Electoral Accountability and Local Public Goods in Ghana

by

Evan P. Morier

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Political Science)

at the

UNIVERSITY OF WISCONSIN–MADISON

2021

Date of final oral examination: 7/14/2021

The dissertation is approved by the following members of the Final Oral Committee:

Rikhil Bhavnani, Associate Professor, Political Science
Scott Straus, Professor, Political Science and International Studies
Alex Tahk, Associate Professor, Political Science
Adeline Lo, Assistant Professor, Political Science

To Emily and Julian

Acknowledgments

This dissertation would not have been possible without the support of a great number of people. First and foremost, I must thank my advisor, Rikhil Bhavnani. Rikhil has supported my academic endeavors since I first visited Wisconsin as a prospective student and has taught me an incredible amount about comparative politics, causal inference, field research, and social science more broadly. He has provided invaluable feedback on numerous drafts of this dissertation and meaningful support in all of my endeavors.

Scott Straus has been an outstanding mentor, teacher, and supervisor during my time at Wisconsin. He helped guide my early research ideas and has consistently offered important feedback and advice. He played an especially important role in helping me acclimate to the program during my first semester and has modeled how to be a caring and hardworking scholar and teacher.

Alex Tahk taught me a great deal about statistics and research methods during several courses at Wisconsin. He fueled my interest in quantitative methodology and exposed me to cutting-edge approaches and insights in the field. He has provided a wealth of knowledge to graduate students and his impact on our skills and knowledge cannot be overstated.

I have also been incredibly fortunate to work with and learn from Adeline Lo. She has taught by example how important doing careful and well-documented research is. She has been generous with her time and is a role model for how to support students.

Emilia Tjernström provided me with an opportunity to collaborate on multiple research projects, including one that introduced me to doing research in Ghana. She was a joyful collaborator and supervisor from whom I learned a great deal about causal inference and research design. Nicole Bonoff also helped introduce me to Ghana and has provided valuable advice throughout my time in graduate school.

My fellow graduate students deserve special thanks. Mike DeCrescenzo was a methodological inspiration and invaluable resource throughout most of my time in Madison. This dissertation owes more to him than nearly anyone outside my committee. Anton Shirikov, Kate Carter, and Molly Minden formed a support network that made surviving the first

years of graduate school possible, as did the other members of our cohort. Many others have provided support and advice over the past six years and I am grateful to have learned and worked alongside them.

I am also thankful to Laura Schechter and the rest of the development research group in the Department of Agricultural and Applied Economics. They welcomed me into their group and provided valuable feedback on my research over the years. I learned a considerable amount about how to conduct and evaluate research of the highest quality from them, for which I am very grateful.

Many other teachers have contributed to my academic success over the years. At Bret Harte Middle School in Oakland, California, I was fortunate to learn from Duane Worm, Malia Dinell, Ann Yamoto, and Peter Mates. At Skyline High School in Oakland, James Richter and Dr. Brown were two of my excellent, inspiring teachers. During my undergraduate at Seattle University, several professors shaped my thinking on important issues of politics and philosophy. Erik Olsen, a graduate of the University of Wisconsin–Madison's political science Ph.D. program, encouraged deep critical thinking on core questions of how political power is organized. And Olufemi Taiwo introduced me to a completely new paradigm for thinking about the relationship between Africa and the West.

I owe much of what I know about Ghana to my friend and research assistant Raymond Annorbah. He also provided excellent research assistance that allowed me to conduct interviews with a diverse sample of Ghanaian voters. But more importantly, I always looked forward to returning to Ghana for fieldwork in large part because it meant I got to see a great friend again. *Medaase*, Kobby.

Much of my data collection in Ghana would not have been possible without the generosity of countless government employees across the country. One person at the Electoral Commission of Ghana (who must remain anonymous) went out of their way to locate several important pieces of data for me. Another person there provided me with the names and cell phone numbers of the regional directors of the electoral commission, which ensured I would be able to get the data I was looking for when I traveled to Ghana's regional capitals. Many others volunteered their time and efforts to help me and I could not be more grateful to all of them.

I would not be where I am now without the love and support of my parents. Early on in life they helped instill in me a love of learning and have always encouraged me to aim high and achieve my goals. They continue to be an endless source of advice and inspiration.

Finally, I am grateful for the love and support of my partner Emily. She has kept me grounded and taken on a significant portion of the work caring for our son, Julian, which has allowed me to complete this dissertation. I am so thankful to have her by my side as

we move onto the next chapter of our lives.

Contents

List of Tables

List of Figures

Abstract

Do voters reward incumbent politicians for public goods provision? While it seems reasonable to assume they would, existing research has produced some contradictory results. In this dissertation, I take up the question in the context of two intergovernmental transfer programs that funded public goods provision by local governments in Ghana. I advance a theory to explain why and when voters will reward incumbents for public goods provision. The main insight from the theory is that voters' preferences matter and that they condition electoral responses to public goods provision. I show that voters reward incumbent presidents for local public goods provision, which is rational because the president appoints all local government executives (mayors). But I do not find strong evidence that voters reward local government representatives. I find moderate evidence in support of the theory that the effect of public goods provision on support for incumbents depends on voters' preferences. I contribute to the literatures on voting behavior and democratic accountability and provide new theory and evidence on the effect of public goods on support for incumbents. Methodologically, this dissertation is, to the best of my knowledge, the first political science research to apply advanced methods of small area estimation to the study of African politics.

Chapter 1

Introduction

Incumbent political leaders are faced with a variety of demands. First, they must maximize their chances of winning re-election by maintaining the loyalty of the powerful party faithful. In many developing countries, this includes political patronage like jobs for well-connected individuals and even preferential government contracts from which graft may be extracted. But they also have to run the government well enough to avoid the serious dissatisfaction that can prevent their re-election, as Ghana's former president John Mahama learned the hard way in losing his presidential re-election bid in 2016. Focusing too much on rewarding supporters, especially through corrupt practices, can leave incumbents vulnerable on the question of competency.

The question remains, though, how much political benefit incumbents receive from providing a key policy output: public goods. In most of the developing world, public goods are chronically under provided. Good governance in these contexts therefore involves, in part, providing public goods. But if voters do not reward public goods provision with a substantial electoral boost, incumbents may be incentivized to provide just the bare minimum, while diverting substantial resources to other sectors.

This dissertation asks two core questions: What is the relationship between policy outcomes and voting behavior? And, how do voters interpret and evaluate the performance

of their elected officials? I examine these questions in the case of Ghana, an established democracy in Africa that is an ideal context in which to study such issues of democratic accountability. Drawing on a a variety of data sets I collected during fieldwork in the country, as well as interviews with voters and bureaucrats,[1] I explore whether voters reward incumbents for providing local public goods. I also investigate the role of voter preferences in conditioning any effect of public goods provision on support for incumbents.

In developing countries, there is an emphasis on clientelism and ethnicity as explanations for why voters vote the way they do. These theories argue that voters want club goods (local public goods like schools, for example) and private goods (e.g jobs in government or cash payments for votes), and that they either trade their votes for goods in a *quid pro quo* arrangement or vote for a co-ethnic who they believe will deliver them. But assuming that voters just want goods delivered without considering the *types* of goods they want assumes that all goods are equally valued by voters. Instead, I argue that voters care about which types of goods/policy benefits they receive and evaluate incumbent performance based on how realized policy outcomes align with their preferences. This is a form of performance voting that is more sophisticated than many would assume could exist for voters in developing countries.

Performance voting is important because it is thought to produce better governance. When voters condition their votes on how well their elected officials perform, it should motivate politicians to perform better. But existing work on this subject generally considers performance along one dimension. For example, performing well could mean producing higher economic growth (e.g Duch and Stevenson 2008) or more disaster relief (Cole, Healy, and Werker 2012). I argue that performance can also be conceptualized as a more complex, multidimensional concept. Voters have preferences over several types of policies, so it is plausible that they would condition their votes on more than one of them. For a

[1]Interviews were approved by the Institutional Review Board (IRB) at the University of Wisconsin–Madison (ID: 2019-0716).

politician to perform well they need to take into account the relative importance of several policy priorities to their constituents when allocating resources under a budget constraint. More responsive governments are also better performing ones.

I ask two related research questions on accountability in Ghana. First, do voters in Ghana take policymaker performance into account when at the ballot box? In particular, does public goods provision, which is sorely needed across Ghana, increase support for incumbents? This is an important link in the accountability chain; if voters do not condition their votes on policymaker performance then it would be difficult for voting to induce better performance. I investigate this question using data on two intergovernmental transfers in Ghana which are earmarked for public goods provision. Using fixed effects, difference-in-differences, and regression discontinuity approaches, I estimate the causal effect of these transfers on support for the incumbent president and for incumbent local assembly members.

Second, does the effect of public goods provision on support for incumbents depend on how much voters prioritize public goods over other types of public policies? If voters who care more about public goods are more willing to reward incumbents who provide them, it would suggest that voters are rather sophisticated and that incumbents could be induced to govern in accordance with the preferences of their constituents. To answer this question, I measure preferences using Afrobarometer data and a modified multilevel regression with post-stratification (MRP) approach to generate more accurate district level estimates. I then use a fixed effects panel setup to test whether the effect of public goods provision on support for incumbents depends on this measure of citizen preferences towards public goods.

I find that public goods funded by an intergovernmental transfer that all districts in Ghana were eligible for had a positive effect on support for the incumbent president, but only in the second period of the analysis. I suggest that there was not a measurable effect in the first period because districts were unable to quickly take advantage of the new programs

funding and turn it into public goods projects. I do not find evidence, however, that public goods funded by this transfer mechanism increased support for local government assembly members. This could be due to aggregation in the data, as I do not observe which single-member districts receive the public goods projects. So support for local incumbents could increase in the areas where projects were built but remain unchanged or even decrease in areas with fewer or no public goods projects.

Examining a second intergovernmental transfer for public goods that was only available to more urbanized districts in Ghana, I do not find statistically significant evidence that transfers increased support for incumbents. However, estimates are nearly uniformly positive across specifications, showing a pattern that provides suggestive evidence that a positive effect may be present. The results are comparable for both presidential and local elections.

Finally, I find some suggestive evidence that preferences for public goods condition the effect of public goods provision on support for presidential incumbents. In districts where more people cited infrastructure as one of the most important problems facing the country that the government should address, there is evidence for a stronger effect of public goods provision on support for the incumbent president. Although this result does not cross conventional levels of statistical significance, I show that about 91% of simulated estimates are positive. Given attenuation that is likely present due to the model-based measure of preferences I use, this provides reason to believe that preferences matter for how public goods provision affects support for incumbents.

I make several contributions to our knowledge of accountability and voting behavior in this dissertation. First, I add theory and evidence to the literature on the determinants of voting behavior. While there is a large body of research on voting behavior, it has tended to emphasize a few explanations and has not adequately explained variation, especially in Africa. Second, I add to existing work on the effects of public goods and service provision on support for incumbents. While some studies have found that they can increase support for

incumbents, recent studies have found that they can actually lower their electoral support. The existing literature reveals that we don't know enough about whether and under what conditions public goods and service provision can help incumbents electorally.

Finally, I estimate how much voters' preferences over different types of policies affect their evaluations of policy outcomes. This is distinct from existing studies of preferences and voting in developed democracies, particularly the US, where voters' are placed on an ideological scale to represent their preferences (e.g. Canes-Wrone, Brady, and Cogan 2002). I argue that in developing countries, though not ideological, voters' preferences over different types of policies are multidimensional and can be placed on scales similar to those used to measure ideology. Political competition in many developing countries (including most African countries) revolves around valence issues like corruption and economic performance rather than ideological concerns like the role of government in the economy.[2]

Existing research that does consider voters' preferences over different types of goods usually distinguishes between private, public, and club goods. For example, Nathan (2019b) makes the distinction between preferences for universalistic public policies (public goods that benefit a large proportion of citizens) and for particularistic benefits like cash handouts and local public goods that are targeted to a relatively small number of people. But he does not emphasize preferences over different types of universalistic and particularistic policies. This distinction is important because, as I argue, voters are not indifferent towards the specific universalistic policies that governments enact.

Methodologically, I apply statistical techniques for small area estimation to a new context. While multilevel regression and post-stratification (MRP) is commonly applied to American public opinion, it has not been used much if at all in Africa. To the best of my knowledge, this is the first application of an MRP-based method to African public opinion

[2]See e.g. Bleck and van de Walle (2011), who argue in part that parties seek to avoid taking positions on controversial issues and instead focus on valence issues where there is broad agreement.

in the political science literature. MRP methods are particularly promising in data scarce developing country contexts because they allow for smaller resolution data using national surveys. Further, I utilize a new and improved version of MRP called BARP (Bayesian additive regression trees with post-stratification) proposed by Bisbee (2019). This replaces the multilevel regression with a machine learning model for more flexible and accurate prediction of opinion.

This research also has implications for the political economy of development and foreign aid. The DDF program is jointly funded by the Government of Ghana and several foreign donors with the goal of increasing the capacity of local governments. The UDG was funded entirely by the World Bank, also with the goal of increasing local government capacity.[3] It would be useful for donors to know whether their governance programs help incumbents to stay in power, particularly in contexts where there is already a significant incumbent advantage. If these programs help poorly-performing incumbents avoid electoral defeat, as has been theorized for foreign aid more broadly, they could reduce the quality of governance in the long-run.

1.1 Plan of the Dissertation

In Chapter 2, I review existing research on voting for public goods provision and outline a theory that seeks to partially explain some surprising results from other studies. In particular, I argue that voter preferences have not been adequately taken into account in existing studies. Voters shouldn't be expected to reward public goods provision if they didn't want the particular goods they end up receiving, or if they preferred government resources to go to something entirely different from public goods provision. This could help explain why voters have been observed to either punish or fail to reward incumbents

[3]The UDG program was actually a part of a broader World Bank project called the Local Government Capacity Support Project.

for public goods or service provision in some contexts.

Chapter 3 discusses the choice of Ghana as the context in which to study public goods provision and voting behavior. Ghana's democratic bona fides are an important reason to study democratic accountability processes like performance voting there. Ghana also has a strong presidential system, like most other countries in Africa, but its stable, two-party political competition makes it more comparable to the United States context than other countries with more fluid party systems.

In Chapter 4 I introduce one primary focus of the dissertation: an intergovernmental funding mechanism called the District Development Facility (DDF) that provided funds to local governments in Ghana for use on public goods projects. I estimate the effect of DDF transfers on support for both the incumbent president's party and incumbent local government assembly members. Using two period panel data, I show that DDF funds led to greater support for the incumbent president's party in one period of the analysis. I suggest that this can be explained by delayed implementation of DDF-funded projects at the start of the program. However, I do not find evidence of an effect on local government incumbents, potentially because of their limited role in public goods provision.

I introduce a second intergovernmental funding mechanism known as the Urban Development Grant (UDG) in Chapter 5. The UDG program was entirely funded by the World Bank and only provided funds to more urbanized districts in Ghana. Although the per-district transfers under the UDG were larger than they were under the DDF, I do not find evidence of an effect of UDG funds on support for either presidential or local incumbents. This may be due to the length of time needed to complete the average UDG project, though I cannot determine the cause with certainty.

In Chapter 6 I test my theory that the effect of public goods provision on support for incumbents depends on voters' preferences. I combine survey and census data in a machine learning-based modification to multilevel regression with post-stratification (MRP) to estimate voters' preferences at the district level. Despite results that do not reach

conventional levels of statistical significance, the results provide some support for the theory. Chapter 7 contextualizes the results in the broader literature and concludes.

Chapter 2

How Public Goods and Preferences Affect Voting

In this chapter I propose a theory to explain how public goods and service provision affects support for incumbents and variation in this effect. While an increase in such provision will often earn incumbents more votes, I argue that an overlooked determinant of incumbent support is the degree of alignment between the distribution of types of public goods and services provided and those preferred by voters. I argue that voters have multidimensional preferences and care about the relative quantity of goods provided. That is, they don't simply care about the overall value of goods they receive but how much they receive of different types of goods. It is this information that they often take into account when evaluating the performance of incumbent politicians.

Throughout this chapter, I discuss whether and under what circumstances voters use information about incumbents' past performance in their voting decisions. I do not intend to claim that voters will always vote retrospectively or that it is always rational for them to do so. I simply posit that retrospective voting is, at least under some realistic conditions and given some assumptions, rational. I rely on this modest claim as I develop a theory of voting for public goods provision.

I first discuss theories of democratic accountability and performance voting. I then outline why we should expect a connection between transfers for public goods provision and support for incumbents, arguing that such a phenomenon is consistent with a rational model of voting behavior. Finally, I present a theory explaining how voter preferences can condition this effect.

2.1 Democratic Accountability, or Why Performance Voting Matters

Theories of democratic accountability suggest that institutions of democracy, particularly voting, have important beneficial effects on political outcomes. Because politicians depend on voters' support to get elected, voters can elect good candidates and remove those who perform poorly. A large body of work in political science and economics examines accountability in a principal-agent framework. The political agency models on which this work rests posit that elected officials act as agents of voters (Besley 2006). Regular elections serve as the formal link between voters and politicians, creating an accountability relationship which allows the former to induce good performance from the latter (Przeworski, Stokes, and Manin 1999). By controlling who gets elected, voters are able to influence what elected officials do in office. As with other principal-agent models, the principals (voters) are tasked with monitoring the behavior of their agents (politicians) in order to reduce the problem of moral hazard (e.g. incumbents shirking). They also need to have information about incumbents and challengers to solve the adverse selection problem (they do not know the politicians' types).

A central prediction of political agency models is that information about incumbents is crucial for voters to be able to monitor and induce good performance from them. Voters only observe a noisy signal of an incumbent's actions, so incumbents will generally perform better as the amount of noise decreases (Ashworth 2012). The importance of

information for political accountability has also received strong support in empirical work. Pande (2011) reviews experimental work on the effect of information on accountability in developing countries, which has shown that voters are receptive to information about politician performance and vote based on it. Gottlieb (2016) argues that voters need relative performance information to properly evaluate government performance and finds that Malian voters who receive such information are more likely to sanction poorly performing incumbents. Besley and Burgess (2002) find that state governments in India are more responsive in their provision of natural disaster relief where newspaper circulation is higher. Access to information via the media can also induce mayors in Brazil to appropriate fewer public resources (Ferraz and Finan 2011). However, a coordinated series of experimental interventions found no discernible effect of information on voter turnout or support for the incumbent (Dunning et al. 2018).

One approach to studying accountability has been to identify the effects of various government policies, including the provision of public goods and services, on support for incumbents. The logic to this strategy is that the voters who benefit from government policies should be less likely to sanction them electorally in the future.[1] When this occurs, it is taken as evidence of a functioning accountability process. For example, Harding (2015) finds that greater improvement in road quality in Ghana is associated with increased incumbent vote share. The economic voting literature has identified positive relationships between economic performance and incumbent support in multiple contexts (Duch and Stevenson 2008).

However, there are also studies which show a negative relationship between apparently beneficial policies and support for the incumbent (and vice versa), pushing the literature to better account for variation in voter responses to government policies. De Kadt and Lieberman (2017) unexpectedly find that increased service provision led to lower support for the incumbent government in South Africa, which cannot be explained by existing theory.

[1] Assuming that voters in a particular context use past information in their voting decisions.

The authors suggest that increased service provision may reveal how corrupt the incumbent government is. That is, what appears to be a positive action in terms of increasing aggregate voter welfare actually reveals that the government is not a "good type." Blattman, Emeriau, and Fiala (2016) show that a government anti-poverty program in Uganda caused an increase in opposition support among recipients, potentially because their increased wealth freed them from clientelistic relationships with the government. Bhavnani and Copelovitch (2018) find that in areas of India hit hardest by the negative effects of demonetization, voters maintained their support for the incumbent BJP, but in areas where the negative effects were lessened by the presence of more bank branches, the BJP lost support. A convincing explanation they offer is that voters supported demonetization despite being negatively affected by it because they thought it would bring the country future benefits, underlining that voters' preferences cannot always be accurately inferred through their immediate, personal, financial self-interest.

In addition, as outlined above, it is not clear that rational voters will always vote based on past information. Even when voters are retrospective they do not evaluate all incumbent actions over the previous term uniformly. Voters in both developed and developing countries exhibit a strong recency bias when evaluating incumbents at the ballot box (Cole, Healy, and Werker 2012; Achen and Bartels 2016, 175). The positive association between government spending and support for the incumbent was also found to attenuate within a few months in the case of mosquito bed net distribution in Tanzania (Croke 2017). Recency bias may also partially explain why Blattman, Emeriau, and Fiala (2016) find that a government anti-poverty program in Uganda caused an increase in opposition support among recipients: because they measure the effect of the program four years after it was implemented, they may have missed any incumbent support effects (with the increase in opposition support a result of some other intervening variable caused by the program).

Even when politicians are actually involved in creating or implementing a program citizens are often unable to properly assign them credit. Harding and Stasavage (2013)

argue that governments take policy actions that are more directly attributable to them due to this problem. Some insight into understanding how voters assign credit for targeted spending come from Pop-Eleches and Pop-Eleches (2012), who examine the political effects of a government computer coupon distribution program in Romania. They find that recipients of the program assigned credit for it to the party which implemented the program rather than the one which created it. Further, recipients become more trusting of the local politicians who administered the program but not the national politicians who adopted and financed it. This suggests that recency and proximity of government actions to voters increase their probability of being attributed to them rather than more distant (in time or space) government actions.

There are also documented cases where politicians were rewarded for programs they did not have any role in implementing. Cruz and Schneider (2014) show that mayors in the Philippines were able to claim credit for World Bank funded development projects, increasing their probability of re-election. Also studying mayors in the Philippines, Labonne (2013) finds that a conditional cash transfer (CCT) program implemented by the national government increased the vote shares of mayors. But Imai, King, and Velasco Rivera (2017) do not find any effect of CCT implementation on incumbent vote shares in two large, randomized CCTs in Mexico.[2]

Another explanation for seemingly irrational voters is that they are in fact irrational. Whether or not voters are rational is an important question because irrationality would distort politicians' incentives to perform well because such performance would not necessarily be rewarded at the ballot box. Research suggesting that voters are irrational has been common in the literature on voter responses to seemingly unrelated (to politics) events. Achen and Bartels (2004) find that voters seem to punish incumbents for shark attacks. Achen and Bartels (2016) make the broader argument that voters are unable to

[2]This study also shows that previous work demonstrating a positive effect of the *Progresa* CCT on incumbent vote share (De La O 2013) was driven by data coding errors or unconventional modeling choices.

make rational, informed decisions due to their punishing incumbents for events outside their control. However, there are reasons that rational voters might behave in these ways. In the case of disasters, voters might punish incumbents for poor relief efforts and reward them for vigorous ones, a proposition with strong empirical evidence (Cole, Healy, and Werker 2012; Bechtel and Hainmueller 2011). There is also theoretical work that shows how rational voters can still respond to seemingly unrelated events, e.g. by using them to infer information about an incumbent's type (Ashworth, Bueno de Mesquita, and Friedenberg 2018). In general, it is difficult to rule out rational reasons for much observed voting behavior, especially under conditions where information acquisition is costly to voters.

Existing explanations for variation in performance voting outside of Africa include partisanship,[3] voter sophistication, and issue salience (Vries and Giger 2014), and education/wealth (i.e. modernization theory, Lipset and Man 1960). It is also important to note that even among rational voters, it is likely that different voters will value different political outcomes and reward different types of incumbent behavior. Some of these voting rationales may be obscured in observable data used by researchers. For example, voters in China value retributive justice and reward politicians who punish corruption and other wrongdoing (Tsai, Trinh, and Liu 2017). Signals such as these about incumbent type can be important to voters but not necessarily observable by researchers, at least unless one is actively looking for them.

