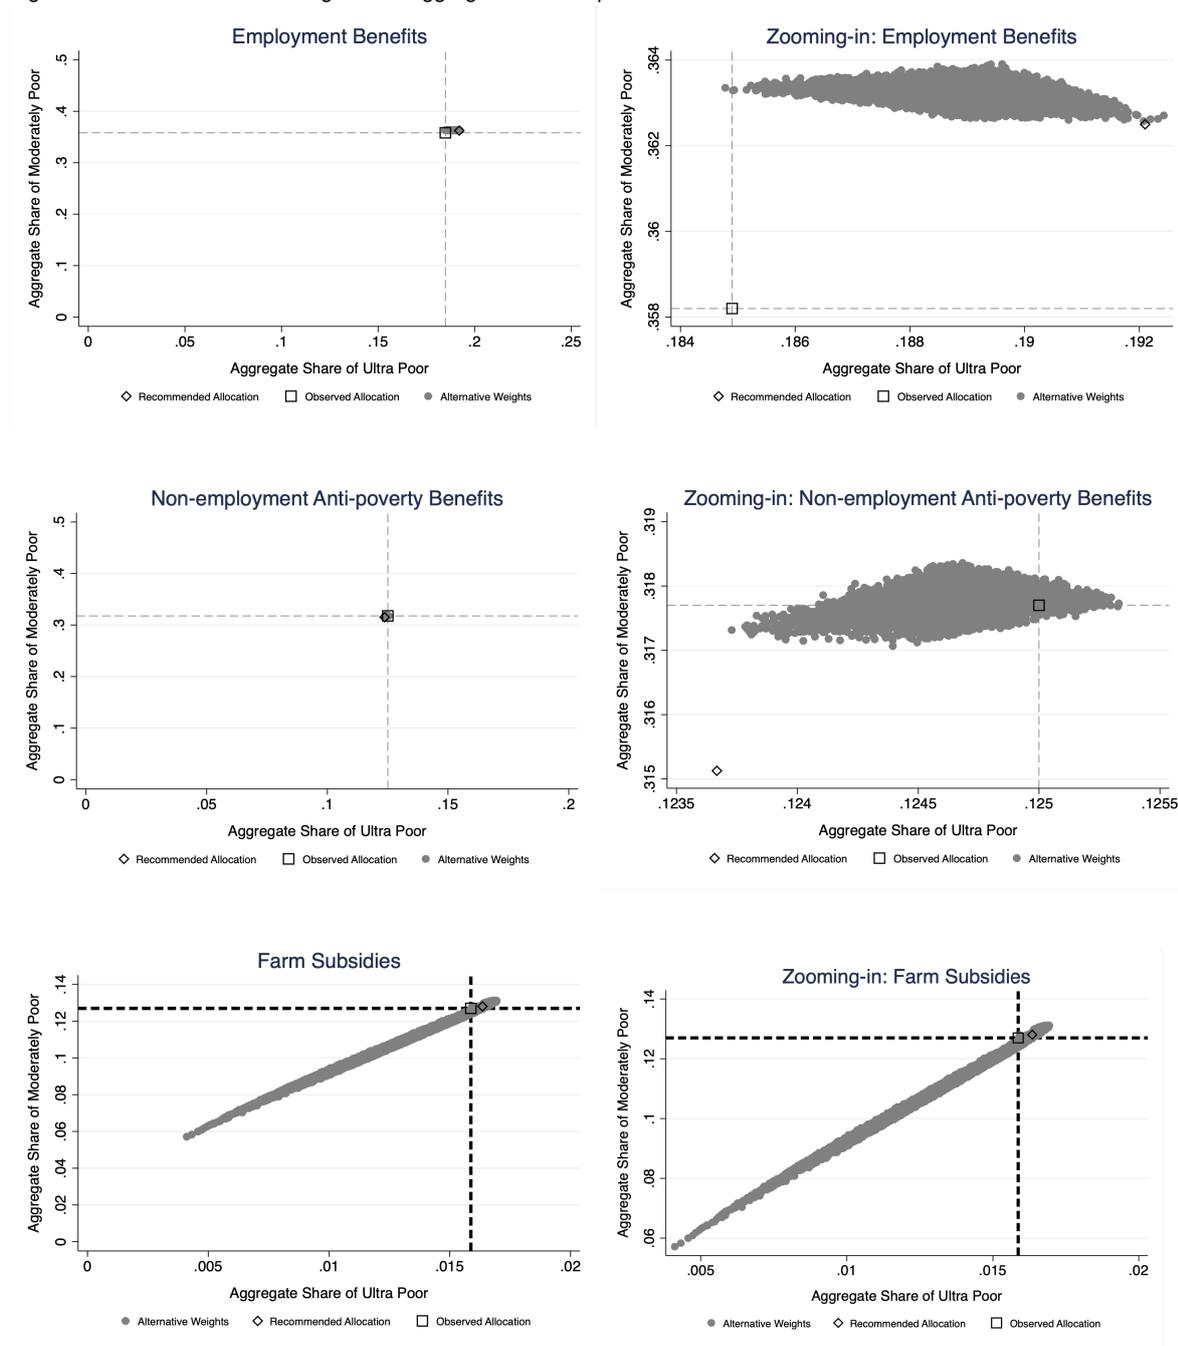
Source: authors' calculations based on survey data.