[3] Attributing credit for good outcomes (e.g. economic performance) has been shown to vary by party identification, with the implication that party identification therefore affects the propensity of voters to use that information to inform their voting decisions. See Healy and Malhotra (2013, 291–293) for a review of this argument.

2.2 How Public Goods Provision Affects Support for Incumbents

It seems straightforward that the provision of public goods, especially in Africa and other developing regions, could improve incumbents' popularity and electoral support. Public goods are chronically under provided in developing countries and voters seem to want more public goods provision.[4] Theories of retrospective voting predict that incumbents who perform well will be rewarded in subsequent elections (Key et al. 1966; Barro 1973; Ferejohn 1986; Fiorina 1981)[5]. Because voters in many developing countries, especially in Africa, want the government to provide more public goods like infrastructure (Afrobarometer 2019), it stands to follow that increases in public goods provision would lead to greater electoral support for incumbents. Theories of clientelism similarly predict that an increase in more targeted public goods provision (i.e. local public goods or club goods) would result in increased support for incumbents, as voters and politicians are thought to engage in a quid pro quo arrangement. Although some voters may prefer private goods to public goods, this does not preclude the provision of public goods being electorally beneficial to incumbents because many voters *do* want the government to provide public goods. In fact, there is evidence that voters prefer candidates for parliament who promise to use the majority of their constituency development funds on public goods rather than private goods (Ofosu 2021).

But not only is there scant evidence for the relationship between public goods and support for incumbents,[6] there is evidence that the provision of public goods and services

[4]Many respondents in Afrobarometer surveys identify infrastructure, a classic public good, as one of their top priorities for the government to address.

[5]It should also be noted that in some models of voting behavior and some equilibria therein, voters are not expected to vote rationally. But given the totality of the literature, it seems reasonable to expect that voters will use past incumbent performance in their voting decisions in at least some cases.

[6]Harding (2015) is one of the few studies to show a positive relationship.

can lower support for incumbents. In South Africa, De Kadt and Lieberman (2017) find that areas that saw increased public service provision showed decreased levels of voting for the incumbent ANC. In West Africa, Briggs (2018) finds that localities close to aid projects saw a decrease in support for incumbents.[7] These surprising results emphasize the need for more developed theory in this area.

One mechanism through which it has been suggested that public goods might negatively impact support for incumbents is by freeing voters from their reliance on the incumbent. To explain a negative association between service provision and voting for the incumbent ANC, De Kadt and Lieberman (2017) argue that people vote based on their expectations of future service provision from the ruling party. According to the theory, voters in places that do not have services vote for the ANC in hopes that they will be rewarded for their loyalty with service provision, which they consider the primary platform of the ANC. On the other hand, those in areas where services are already provided do not expect to receive future benefits from the ANC, presumably because they expect the ANC to focus on service provision in areas that do not have services yet. Blattman, Emeriau, and Fiala (2016) find evidence for a related effect, where increases in income free Ugandan voters from their reliance on clinetelistic handouts from the ruling party.

These two cases suggest a model of voting behavior that is forward-looking rather than retrospective; once voters get something from the government, they no longer support it. However, elections are repeated strategic interactions, not one-shot games, and these accounts fail to address the concern that the ruling party could punish voters who defect after receiving goods or services. Presumably, there are always going to be things that voters want from the government (e.g. new public goods or maintenance of existing ones), so withdrawing support after receiving just some of them is difficult to explain. The above studies also focused on South Africa, a dominant-party system, and Uganda,

[7]Though the dynamics of aid projects are probably different from those of government projects, many voters do not know who funds or implements aid projects (Baldwin and Winters 2017) and incumbents can claim credit for them (Cruz and Schneider 2014).

a competitive authoritarian one. The logic of the theory might apply in contexts where the current government is the only real contender for political power. But the dynamics may work differently where there are two competitive parties in national elections, as in Ghana. Voters have real options and their votes are more important for deciding elections, providing a stronger incentive to reward incumbents who perform well. That incentive isn't as strong where the outcome of national elections is a *fait accompli*.

Previous work has identified policy attribution as an important requirement for electoral accountability (Harding and Stasavage 2013). Voters must be able to attribute credit for a policy outcome to a politician or level of government in order to reward or punish them for the outcome. Then, for electoral accountability for public goods provision to hold, voters must be able to attribute credit for it to at least one elected official or body with control over the outcome.

Harding (2015) argues that decentralization and multiple foreign funding sources (both of which are present in the context of local public goods provision in Ghana) can make it more difficult for voters to correctly attribute responsibility for policy outcomes, leading to "limited" levels of accountability. The idea here is that where more actors are responsible for outcomes it is more difficult for voters to assign credit correctly than in a situation where only one actor is completely responsible for an outcome.

However, the implication that voters will not be able to hold elected officials accountable in situations of joint responsibility seems weakly supported. Indeed, there are reasons to expect that, at least in some cases, voters and politicians would be able to compensate for the difficulty of assigning credit to multiple actors. Incumbents can engage in credit claiming activities to show voters that they are (at least partly) responsible for an outcome.[8] The extent to which rewarding various actors for a policy outcome is rational will of course depend on the degree to which they influenced it.

[8]This can even result in undeserved credit claiming, where incumbents who had nothing to do with an outcome are nevertheless able to convince voters that they did (Cruz and Schneider 2014).

Where multiple actors are jointly responsible for policy outcomes and multiple actors claim credit for those outcomes, it seems rational for voters to attribute responsibility jointly across multiple actors. If a voter who is pleased with new local public goods provision in her community wants to reward those responsible for it, she might rationally increase her probability of voting for the incumbent president *and* her incumbent local government representative. Other factors will likely also influence her vote choices, so the weight of this one project might be small. But that does not mean her vote is not affected by the project. Further, just because the source of funding for the project might have come from a number of donor funds, the central government, or the local government, it does not mean it is irrational to reward the local and national incumbent for a job well done because they still had some influence over the outcome. It could be argued that in this scenario voters would be giving too much credit for the project to incumbents given that the funding came from elsewhere. But in any case where incumbents have some influence over the outcome, it is rational to assign them some credit. If voters do not know the source of funding for a given local public goods project (which is likely the case for most Ghanaian voters), it may be irrational to expend time and effort to discover the funding source in order to condition one's vote on that information.

If transfers were not used for public goods provision but instead directed towards private goods like cash payments to political supporters or stolen, would it still be rational for voters to reward incumbents for those transfers? Because the national government has a role in securing funding from donors, it would be rational to reward the president's party in presidential elections. But, unless the funds were directed to assembly members who were able to leverage them for other goods that their constituents wanted, it would be irrational (and difficult to explain empirically) for voters to reward them for transfers in local elections because they did not play a role in securing the funds.

Another possibility is that voters in the area immediately surrounding a public goods project vote for the incumbent at higher rates but those in nearby areas that did not receive

public goods projects (or received fewer/less desirable projects) punish the incumbent. This is a version of a relative deprivation hypothesis, where it is only upon seeing a neighboring community receive an increased investment in public goods that voters become unhappy with the level of investment they are receiving. Voters may even resent the government if their own community receives a public goods investment if observing it gives them negative information about the state's capacity. For example, voters who learn that the government is capable of a certain level or type of public goods investment may resent the fact that the previous public goods were not up to the same standards. This is more likely in the case of new funding because it increases funding for public goods provision and voters may not know that local governments did not have access to as much funding before. This increase in public goods provision could cause voters to (potentially incorrectly) update their estimates of how corrupt or poorly performing the government was before the increase.

The level of swing voting in a district should also affect the relationship between public goods and voting behavior. In highly partisan contexts where voters are reliable supporters of a given party, their vote choices in national (partisan) elections will probably not be greatly affected by public goods provision. For these voters, the more likely course of action in the face of seriously unsatisfactory performance by their preferred party is abstention from voting. This would of course still lead to changes in aggregate vote share for the incumbent, assuming the opposition party's voters still turn out to vote. But the impact of a single voter abstaining is half as large as switching to the other party would be. Therefore, the effect of public goods provision on voting in partisan elections should be magnified where there are more swing voters.

The above discussion sets out reasons that voters may either punish or reward incumbents for the provision of public goods. Because voters might respond differently depending on these moderating variables, it is difficult to specify a hypothesis of the average effect of public goods provision on support for incumbents. It is possible that the

average effect is near zero, either because other factors dwarf local public goods provision in presidential elections or because positive effects in some districts are canceled out by negative effects in other districts. However, I present a hypothesis that the average effect will be positive based on the totality of the theory and the particulars of the Ghanaian case (discussed further in Chapter 3. I propose another explanation for heterogeneity in voter responses to public goods provision in the next section.

2.3 Multidimensional Preferences

It is often implicitly assumed in models of distributive politics, especially in developing countries, that voters have similar preferences for all types of goods and services. Where politics are dominated by clientelism, votes are thought to be exchanged for various handouts, goods, and services, but the types are generally not explicitly considered. In contexts where ethnicity strongly predicts voting behavior, it is argued that people vote for co-ethnics or ethnically-aligned parties[9] because they expect to receive patronage or other benefits from them (e.g. Bates 1983). But, again, the exact types of goods and services are unspecified.

But, even in developing countries, voters' preferences are more complex than simply wanting more goods distributed to them; they have preferences over the *types* of goods they receive. That is, they are not indifferent between different types of public goods that could be allocated to them by the incumbent government. In areas with a poor road network, voters' demands for road improvements relative to other types of public goods should be greater than in areas with better roads.

So, if voters who want the government to invest in school buildings instead observe the government building latrines, it signals that the government is not being responsive

[9]Ethnic voting can take the form of voting for parties that are expected to govern with favoritism towards certain ethnic groups rather than voting for co-ethnic candidates. In these contexts, voting for an ethnically-aligned party means voting for the one which is commonly thought to favor one's own ethnic group.

to their preferences (or at least that the government doesn't share their preferences), which can lead to electoral sanctioning. Looking at it another way, different voters have different preferences over the types of goods the government provides and this can help explain heterogeneity in the electoral effects of public goods provision. Indeed,there is evidence of such heterogeneity in voters' responses to educational investments made by local governments in Brazil (Bursztyn 2016).

I argue that voters do vote rationally and retrospectively but that their evaluations of incumbents is less straightforward than has often been assumed. Rather than simply rewarding incumbents for any public goods or service provision, voters evaluate politicians based on how responsively they govern. Voters' calculations on election day, whether they follow a sanctioning or selection model, likely take into consideration how closely the distribution of types of public goods and services provided to their locality matches what they feel should have been provided.

I present a visual example in Figure 2.1. Let us assume that citizens in a given district have identical preferences over two types of public goods: healthcare and education. They are indifferent over the specific proportions of goods they receive along their indifference curve. The district government decides how much healthcare and education to provide in the district, limited by its budget constraint. The most responsive government would allocate the mix of goods represented by point θ, i.e. the point where the budget constraint intersects the constituents' indifference curve.

But a less responsive government may over-provide education relative to healthcare. If it still spent up to its budget constraint, it would provide the mix of goods at point τ'. While this is a situation with no loss from corruption, in contrast to τ, it does not meet the constituents' minimum demands for healthcare. Let us suppose that there is an exogenous increase in education goods provision (e.g. from a central government grant), moving the provision from τ' to τ''. Although constituents prefer more education goods than less, they do not benefit from this increase as much as they would if they had received more

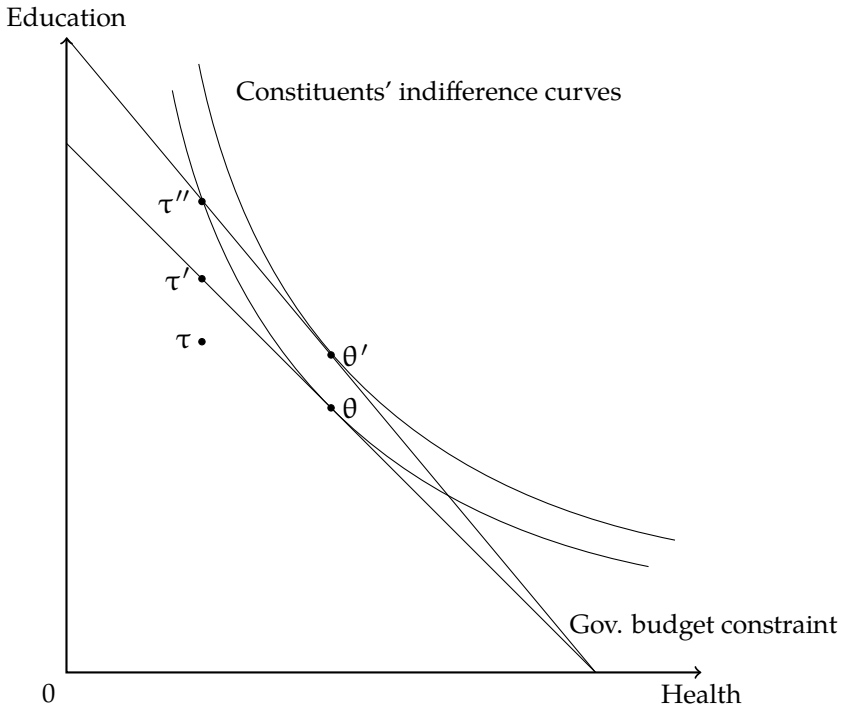

Figure 2.1: Indifference Curve and Budget Constraint

healthcare instead.

Although this simple formalization helps to demonstrate the preferences and available choices for constituents and government respectively, it does not capture what happens after the government chooses a mix to provide: constituents then choose whether to re-elect the incumbent or vote for the challenger. These voters can tell whether the mix of goods provided corresponded with their relative preferences over them. Because the mix of goods provided contains information about how well the government knows the preferences of its constituents, it is rational for voters to use this information when deciding whether to re-elect the incumbent. This leads to the following theoretical prediction: where the distribution of types of public goods provided is more closely aligned with constituents' preferences over those types of goods, incumbents will receive higher vote shares.

Urbanization and Support for Incumbents

Another important variable on which performance voting depends is whether a constituent (or district) is in an urban or rural area. Urbanites in Africa tend to be better informed about politics and less attached to political parties, both of which should make retrospective voting more likely. As noted in the previous section, models of democratic accountability and some empirical results emphasize the importance of information for retrospective voting (Pande 2011). If urbanites have more and better information about government performance, they will be more likely to reward the incumbent party in the event a policy positively impacts them. Rural voters, on the other hand, tend to be more closely associated with a political party, potentially due to vote buying and other clientelistic mobilization strategies (Kuenzi and Lambright 2011). Being more closely aligned with a political party will make these voters less likely to change their vote based on policy choices that negatively impact them.

By contrast, the expectations of urbanites are higher and therefore these voters are harder to placate. Urban voters often live in areas that already have paved roads and other basic public goods. At least, they do at a much greater degree than rural voters. New public goods projects in urban areas must be above a certain quality/level of effort to impress voters enough to affect their vote choice. But in rural areas, where the existing stock of public goods tends to be much smaller and of poorer quality, a relatively cheap or low quality project can be enough to produce an electoral benefit for the incumbent. There is also a lot more going on in urban areas, so many voters might not even be aware of some local public goods projects despite knowing a lot about national politics. These factors may lead to less observed performance voting in the case of local public goods projects and could outweigh any countervailing effects discussed above.

Thus, theory presents us with two opposite predictions for how urbanization affects performance voting. In interviews with Ghanaians, it was more common to hear that the

local development projects I study would not sway voters as much in urban areas. So, while I acknowledge that the effect could in theory go in either direction, I present a hypothesis that performance voting will be lower in districts with higher urbanization.

2.4 Hypotheses

I present several hypotheses based on the theory developed above. First, in line with the conventional wisdom that receiving goods from the government should increase voter's support for incumbents, I propose the following:

Hypothesis 1a: Larger transfers for public goods provision lead to increased support for the incumbent president's party.

Related to this is the proposition that transfers for public goods will increase support for local government representatives. In the case of Ghana it is impossible to link projects with the assembly member in whose constituency the project is located, so I examine whether districts with more funding see average increases in support for incumbent assembly members.

Hypothesis 1b: Larger transfers for public goods provision lead to an average increase in support for incumbent local government representatives.

The effect of public goods provision should depend on voters' preferences. Specifically, the effect will be larger where voters have relatively stronger preferences for public goods projects (compared with other government spending/policies).

Hypothesis 2: Where voters have stronger preferences for public goods provision, transfers for public goods will lead to a greater increase in support for incumbents than in places where voters care more about other aspects of governance.

Finally, because I theorize that urbanites will be less responsive to public goods provision than those in rural areas, I propose:

Hypothesis 3: Where the level of urbanization is higher, the effect of transfers for public

goods provision on support for incumbents will be lower.

Chapter 3

The Context

3.1 Case Selection

Ghana represents a good context in which to study democratic accountability for several reasons. First, and most importantly, Ghana is a democracy. Many studies on democratic accountability in Africa concern countries that are not true democracies. A particularly common setting for such studies is Uganda, where President Yoweri Museveni shows little sign of relinquishing the office he has held since 1986. Yet a plethora of research on democratic accountability has been conducted there in recent years (e.g. Dunning et al. 2018; Grossman, Humphreys, and Sacramone-Lutz 2014; Blattman, Emeriau, and Fiala 2016; Carlson 2015). The lessons learned from those studies must be taken with a grain of salt when applied to democratic contexts.

Ghana, on the other hand, is a liberal democracy known for its relatively long experience with democratic elections compared to most other countries in Africa. The current era of multiparty democracy was inaugurated in 1992, with free and fair elections taking place every four years subsequently.[1] There have been three peaceful electoral transfers of

[1]There is some disagreement over how credible the 1992 elections were. Jeffries and Thomas (1993, pg. 331) note that international observers declared the election mostly free and fair, stating that voting irregularities were not serious enough to impact the result. But Gyimah-Boadi (1999) criticizes the administration

power from one party to another in the current era (2000, 2008, and 2016). The external validity of research on accountability in Ghana should be relatively high in regards to other democracies.

In addition, it is not as if Ghana can only teach us about African or developing democracies, as its political context shares some features commonly found in older democracies. For example, its party system is highly institutionalized. Ghana's politics are dominated by two main parties, the New Patriotic Party (NPP) and the National Democratic Congress (NDC), which have been the only two parties that have come close to winning the presidential election in the current era of multiparty democracy.[2] In contrast, most other African democracies have less stable party systems (Riedl 2014, 41), so studies of democratic accountability in them are fundamentally different (e.g. there should be lower accountability without parties to reward/punish).

Although some features of Ghana's political system are similar to non-African democracies, there are aspects of it that are more typical of African ones. For one, Ghana has a strong presidential system, where the majority of policymaking authority rests in the executive. There is a parliament, though it is mostly a rubber stamp for the president's policies.[3] Because most Ghanaians vote for the presidential and parliamentary candidates from a single party (straight ticket voting),[4] and because elections for president and parliament are held concurrently, the president's party has always held a majority in parliament.[5]

---

of the election, saying that the rushed preparations and other challenges weakened the credibility of the process. In any case, the election was not seen as legitimate by the majority of the opposition, leading to a boycott of the parliamentary elections the following month.

[2]While there are several minor parties, they never garner more than a few percentage points in the presidential election and a few seats in parliament.

[3]Though parliament does provide the minority party enough institutional power to level critiques and provide some oversight of the executive branch.

[4]Though a large minority of voters (known colloquially as skirt and blouse voters) split their tickets each election year.

[5]Though the 2020 elections produced an even split in parliament, with the re-elected president Akufo-Addo's NPP having a one seat majority after including an independent candidate who will caucus with them.

There are also research design considerations that make Ghana a good case for studying democratic accountability. There are two funding mechanisms which produce pronounced variation in local public goods provision. Additionally, these mechanisms rely on transparent criteria for qualification (i.e. local government performance measures) as opposed to a situation where the incumbent president was able to completely control allocations to local governments. Both funding mechanisms were partially or wholly funded by international donors, leading to more stringent requirements and restrictions for qualifying for funds and putting them to use on public goods projects. These factors should lead to much less political influence in which local governments provide more public goods, leading to less of a concern of confounding bias resulting from political factors causing both levels of public goods provision and electoral outcomes for incumbents.

3.2 Local Government in Ghana

Ghana's system of local government traces its earliest origins to pre-colonial period, when traditional authorities (chiefs) ruled many localities in what is now Ghana. The British engaged in indirect rule, allowing existing chiefs substantial control over localities. The British also "invented" chieftaincy in many other areas, installing chiefs where there had been none previously (Nathan 2019a). Several changes to this system were implemented in the post-independence period, most importantly shifting power away from traditional authorities, culminating in the Local Government Act of 1993 that established the current system.

As of today, Ghana is divided into local governments known as metropolitan, municipal, and district assemblies (MMDAs). Metropolitan assemblies are found in the largest urban areas and have populations greater than 250,000. Municipal assemblies tend to contain at least one large town and must have at least 95,000 people. District assemblies are more

But the speaker of parliament is from the NDC, so some power sharing will take place.

rural and have populations under 75,000. MMDAs are often collectively referred to as districts for convenience (a practice which I adopt here). Districts are roughly analogous to counties in the US system, especially outside of the largest metropolitan areas. In most of the country, a district can cover a wide area composed of many towns and villages (though the largest urban areas are composed of multiple districts each). Some districts are exactly congruous with a parliamentary constituency, while others contain multiple constituencies. No district contains only part of a constituency.

Districts are only permitted to collect a few types of minor taxes, making internal revenue generation a perennial challenge. They rely on transfers from the central government for the vast majority of their operating funds. The primary intergovernmental transfer mechanism is the District Assemblies Common Fund, through which the central government is constitutionally required to disburse 5% of national revenue to districts. The precise share given to each district depends on a frequently-changing formula determined by parliament. Another transfer fund is the donor initiative District Development Facility (DDF), discussed in detail in Chapter 4.

Each district is governed by a chief executive and an assembly. The municipal, metropolitan, and district chief executives (MMDCEs[6]) are appointed by the president of Ghana and have a great deal of influence over the districts' priorities (e.g. which infrastructure projects to undertake). At least 70% of the assembly members are elected in single member districts, while at most 30% are appointed by the president. Assembly elections are officially non-partisan, though candidates' political loyalties are often known in the community. So, the majority of voters I interviewed were able to tell me the political party of their assembly member. Assembly members are not paid a salary or stipend for their service, though some have received motorbikes (ostensibly to enable them to more easily connect with their constituents).

[6]Districts have district chief executives, or DCEs, while municipalities and metropolis have MCEs (municipal or metropolitan chief executive).

Although MMDCEs are the most influential actors in determining which projects get funded, the assembly also has some formal and informal influence (Williams 2017). This system allows for some electoral input by citizens while maintaining substantial central government control over local governments. Partial decentralization of power like this is common in Africa, where leaders engage in rhetoric that emphasizes the importance of decentralization to development outcomes but authorities have been hesitant to relinquish too much power. Ghana was even classified as being on the high end of political decentralization by one study (Ndegwa 2002).

Public support for further decentralization in Ghana is strong, with 69% of Ghanaians supporting direct elections of MMDCEs over presidential appointment (Afrobarometer 2017). Ghanaian voters I interviewed in the Central Region of Ghana also tended to support direct elections of chief executives.[7] However, legislation to enact such a reform was scrapped in December 2019 by the incumbent NPP president Nana Akufo-Addo. This came as he was forced to cancel a referendum which would have allowed political parties to sponsor candidates for chief executive and district assembly seats. The opposition National Democratic Congress (NDC) party came out against allowing political party involvement in local elections, which made the referendum's success much less likely. And many voters I interviewed were not favorable towards political party participation in local elections.[8] One reading of Akufo-Addo's actions is that without partisan participation in local elections, he was unwilling to cede his power to appoint chief executives because that could have allowed independent candidates to win seats that his NPP would have otherwise been able to capture by sponsoring their own candidates. At this point the future of decentralization

[7]I asked voters the question "Should district chief executives be directly elected by the people?" to which the vast majority responded yes. Their reasons included poor performance by appointed DCEs and that the president could protect poorly-performing DCEs.

[8]I asked voters: "Do you think district assembly elections should be non-partisan?" and most agreed that they should remain non-partisan. They tended to not be confident in assembly members and DCEs from different parties working together, though one respondent noted that even in officially non-partisan elections candidates have their preferred parties.

in Ghana remains murky.

The creation of new districts has occurred every four to six years since 2004, often in election years.[9] Ayee (2013) argues that this process is political, in part because it allows the incumbent president to deliver jobs to political allies in the newly created districts. There are other benefits to areas that become part of new districts. They see increases in investment due to the construction of offices and other buildings in the new district capitals. And since a portion of intergovernmental transfers is allocated equally to all districts, these areas receive more funds for local public goods provision than they would have as part of their parent district.

Politics in Africa and Ghana

Although many Africans had high hopes for self-governance upon gaining their independence, they were soon disappointed by the corrupt and authoritarian regimes that dominated the continent. It wasn't until the early 1990s that the tide would turn in a significant way towards democracy on the continent, with numerous countries transitioning to democratic systems (including Ghana) (Bratton and Van de Walle 1997). Nevertheless, democracy has not delivered everything African citizens have hoped for. Corruption has remained a major concern across the continent, and many Africans remain dissatisfied with their political leaders (see e.g. Bratton, Mattes, and Gyimah-Boadi 2005, 232-239). More recently, several democratic states have faltered, experiencing coups, leaders clinging to power longer than they were allowed, or restrictions on the opposition.[10]

Regardless of a country's regime type, African politics is often characterized as being dominated by ethnic voting and clientelism. The role of ethnicity in African politics has been emphasized by scholars for decades (e.g. Young 1976; Bates 1983; Horowitz 1985). A

[9]The number of districts increased from 138 to 170 in 2008, 170 to 216 in 2012, and 254 in 2018. Presidential/parliamentary elections took place in 2008 and 2012.

[10]The countries I have in mind include Mali, Cameroon, Benin, and Côte d'Ivoire.

robust finding that continues to find support in empirical work today is that Africans in many countries tend to vote for politicians of the same ethnic group (Dunning and Harrison 2010; Arriola, Choi, and Gichohi 2017). Even studies which argue that ethnicity is not the only factor at play find that it exerts a strong influence on vote choice (Bratton and Kimenyi 2008; Bratton, Bhavnani, and Chen 2012). Ethnic voting appears to occur for instrumental reasons: people vote for co-ethnics because they expect to receive some sort of material benefits from them (e.g. Bates 1983). Politicians cannot credibly commit to post-election actions and ethnicity is the best available signal of a candidate's likelihood of "delivering the goods" once elected (Posner 2005). Emphasizing this finding and refining it, Ichino and Nathan (2013) show that Ghanaians who live in areas dominated by a non-co-ethnic party are more likely to vote along with their neighbors rather than for their own co-ethnics' party.

Clientelism, the other major theme in work on African politics, is the direct exchange of goods or services for votes (Kitschelt and Wilkinson 2007, 2).[11] Clientelism is thought to be harmful to democratic accountability because it is a system where voters surrender their ability to vote based on incumbent quality, performance, etc. (Hicken 2011), instead agreeing to vote according to the wishes of their patron (or their brokers, e.g. Stokes et al. 2013). This accountability relationship is thus reversed: voters are held to account by politicians, who withhold goods and services if they do not receive enough electoral support. Contrast this with the democratic accountability perspective, which theorizes that politicians are accountable to voters, who condition their votes on the politicians' policy performance.

There is, however, evidence that African voters care about policy performance and do not simply vote based on who gives them handouts. Harding (2015) shows that changes

[11]Clientelism is a broad term that includes patronage, or the exchange of votes for access to state resources (Hicken 2011). Although patronage was once emphasized more than clientelism in the study of African politics (e.g. van de Walle 2007), I refer to clientelism because it is the umbrella term for a targeted disbursement of goods to voters contingent on electoral support.

in road quality in Ghana are positively associated with changes in incumbent vote share. Weghorst and Lindberg (2013) argue that MP performance influences Ghanaian support for incumbent MPs. However, their survey-based research design cannot rule out an omitted variable influencing both voters' evaluations of incumbent performance and their propensity to vote against them. Other results claiming similar findings suffer from other empirical problems (e.g. Bratton, Bhavnani, and Chen 2012, which simply asks voters what motivates their decisions).

Recent, ongoing work in Ghana examines how local need affects public goods provision. Ichino, Williams, and Wibbels (2018) leverage a large data set of local government development projects to examine, in part, whether areas within districts that have a greater need for development receive more projects. Although they do not find a relationship between need and projects, there are areas where their analysis could be improved. First, because it is preliminary, they only examine three districts. Second, one of the two measures of need they use (the share of households with electricity) is not closely related to public goods that local governments provide, as the national government is charged with electricity provision. I improve on the above study by including districts from all of Ghana and measuring preferences in ways that more closely correspond to potential local government outputs.

3.3 The Importance of Public Goods Provision in Ghana

Local governments in Ghana are tasked with providing a great deal of local public goods in their jurisdictions. Typical projects include building classrooms, public latrines, drainage/irrigation, and repairing minor roads. But these same governments are only allowed to collect a few minor taxes, so their internally-generated funds (IGF) are perennially too low to fund more than a couple projects (IGF makes up only about 17% of districts' revenue (Debrah 2014)). Their main source of funds is the District Assemblies Common Fund (DACF), a

formula-based intergovernmental transfer allocated to all districts. The DACF distributes 5% of national revenue among the various districts according to an ever-changing formula, amounting to about 40% of districts' revenue.

Another major source of funds is the "donor-partners" who give money to districts, e.g. the World Bank and bilateral aid donors. The two programs I focus on in this dissertation, the DDF and UDG (discussed in more detail in Chapters 4 and 5), are examples of donor-partner funding. These funds make up around a third of district's revenue, though there is significant variation over time and across districts depending on funding levels, district demographics, etc. Despite these various sources of funding, public goods are so desperately needed across most of the country that not even all of the most basic priorities are implemented in a given district. Local governments are therefore tasked with prioritizing some projects and localities over others.

Interviews with development planning officers and other district bureaucrats show that the needs and preferences of different communities are taken seriously, at least at the early stages of planning. Every three years, districts produce a medium term development plan which outlines the priorities and projects it will focus on for the next three years. In preparation for this report, a district holds meetings with stakeholders in communities across a given district. Invitees include the local traditional authority (a chief or their representative), trade associations, churches, and, importantly, the local assembly member(s). Attendees are asked about the main priorities for investment in their community and their preferences are aggregated to produce a ranking of projects. The aggregation process may depend on the district, but in one case a district official described having attendees break into groups to brainstorm and rank projects, after which the top projects from each group were listed and voted on by all attendees.

These stakeholder meetings allow assembly members and chief executives to learn about the preferences of their constituents. This is important because it establishes one link between what public goods voters want and what their government provides. Of course,

there are other ways that incumbents can learn about voters' preferences in Ghana. People are not shy about criticizing the government's performance and voicing their preferences. This occurs frequently on social media and call-in shows on TV and radio programs. At a very local level, responsive assembly members discuss neighborhood issues with their constituents on an ongoing basis.

3.4 Voting Behavior in Ghana and Africa

It is important to discuss the extant literature's treatment of voting behavior in Ghana, as I seek to build on it rather than supplant it. I begin by discussing the role of ethnicity in structuring vote choice, a frequently studied and cited factor in both Ghana and many other African countries. Next, I review how Ghana's high levels of partisanship make vote patterns highly predictable. Finally, I discuss existing research on performance-based voting in Ghana.

Ethnicity

There is some disagreement in the literature over the main determinants of voting behavior in Ghana. The role of ethnicity has been emphasized by some scholars (e.g. Ichino and Nathan 2013; Nathan 2016) and is a common source of voting preferences cited by Ghanaians. Some scholars have even claimed that ethnic voting in Ghana is as high as it is in Kenya, where ethnicity is highly predictive of vote choice. It is true that members of Akan ethnic groups tend to vote for the NPP while those of the Ewe, and northern ethnic groups are more likely to choose the NDC.[12] However, a closer look at the data and empirical work to date reveals that the role of ethnicity in Ghanaians' voting behavior has been exaggerated.

[12]It has been argued that the Ga ethnic group also tends to favor the NDC. However, Afrobarometer survey data suggest that the Ga are not reliably pro-NDC, being roughly evenly split between the NDC and NPP leading up to the 2016 election.

The way ethnic voting is measured in studies of Ghana's politics is one reason it has been overestimated. Ethnic voting is often thought of as voters choosing candidates who share their own ethnic origins. But Nathan (2016) operationalizes ethnic voting as voting for one's "ethnically-aligned party." But it is very difficult to tell whether voters who vote for their ethnically-aligned party do so due to shared ethnicity with many of the party's candidates, expectations of ethnically-targeted future benefits from the party, the performance of the candidate/party, or simple partisanship. Further, although Ghana's two main parties are colloquially thought of as ethnically-aligned, the reality is more complicated. For example, even though the NPP is the party associated with the Akan group of ethnicities, the second and third most senior members of the NDC were Akans in 2019. And some ethnic groups do not even have an "ethnically-aligned party."

One way of measuring ethnic voting systematically is by measuring the proportion of people whose votes can be predicted solely based on their ethnicity. I carry out this exercise using data from Round 6 of the Afrobarometer survey conducted in Ghana. I first count the number of voters who voted with their ethnically-aligned party, using the ethnic group-party pairs defined by Nathan (2016). I then drop all respondents who did not answer the vote intention question. The denominator is then the total number of respondents who declared a vote intention (for any candidate from any party, including independents and minor parties): 1,563. Of those voters, 1,018, or 65%, said they would vote for their predicted party. Akans are the largest group and also the most inclined to vote ethnically in this sample at 76% for the NPP. Seventy-two percent of Ewe respondents said they would vote for the NDC. But the northerners and Ga were split much more evenly, despite being considered NDC partisans.

Another way to measure ethnic voting is by explicitly fitting a predictive model of vote choice using solely ethnicity information. First, I split the Afrobarometer Round 6 data for Ghana into training and test samples to get an out of sample measure of accuracy. Next, I fit a logistic regression model on the training data, regressing vote choice (vote NDC or

vote NPP) on ethnicity dummy variables. Finally, I apply this fitted model to the held out test data to measure its predictive accuracy. I find that the model correctly predicts about 77% of votes based on voter ethnicity alone. However, this model only predicts vote choice, not whether someone votes for their ethnically-aligned party. Because survey respondents did not necessarily tend to vote with their ethnically-aligned party, this predictive model was more accurate than simply counting how many did vote with their ethnically-aligned party.

This exercise shows that while ethnic voting is certainly common in Ghana, it is not as widespread as one might expect. If there were no ethnic voting whatsoever, people from a given ethnic group would vote for each of the parties at the same rate (conditional on party performance and other factors that would affect the parties' popularity with all voters). So, for instance, Ewes would vote for the NPP and NDC at about 50% each. Knowing a voter is Ewe would give us no information about their vote choice, so if one predicted that Ewes would vote for the NDC or NPP, they would only correctly predict about 50% of votes. If ethnicity perfectly predicted vote choice, of course, we would be able to predict 100% of voters' choices. Therefore, we can think about ethnic voting potentially ranging from about 50% to 100% in Ghana. The above estimates of 65% and 75% are 30% and 50% of the way to 100%, indicating modest but not overwhelming ethnic voting.

To further benchmark this level of ethnic voting, I conduct a similar prediction exercise for the United States. I use data from the American National Election Studies' 2016 pre-election survey to predict support for candidate Trump. The dependent variable equals 1 for respondents who say they will vote for Trump, 0 for those who say they will vote for Clinton, and NA otherwise. The independent variable is a categorical race variable with possible values of Asian, Black, Hispanic, Native, White, and other. I fit a logistic regression model using a training data sample (85% of the data set) and evaluate its predictive accuracy using a test set (15% of the data set). The model correctly predicts 61% of votes using race alone, somewhat lower than the 75% figure obtained for Ghana (though

the Ghana data has more detailed ethnic categories, which could account for some of the increased precision). This indicates that ethnic voting in Ghana is only somewhat more prevalent than it is in the US, a country where race/ethnicity is important in structuring vote choice but far from deterministic.

Other Factors

Voting behavior in Africa is not immune to the effects of partisanship on structuring long-term voting patterns. Driscoll (2015) argues that partisanship in Ghana is so strong that changes in aggregate voting over time are mostly due to differential partisan turnout rather than swing voters changing their votes. While this likely understates the importance of swing voting, it is true that most voters in Ghana will only ever consider voting for a single party.

It is also instructive to consider work on African voting behavior writ large rather than solely focusing on Ghana. In a cross-national study based on the Afrobarometer survey, Bratton, Bhavnani, and Chen (2012) emphasize the role of partisanship in voting, though it is unlikely that their work produced an accurate and unbiased estimate of the effect of partisanship.[13] Still, their study raises the possibility that some African voters are just as strongly partisan as voters in other regions. Indeed, this is consistent with the case of Ghana, where most voters choose the same party each election. But it is difficult to pull apart the effect of voting based on ethnicity and partisanship and there does not appear to be work that has done so yet in Africa. Nevertheless, partisanship is popularly understood to be quite high in Ghana. Although there are some ethnic voting patterns, there are reasons to believe that partisanship exerts strong effects on voting behavior independent of ethnicity.

Other work in Ghana suggests that the performance of the national government matters

[13]They estimate the probability of a voter supporting the ruling party based on whether she felt close to the ruling party, which cannot rule out reverse causation and omitted variables bias.

for presidential elections. Harding (2015) shows that changes in road conditions are positively associated with changes in vote share. Weghorst and Lindberg (2013) also provide some evidence that voters care about performance when choosing whom to vote for. But it remains unclear how much performance matters compared to ethnicity and partisanship and whether performance voting is increasing or decreasing over time. If Ghana's democracy is maturing and voters are becoming more discerning, learning the policies and performance of the two main parties over time, it would suggest an increase in performance voting over time. But if partisanship is increasing over time, voters may pay less attention to performance. Encouragingly, the opposition NDC gained seats in parliament despite losing the presidential election in 2020, suggesting an increase in split-ticket voting and therefore less reliance on partisanship for some voters.

3.5 Theoretical Considerations and the Ghanaian Context

The case of local public goods provision in Ghana presents some factors which influence the theoretical predictions for the political effects of public goods provision. Perhaps most importantly, there are multiple levels of government involved in local governance outcomes, making it difficult to predict which level of government voters should hold accountable for the performance of their district assembly. In Ghana, local governance outcomes are determined jointly by the national government, controlled by the president and their party, and local governments, headed by presidential appointees but with elected assemblies. The national government exerts influence over outcomes in local governments through several channels. For one, changes in the allocation formula for a major intergovernmental transfer (the District Assemblies Common Fund or DACF) are proposed by an administrator who is appointed by the president and must be approved by Parliament, which has always been controlled by the president's party. In addition, the national government signs agreements with donors to provide additional funding to local governments, i.e. through the District

Development Facility (DDF) and Urban Development Grants (UDG).[14]

Finally, and perhaps most importantly, the president appoints the chief executive (mayor) for all local governments in the country. The chief executive plays a major role in administrative and policy outcomes in assemblies, e.g. by determining which projects get funded. There is also an indirect channel through which chief executives affect policy outcomes: administrative performance affects how much funding for public goods projects is available to a given assembly through the DDF and UDG funding programs. But because voters do not elect their local chief executives, they have limited ability to hold them to account for their performance. Informally, voters have staged protests against their local chief executive due to poor performance or corruption with the aim of convincing the president to appoint a replacement. But the only formal accountability mechanism for chief executives is through voting in presidential elections.

Importantly, voters are aware that the president is responsible for appointing local chief executives. The media covers appointments of chief executives by the president, protests against unpopular executives, and some stories of local advocacy for particular individuals to be appointed in the future. Chief executives are also visibly associated with the president's party, e.g. by wearing the party's colors or lauding the incumbent president's efforts during public appearances.

But there are also reasons to expect that intergovernmental transfers for public goods provision affect voting behavior more in local than presidential elections. Most importantly, there are a multitude of issues that are salient in a given presidential election, while there are presumably fewer for district elections. The national government is responsible for many more outcomes than local governments, even in rural areas. This would suggest that local public goods provision has a larger effect in district elections than in presidential elections.

Even though intergovernmental transfers are channeled to officially non-partisan local

[14]The national government also contributes its own matching funds to the DDF program.

governments, there are reasons to believe that the funds may be partially intended to increase the incumbent party's electoral prospects. Banful (2011) finds evidence that the main intergovernmental transfer in Ghana (the DACF) is targeted to districts that had lower vote margins in the previous year's presidential election, which they suggest is evidence of targeting to swing voters. This implies that incumbents believe increasing transfers will lead to higher vote margins for their party in national elections, so it is not inconceivable that they do.

While DDF projects are relatively small scale, they are large and expensive enough that most assembly members will not see one carried out in their electoral area in a given term. When looking at public goods provision and support for incumbents at the district level, which aggregates many electoral areas, it may be difficult to detect effects. Even if DDF projects increase support for the assembly members in whose electoral areas the projects are built, they could have null or negative effects on support for assembly members in non-project districts. Thus, I expect that estimated effects at the local level to underestimate the true effect of DDF projects on incumbent support.

A difference between many developing countries versus developed ones is the relative absence of ideological content in policy preferences and party's governing priorities. There is no meaningful left-right continuum in most African countries, including Ghana. For instance, Ghana's supposedly pro-business party (the NPP) abolished school fees for secondary schools in Ghana, not the party that is sometimes thought of as more left wing (the NDC). Further, when looking at the local level, ideology is less relevant in any context. The local issues that matter most to voters usually relate to the quality of public goods and services, even in a more ideologically-divided context like the US. All this is to say that when voting in local elections or based on local governance outcomes, voters' preferences in most African countries are not likely to contain much, if any, ideological content.

While ideology shouldn't play a role in preferences over local governance in Ghana and many voters say they would be happy with any sort of public goods provision, voters

do want certain goods and services more than others. This can be seen in public opinion surveys, for example, where there is a diversity of responses to the question of what is the most important problem facing the country that the government should address. Voters also express opinions about the relative importance of different public goods projects in community meetings held in preparation for districts' medium term development plans.

I argue that the main source of these preferences is the availability and condition of local public goods. In interviews, Ghanaian voters often referenced the existing condition of public goods when describing their preferences for future public goods provision. There is a substantial amount of internal migration and travel in Ghana, so voters are aware of how their communities' public goods compare to those in other communities. Thus, voters are aware of the condition of public goods and express preferences for new goods at least in part based on this knowledge.

Chapter 4

Rewarding Public Goods Provision: The District Development Facility

4.1 The District Development Facility (DDF)

The District Development Facility (DDF) program in Ghana represents a unique opportunity to study the electoral effects of public goods provision. The DDF is a special intergovernmental transfer whose funds are required to be used for infrastructure projects. It is co-funded by Ghana's central government, which contributes about half the funds, and several governmental aid donor agencies. The agencies participating have changed over time, but have included AFD (France), DANIDA (Denmark), and KFW/GIZ (Germany). There were seven rounds of assessments with subsequent disbursements in the period under study (2008-2016), referred to as FOAT I–FOAT VII after the name of the assessment used to evaluate districts (the Functional and Organisational Assessment Tool).

DDF funds are only available to local governments that meet minimum administrative performance criteria before receiving funds. These criteria have changed over time, but have centered on planning, accounting, procurement, and capacity. For example, in all years districts were required to have no audit reports that indicated significant financial

irregularities. Another requirement was that district assemblies actually had to convene for meetings and take minutes. Other requirements have included preparing an annual action plan detailing planned policy activities and submitting progress reports on the implementation of annual action plans.

The specific formula that determines the allocation of funds to eligible districts has also changed over time. The Ministry of Local Government and Rural Development (MLGRD) was unwilling to provide documents showing exactly how the formula was applied in specific years. However, through an interview with the bureaucrat at the ministry responsible for implementation of the program I learned about the formula's components and how they changed over time. In all rounds, the funds to be disbursed were allocated based partly on the performance of local governments on a variety of indicators related to planning, administrative capacity, transparency, and other administrative performance criteria. The proportion of funds allocated based on performance increased from 40% in FOAT I to 50% in FOAT II, then to 68% for the remaining rounds. The other factors used for FOAT I–FOAT III were district population and land area, with some portion of funds allocated equally to all qualifying districts. Beginning with FOAT IV, 5% of the funds were allocated based on the number of poor people living in each district.

Ghana's districts almost exclusively undertake relatively small projects. So, districts build classroom blocks not schools, for example. A photo of one DDF project, paving the lorry station in the Ekumfi District in Ghana's Central Region, is shown in Figure 4.2. A breakdown of the types of projects typically built using DDF funds is shown in Figure 4.1. The data come from Williams (2017), who collected annual progress reports submitted by individual districts that listed ongoing and completed development projects. Although this data is partially incomplete, it provides the best available view of what types of projects are built using DDF funds. Education projects are by far the most common types of projects at around 35% of projects, most of them new classroom blocks (buildings). Water, sanitation, and hygiene (WASH) projects constitute a significant portion as well, though the rate

is about half of education projects. Administrative projects include the construction of buildings for district offices and for housing the district chief executive. Several other categories of projects are less common, including construction and repairs of roads, market stalls, and streetlights.

Figure 4.1: Distribution of DDF Project Types

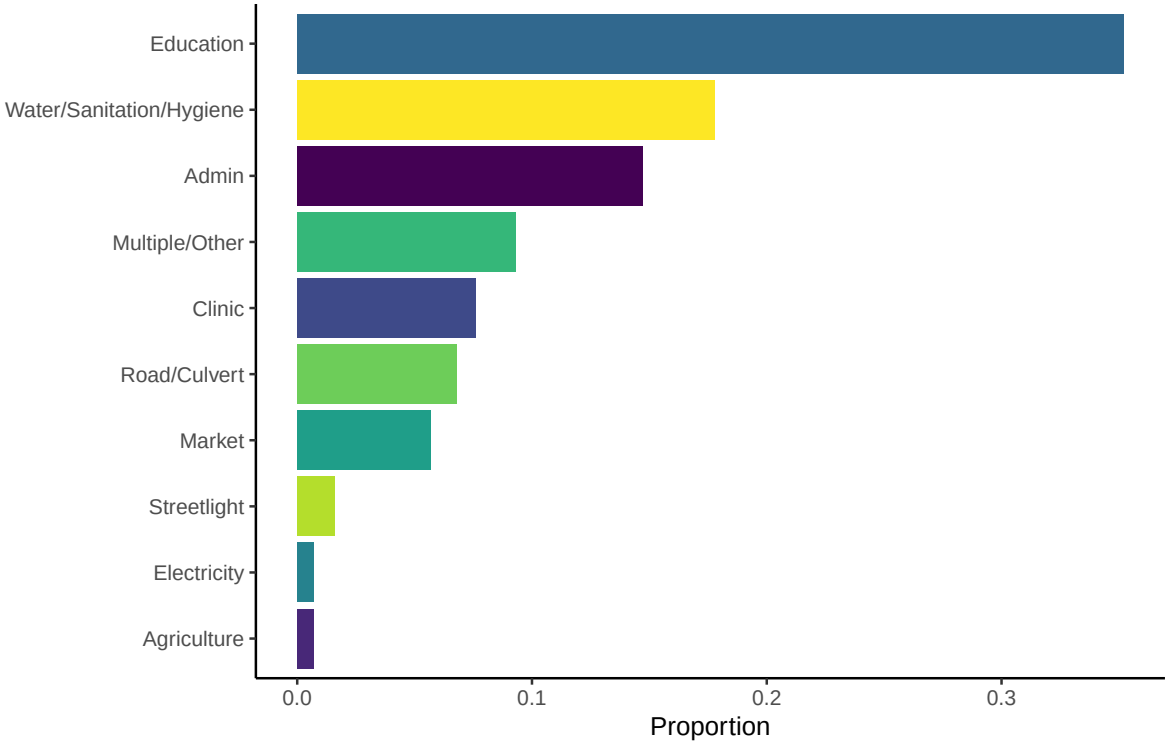

Note: Includes only DDF-funded projects reported by districts in annual progress reports between 2011-2013. List of reported projects was de-duplicated prior to calculations. Original data are from Williams (2017). Projects from other years and from districts that did not report projects are not observed.