4.3 Alternative formula weights

We now examine whether alternative formulae based on changing the weights on GP demographic variables used by the SFC can improve the targeting of benefits to poorer groups compared with that in observed allocations. We consider within-district reallocations of GP budgets, using the set of GP characteristics from equation 1. We draw 10,000 alternative weights from the Dirichlet distribution using a likelihood model with uniform density over each weight in the unit simplex defined by $\sum_i w_i = 1; w_i > 0$ in R^7 .

For each draw, we use proportional scaling to balance the budget and calculate the aggregate share of benefits going to ultra poor and moderately poor households. Figure 8 plots the two groups' aggregate shares implied by each alternative formula. The pair of aggregate shares associated with the observed household allocation is depicted by dashed lines. The horizontal and vertical lines depicting observed allocation partition the graph into four parts. The upper right quadrant depicts the set of weights for which the aggregate share of benefits for both the ultra and moderately poor would be higher in the corresponding formula-based budget than in the observed allocation.

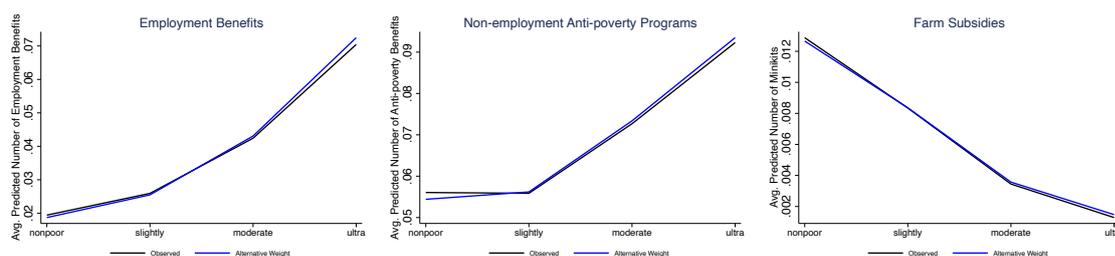
Figure 8: Alternative formula weights and aggregate share of poor households



Source: authors' calculations based on survey data.

The results show that compared to the observed allocation, formula-based budgets with suitably chosen weights different can improve aggregate shares for the two poor groups, but only marginally. These are depicted by the set of weights in the upper right quadrant of the graph. Figure 9 plots the predicted number of benefits for each poverty group if the formula weights had been chosen to maximize the average share of the ultra-poor group. In this case, the quantitative improvement continues to be small. The ultra-poor group's shares of employment and anti-poverty benefits increase from 18.4% to 19.2% and from 12.37% to 12.52%, respectively.

Figure 9: Predicted benefits for alternative weights that maximize the ultra poor share



Source: authors' calculations based on survey data.

Table 12: Aggregate shares under observed and alternative allocations

Group	Demog. share	Employment		Non-emp anti-pov.		Farm subsidy	
		Observed	Alt.	Observed	Alt.	Observed	Alt.
Ultra poor	8.53	18.42	19.24	12.37	12.52	01.45	01.66
Moderately poor	27.56	35.86	36.27	31.47	31.77	12.58	13.03
Marginally poor	38.33	30.48	29.75	33.64	33.84	42.35	42.42
Non-poor	25.58	15.24	14.73	22.53	21.87	43.62	42.88

Source: authors' calculations based on survey data.

5 Conclusion

In summary, observed anti-poverty programme targeting patterns were pro-poor, both within and across GPs in rural West Bengal. Switching to a rule-based financing system based on the State Finance Commission's formula would have reduced the extent of pro-poor targeting. Our calculations indicate that alternative formulae obtained by varying the weights on GP characteristics used in the SFC formula would have improved pro-poor targeting only marginally. Hence, as long as formula based budgets are based on the measures of village need used by the SFC, little improvement in pro-poor targeting can be expected.

The results highlight the need for the state government or the SFC to use more accurate information regarding the distribution of poverty in the event of a transition to centralized budgeting. Village demographics contained in the census are unlikely to be precise enough; they need to be supplemented by more detailed measures of local poverty that are based on disaggregated household surveys. Moreover, these surveys could be used to estimate targeting patterns and the extent to which they differ across regions; these estimates could also be used to fine-tune formulae used to determine budgets. For instance, districts that exhibit greater targeting errors could receive smaller grants.

The results also may reflect the inherent limitations of delegating within-GP targeting to GPs, instead of direct allocation of transfers to households (which would require upper level governments to build a reliable database of proxy means of household poverty for **all** households, combined with the capacity to deliver benefits directly to them). While intra-village targeting of anti-poverty programmes is progressive on average, a large fraction of these benefits (exceeding 40%) were delivered to households that are neither ultra or moderately poor. However this measure of targeting leakage may be an over-estimate, if our measure of household poverty status include measurement error. Future research could be devoted to studying effects of adverse transitory (idiosyncratic or village level weather) shocks on targeting and whether our results continue to be robust when these are incorporated.