The DDF is not the largest intergovernmental transfer in Ghana, [1] but it is advantageous for the present study. Unlike many other intergovernmental transfers, DDF funds must be used to fund local public goods provision. This means that larger transfers result in

[1]It is difficult to estimate the share of the average district budget that DDF funds represent because this data hasn't been previously compiled. In addition, the proportion varies by year. To give two examples: for Mfantseman Municipal assembly, a relatively large, urban district, DDF funds represented 7.7% and 3.8% in 2012 and 2013, respectively. For the more rural Gomoa West district, these figures were 17.1% and 13.6%.

Figure 4.2: A DDF-Funded Public Goods Project

Note: DDF-funded lorry station paving project in Ekumfi District Assembly, Central Region (author's photograph, 2019). Inset is the project's sign, which sits at the edge of the station.

more public goods,[2] which makes transfers a good proxy variable for public goods. This is important because it can be difficult to directly measure public goods provision, especially in Ghana. For example, Harding (2015) only looks at a sample of road condition reports (and only one public good–road improvements) and Williams (2017) can only observe public goods for districts that reported them. By using transfers earmarked for public goods provision, I am not limited by a single type of good or by individual districts' reporting quality.

4.2 Empirical Approach

To measure the effect of public goods provision on support for incumbents, I first estimate the effect of transfers from the DDF program on voting in presidential elections. The

[2]Although at the beginning of the program it appears that it took local governments some time to ramp up construction using DDF funds, there do not appear to be lingering capacity limitations. Indeed, districts often receive fewer funds than they budgeted for and have to postpone projects.

primary threat to identification is the concern that districts that receive larger DDF transfers are systematically different from those that receive smaller ones and that these differences are also correlated with presidential vote share. To address this concern, I estimate a model of the form:

$$y_{it} = \alpha_1 + \beta_1 D_{it} + \gamma_{1t} + \phi_{1i} + \epsilon_{1it} \tag{4.1}$$

where y_{it} is the vote share for the incumbent NDC president in district i in the election concluding term t. The level of DDF transfers is given by D_{it}. The coefficient β_1 captures the effect DDF transfers on incumbent vote share. A term-specific intercept γ_{1t} is included to capture the average difference in vote share between the elections across all districts. A vector of district-specific dummy variables ϕ_{1i} captures the average vote share in individual districts, leaving no possibility for time-invariant district characteristics to bias the estimation of β_1.

The primary identifying assumption here is that there are no time-*variant* characteristics of districts that are correlated with both DDF funding and presidential vote share. One example is if districts that were becoming less supportive of the incumbent anyway were targeted with more DDF funding, perhaps to try to shore up losses. This would exert a downward bias on the estimate of the coefficient β_1. Although this is certainly possible, it seems unlikely that the bias would be very large. Qualification for DDF funds is determined by district performance, and allocation among qualifying districts is done according to a formula. The formula has changed over time, but even if the incumbent government were able to target certain districts through formula changes the differences between different allocation schemes shouldn't produce drastically different patterns of allocation.

It seems even less plausible that districts that were becoming more supportive of the incumbent would be targeted for increased DDF funds since any reward for supporting the incumbent would likely come only after an election, when an area had already delivered

their votes.

For estimating the effect of DDF projects on support for local government incumbents, I take two approaches. In the first, I conduct a cross-sectional analysis using results from the 2015 local assembly elections. This allows me to compare the results using two measures of incumbent support: the proportion of incumbents re-elected of those who ran for re-election, and the proportion of all seats won by incumbents.

I also estimate models using two-period panel data (from the 2015 and 2010 local elections). These models take the form:

$$y_{it} = \alpha_2 + \beta_2 x_{it} + \gamma_{2t} + \phi_{2i} + \epsilon_{2it} \tag{4.2}$$

where y_{it} is the proportion of incumbents who were re-elected in district i in the local assembly elections at the end of term t. Time-invariant district characteristics are captured by the vector ϕ_{2i} and term-level differences in re-election rates across districts are captured by γ_{2t}.

The primary identifying assumption is similar to that for presidential elections, i.e. that there are no time-variant differences among districts that affect both their levels of DDF transfers and the re-election rates of incumbent assembly members.

To test the hypothesis that the effect of DDF transfers on support for incumbents is greater in more urbanized areas, I interact a level of urbanization measure with the DDF funding variable in some cross-sectional specifications. I calculate the urbanization variable using individual-level data from the 2010 Ghana Census (a 10% sample of all individuals in Ghana in 2010 representative at the district level). It is simply the proportion of a district's respondents living in an urban area. I use the Ghana Statistical Service's definition of an urban area here (localities with over 5,000 inhabitants), though it is an imperfect one.

Table 4.1: Summary Statistics–Presidential Voting

	N	Mean	SD	Min	Max
Incumbent (NDC) vote share	339	0.52	0.19	0.10	0.96
DDF transfers (100,000 cedis)	340	20.23	11.43	1.06	115.36
Real DDF transfers	340	23.34	11.91	2.02	126.76
Logged real transfers	340	3.05	0.45	0.70	4.84
Real election year DDF transfers	340	7.15	4.67	0.60	27.35
Real avg. yearly transfers	340	6.97	4.02	0.67	42.25
Logged real avg. yearly transfers	340	1.81	0.52	-0.40	3.74
District population (Thous.)	340	145.08	208.01	47.88	2029.60
Urbanization (prop.)	340	0.36	0.24	0.00	1.00
Median Age	340	18.83	1.96	14.00	24.00
Adult literacy rate	340	0.61	0.21	0.16	0.92
Internal migration (prop.)	340	0.30	0.13	0.05	0.77
Prop HHs engaged in ag.	340	0.64	0.24	0.03	0.96
Open defecation rate	340	0.30	0.30	0.02	0.93

4.3 Data

Summary statistics are shown in Table 4.1. One dependent variable of interest is the proportion of the two party vote in a district earned by the incumbent presidential candidate (the NDC's John Mahama) in the 2012 and 2016 elections. Presidential election data at the constituency level come from the Electoral Commission of Ghana. I use constituency maps provided by (Ichino and Nathan 2013) and district maps from the Ghana Statistical Service to link constituencies to districts. Districts and constituencies are mostly congruous, and in most of the cases where they aren't a district is made up of exactly two or more constituencies. In these cases the NDC's two-party vote share is calculated as the sum of NDC votes in all constituencies divided by the sum of the NDC and NPP votes in all constituencies. For any case where a constituency was split across more than one district, all overlapping districts and constituencies are agglomerated.

I also operationalize support for incumbents based on local (District Assembly [DA]) election results. These results come from regional offices of the Electoral Commission of Ghana (EC), the Accra headquarters of the EC, and the Ghana Law School Library (a

detailed description of data cleaning and processing is in Appendix A). For cross-sectional specifications based on the 2015 DA elections, I calculate the incumbent re-election rate for all electoral areas (EAs) in a district where the incumbent stood for re-election. I also calculate the mean vote share for incumbents who stood for re-election. These measures are highly correlated ($\rho = .74$). The proportion of the two-party vote earned by incumbents cannot be used here because DA elections are officially non-partisan, so any candidate party affiliation is not recorded.

For the DA elections in 2010, I was unable to find detailed election results that included the names of each candidate and their vote totals. I therefore operationalize support for incumbent local assembly members in panel specifications as the proportion of EAs in a district in which an incumbent was re-elected, using the total number of EAs in a district as the denominator rather than the number of EAs where an incumbent ran. The disadvantage of this measure is that it may understate true incumbent support, especially in districts where more incumbents decided against running for re-election for reasons unrelated to their level of support among their constituents. The advantage of this measure, though, is that it does not suffer from bias related to endogenous retirement of assembly members. When the denominator used to calculate incumbent re-election rate is the number of incumbents who ran, true support for incumbents may be exaggerated, especially in districts where incumbents decided against running for re-election due at least in part to their self-perceived electoral weakness. For the 2015 DA elections I can compare the two measures, which are highly correlated ($\rho = .71$).

The relevant independent variable is the amount of transfers to a district through the DDF. I measure this in hundreds of thousands of real Ghana cedis, a sum roughly equivalent to $50,000 in 2013. Due to consistent inflation and currency depreciation, I use real transfers (deflated using World Bank GDP deflator data). I use the log of transfers in some specifications since there are some extreme values on the higher end of the distribution. Data on transfers from the central government to the districts through the DDF program

were obtained from the Ministry of Local Government and Rural Development in Accra.

Ghana has seen the number of local government units increase substantially over time, from 138 in 2006 to 170 in 2010 to 216 in 2016. Because DDF transfers began in 2009, I conduct the analyses at the pre-2010 split level (when there were 170 districts). Transfers to any child districts are aggregated at the level of the parent district by simply summing them. Vote shares are calculated at the aggregated level as well. The number of electoral areas (the single member districts for which assembly members are elected) also increased between 2006 and 2010. As a result, it is difficult to determine the proportion of assembly seats in a district in 2010 that were occupied by an incumbent and the proportion that were for newly created electoral areas. To address this, when calculating support for local incumbents, I take the total number of incumbents who were re-elected and divide by the number of seats that were present in the district before the new ones were created.

4.4 Results

Presidential Elections

I begin by presenting results from cross-sectional regressions of incumbent support on DDF transfers. With the caveat that these estimates should not be interpreted as causal, Table 4.2 shows the relationship between DDF transfers and incumbent presidential vote share separately for the periods ending in the presidential elections of 2012 and 2016. Although the signs of coefficients are the same for the 2012 and 2016 elections in columns 1-5, the estimates are smaller and standard errors larger for 2012 (few terms are statistically significant in 2012). The relationship between transfers and incumbent vote share is also presented graphically in Figure 4.3.

Columns 1-4 are bivariate regressions using different operationalizations of the dependent variable. Column 1 uses the total amount of DDF transfers received in the electoral

term[3] and shows a positive relationship that is significant for 2016 but not 2012. Column 2 uses logged DDF transfers to deal with very large and small values, with similar results. Column 3 uses an annual average of DDF transfers received in an electoral term which also shows a positive relationship that is only significant in 2016. In column 4, only transfers received in the year of the presidential election are counted and the magnitude of the effect is close to zero and not significant in 2012 or 2016. In column 5 I include region fixed effects in a regression using logged transfers. The magnitudes of the coefficients on transfers are about half as large as in column 2 but the result remains significant in 2016.

My preferred specification for the cross sectional analysis is column 6, which includes region fixed effects and controls. Here, the results are still positive and significant in 2016 but the coefficient estimate is close to zero in 2012. In 2016, a one percent increase in log DDF transfers is associated with a .064 percentage point increase in incumbent vote share. For a result that is more intuitive, I use this model to predict incumbent vote share for hypothetical districts at the 25^{th} and 75^{th} percentiles of logged DDF transfers, holding all other variables constant. This difference in transfers is about 760,000 cedis or $380,000 in 2013,[4] and it is associated with a 2.74 percentage point higher incumbent vote share. A three classroom block (a common DDF project) would cost about 100,000–130,000 cedis in 2013, so in theory each additional such project is associated with about a 0.4 percentage point increase in incumbent vote share. The median district received about 1,700,000 cedis in the second period of the analysis (leading up to the 2016 election) and the average absolute change in incumbent vote share from 2012 to 2016 was 5.5 percentage points, so this is a moderate effect size and is substantively significant.

I find mixed evidence for the question of how urbanization conditions the effect of DDF transfers. For the 2012 period, the coefficient on the interaction term in column 7 is positive and significant, suggesting that, while the average effect may have been zero, there was a

[3] All measures using cedis are in hundreds of thousands of constant cedis.

[4] Ghana's GDP per capita was $2,345 in 2013, or 4,690 cedis.

Table 4.2: Cross-Sectional Regressions of Incumbent Vote Share on DDF Transfers

	Panel A: 2012						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
DDF Transfers	0.001						
	(0.001)						
Log DDF Trans.		0.027			0.011	0.005	−0.058
		(0.031)			(0.020)	(0.017)	(0.031)
Avg. DDF Trans.			0.002				
			(0.003)				
Elec. Yr. DDF Trans.				0.001			
				(0.003)			
Prop. Pop. Urban						−0.110*	−0.561**
						(0.055)	(0.195)
Log DDF*Urban							0.146*
							(0.061)
Intercept	0.525***	0.458***	0.525***	0.535***	0.271***	0.580**	0.792***
	(0.034)	(0.101)	(0.034)	(0.035)	(0.068)	(0.209)	(0.224)
Mean of D.V.	0.546	0.546	0.546	0.546	0.546	0.546	0.546
Region FEs	No	No	No	No	Yes	Yes	Yes
Controls	No	No	No	No	No	Yes	Yes
N	170	170	170	170	170	170	170
Adjusted R^2	-0.003	-0.001	-0.003	-0.005	0.659	0.793	0.799

	Panel B: 2016						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
DDF Transfers	0.003*						
	(0.001)						
Log DDF Trans.		0.127***			0.052*	0.064*	0.072
		(0.037)			(0.025)	(0.029)	(0.041)
Avg. DDF Trans.			0.012*				
			(0.006)				
Elec. Yr. DDF Trans.				0.004			
				(0.010)			
Prop. Pop. Urban						−0.118*	−0.053
						(0.057)	(0.254)
Log DDF*Urban							−0.025
							(0.093)
Intercept	0.428***	0.119	0.428***	0.472***	0.116	0.295	0.278
	(0.031)	(0.107)	(0.031)	(0.043)	(0.073)	(0.230)	(0.239)
Mean of D.V.	0.486	0.486	0.486	0.486	0.486	0.486	0.486
Region FEs	No	No	No	No	Yes	Yes	Yes
Controls	No	No	No	No	No	Yes	Yes
N	169	169	169	169	169	169	169
Adjusted R^2	0.020	0.061	0.020	-0.005	0.662	0.787	0.785

*p < .05; **p < .01; ***p < .001 Note: The dependent variable in all specifications is the incumbent party's (NDC's) vote share in the 2012 or 2016 presidential election. In column 1, DDF transfers is the total amount transferred in that presidential term (in real hundreds of thousands of cedis, appx. $50,000). In column 3, DDF transfers are computed as the annual average for the four years up to and including the election year. Controls include population, Akan population, prop. of households engaged in agriculture, adult literacy rate, prop. of residents who are internal migrants, and open defecation rate.

Figure 4.3: DDF Transfers and Support for Incumbent President

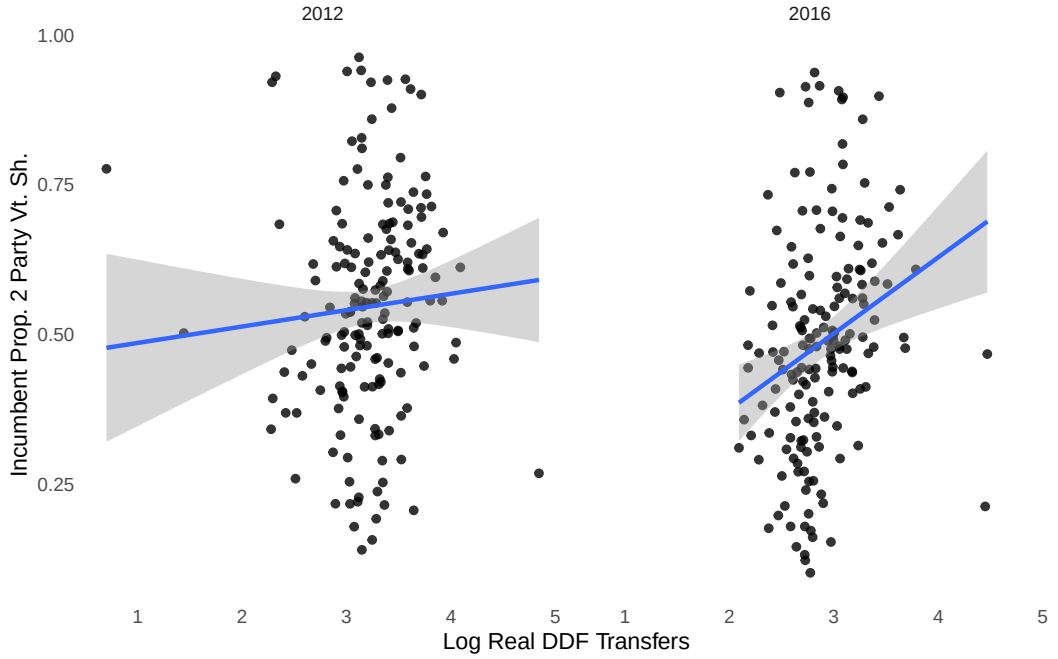

Note: The independent variable is the log of real DDF transfers. The dependent variable is the incumbent NDC's share of the two party vote. Each point represents one district. The left panel depicts transfers from 2009-2012 and 2012 election results, while the right panel shows transfers from 2013-2016 and the results of the 2016 election. Lines are from bivariate regressions and shading represents 95% confidence intervals of the slope of the regression lines.

larger effect in more urbanized districts. However, for the 2016 election the interaction of logged transfers and the proportion of the population living in an urban area is negative but not significant. These results could be due to more urbanized districts having higher administrative capacity and therefore being more quickly able to convert transfers into projects before the 2012 election. Though transfers began in 2009, it is unclear how quickly districts actually began construction on DDF-funded projects. By 2016, less urbanized districts could have developed better capacity and been able to build more DDF-funded projects with their transfers, resulting in a smaller difference between them and more urbanized districts.

Results from the pooled (two period) regressions described in Section 4.2 are reported in

Table 4.3: Support for Incumbent President and DDF Transfers–Pooled

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log DDF Trans.	0.084***	0.075***	0.004	−0.004			
	(0.022)	(0.008)	(0.007)	(0.007)			
DDF Transfers					−0.0001		
					(0.0003)		
Log Avg. DDF Trans.						0.004	
						(0.007)	
Log Elec. Yr. DDF Trans.							−0.003
							(0.006)
Elec. Yr. 2016			−0.058***	−0.131***	−0.059***	−0.056***	−0.062***
			(0.003)	(0.022)	(0.003)	(0.005)	(0.006)
Log DDF*Elec. Yr. 2016				0.025***			
				(0.007)			
Intercept	0.261***	0.350***	0.580***	0.605***	0.593***	0.584***	0.598***
	(0.069)	(0.037)	(0.027)	(0.027)	(0.018)	(0.022)	(0.022)
Mean of D.V.	0.516	0.516	0.516	0.516	0.516	0.516	0.516
District FEs	No	Yes	Yes	Yes	Yes	Yes	Yes
N	339	339	339	339	339	339	339
Adjusted R^2	0.037	0.954	0.983	0.984	0.983	0.983	0.983

*p < .05; **p < .01; ***p < .001

Note: The dependent variable in all specifications is the incumbent party's (NDC's) vote share in the 2012 or 2016 presidential election. In columns 1-4, the primary independent variable of interest is the log of the total amount of DDF transfers in a district-presidential term observation (in real hundreds of thousands of cedis, appx. $50,000). In column 3, raw (not logged) transfers are used. In column 4, DDF transfers are computed as the annual average (transfers divided by number of years in which there were transfers to any district) for the four years up to and including the election year. *Elec. Yr. 2016* is a dummy variable that equals 1 for the 2016 election and 0 for the 2012 election.

Table 4.3. The coefficient on the log of DDF transfers is positive and significant in columns 1 (without any fixed effects) and 2 (with district but not term fixed effects). Adding term fixed effects (a dummy variable for the term ending with the 2016 election) substantially reduces the magnitude of the coefficient on log DDF transfers and it loses statistical significance. Column 3 would suggest that a one percent increase in DDF transfers is associated with only a .004 percentage point increase in the incumbent NDC's vote share. This is a relatively small effect size.

To formally test whether the effect of DDF transfers is different for the 2012 and 2016 elections, I interact the *DDF Transfers* variable with an election year dummy variable. As shown in column 4 of Table 4.3, the coefficient on this interaction term is positive and

statistically significant, suggesting that the effect of DDF transfers on incumbent presidential vote share was larger in 2016 than in 2012. As mentioned above, it is unclear how quickly districts were able to build DDF-funded projects after they started to receive transfers in 2009. Many districts' annual progress reports, which list in-progress development projects, do not list DDF-funded projects for several years after the first transfers were received. By the 2013–2016 period, though, they may have developed stronger capacity and familiarity with the process for approving and building DDF-funded projects.[5] This would result in more projects being built and a larger effect on support for incumbents in the second time period in this study.

DDF Transfers and Support for Local Incumbents

Turning to support for incumbent local assembly members, Table 4.4 shows results from four regressions related to incumbent support and DDF transfers from a one period analysis. The independent variable (logged real DDF transfers from 2011–2015) is the same in all four columns. In Column 1, the dependent variable is the proportion of incumbents in a local government who who were re-elected in 2015 (of those who ran for re-election). In Column 2, the dependent variable is the mean vote share of incumbents who ran for re-election in a given government. Both coefficients are small (even negative in column 1) with relatively large standard errors. In Column 3, the dependent variable is the proportion of newly elected representatives that were incumbents. That is, the denominator is the total number of assembly member seats rather than the total number of incumbents who ran for re-election. In this specification, the coefficient is positive and much larger, though still not significant. The adjusted R-squared and F statistic are also much larger than in the models in columns 1 and 2. The magnitude implies that a one percent increase in DDF transfers is associated with a .034 percentage point higher proportion of elected assembly

[5]Projects must be listed on a district's medium term development plan to be eligible for DDF funding, so there may have been a lag between the start of transfers and their approval and construction.

Table 4.4: DDF Transfers and Support for Local Incumbents–2015

	Inc Win Prop	Mean Inc Vtsh	Prop Re-Elected	Prop Incumbs Ran
	(1)	(2)	(3)	(4)
Log DDF Transfers	−0.0033	0.0029	0.0341	0.0773**
	(0.0308)	(0.0133)	(0.0215)	(0.0272)
Intercept	0.5217***	0.4227***	0.1658*	0.2927**
	(0.0995)	(0.0430)	(0.0697)	(0.0879)
Mean of D.V.	0.51	0.43	0.28	0.54
N	142	142	142	142
Adjusted R^2	-0.0071	-0.0068	0.0106	0.0479
F Statistic (df = 1; 140)	0.0112	0.0472	2.5097	8.0931**

Note: The independent variable in all specifications is the log of DDF transfers received from 2011–2015.

members who were incumbents.

The results in Column 4 suggest one possible reason for the difference between the results from Columns 1 and 3: larger DDF transfers are associated with a greater number of incumbents who ran for re-election. A one percent increase in DDF transfers is associated with a .08 percentage point increase in the proportion of incumbents who ran for re-election. It could be that DDF transfers encouraged incumbents to run for re-election because they thought the additional public goods projects gave them a higher probability of winning re-election. If this is the case, the estimated effect in columns 1 and 2 should be biased downwards. The logic here is that DDF transfers could encourage marginal incumbents (potentially poorer quality ones who did not have a high chance of winning re-election) to run for re-election, resulting in lower incumbent win proportion and vote share in areas with more DDF transfers.

The one period analysis above does not include all districts or results from 2010 local elections. Looking at data from both the 2010 and 2015 elections increases statistical power and allows me to control for term effects. However, because the pre-2015 election data only included winners, I cannot examine how many incumbents ran, etc., as I did in Table 4.4. Table 4.5 shows results from two-period analyses of DDF transfers and incumbent

Table 4.5: DDF Transfers and Local Incumbents: Two Period Analysis

	(1)	(2)	(3)	(4)	(5)	(6)
Log DDF Transfers	0.004	−0.017	0.014	−0.012		−0.014
	(0.008)	(0.015)	(0.009)	(0.014)		(0.014)
Elec. Yr. Transfers					−0.003	
					(0.003)	
Elec. Yr. 2015		0.065***		0.059**	0.031	0.006
		(0.016)		(0.019)	(0.022)	(0.110)
Log DDF*Elec. Yr. 2015						0.017
						(0.034)
Lag N. Elec. Areas			−0.001***	−0.002*	−0.002*	−0.003*
			(0.0003)	(0.001)	(0.001)	(0.001)
Lag N. Districts			0.033	−0.003	−0.006	−0.015
			(0.019)	(0.046)	(0.046)	(0.052)
Intercept	0.238***	0.209**	0.236***	0.275**	0.282**	0.296**
	(0.026)	(0.076)	(0.027)	(0.100)	(0.099)	(0.110)
Mean of D.V.	0.251	0.251	0.248	0.248	0.248	0.248
Dist & Year FEs	No	Yes	No	Yes	Yes	Yes
N	266	266	258	258	258	258
Adjusted R^2	-0.003	0.130	0.101	0.223	0.227	0.218

*p < .05; **p < .01; ***p < .001

Note: The dependent variable is the proportion of elected district assembly members who were incumbents in a district-year observation.

assembly member re-election rates.

In column 1 of Table 4.5, the coefficient on log DDF transfers is positive but not significant. Adding district and year (term) fixed effects results in a negative but not significant sign (column 2). Because the number of electoral areas and districts increased between 2006 and 2010, I control for the lagged number of each in columns 3-6. I do this because a few observations are agglomerations of multiple districts (due to an inability to place former electoral areas in current districts for some that split) and because some districts saw a larger increase in the number of electoral areas than others (which would, all else equal, lead to a lower proportion of incumbents elected in the next election). As expected, the lagged number of electoral areas is negative and significant in all specifications. When included in column 3, the magnitude of the estimated effect of logged transfers is much larger, though not significant. Adding fixed effects (column 4) results in a negative and

insignificant coefficient. There is not evidence to suggest that the effect of transfers differs by election year, as indicated by the interaction term in column 6.

4.5 Robustness Checks

One concern with the above analyses is that the DDF funding variable is somewhat skewed. There are some district-year observations with much higher values of this variable than other districts. For example, Accra and Kumasi, the two largest cities/districts in Ghana, received much larger transfers due to their large populations. There are also a few districts with very small transfers in one of the periods. To test the robustness of the results above, I conduct the same analyses but without district-election year observations in which a district received extremely high or low levels of transfers in either period (five district-year observations are dropped). The results of the analyses without outliers, reported in Table 4.6, are broadly similar to those from the full sample. DDF transfers are associated with greater incumbent presidential vote share in 2016 but not 2012.

For the cross-sectional analyses reported in Table 4.2, there remains a concern that omitted variables could be driving the results. I therefore conduct a sensitivity analyses to estimate how important unobservable variables would have to be relative to observables to account for the estimated positive effect found in 2016. Following Altonji, Elder, and Taber (2005) and Nunn and Wantchekon (2011), this ratio can be calculated by comparing results from both a restricted model (no controls) and an unrestricted model (full controls). These two models are shown in columns 2 and 6, respectively, of Table 4.2 (Panel B). This analysis indicates that unobservables would have to have 1.02 times as large an effect as observables to explain the entire positive result. Although this is not particularly large, it is greater than the suggested cutoff of 1. This suggests that unobservable factors are unlikely to explain the entire result.

In the next chapter I introduce another transfer program, the UDG, which was only

Table 4.6: DDF Transfers and Support for the Incumbent President: No Outliers

	2012	2016	Pooled			
	(1)	(2)	(3)	(4)	(5)	(6)
DDF Transfers	−0.001	0.003*	−0.00004		−0.0003	
	(0.001)	(0.001)	(0.0003)		(0.0003)	
Log DDF Trans.				0.007		−0.005
				(0.008)		(0.008)
Elec. Yr. 2016			−0.059***	−0.056***	−0.092***	−0.148***
			(0.004)	(0.004)	(0.009)	(0.026)
DDF*Elec. Yr. 2016					0.002***	
					(0.0004)	
Log DDF*Elec. Yr. 2016						0.031***
						(0.009)
Intercept	0.296	0.119	0.593***	0.570***	0.600***	0.610***
	(0.244)	(0.245)	(0.019)	(0.029)	(0.018)	(0.030)
Mean of D.V.	0.547	0.487	0.517	0.517	0.517	0.517
District FEs	No	No	Yes	Yes	Yes	Yes
Region FEs/Controls	Yes	Yes	No	No	No	No
N	166	165	331	331	331	331
Adjusted R^2	0.696	0.704	0.983	0.983	0.984	0.984

*p < .05; **p < .01; ***p < .001

Note: The dependent variable in all specifications is the incumbent party's (NDC's) vote share in the 2012 or 2016 presidential election. Columns 1 and 2 are cross-sectional regressions for 2012 and 2016, respectively. Columns 3-6 are pooled regressions across both election years. District-year observations with extreme values of DDF transfers are dropped.

available to more urbanized districts. In comparing the two programs, it is useful to test whether there are differences between these two sets of districts. One way I do this here is by restricting the *DDF* analysis to those districts that were eligible for UDG transfers. The results of this analysis are presented in Table 4.7 and are rather similar to the results for all districts reported earlier in the chapter. Importantly, in Column 6 the interaction between logged DDF transfers and the 2016 election year dummy variable is positive, significant, and of a similar magnitude to the main result (.022 here vs. .025 in the corresponding model shown in Column 4 of Table 4.3).

Table 4.7: DDF Transfers and Support for the Incumbent President: UDG-Eligible Districts

	2012	2016	Pooled			
	(1)	(2)	(3)	(4)	(5)	(6)
DDF Transfers	−0.0001	−0.0004	−0.0001		−0.0003	
	(0.001)	(0.001)	(0.0003)		(0.0003)	
Log DDF Trans.				−0.005		−0.016
				(0.008)		(0.008)
Elec. Yr. 2016			−0.061***	−0.061***	−0.071***	−0.130***
			(0.004)	(0.004)	(0.007)	(0.022)
DDF*Elec. Yr. 2016					0.0004	
					(0.0002)	
Log DDF*Elec. Yr. 2016						0.022**
						(0.007)
Intercept	0.997*	0.914*	0.532***	0.544***	0.526***	0.572***
	(0.369)	(0.371)	(0.022)	(0.035)	(0.021)	(0.033)
Mean of D.V.	0.5	0.44	0.47	0.47	0.47	0.47
District FEs	No	No	Yes	Yes	Yes	Yes
Region FEs/Controls	Yes	Yes	No	No	No	No
N	46	46	92	92	92	92
Adjusted R^2	0.805	0.816	0.991	0.991	0.991	0.992

*p < .05; **p < .01; ***p < .001

Note: The dependent variable in all specifications is the incumbent party's (NDC's) vote share in the 2012 or 2016 presidential election. Columns 1 and 2 are cross-sectional regressions for 2012 and 2016, respectively. Columns 3-6 are pooled regressions across both election years. Only districts that were eligible for UDG funds are included.

4.6 Discussion

These results suggest that incumbents can receive an electoral boost from public goods projects undertaken by local governments. However, I find evidence to suggest that in Ghana the president's party is rewarded for this public goods provision but not local assembly members. There are a couple of reasons that this might be the case. First, in reality assembly members do not have much if anything to do with whether their district qualifies for DDF funding (and do not have any control over the existence of the DDF funding mechanism writ large). So it would be rational for voters to not reward their assembly members for DDF projects in the district as a whole. Second, assembly members who represent the electoral areas where projects are actually built may receive a boost in

support but not other assembly members (or may even be punished for not securing a project for their own constituents.) Evidence from interviews suggests that many actors try to claim credit for the provision of these public goods projects, including the president (or his party) via proxies like the district chief executive, the Member of Parliament, and the assembly member. So if assembly members do not see an electoral boost from DDF projects it is not for lack of trying.

Of course, the president appoints district chief executives and one third of a district assembly's members. So it is not irrational to reward the national leader for projects undertaken by the local government (as it might be in other contexts, like a federal system). And although it is unlikely that voters know the source of funding for various projects and how they come about, the national government does control whether these funding mechanisms get implemented (and in the case of the DDF, provided half of the funds for it).

4.7 Conclusion

In this chapter I tested how public goods provision affects support for presidential and local government incumbents using the case of the District Development Facility (DDF), an intergovernmental transfer funded jointly by the Government of Ghana and a group of bilateral foreign donor agencies. Using a two-period panel data set and a variety of estimation approaches, I find evidence that the program caused an increase in support for the incumbent president in the 2016 election. However, I do not find strong evidence that the program affected voting in the 2012 presidential election or in local elections.

In the next chapter, I examine the same general research questions using the case of the Urban Development Grant (UDG), another intergovernmental transfer with some important differences for estimation and interpretation.

Chapter 5

The Urban Development Grant: Incumbent Support and Funding for Public Goods

The Urban Development Grant (UDG) was a unique program that funded public goods provision by local governments in Ghana. Its funding was only available to urban local governments that scored above an eligibility threshold, allowing for a comparison incumbent support for governments that barely qualified to those that barely did not qualify.

In this chapter I draw on a variety data sources, including interviews with local government officials involved in the planning and implementation of public goods projects in three local governments of Ghana.[1] These officials were able to provide insights into the administrative procedures involved with building public goods projects with UDG funding, as well as some of the political implications of public goods provision.

[1]The specific governments were chosen so that I would visit one (rural) district assembly, one (partly urban) municipal assembly, and one (highly urbanized) metropolitan assembly. I do not disclose the specific local governments in which I conducted interviews to protect the anonymity of those who spoke to me.

5.1 Context: The Urban Development Grant and Public Goods Provision

As part of the broader Local Government Capacity Support Project (LGCSP), the World Bank funded the Urban Development Grant (UDG) to provide supplementary funding to local governments in Ghana with relatively high rates of urbanization.[2] Between 2012 and 2018, over $120 million was transferred to local governments in Ghana under the UDG, enabling the construction of 786 public goods projects.[3]

The UDG was similar in many ways to the District Development Facility (DDF, discussed in Chapter 4). Both programs only granted funds to districts that met certain administrative performance benchmarks. Districts also had to comply with donor requirements under both programs when actually using the funds on projects, such as competitive procurement procedures. The main difference between the programs was that UDG funds were only available to more urbanized districts, whereas DDF funds were open to all districts in Ghana. In addition, the DDF transferred over 687 million cedis between 2009 and 2016, whereas the UDG transferred just over 325 million cedis from 2012 to 2016.[4] This actually resulted in much larger per-district transfers under the UDG because only 46 districts were eligible (vs. 170 until 2012 and then 216 for the DDF).

Another potentially important difference between the DDF and UDG was that districts could use DDF funds for construction of administrative buildings like district offices, whereas this was not allowed under the UDG. If voters do not observe or value administrative construction as much as they do other types of public spending, this could result in DDF funds having a smaller effect on support for incumbents than UDG funds. However,

[2]Eligible local governments were those with metropolitan or municipal status when the program began (46 out of 170). District assemblies, which are more rural, were not included in the program.

[3]Some projects include several sub-projects, so the effective number of projects is actually greater.

[4]Both programs continued after 2016, but I limit my analyses to the period leading up to 2016 because the 2016 presidential election is the last one for which I have data.

other differences between the programs, both observed and unobserved, make it hard to confidently predict this. An obvious observed difference is that the districts that were eligible for each program differed. This can be accounted for by comparing the effect of each program for the set of districts that was eligible for UDG funds, which I do in the previous chapter in Table 4.7. The main result for DDF transfers is nearly unchanged when restricting the analysis to districts that were UDG-eligible, suggesting that the sample of districts is not what is driving differences in estimated effects. But the precise type and scale of projects could differ between the programs as well due to the UDG's larger per-district transfers, which cannot be observed with available data. Any differences in the time from receipt of funds to project construction could also affect the proportion of transfers converted into projects in a given period.

Because UDG funds had to go towards public goods provision, I use UDG qualification/disbursements as proxy variables for public goods provision. There are reasons to believe that the vast majority of the funding did indeed go to public goods provision. Local governments were required to submit work plans and other documentation to the Ministry of Local Government and Rural Development (MLGRD) for approval before they even began projects, as well as at later stages of projects. The World Bank also reviewed the procurement processes of 20 local governments after the project was completed and found no evidence of "ineligible expenditures" (World Bank 2018). Requirements are more stringent in general with World Bank and other donor-funded projects, e.g. the DDF requires projects be budgeted to completion (Williams 2017).

The most common project types included classroom blocks, health clinics, road repairs, police stations, public toilets, and boreholes for drinking water (World Bank 2018). For example, Figure 5.1 shows a completed borehole with tank built by the Hohoe Municipal Assembly in the Volta Region. This particular water tower was one of 20 built by the municipal assembly in various communities under a single administrative project costing nearly 600,000 cedis (about $165,000 in 2014 when the project was undertaken). Of course,

the municipal assembly's name is printed on the front of the support wall but not the funder's (i.e. the World Bank).

Figure 5.1: A UDG-Funded Public Goods Project

Note: UDG-funded borehole with water storage tank in Hohoe Municipal Assembly, Volta Region (author's photograph, 2019). This was part of a larger project that built 20 water towers in the assembly.

However, the funding was not guaranteed to every urban local government. As discussed further below, governments had to meet a handful of minimum administrative conditions (e.g. no adverse audit report) and score above the national average on an index of broader set of administrative performance measures.

5.2 Data and Empirical Strategy

The primary threat to identifying the effect of public goods provision on support for incumbents is that an omitted variable could bias the estimate. If such a variable were positively correlated with both public goods provision and support for incumbents, the estimated effect would be upwardly biased. In the case of the UDG, better performing governments should be more likely to qualify for UDG funds and be rewarded by voters for their good performance independent of whether or not they actually receive funds and provide public goods. In other words, governments that provide more public goods could simply be popular with voters for other reasons.

To address this concern, I use a regression discontinuity design, which takes advantage of a discontinuity in the UDG's qualification requirements to make causal inferences about the programs effects. But unlike the typical RD setup in which the running variable can take values arbitrarily close to the cutoff, the running variable here is discrete. This is perhaps most easily seen in Figure 5.6, where multiple observations take the same values of the running variable. This bunching makes the primary identifying assumption of typical RD designs–continuity of the conditional expectation function at the cutoff–implausible.[5]

I therefore proceed with a slightly stronger (but still plausible) assumption: the assignment of local governments to treatment or control is as-if random within a narrow bandwidth around the cutoff. This follows the recommendation of Cattaneo, Idrobo, and Titiunik (2018), who argue that this assumption is more plausible than the continuity assumption for discrete running variables. There are two sources of randomness in the UDG qualification system that provide support for this assumption. First, the cutoff is based on the national average of performance scores (the running variable), which is only determined after all local governments have been evaluated. In addition, the eventual location of the cutoff is difficult to predict because it is subject to the performance of many

[5]Lee and Lemieux (2010) note the continuity assumption for ordinary RDDs, while Cattaneo, Idrobo, and Titiunik (2019) question its plausibility for the case of a discrete running variable.

governments.

Second, the scores of individual governments are subject to some random variation each year. This can be seen in Figure 5.2, where performance scores are plotted over time by region.[6] Scores for individual governments vary over time and a government's precise score cannot be predicted with precision in a given round. It is important to note that performance scores are not completely random. However, the primary RD assumption is *local randomization* in a small bandwidth around the cutoff. All that is required is that a local government's *precise* score is subject to some random variation, which seems plausible given the bounciness of scores evident in Figure 5.2. Therefore, due to the unpredictability of the cutoff's location and the random variation in performance scores, it is likely that in the precise area around a given year's cutoff score the assignment of local governments is as-if random.

I conduct several tests using data only from UDG-eligible districts. First, I estimate how qualification for UDG 1 affected changes in NDC vote share from 2008 to 2012. UDG 1 was the only round of UDG funding disbursed prior to the 2012 election and all eligible governments received their allocation of funds at the same time. Second, I estimate the effect of qualification for UDG 3 on support for local incumbents in the 2015 assembly elections. The only round of UDG funding disbursed in 2015 was UDG 3, so this test should pick up any effects of election year disbursements on local elections. These analyses are under-powered, but should provide some evidence of the effect of UDG funds on support for incumbents and inform future analyses.

Combining Urban and Rural Districts

The tests described above are unbiased but also under-powered, as the number of observations within a small bandwidth of the cutoff ranges from 4 to 31 depending on the

[6]The sample here is restricted to local governments that *ever* scored within three points of the qualification cutoff in a UDG round, which is closer to the sample used in analyses.

Figure 5.2: Variation in Performance Scores by Region and UDG Round

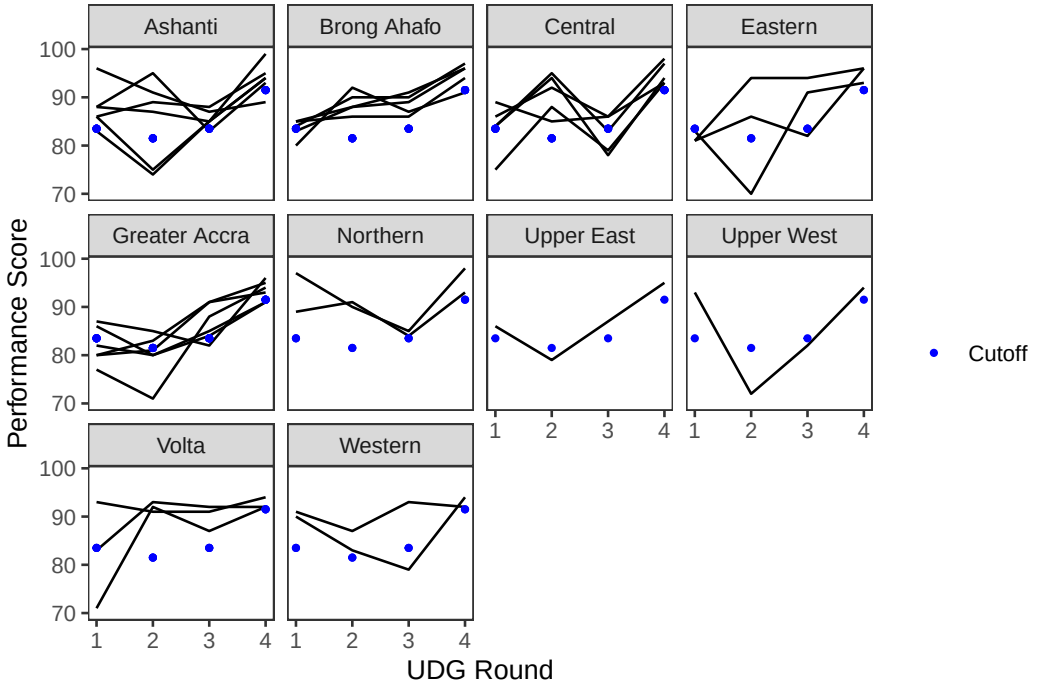

Note: Each line represents one local government and each plot one region. The sample is restricted to local governments that *ever* scored within 3 points of the cutoff in a round of UDG funding. The x-axis is the round of UDG funding and the y-axis is the performance score. Blue points are placed at the qualification cutoff for a given round (not region-dependent).

specification. To gain precision in the estimated effect of UDG funding on support for incumbents, I conduct a difference-in-differences analysis which uses data from all districts in Ghana, including those that were ineligible for UDG funding (i.e. the mainly rural districts). This increases the number of observations from 46 to 170. Although the rural districts were not eligible for UDG funds, they were still evaluated on the same scale of administrative capacity (FOAT) that the UDG-eligible ones were. This allows for the comparison of the difference in the NDC's vote share in those districts just above and just below the UDG round 1 funding discontinuity, in urban versus rural districts. I compute this quantity for both the 2008 (pre-UDG) and 2012 presidential elections, then subtract the latter from the former to get a triple difference-in-differences estimate.

I estimate models of the form:

$$
\begin{aligned}
y_{it} = \alpha_1 + \beta_1 \mathbf{U}_i \cdot \mathbf{A}_i \cdot \mathbf{P}_t + \lambda_1 \mathbf{S}_i \\
+\gamma_1 \mathbf{U}_i \cdot \mathbf{A}_i + \delta_1 \mathbf{U}_i \cdot \mathbf{P}_t + \zeta_1 \mathbf{A}_i \cdot \mathbf{P}_t + \kappa_1 \mathbf{U}_i + \eta_1 \mathbf{A}_i + \upsilon_1 \mathbf{P}_t + \epsilon_{1it}
\end{aligned}
\tag{5.1}
$$

where $\mathbf{U}_i$ is a dummy variable that equals 1 for the (urban) UDG-eligible districts. A second dummy variable, $\mathbf{A}_i$, equals 1 for districts above the FOAT score that was the eligibility threshold for UDG-1. The post period, i.e. the election year after UDG-1 had begun disbursements, is represented by the dummy variable $\mathbf{P}_t$. The coefficient of interest β_1 captures the effect of UDG funding for urban districts above the eligibility threshold in the period after UDG funding started. I control for the actual FOAT score received by districts, represented by $\mathbf{S}_i$. Districts are indexed by i and terms by t. Only observations within a small distance from the cutoff will be included in these estimates.[7]

This empirical strategy can be viewed as a combination of a regression discontinuity and difference-in-differences design. The use of both urban districts just below the eligibility threshold and rural districts above and below the urban districts' eligibility threshold creates a convincing set of counterfactual observations. The strategy can also be looked at as a difference-in-differences design where the estimate is computed separately for observations just above and below the cutoff and then pooled. This makes the parallel trends assumption more plausible since there should not be substantial differences in treatment and control governments conditional on other covariates (i.e. urban/rural status and the exact score on the running variable).

[7]Because optimal bandwidth selection methods are designed for purely continuous running variables, I will run robustness checks using different bandwidth sizes.

Data

The Urban Development Grant transferred over $120 million for public goods provision to qualifying local governments over six rounds (UDG 1 through UDG 6[8]) during the years 2012-2018. These transfers represented substantial proportions of eligible districts' budgets.[9] To be eligible for a round of UDG funding, governments first had to fulfill a handful of criteria called minimum conditions in their FOAT assessment[10] for the corresponding assessment year (the same conditions used to determine eligibility for DDF funds). These varied slightly by year, but generally included administrative criteria such as the absence of an adverse audit report on financial irregularities, holding regular meetings, and preparing an annual action plan. Every year required the absence of an adverse audit and for assemblies to hold meetings and take minutes.

Second, governments had to score above the national average on the FOAT performance score. This score was determined by looking at a wide variety of administrative indicators related to assemblies' management, organization, and capacity/effectiveness. Because the national average was calculated based on all local governments (including rural ones), it was theoretically possible for all urban governments to score above the average and qualify for UDG funds in any given year– though this never occurred.[11] Instead, there was a discontinuity in qualification, with some urban governments scoring above and some below the threshold at the national average.

I create a standardized running variable where the cutoff value equals zero in all UDG

[8]Ghana's government referred to the final round as UDG-Supplement rather than UDG 6, but I use the latter term for simplicity.

[9]Although the data required to calculate the proportion of all eligible districts' budgets represented by UDG funds has not been collected/processed, for the district that received the median level of UDG transfers from 2013–2016 (Agona West), these funds were 25% of total income in 2016.