A number of further qualifications are in order. We focused entirely on questions of vertical distributive equity in the allocation of private benefits and abstracted from many other welfare-relevant dimensions. Politically manipulated variations in GP budgets result in horizontal inequity—that is, unequal treat-

ment of different GP areas that cannot be defended on normative grounds—and reduce the legitimacy of incumbent parties. Moreover, focusing on pro-poor targeting alone ignores possible under-provision of public goods and reduced political competition that have been alleged by many scholars to be pernicious consequences of clientelism. Assessing the empirical relevance of these concerns constitutes an important and challenging agenda for future research and policy experimentation.

References

- Alatas V., A. Banerjee, R. Hanna, B. Olken, and J. Tobias (2012). 'Targeting the Poor: Evidence from a Field Experiment in Indonesia'. *American Economic Review*, 102: 1206–40. <https://doi.org/10.1257/aer.102.4.1206>
- Azulai, M. (2017). 'Public Good Allocation and the Welfare Cost of Political Connections: Evidence from Brazilian Matching Grants.' Working Paper. London: Institute of Fiscal Studies.
- Bardhan P, D. Mookherjee, and M. Parra Torrado (2010). 'Impact of Political Reservations in West Bengal Local Governments on Anti-Poverty Targeting'. *Journal of Globalization and Development*, 1(1): art. 5. <https://doi.org/10.2202/1948-1837.1025>
- Bardhan P., and D. Mookherjee (2012). 'Political Clientelism and Capture: Theory and Evidence from West Bengal'. working paper. Boston: Institute for Economic Development, Boston University.
- Bardhan P, D. Mookherjee, M. Luca and F. Pino (2014), 'Evolution of Land Distribution in West Bengal 1967–2003: Role of Land Reform and Demographic Changes'. *Journal of Development Economics*, 110: 171–90. <https://doi.org/10.1016/j.jdeveco.2014.02.001>
- Bardhan P., S. Mitra, D. Mookherjee, and A. Sarkar (2015). 'Political Participation, Clientelism, and Targeting of Local Government Programs: Results from a Rural Household Survey in West Bengal, India'. In J.-P. Faguet and C. Poschl (eds), *Is Decentralization Good For Development? Perspectives from Academics and Policy Makers*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198737506.003.0013>
- Bardhan P., S. Mitra, D. Mookherjee, and A. Nath (2020). 'How Do Voters Respond to Welfare Programs Vis-a-Vis Infrastructure Programs? Evidence for Clientelism in West Bengal'. Working paper. Boston: Institute for Economic Development, Boston University.
- Dey, S., and K. Sen (2016). 'Is Partisan Alignment Electorally Rewarding? Evidence from Village Council Elections in India'. IZA Discussion Paper 9994. Bonn: Institute of Labor Economics. <https://doi.org/10.2139/ssrn.2835536>
- Dreze, J., R. Khera, and A. Somanchi (2020). 'Balancing Corruption and Exclusion: A Rejoinder'. Ideas for India, 28 September 2020. Available at: <https://www.ideasforindia.in/topics/poverty-inequality/balancing-corruption-and-exclusion-a-rejoinder.html> (accessed 11 June 2021).
- Finan, F., and M. Mazzocco (2020). 'Electoral Incentives and the Allocation of Public Funds'. Working Paper. Berkeley: Department of Economics, University of California. <https://doi.org/10.1093/jeea/jvaa055>
- Holland, A. (2017). *Forbearance as Redistribution: The Politics of Informal Welfare in Latin America*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/9781316795613>
- Levitt, S., and J. Snyder (1997). 'The Impact of Federal Spending on House Election Outcomes'. *Journal of Political Economy*, 105(1): 30–53. <https://doi.org/10.1086/262064>
- Mansuri, G., and V. Rao (2013). *Localizing Development: Does Participation Work?* Washington, DC: World Bank Publications. <https://doi.org/10.1596/978-0-8213-8256-1>

- Mookherjee, D. (2015). 'Political Decentralization'. *Annual Review of Economics*, 7: 231–49. <https://doi.org/10.1146/annurev-economics-080614-115527>
- Shenoy, A., and L. Zimmerman (2020). 'Political Organizations and Persistent Policy Distortions'. Working paper. Santa Cruz: Department of Economics, University of California.
- State Finance Commission (2008). *Report of the Third State Finance Commission, West Bengal 2008*. Kolkata: Government of West Bengal.
- Stokes, S. (2005). 'Perverse Accountability: A Formal Model of Machine Politics with Evidence from Argentina'. *American Political Science Review*, 99(3): 315–25. <https://doi.org/10.1017/S0003055405051683>.
- Stokes, S., T. Dunning, M. Nazareno, and V. Brusco (2013). *Brokers, Voters, and Clientelism: The Puzzle of Distributive Politics*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9781107324909>
- Tarquinio, L. (2020). 'The Politics of Drought Relief: Evidence from South India'. Working paper. Boston: Department of Economics, Boston University.
- World Bank (2004). *World Development Report 2004: Making Services Work for Poor People*. Washington, DC: World Bank. <http://hdl.handle.net/10986/5986>