[10]The Functional and Organisational Tool (FOAT) is an assessment tool used to evaluate the administrative performance of each of Ghana's districts.

[11]In the fifth and final year of the UDG program, all urban governments were given funds due to the World Bank's need to disburse money by a certain deadline, though FOAT assessments were not used to determine eligibility.

Figure 5.3: UDG Disbursements by Year and Round

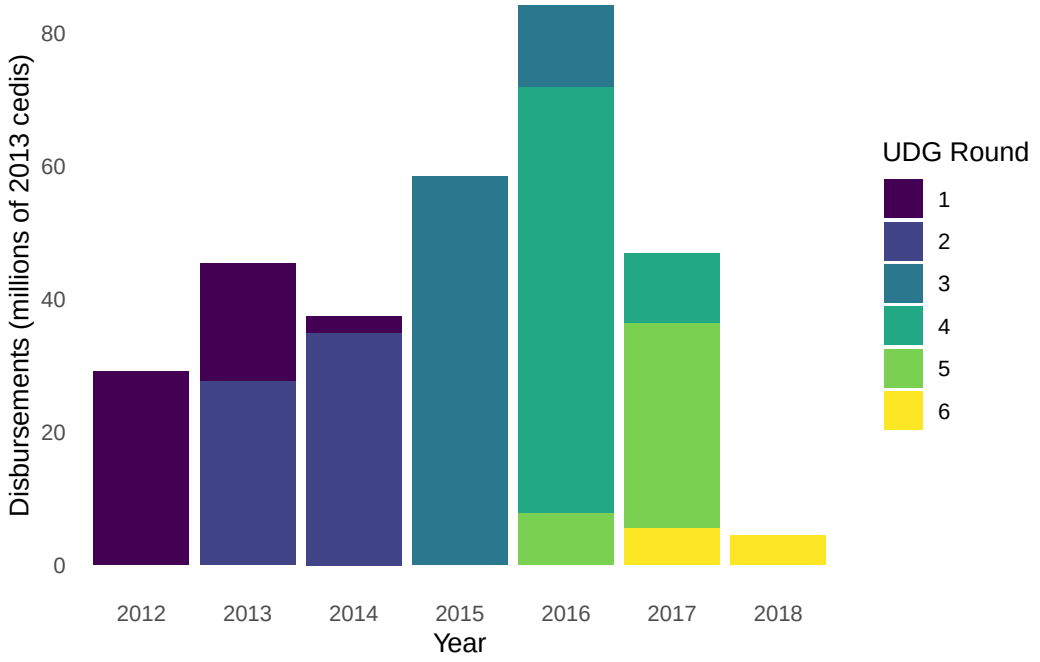

Note: Raw disbursements were converted to real disbursements to adjust for inflation over time.

rounds that used scores to determine eligibility. I recode the original running variable (the FOAT performance score in a given round of UDG) by subtracting the qualification threshold score for a given year less 0.5 from each score in that year. So, in the first year of the UDG program when the qualification score was 84, a score of 83 becomes -0.5 and an 84 becomes 0.5. This places the cutoff in qualification at the midpoint of the highest possible non-qualifying score and the lowest possible qualifying score. The discontinuities in funding for each UDG round are easily visible in Figure 5.5.

There are several potential issues that must be addressed when using a regression discontinuity design. For one, there is a concern that sorting might occur at the discontinuity, when units with higher ability are able to manipulate their score to be just above the cutoff when they are close. This would be the case if the better-performing governments near the discontinuity were able to inflate their scores to barely qualify for funds. This is not a

serious problem in the case of the UDG because local governments were not able to predict the national average score (i.e. the discontinuity) in any given year, preventing precise manipulation of the running variable.

Empirically, we can examine a histogram of the running variable to check for evidence of sorting. An observable implication of sorting is that there will be a higher number of observations just above the cutoff than just below it. Figure 5.2 shows that there are an equal number of observations at the scores just above and below the cutoff (three). While there are slightly more observations above the cutoff at scores a bit farther from the discontinuity, there is no evidence of a discontinuous change in density at the cutoff.

Figure 5.4: Distribution of the Running Variable: UDG 1

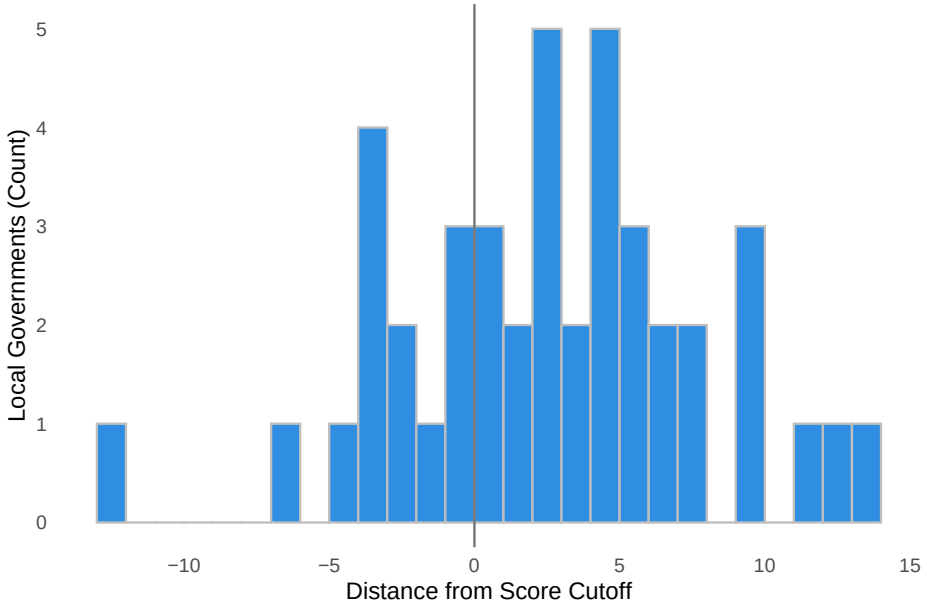

Note: Performance score scaled to have discontinuity at zero.

To formally test for sorting at the cutoff, I follow Cattaneo, Idrobo, and Titiunik (2018), who suggest using a binomial test in the case of a discontinuous running variable rather than the standard test (McCrary 2008) used when the running variable is continuous. This binomial test compares the number of observations just above the cutoff to the number just below it within a given bandwidth, producing a p-value to quantify evidence against

Table 5.1: Sorting Tests for UDG 1 Qualification

Bandwidth	1	2	3	4	5	6
P-Value	1	1	0.45	0.83	0.34	0.15

Note: P-values are from binomial tests of the null hypothesis that the number of observations above and below the cutoff for a given bandwidth is equal.

the null hypothesis that the number of observations on either side is equal. I conduct this test for bandwidths of width 1 through 6 and present the results in Table 5.1. No p-values are below 0.5 and the only one that is close (.15) is for a bandwidth of 6, where the local randomization assumption is less plausible.

Another concern is potential non-compliance, where units just below the cutoff nevertheless receive treatment (or units just above the cutoff do not take up treatment). In the case of the UDG, non-compliance is zero. No local governments below the national average in a given year received UDG transfers from that year's pool of funds, as seen in Figure 5.5.

I use both presidential election data and local election results to test how UDG funding affects support for incumbents (these data collection efforts are described in detail in Chapter 4 and Appendix A). Since the NDC held the presidency in the lead up to both the 2012 and 2016 elections, the NDC candidate in both of those elections (John Mahama) was the incumbent. I calculate the change in vote share from one election to the next, which uses the NDC vote share from 2008 and 2012 as the baseline for the 2012 and 2016 presidential elections, respectively.[12] Change in vote share is ideal for the present set-up because it adjusts for baseline support for the NDC in each district, which varies considerably. The small sample size makes estimates of the effect on vote share itself imprecise, whereas estimates of the effect on change in vote share are less so. For the combined rural and urban district analysis, I use the level of NDC support in 2008 and 2012 as the dependent variable

[12] Although the NDC wasn't the incumbent in 2008 so the measure isn't strictly speaking the change in incumbent vote share, Ghana's entrenched two-party system and relatively stable support for each party over time makes the party's own previous support the relevant comparison.

Figure 5.5: UDG Disbursements and Performance Scores

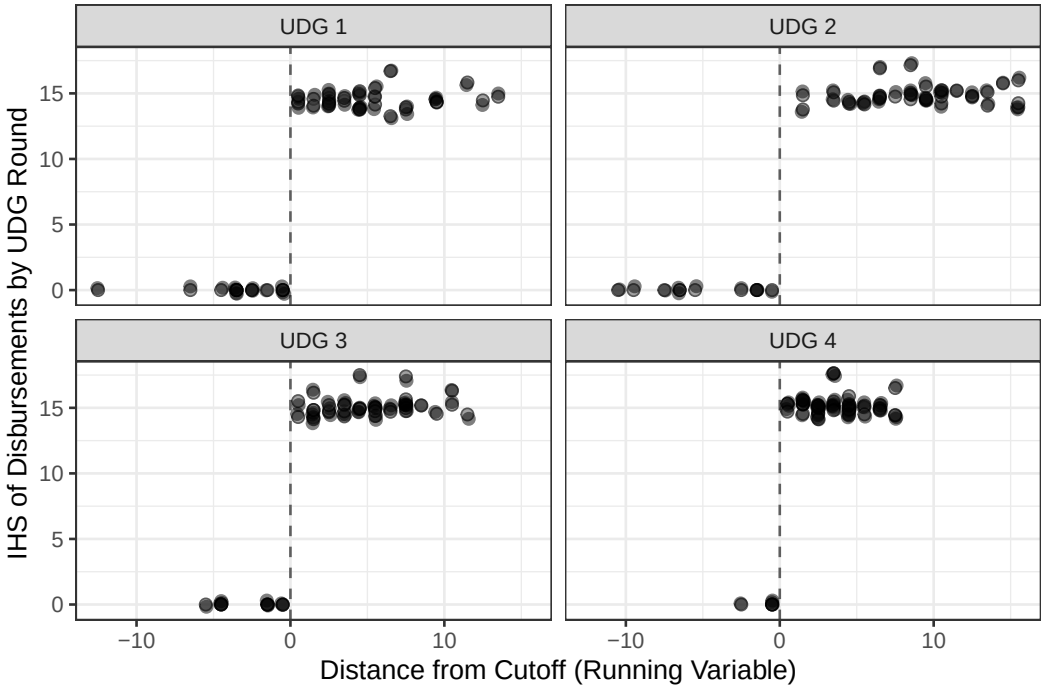

Note: Each small point represents one local government. Points are slightly jittered for better visibility. The y-axis is the inverse hyperbolic sine transformation of total disbursements under a given round of UDG funding. Vertical lines are at zero, the cutoff score at which point local governments become eligible for funds.

In addition to support for the incumbent president, I also measure support for incumbent local government representatives (assembly members). Two-thirds of assembly members in each local government are elected in single member electoral areas with plurality (first past the post) voting. One third of assembly members, as well as the local government's chief executive, are appointed by the president so their electoral support is not directly measurable with local election results. I therefore focus on elected members.

Because funding for public goods varies at the local government level, I summarize support for incumbent local government assembly members for each district.[13] One measure is the proportion of incumbents who were re-elected in each district, which uses the

[13]There are no maps of electoral areas, so it would be challenging to match public goods provision in an electoral area to support for that area's incumbent assembly member even if I had data on sub-district public goods provision.

Table 5.2: Balance Tests for UDG 1 Qualification

Variable	Did not Qualify	Qualified	Difference	P-Value
NDC Voteshare 2008	0.511	0.47	-0.0414	0.698
Population (2010)	138000	153000	15300	0.669
Prop. Pop. Urban	0.673	0.493	-0.181	0.0375
Mean Age	25.8	23.9	-1.87	0.0139
Literacy Rate (Adult)	0.858	0.706	-0.152	0.00118
Prop. Agric. HHs	0.342	0.503	0.161	0.149
Prop. HHs Open Defecating	0.103	0.173	0.0706	0.209
Number of Obs.	6	10		

Note: Includes observations within three points of the cutoff (BW = 3)

number of incumbents who ran for re-election as the denominator.

5.3 Results

I begin by estimating the effect of the first round of UDG funding on the change in the incumbent president's vote share from 2008 to 2012. (The first disbursement of this round, constituting half of funds disbursed under the first round of UDG funding, occurred just over two months before the presidential election in 2012.) The raw data and binned means for a small bandwidth are shown in Figure 5.6. Districts just above the qualification cutoff showed slightly greater changes in incumbent vote share between 2008 and 2012. Although the estimated effect of qualifying for funding is positive, there is some uncertainty in the estimation due to the small number of observations, so the difference is not statistically significant.

Figure 5.7 shows how the estimated effect of qualifying for UDG funds on support for incumbents is affected by the choice of bandwidth.[14] Results are shown for bandwidths

[14]With most approaches to regression discontinuity analysis, a bandwidth must be chosen so that only observations close to the cutoff are compared to one another. Smaller bandwidths reduce bias but increase variance due to smaller sample sizes. Regression discontinuity estimates tend to be sensitive to changes in the bandwidth, so it is generally good practice to report results for multiple bandwidths.

Figure 5.6: Effect of UDG 1 on Change in Incumbent President's Vote Share

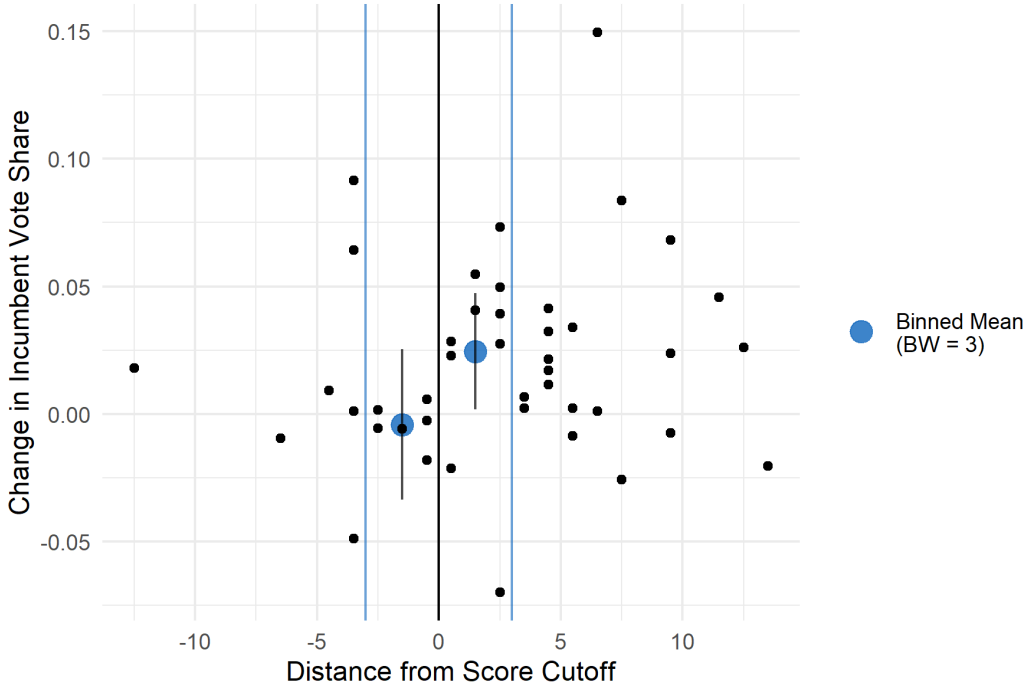

Note: Each small point represents one local government. The two large points represent means of the outcome variable on either side of the discontinuity for observations within 3 points of the cutoff. The dependent variable is the change in the incumbent president's vote share from the 2008 election to the 2012 election. The running variable is the distance from the qualification threshold in the first round of UDG funding. Error bars are 95% prediction intervals.

from 1 (the lowest possible distance from the cutoff) to 6.[15] Coefficients are uniformly positive, ranging from .012 (bandwidth = 6) to .03 (bandwidth = 1). An effect of .014 (the median) would imply that the UDG program resulted in a 1.4 percentage point *larger* increase in NDC vote share for qualifying local governments than for non-qualifying ones.

Note that because the running variable is not purely continuous, the continuity assumption for traditional regression discontinuity approaches is not plausible, so I follow Cattaneo, Idrobo, and Titiunik (2018) in treating the setup as a randomized experiment around the cutoff. This means that the models are simple bivariate regressions of the outcome on an indicator for being above the cutoff, restricted to observations within a

[15]Seventy-two percent of observations are within 6 points of the cutoff.

Figure 5.7: Bandwidth Size and Effect of UDG 1

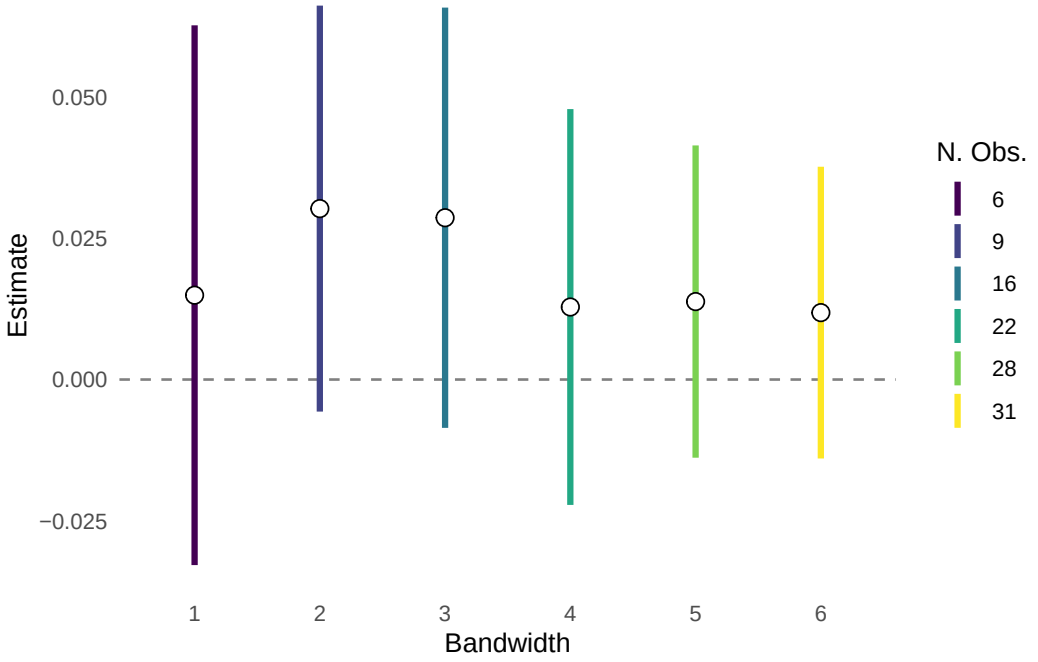

Note: Coefficients are from regressions of change in NDC vote share on the UDG treatment indicator for each bandwidth. Only observations within the bandwidth distance to the cutoff are included in each regression's sample. The dependent variable is the change in the incumbent president's vote share from the 2008 election to the 2012 election. The running variable is the distance from the qualification threshold in the first round of UDG funding.

given bandwidth around the cutoff. So, I do not estimate models or report results with the familiar RD statistics like the order of the forcing variable or the shape of the kernel. The point estimates and confidence intervals reported in Figure 5.7 contain all of the relevant results for the models I run.

UDG 3 and Local Elections

Results of the analysis of UDG 3 qualification and local incumbent support are presented in Figures 5.8 and 5.9. Figure 5.8 shows the typical RD plot, with the running variable on the x-axis and the outcome (the proportion of elected assembly members who were incumbents) on the y-axis. I use a bandwidth of two to calculate the means on either

side of the cutoff because there are no districts that scored either 3 or 4 points below the cutoff. Comparing local governments just above the qualification threshold to those just below it, UDG 3 qualification appears to have a positive effect on incumbent re-election rates. Formally, districts just above (within 2 points of) the cutoff saw a seven percentage point higher rate of incumbents among those who were elected in 2015 than districts just below (the standard deviation of this outcome measure is 12 percentage points). However, potentially due to the small sample size, this result is not statistically significant.

The confidence intervals for effect estimates include zero for various bandwidths, as shown in Figure 5.8. Results using the alternative measure of incumbent support (the proportion of incumbents who ran for re-election who were re-elected) are similarly positive but not significant.

Figure 5.8: Support for Local Incumbents and UDG 3

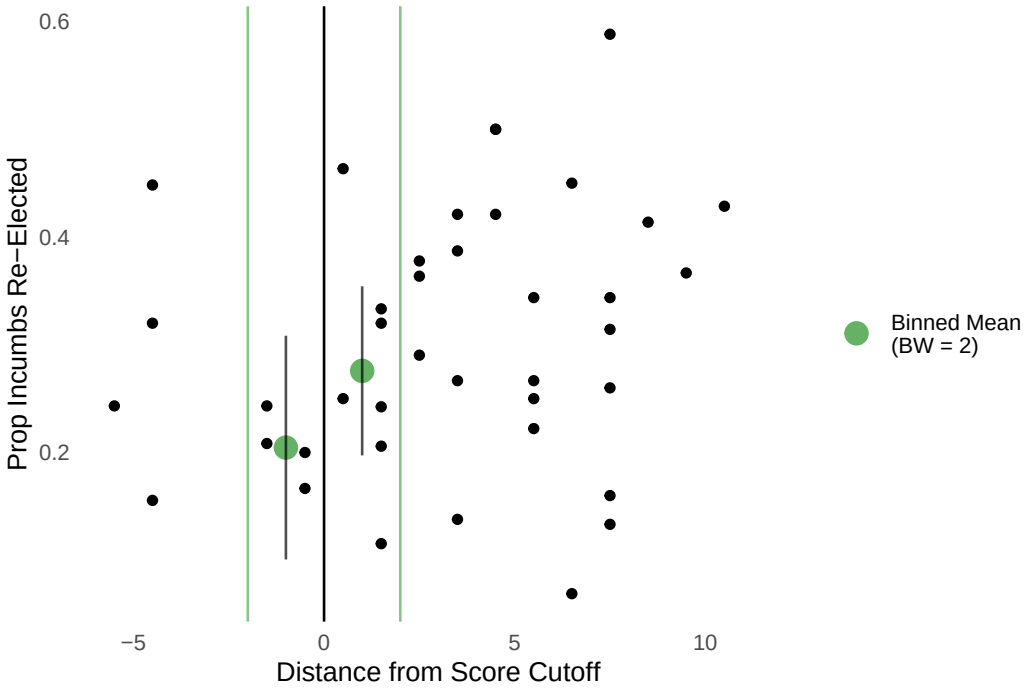

Note: The dependent variable is the proportion of elected assembly members in 2015 who were incumbents. Each black point is one local government, while the large green circles are binned means from just above and below the threshold. Error bars are 95% prediction intervals.

Figure 5.9: Bandwidth Size and Effect of UDG 3

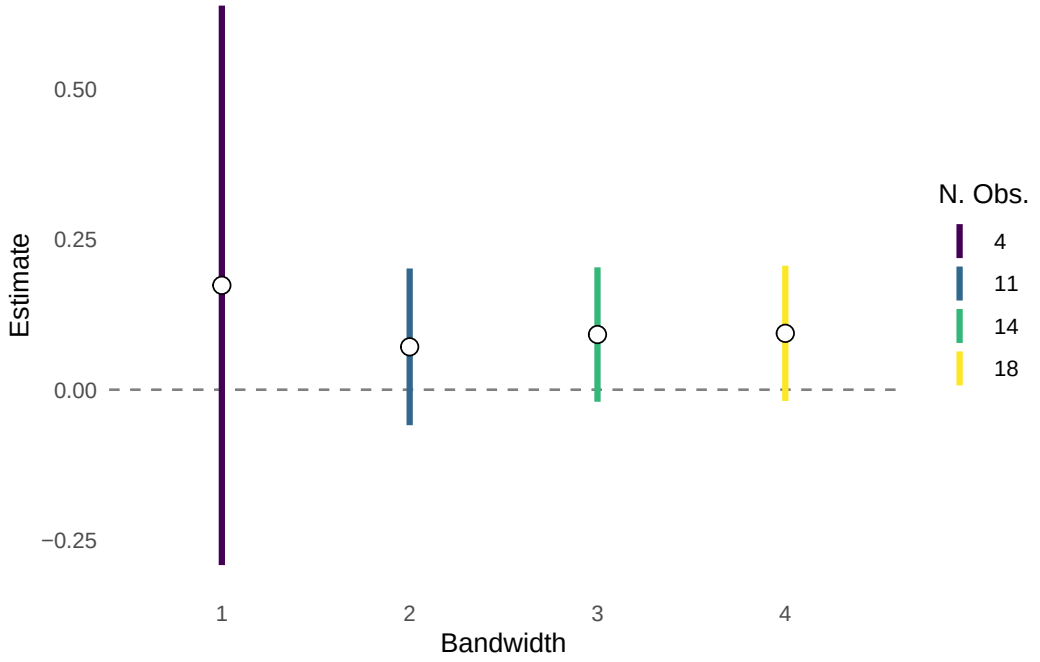

Note: Coefficients are from regressions of incumbent win proportion on UDG qualification in small bandwidths around the qualification threshold.

5.4 Difference-in-Differences–Urban and Rural Districts

Results from the difference-in-differences model described in Section 5.2 are reported in Table 5.3. The coefficient on the triple interaction term represents the effect of scoring above the eligibility threshold for UDG in an urban district that could get UDG funds in the period after UDG funds were first disbursed. This term is positive, which would indicate a positive effect of UDG funds, but the standard error is quite large and the estimate is not statistically significant. Curiously, none of the other coefficients are significantly difference from zero, either.

One possibility for the above null results is that districts did not have enough time to begin building projects with the funds they received before the 2012 election. That first transfer was in September, while the election was in December. However, additional tests using the 2016 election as the dependent variable did not produce dramatically different

Table 5.3: Difference-in-Differences: Urban and Rural Districts

Urban*Post*Above	0.033
	(0.119)
Mean of D.V.	0.52
N	166
Adjusted R^2	0.042

$^{*}p < .05$; $^{**}p < .01$; $^{***}p < .001$

Note: The dependent variable is the incumbent party's (NDC's) vote share in the presidential election. The sample includes all districts in Ghana that met the minimum conditions for DDF/UDG funding in 2010 (FOAT IV) and whose performance scores (the running variable) were within 3 points of the UDG eligibility threshold for UDG-1. There are two observations for each district: one in 2008 (pre-UDG) and one in 2012 (post-UDG, i.e. after the first transfer).

results. It may instead be an issue with the sample size, as the triple interaction approach splits the data into many different cells, many of which do not have more than a few observations.

5.5 Robustness Checks

As a robustness check on the above results, I set up the data in the same way as the DDF data from the previous chapter. That is, I create a two-period panel of district–presidential election term observations. The dependent variable is the incubment NDC's vote share and the independent variable is the inverse hyperbolic sine (IHS) transformed, total real transfers to a district during that term. The IHS transformation is used here because there are a few extremely high observations and many zeroes. It accomplishes the same goal as a log transformation but allows for zeroes in the data. Results are reported in Table 5.4.

Column 1 shows a bivariate regression of incumbent support on transfers across both periods. The sign of the coefficient is negative but not significant. In Column 2 I include district fixed effects and the coefficient on transfers remains nearly the same but becomes significant. In Column 3 I also include election term fixed effects (a dummy variable for the 2016 election year) and the coefficient on transfers becomes close to zero and not significant,

Table 5.4: Support for Incumbent President and UDG Transfers–Pooled

	(1)	(2)	(3)	(4)
IHS UDG Trans.	−0.017	−0.018***	0.002	0.002
	(0.011)	(0.002)	(0.003)	(0.003)
Elec. Yr. 2016			−0.064***	−0.068**
			(0.008)	(0.020)
IHS UDG*Elec. Yr. 2016				0.001
				(0.005)
Intercept	0.511***	0.547***	0.510***	0.509***
	(0.038)	(0.020)	(0.013)	(0.014)
Mean of D.V.	0.462	0.462	0.462	0.462
District FEs	No	Yes	Yes	Yes
N	92	92	92	92
Adjusted R^2	0.013	0.977	0.991	0.991

*p < .05; **p < .01; ***p < .001

Note: The dependent variable in all specifications is the incumbent party's (NDC's) vote share in the 2012 or 2016 presidential election. The primary independent variable of interest is the IHS of the total amount of UDG transfers in a district-presidential term observation (in real hundreds of thousands of cedis, appx. $50,000). *Elec. Yr. 2016* is a dummy variable that equals 1 for the 2016 election and 0 for the 2012 election.

but positive. Finally, Column 4 tests whether the effect of transfers differs according to the election term by interacting the election year dummy variable with the transfers variable. This interaction effect is very small and not significant.

5.6 Discussion

The results presented above follow a pattern: estimates of the effect of UDG funding on support for incumbents tend to be positive but not statistically significant. The tests, while low in bias, are not very highly powered. While I cannot definitively conclude that UDG funding positively affected support for presidential or local incumbents, the evidence taken together suggests that a positive effect could be present.

5.7 Conclusion

In this chapter I explored the effects of public goods provision on support for incumbents in the case of the Urban Development Grant (UDG), an intergovernmental transfer funded by the World Bank. The UDG was only available to relatively urbanized districts in Ghana, which had to also score well on an administrative performance measure in order to qualify for funds. This lent itself to a regression discontinuity approach to estimate its causal effect on both presidential and local incumbents.

I do not find strong evidence that the program had an impact on support for either national or local incumbents. However, a pattern of positive coefficients emerged across the various models, suggesting that there may have been a positive effect that was small or simply hard to detect given the sample size of the study.

Looking back at the past two chapters, there is consistent and sometimes significant evidence that public goods provision can help incumbents electorally. A remaining question relating to mechanisms is whether these effects depend on the preferences of voters living in each district. It could be that districts where voters do not place a high importance on public goods provision had a smaller effect than those where voters do highly value public goods. This could help explain some of the inconclusive results above. I address the role of preferences in the next chapter.

Chapter 6

Preferences, Public Goods Provision, and Electoral Outcomes

6.1 Introduction

Does the relationship between public goods provision and electoral outcomes depend on voter preferences? In Chapter 2 I outline a theory of how intergovernmental transfers earmarked for public goods provision could influence voting behavior. In particular, I suggest that the effect of these transfers on voting depends on what sorts of public goods projects people actually want. I emphasize that voters have multidimensional preferences over public goods provision, in contrast to a simpler model of voting in developing countries where voters can be swayed by any sort of material transfers from the government. In this chapter, I test whether Ghanaian voters' preferences condition the effect of transfers on support for incumbents.

To test this hypothesis, I use data from Afrobarometer surveys in Ghana to construct a measure of the average importance of public goods provision to Ghanaians in each district in the country. Because of small sample sizes in most districts, I use a machine learning-based method similar to multilevel regression with post-stratification (MRP) to produce

more accurate district-level estimates. MRP is a model-based estimation method that has been shown to produce more accurate estimates of public opinion for small geographic areas (Lax and Phillips 2009). It involves fitting an individual-level predictive model on public opinion data, predicting responses for thousands of different demographic profiles, then weighting those predictions in each geographical unit by the particular demographic profile's share of that unit's population.

To generate estimates of public opinion on the importance of public goods provision at the district level in Ghana, I use an improved version of MRP called BARP, which stands for Bayesian additive regression trees (BART) and post-stratification (Bisbee 2019). BARP replaces the multilevel model in the first stage of traditional MRP with a machine learning model. That model, BART (Chipman, George, and McCulloch 2010), is highly flexible, so it can make more accurate predictions when the precise functional form of a relationship is unknown. It can also include complex interactions as the data demand. BART can deal quite easily with missing data and its regularization makes overfitting[1] highly unlikely (though it is still important to take steps to prevent it, as I do below). And because it is an approximate Bayesian method that employs Markov chain Monte Carlo (MCMC) methods, propagating predictive uncertainty into regression models using draws from the posterior distribution is relatively straightforward. The post-stratification (weighting) step in BARP is the same as in MRP.

Once the district-level opinion estimates are generated, I take the $1,000$ posterior estimates for each district and use them to fit $1,000$ regression models. This ensures that the uncertainty inherent in a model-based estimate is conveyed in the regression analysis. And because each of the $1,000$ estimated regression coefficients has its own uncertainty estimates, I use simulation to propagate that uncertainty into the final estimates.

The analysis in this chapter makes several contributions to existing work. First, it applies

[1]Overfitting occurs when a model is too flexible. It learns the idiosyncrasies of a given data set to predict its outcomes with high accuracy, but performs very poorly on unseen data. I discuss overfitting further in section 6.3.

a regression and post-stratification method to an African country using Afrobarometer data to improve small area estimation. This has been rarely, if ever, done before, as most applications of MRP have been in the US (with a few in Europe and elsewhere). Second, it takes several specific steps to prevent and detect overfitting, which is not standard practice in the MRP literature. This is an important step when doing any predictive modeling, which the multilevel regression portion of MRP is. Third, the model pools information across time as well as space, a relatively new development in the standard MRP literature. Finally, it tests a novel theory explaining how public goods provision can affect support for incumbent politicians, an important link in the accountability and responsiveness process.

6.2 Data

To test the theory that voter preferences condition the effect of intergovernmental transfers (and by extension, public goods provision) on support for incumbents, I combine several sources of data. I use Ghanaian government data on transfers that are earmarked for public goods provision under the District Development Facility (DDF) program, discussed in detail in Chapter 4.

To measure voter preferences, I use data from rounds 2-6 of the Afrobarometer survey conducted in Ghana. In these rounds, respondents were asked for their view on the most important problems facing the country that the government should address. They could give up to three responses free form, which were later coded into about 35 categories of responses. I use these responses to create a dummy variable that indicates whether a respondent mentioned infrastructure/roads in any of their three responses. The infrastructure variable should capture preferences for physical infrastructure, which is broad and includes the majority of DDF project types. I also consider an education variable, as education projects are the most common types built with DDF funds. This variable is constructed in the same way as the infrastructure variable: any respondent who mentions

education as one of the most important problems facing the country is coded as supporting education.

About 24% of respondents in the sample I use mentioned infrastructure as one of the most important problems. The percentage in any given year of the survey ranged from 14% to 32%. For education, about 36% of respondents cited it as one of the most important problems. The figures by survey year ranged from 30% to 42%.

Using the most important problem question is not without its drawbacks. Wlezien (2005) argues that this question may not actually capture the salience of a given issue very well. A primary critique is that the importance of other issues affects how often the issue of interest is mentioned. However, this is not necessarily a problem for my analysis. I *do* want to capture the relative importance of infrastructure/public goods, so it is acceptable if this is affected by the importance of other issues.

The question is also not specific about what types of infrastructure are most important to respondents. Responses are given free form and later, any response falling into the category of roads or infrastructure is grouped into a single value. So, respondents who view roads as being very important but not other types of public goods are treated the same as ones who view school buildings as important but not roads. However, DDF funds were used on a variety of public goods projects, including road improvements, so a collapsed indicator is probably the best way to measure preferences such a variety of types of public goods.

Post-stratification (discussed below) requires census data, and I use micro-level data from Ghana's 2010 Population and Housing Census. The data set contains a 10% sample of individuals living in the country in 2010. I subset the data to exclude children and foreigners, which matches the sample universe of the Afrobarometer surveys (and is the population of interest, i.e. voting-aged citizens).

6.3 Methods

Testing whether voters have multidimensional preferences, and whether those preferences affect their voting behavior, is not without challenges. For one, voters might not think of their preferences in these ways. In interviews I conducted in Ghana's Central Region, voters overwhelmingly said they would be satisfied with any sort of development project that the government brought to their locality. But, they also were able to give their top priorities for investment, suggesting that they would indeed prefer one sort of project to another.

Another complication in testing observable implications of the theory is the distinction between what voters want and what they need. In some cases there may be no differences between the two, but in others people may want investments in a given type of good but objectively have a greater need for another type. For example, a district that has a very high student to classroom ratio would appear to need more classrooms built, but its citizens may state that health infrastructure is a higher priority for investment. Ideally, researchers interested in citizen preferences could measure what people want, potentially using measures of need as alternative measures/robustness checks. But it is difficult to measure what people want at a highly localized level in developing countries because of statistical power issues in public opinion surveys. Respondents in every administrative unit of a country are not interviewed every time a survey is conducted and the units that are surveyed do not have very large sample sizes (some districts in Ghana have few as 8 voters when included in an Afrobarometer survey). Even if a researcher were conducting their own survey it would be logistically and financially challenging to interview respondents in every administrative unit of a country.[2]

Testing the theory involves a multi-step methodological process. First, I use a modified multilevel regression and post-stratification (MRP) approach to estimate the importance of

[2]There were 216 districts in Ghana as of 2016.

public goods provision to Ghanaians in each district in Ghana at two points in time (before the 2012 and 2016 presidential elections). I then regress presidential vote share on the interaction between this public goods public opinion measure and the earmarked public goods transfers received by a district. Because the public opinion measure is a model-based estimate, I propagate its uncertainty into the regression model by running 1,000 separate regressions, one for each estimate of the opinion measure produced by the estimation model. To propagate the uncertainty in each regression's estimate, I then simulate 1,000 draws from a multivariate normal distribution with the mean and variance from each regression's coefficients. This produces one million estimates of the model parameters, where the mean and standard deviation are analogous to a coefficient and standard error estimate from an OLS regression. The interaction between transfers and the public opinion measure is the primary parameter of interest. Each step is described in more detail below.

Multilevel Regression with Post-Stratification (MRP)

When researchers want to measure public opinion in geographical areas but there are only a few respondents from some of them, simple averages (sometimes called the disaggregation method) will be unreliable due to high variances. To help alleviate problems caused by small sample sizes, researchers often use multilevel regression with post-stratification (MRP). This approach uses information from all respondents in a survey, as well as census data, to increase the precision of parameter estimates for small geographic units. In this chapter, I use a modified (but very similar in purpose and effect) MRP method called BARP (discussed later) to achieve more accurate estimates with less data.

MRP is useful for estimating public opinion at the level of relatively small geographic areas. In such cases, many units will have few observations in nationally-representative polls/surveys. Using a simple mean as the estimate for units with few respondents will produce inaccurate estimates because the sample sizes are small and often unrepresentative of the broader unit. For Afrobarometer surveys in Ghana most districts only have eight

respondents polled per survey wave, with many districts skipped entirely in a given wave. Sampling is done by clusters (sampling points are chosen at random and then eight respondents in that immediate area are surveyed), further reducing the representativeness of district-level samples. Estimates would also be unavailable for districts that are not surveyed in a given wave.

In the first step of MRP, a multilevel model is fit to predict individual-level opinion (usually, whether a respondent provides a certain response to a given question). The multilevel model pools information from respondents across the entire sample based on demographic and district-level predictors. This is achieved by fitting a district-level and an individual-level model, a procedure commonly referred to as fitting a random intercepts/random effects model. The district-level intercepts are modeled using geographic predictors (like the literacy rate in the district), which improves their accuracy. Individual-level predictions are then based on these district intercepts and the individual-level demographic predictors in the model.

In the second step of MRP, the multilevel model is used to generate predictions for every type of respondent in the population. These types (also referred to as cells) are defined by the unique combination of demographic variables and geographic units that are included in the model (limited by those that are available in both the poll/survey data and the census data in that country). So, for one example type/cell in Ghana, the model would predict opinion for Akan men, aged 18-25, without a formal education, living in an urban area of the Mfantseman district. To obtain district-level predictions, the individual-level prediction for each cell in a district is weighted by its share of the population in that district (as determined by census data) and then all weighted predictions in a district are summed.

MRP improves estimates of public opinion for small geographic areas in two ways. First, the multilevel model takes advantage of the nested structure of the data, partially pooling information across districts. This means that predictions for districts with fewer respondents will be pulled more towards the predictions for similar districts (for example,

districts in the same region or with similar levels of literacy). Second, post-stratification corrects for the major issues associated with non-representative sampling by weighting predictions according to their share of a district's population.

Bayesian Additive Regression Trees with Post-Stratification (BARP)

Although MRP can improve small area estimation of public opinion, it has some drawbacks. MRP can be quite accurate when the model is correctly specified, but in practice that is often an assumption researchers should be hesitant to make. This is especially true when estimating subnational opinion in contexts outside the US, as MRP has been applied much less often than in the US and it is unlikely that there will be an accepted consensus on the appropriate covariates to include in a model to render it correctly specified. This means that MRP might not produce the most accurate estimates outside of the US context. A more practical consideration is that complex multilevel models can be difficult to fit due to issues with convergence.[3] A model that doesn't converge cannot be used, so MRP wouldn't even be usable in some situations.

A recent development has solved these problems and given researchers the ability to produce even more accurate estimates of subnational opinion than before: replacing the multilevel model in MRP with a machine learning model (Bisbee 2019). In contrast to a multilevel model, the flexibility of machine learning models means they do not make assumptions about the functional form of relationships. Machine learning models can also allow for more complex interaction effects than traditional multilevel regression models used in MRP. Essentially, these models do not require the researcher to correctly specify the model in order to produce accurate predictions. Machine learning models can learn to ignore irrelevant predictors and include complex interaction effects as long as there is information contained in the predictors fed into the model.

[3] I began the analyses for this chapter by trying MRP, both the maximum likelihood and fully Bayesian approaches. However, the models would not converge using either method.

Using machine learning and post-stratification has been shown to produce more accurate estimates of subnational opinion than MRP. Bisbee (2019) tests a variety of machine learning methods and finds that Bayesian additive regression trees (BART, or BARP when combined with post-stratification) is best at reducing mean absolute error of (US) state-level opinion estimates. Because MRP and BARP do not appear to have been applied in an African context, I conduct a validation exercise and report the results in Appendix B. The exercise shows that BARP can produce very accurate estimates in a sample the size I am working with, greatly outperforming simple disaggregation.

Replacing the multilevel model with BART specifically has several advantages. Protection against overfitting is baked into the BART model, which includes a prior that discourages individual regression trees from being assigned too much weight and from growing too large.[4] In addition, BART produces samples from a posterior distribution, which provide a natural way to propagate uncertainty into the second stage estimation. And using BART for this application has a proven track record of success, at least in some applications (Bisbee 2019).

A further strength of BART is that its best implementation in R (bartMachine Kapelner and Bleich 2013) natively handles missing data (Kapelner and Bleich 2014), which is not the case for the random forest, logistic regression, or ensemble methods. This is important because there is a nontrivial amount of missingness in the Afrobarometer data I use. Throwing away observations with missingness in any predictor would substantially reduce the number of observations, with lower model accuracy the result. The bartMachine package handles missingness by allowing individual decision trees to create splitting rules based on whether or not an observation's value for a variable is missing. This allows the model to learn whether missingness is important for a given variable and produce better predictions overall without the need for imputation or listwise deletion.

One downside to the model complexity of machine learning methods is the risk of

[4]This may not sufficient to prevent overfitting, so I take additional steps outlined below.

overfitting, where a very complex model learns to predict observations from a given data set very well but performs poorly in out of sample predictions. When fitting machine learning models, it is important to take steps to prevent overfitting and detect it when it does occur. There are several ways to prevent or reduce overfitting. Many models include some sort of regularization, which prevents individual parameter values from having too much influence in the model. Another approach is k-fold cross-validation, where a model is trained several different times using different subsets of the training data and evaluated on different validation sets from the training data. When repeated using different model hyperparameters, this allows the researcher to choose those that produce the most accurate predictions for validation data on average. Hyperparameters that lead to overfitting will not perform well because their accuracy on the validation sets will be lower than on the training sets. Researchers should also reserve a portion of their data for testing their finished model, and that data should not be used in any part of the model training process. If the model is much more accurate on the training set than this final test set, there is strong evidence of overfitting.

Although the BART model's prior regularizes and helps to prevent overfitting, it likely does not completely eliminate the problem. I therefore augment the native regularization of BART in two ways. First, I randomly split the Afrobarometer survey data into training and testing sets. I only use the training set when fitting (training) the BART model, while the test set is reserved for a final application of the trained model to test its overall accuracy on unseen data. By only evaluating the model on data it has never seen before, I can detect whether it was overfit to the training data. I include 95% of the data in the training set and 5% in the test set in order to maximize the amount of data seen by the model.[5]

The second step I take to prevent overfitting it using 5-fold cross validation to choose

[5]This is a higher proportion of training data than the conventional 80/20 or 90/10 split. Since the predictive accuracy of the individual-level model is not the end goal (more accurate small-area estimation is), I am willing to accept higher variance in the test set accuracy rate. The test set mainly serves as a final test to ensure that out of sample accuracy is not dramatically different from training accuracy.

the BART model's hyperparameters. First, I randomly split the training data into 5 equally-sized partitions. I then run (train) the BART model on a data set comprised of the first four partitions several times, each time with a different combination of hyperparameters.[6] I test the accuracy of a given model fit by computing the accuracy rate of the trained model on the fifth partition (the "holdout" or "validation" set). I repeat this using each of the other possible combinations of four partitions as training data. Then, for each set of hyperparameter values, I compute the average accuracy rate (area under the ROC curve) across the five training/validation splits. I select the combination of hyperparameters that performed best and then re-fit the BART model on the entire training data set (all five partitions, i.e. without a validation set).

Once these steps are complete, the BART model remains unchanged. I apply the model to the 5% of the original data that were not used for model fitting (the test set) to obtain an estimate of the model's out-of-sample accuracy at the individual level. I also fit the model on the training data to calculate its in-sample accuracy. As long as the model's accuracy isn't drastically higher on the training sample than the test sample, overfitting is not great enough to be a cause for concern.

After the individual-level BART model has been fit, I use it to generate predictions for all demographic cells in the data (e.g. ethnically Ewe men with no formal education, aged 18-25, living in an urban area of the Gomoa East district). I do this separately for 2008 and 2014 to get estimates that vary over time for a panel analysis. I chose these years for two reasons. First, I can only generate estimates for years that Afrobarometer did surveys in Ghana that asked the question I use and for which geocoded data was released (2002, 2005, 2008, 2012, and 2014). Second, I want estimates of public opinion on the importance of public goods provision from before an electoral term's public goods have been provided. So, for the 2012 election (first period in the panel analysis), I use predictions from 2008 because

[6]The hyperparameter combinations were chosen at random, with possible values of each hyperparameter chosen from suggestions in Chipman, George, and McCulloch (2010) and Kapelner and Bleich (2013)

the (May) 2012 figures may have already been affected by the public goods provision I measure for that time period. For the 2016 election (second period), I use the 2014 data because the 2012 data is from before the 2012 election and would likely be too out of date by the time public goods provision in the lead up to the 2016 election was taking place.

Finally, the post-stratification step combines and weights these predictions based on the underlying population figures in each district. For example, the predicted proportion for a demographic group that makes up 5% of a district's voting age citizen population receives a 5% weight when calculating the district-level estimate. The general equation for this step is:

$$\hat{\theta}_{it} = \sum_{g=1}^{n} \hat{\gamma}_{git} \cdot \rho_{gi}$$

where $\hat{\theta}_{it}$ is the estimated proportion in district i at time t. The estimated proportion of interest in group g is given by $\hat{\gamma}_{git}$ and the proportion of a district's population composed of people in group g is ρ_{gi}.[7]

Propagating Uncertainty

It is not sufficient to simply plug the district opinion estimates obtained by BARP into a regression because BARP (and MRP) estimates are model based and uncertain. To propagate the uncertainty from the BARP estimation process into subsequent regression models, I use the MCMC samples produced during BART model fitting. These samples are draws from the posterior distribution, so they reflect both the central tendency of estimates of district level opinion derived from them and the uncertainty in those estimates. I conduct the post-stratification step for each iteration of the BART algorithm performed after the

[7]The proportion of a district's population made up of people in a given group does not vary over time because it is calculated using census data, which is only collected every ten years. While an approximation exist in the US context for non-census years (i.e. the ACS), such a survey does not exist in Ghana.

burn-in period, resulting in $1,000$ estimates for each district.[8] These estimates are closer together for districts in which there is less uncertainty (e.g. those with larger numbers of observations).

I then run the regression models discussed below once for each of the $1,000$ samples from the posterior. In each regression model i, the estimate for a given district is that produced by the ith iteration of the BART algorithm. The resulting $1,000$ regression coefficients therefore reflect the uncertainty in the BART estimates.

However, the question remains how to calculate a single set of standard errors and confidence intervals for the regression parameters across these $1,000$ models. One approach would be to simply calculate the standard deviation and 95% credible interval for the $1,000$ regression coefficients. But this does not account for the uncertainty in each coefficient from each regression. So, for each regression model i, I take the further step of simulating 1,000 draws from a multivariate normal distribution with vector of means given by the vector of regression coefficients for model i and variance-covariance matrix that of the same regression model. This ensures that the uncertainty in coefficient estimates for each regression can be reflected in a single standard error/confidence interval for each parameter. These are calculated by taking the standard deviation of the $1{,}000 \times 1{,}000$ coefficients and the 2.5^{th} and 97.5^{th} percentiles of their distribution.

Identification and Estimation

The individual regression models I fit are similar to the models of presidential voting and DDF transfers in Chapter 4. I estimate two separate single period models, both with and without controls, where I interact district-level opinion on the importance of public goods with the DDF transfers variable. I also estimate two-period panel models with DDF transfers in a given period interacted with the public opinion estimates, including district

[8]The best combination of hyperparameters for BART included $1,000$ as the number of iterations.

and time fixed effects.

In order to interpret the effect estimates in the cross-sectional models as causal, we must assume that there are no confounding variables that affect both DDF transfers and presidential vote share after including control variables. The identifying assumption in the fixed effects models is that there are no *time-variant* variables omitted from the model that affect both DDF transfers and presidential vote share. Because I am only interested in how the effect of DDF transfers depends on different levels of a baseline covariate (public opinion towards public goods), an identification strategy for that covariate is not necessary (Keele and Stevenson 2020).

6.4 Results

I begin by discussing the results of the BARP estimation of district-level public opinion towards public goods provision. I then present results of the regression models using these estimates.

BARP Estimates

Figure 6.1 shows the estimates of the importance of infrastructure for all 170 districts for one year (2008), ordered from least to greatest. Even though BARP estimates will be pulled closer to the mean due to model regularization, we still observe considerable variation in the importance of infrastructure across districts. The lowest estimates of the proportion of citizens who think infrastructure is one of the most important problems are in the single digits, whereas the highest approach 50%. Uncertainty is somewhat high for most districts, however, with 95% credible intervals around 20 percentage points. This demonstrates the importance of propagating measurement uncertainty into regression models.

The BARP and disaggregation estimates are compared in maps of Ghana's districts in Figure 6.2. The left panel shows the BARP estimates of the proportion of adult citizens

Figure 6.1: Importance of Infrastructure–BARP Estimates

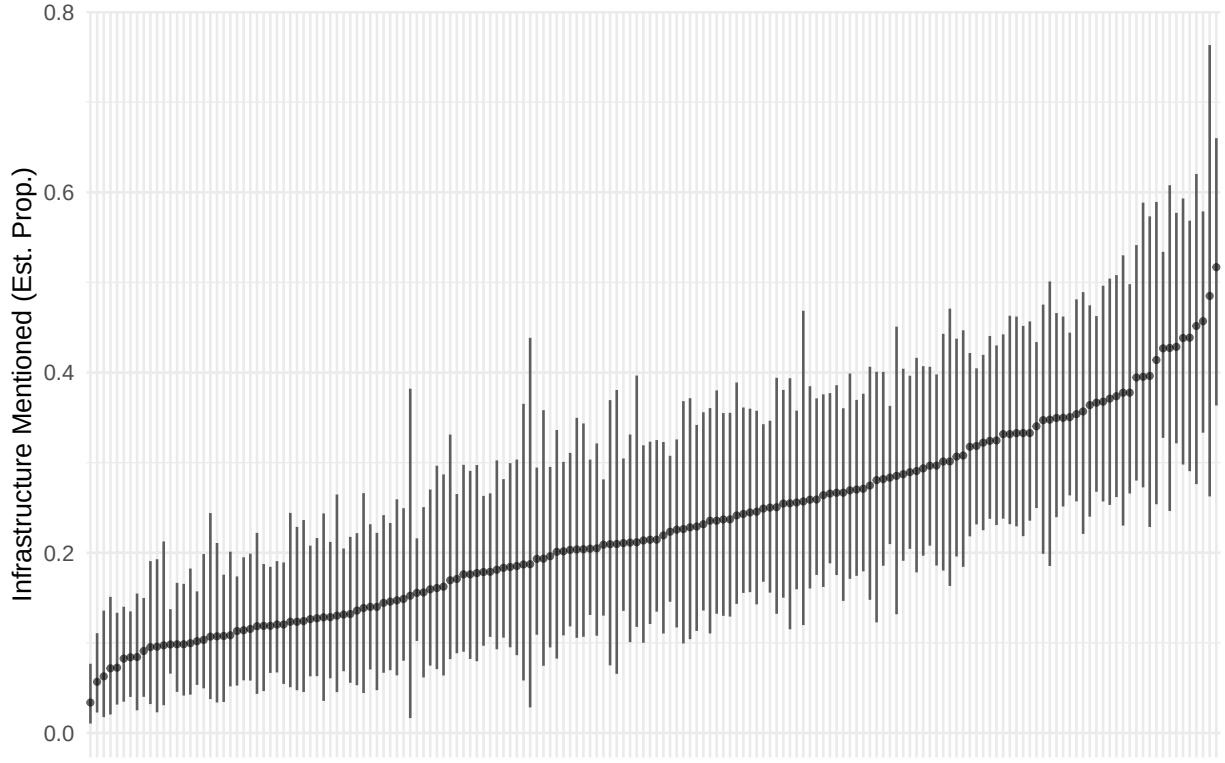

Note: Each point is one district's estimate from 2008, ordered from least to greatest. Error bars are 95% credible intervals.

in 2008 who would infrastructure as one of the three most important problems facing Ghana that the government should address, while the right panel shows the same for disaggregation. The first thing to notice is that the disaggregation approach cannot produce estimates for districts in which Afrobarometer did not survey any respondents in 2008. BARP, on the other hand, produces estimates for all districts in 2008 (even the two districts for which there are no respondents in any year). Second, there are many extreme values in the disaggregation map (due to the small sample sizes in many districts). The BARP map, though, shows few extreme estimates. This is because BARP pulls district estimates towards their expected values (conditional on predictors and population distributions), especially for those with few respondents in the data. Both of these differences are due in large part to the ability of BARP to pool data across all survey waves/years.

Figure 6.2: BARP vs. Disaggregation–Mapped

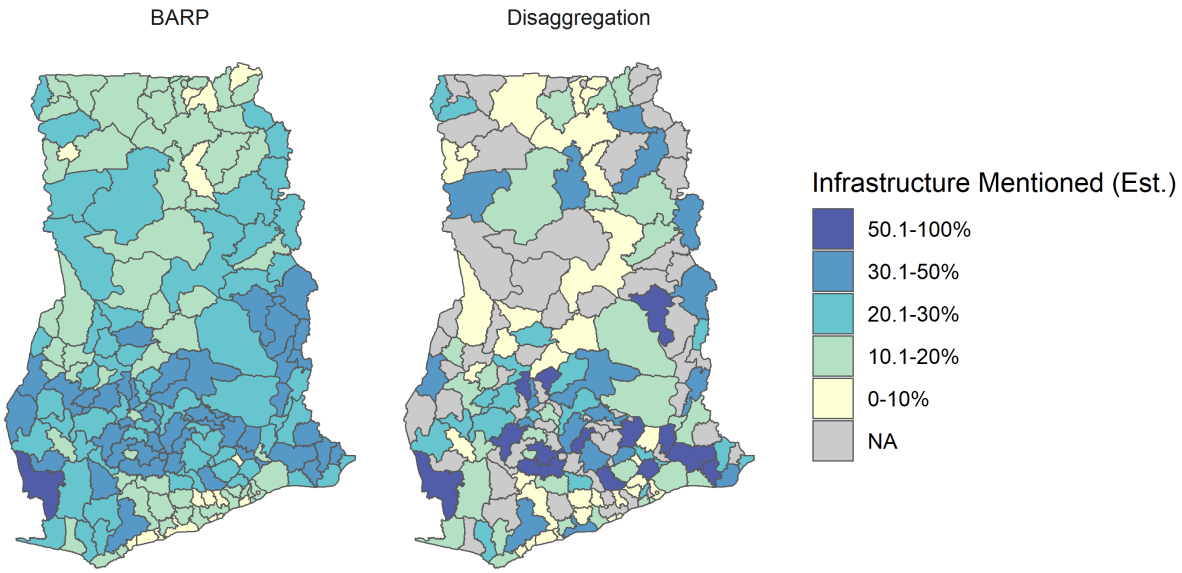

Note: Maps show estimates of the proportion of adult citizens in 2008 who would cite infrastructure as one of the three most important problems facing Ghana that the government should address. Missing values in the disaggregation plot appear because no respondents in those districts were surveyed in 2008, one disadvantage of that approach.

Validity of the Measure

Theoretically, the measure of preferences for infrastructure/public goods I present seems reasonable. As more survey respondents in a district say that infrastructure is an important problem facing the country, the importance of infrastructure to the average voter in that district should be higher. Further, this seems to be the best way to measure preferences with the data that is currently available for Ghana (and presumably many other countries).

However, it is also worth engaging in a data-driven examination of the validity of this measure. Although, to the best of my knowledge, there are no other survey questions asking people their preferences over infrastructure/public goods provision, there are

district-level measures with which preferences should be correlated. For instance, more urbanized districts should have more and better existing public goods, which would lead to people placing relatively less importance on additional infrastructure spending. I calculate urbanization as the proportion of the population in a district living in an urban area using micro data from Ghana's 2010 census.

I compare both the disaggregated (simple averages) and BARP measures of preferences for infrastructure from 2008 and 2012 to the proportion of the district population that lives in an urban area in 2010. Both versions of the preference measure are negatively correlated with urbanization. For 2008 the correlation with the disaggregation measure is –.31 and with the BARP measure –.44. For 2012 the correlations are –.07 and –.14.

Other variables that should be correlated with preferences for public goods are the quality and quantity of existing public goods. I test whether this is the case in the context of education public goods, using data from Ghana's Education Management Information System (EMIS). First, I take the average number of primary and junior high school classrooms in need of repairs for each district for the 2011-2012 school year.[9] I compare this to the survey-based measures of preferences. Here, we also see the expected relationship: districts where schools are in greater need of repairs score higher on the importance of infrastructure measures. The correlation with the BARP measure is .15 and with the disaggregated measure .02.

However, the measure of the quantity of education public goods—the pupil-classroom ratio—is negatively correlated with the importance of infrastructure measures. This is surprising because one would think that in districts with more pupils per classroom the importance of infrastructure would be higher. However, the relationship becomes close to zero when estimated using regression with a full set of controls and region fixed effects included. The urbanization and classrooms needing repair measures, though, tend to

[9]This is when one of the panel observations for the preference measure occurs (2012). I do not use the earlier one (2008) because the school data is reported at the previous district level for that year (when there were only 138 districts instead of the 170 I use in my analyses).

maintain their observed relationships with the infrastructure preference measures when using regression with the same set of controls.

This evidence taken together lends support to the validity of the proposed measure for capturing preferences for infrastructure.

Regression Models

I report regression results in Table 6.1. While the coefficient of interest (the interaction between public opinion towards infrastructure and the log of DDF transfers) remains positive across specifications, it does not reach conventional levels of statistical significance. In Model 1, the coefficient magnitude is highest (without any fixed effects), but the standard error is also quite high. Adding district fixed effects in Model 2 reduces the coefficient to nearly zero, though the standard error shows the estimate is quite imprecise. In Model 3, I include both district and election year (time) fixed effects. The coefficient of interest is positive with moderate magnitude and its distribution is visualized in Figure 6.3. A test of differential effects in 2012 vs 2016 does not reveal a strong relationship, with a negative but not precisely estimated or significant coefficient on the triple interaction term.

I further explore the results of Model 3 (both fixed effects) in Figure 6.3, which shows the distribution of the coefficient of interest (the interaction between opinion towards infrastructure and DDF transfers). These values come from simulating 1,000 values of the coefficients for each of the $1,000$ regressions I fit (one regression for each value of the infrastructure opinion variable from the BART posterior samples). The simulated coefficients are draws from a multivariate normal distribution with means and standard deviations taken from each regression model estimate. Although the 95% interval crosses zero, the bulk of the distribution (about 91%) rests above zero, providing some evidence that the effect of DDF transfers on incumbent presidential vote share depends on how important infrastructure is to people in a given district.

Table 6.1: Support for Incumbent President, Preferences, and DDF Transfers

	Model 1	Model 2	Model 3	Model 4
Infrastructure Opinion	-1.154	-0.237	-0.211	-0.293
	(0.688)	(0.200)	(0.123)	(0.180)
Log DDF Trans.	-0.010	0.043	-0.012	-0.021
	(0.061)	(0.022)	(0.013)	(0.016)
Elec. Yr. 2016			-0.053	-0.167
			(0.005)	(0.067)
Infr. Op. X Log DDF Trans.	0.284	0.002	0.054	0.067
	(0.228)	(0.065)	(0.041)	(0.056)
Infr. Op. X Log DDF Trans X Elec. Yr. 2016				-0.039
				(0.081)
Intercept	0.632	0.488	0.636	0.675
	(0.192)	(0.075)	(0.044)	(0.052)
District FEs	N	Y	Y	Y
Mean of DV	0.516	0.516	0.516	0.516
N	339	339	339	339

Note: One thousand regression models were run with each column's specification. For each regression, 1,000 draws of the coefficients were then simulated by taking random draws from the multivariate normal distribution. Coefficients are then means of one million parameter values and standard deviations are reported in parentheses. The dependent variable in all specifications is the incumbent party's (NDC's) vote share in the 2012 or 2016 presidential election. Elec. Yr. 2016 is a dummy variable that equals 1 for the 2016 election and 0 for the 2012 election. The remaining (double) interaction terms from Model 4 are not shown.

6.5 Discussion

I find modest evidence in support of the theory that the effect of public goods provision depends on voters' preferences on the importance of infrastructure. Although the main result does not cross the conventional threshold for statistical significance, that may be a rather high bar in this case. Given the level of regularization involved in smoothing public opinion estimates, the effect estimate is likely to be substantially attenuated, as it would be with MRP (Caughey and Warshaw 2019).

In regressions where uncertainty is not propagated (not shown), point estimates are similar to those reported in Table 6.3. But standard errors and p-values are smaller. This demonstrates the importance of propagating uncertainty when plugging model-based

Figure 6.3: Results: Distribution of Coefficient on Interaction Term

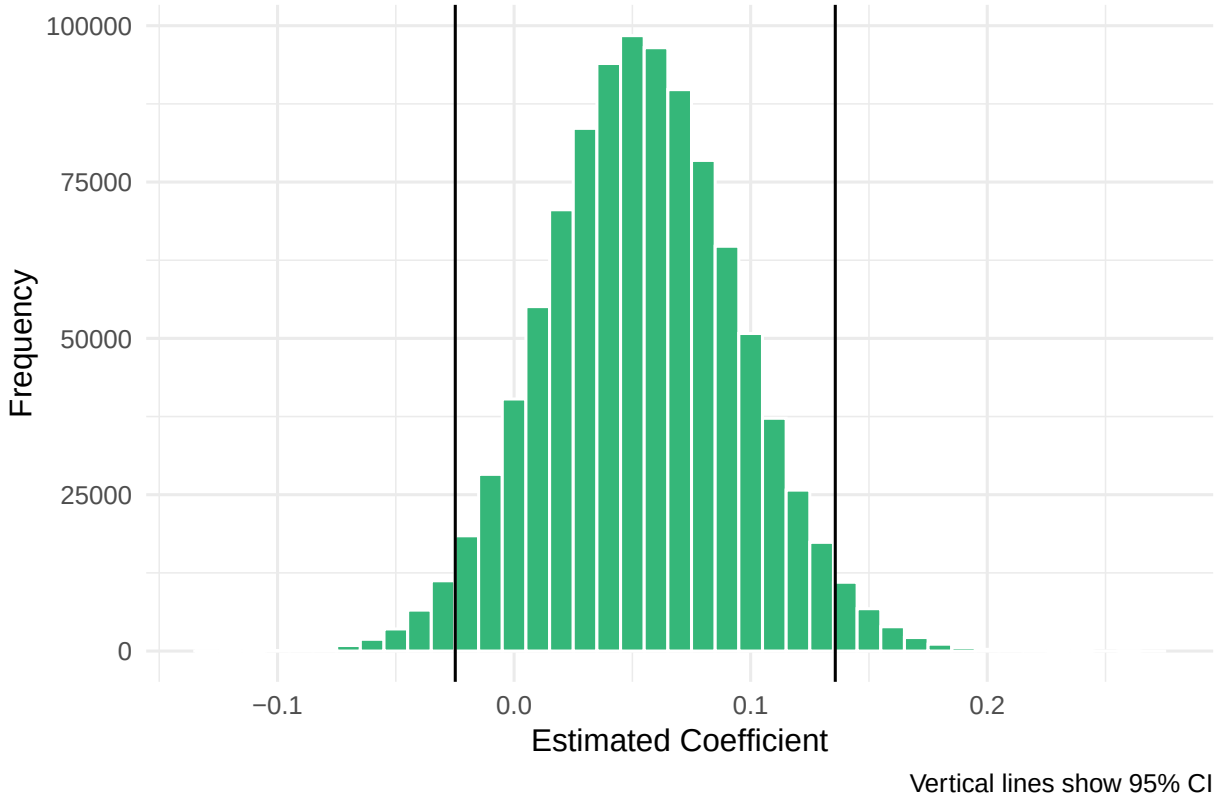

Note: Coefficient estimates are simulated from 1,000 regressions, one per iteration of the BART model. The specification is Model 3 from Table 6.1, where the dependent variable is incumbent presidential vote share and the coefficient of interest (plotted here) is the interaction between public opinion towards infrastructure and DDF transfers. District and year fixed effects are included.

estimates into regressions. Without doing so, I would risk overstating the confidence of the results.

A more complex approach to propagating uncertainty would involve weighting the regressions in the second stage of the process based on their likelihoods. This would allow information from the second stage to influence the first stage estimations by placing greater weight on the opinion estimates that produce models with higher likelihoods. I conduct a preliminary test to determine whether this approach would substantially change the results reported above. For Model 3 in Table 6.1, the coefficient of interest (on the interaction term in row 4) is nearly identical when using the weighting approach (.0542, vs .0541 in the unweighted result reported above). Because weighting does not result in a substantial

change to the mean of the parameter estimate, I do not undertake the more complex task of calculating weighted standard errors or confidence intervals.

As an alternative measure of voter preferences, I use a variable that captures whether a survey respondent cited education as one of the most important problems facing the country. However, the individual-level BART model performed quite poorly for this variable. The out of sample performance measured by the area under the ROC curve did not go above .6 during model tuning cross-validation and there was evidence of overfitting in all hyperparameter configurations used. Therefore, I do not continue with prediction and post-stratification for measuring preferences for education at the district level.

As noted above, BART is not the only machine learning method that can be used the replace the multilevel model in in MRP. Bisbee (2019) chose to focus on BART because it performed best overall in the US data he examined. To determine whether other models may perform better on the data I use, I implemented a logistic regression and a random forest model. Neither outperformed BART in predictive performance for the first-stage (individual-level) model, nor were they well equipped to deal with the numerous missing values present in the data. Broniecki, Leemann, and Wüest (n.d.) use an ensemble approach (fitting and averaging multiple machine learning models) and also show improvement over MRP (and over BARP). However, their approach is more complex, as it relies on several underlying models, and getting uncertainty estimates via bootstrapping is prohibitively time-consuming.[10] While BART models take longer to run than some other machine learning methods, uncertainty estimates are easier to calculate because BART uses Markov chain Monte Carlo (MCMC), which produces draws from the posterior distribution that reflect estimation uncertainty.[11] Therefore, while other methods may work better in other contexts, BART was the best choice here.

[10]The authors report in the vignette for their R package that fitting a single model and getting uncertainty estimates took 12 hours.

[11]Using the bartMachine package in R, the MCMC draws are saved in BART model objects, so producing credible intervals is computationally cheap.

6.6 Conclusion

In this chapter I applied an innovative method to test a novel research question: whether the effects of public goods provision depends on the varying preferences of voters. Using Bayesian additive regression trees with post-stratification (BARP), I estimated the average importance of infrastructure to voters in all of Ghana's districts at multiple points in time. I then used these estimates in regression models to explore whether they condition the effect of public goods provision on support for incumbents.

I find modest support for this theory. Across model specifications, the estimate of this conditional effect was positive. Although statistical significance was not reached, this is likely due to the attenuation caused by the regularization of estimates in the BARP method.

These results imply that voters' preferences are an important factor to consider when studying democratic accountability and performance voting. Researchers cannot assume that all voters actually place a high priority on the public goods and services examined in a particular study. Counter-intuitive findings, where voters appear to not reward or even punish governments for policies one might expect them to reward, should be scrutinized to determine whether the voters who *do* favor those policies are more likely to reward incumbents for them.

Chapter 7

Conclusion

A crucial aspect of democracy is the relationship between what voters want and what governments actually provide. If incumbent leaders do not govern according to the wishes of their constituents, it would lead us to question how well democracy is actually functioning. But if incumbents do provide the kind of governance their constituents want only for voters to fail to reward them electorally for doing so, it would discourage future incumbents from governing in accordance with popular will.

In this dissertation I have addressed these issues of democratic accountability and performance voting by looking at the effects of public goods provision on support for incumbent politicians in Ghana. I show that public goods provision, which is high on the list of citizen priorities across the developing world, can increase support for incumbent politicians. I also find some evidence that where more voters prioritize government action on infrastructure, incumbents are rewarded for public goods provision to a greater extent. However, voters do not appear to reward public goods provision in all contexts or reward all levels of government equally.

Despite these results that suggest a positive benefit for incumbents to transfers for public goods provision, the current incumbent president Nana Akufo-Addo has reduced the amount of transfers for local governments. Transfers under both the DDF and the

constitutionally-established district assemblies common fund (DACF) were cut in recent years. If these transfers for public goods provision increase support for incumbents, why don't national incumbents increase the provision of funds to local governments even more rather than cutting them?

The main reason I propose is that these local public goods projects are not the most efficient ways to increase national parties' support. When public goods projects are implemented by local governments, as the ones examined in this dissertation were, the president must share some of the credit with the local government. Even though district chief executives (DCEs) are appointed by the president and are loyal to his party, voters may assign some credit to the individual DCE rather than her party. Further, if local assembly members receive any electoral boost from these projects, the credit is spread even thinner. Contrast this with projects implemented by a national government ministry or other agency. In those cases, the president can claim nearly full credit for the implementation of the project.

It may also be necessary for incumbents to publicize their efforts to provide public goods, e.g. through the use of information campaigns to inform voters about them. Presumably, these funds would come from campaign war chests. But as Cruz et al. (2018) find, these information campaigns may be costly on a per-vote basis compared to clientelistic strategies like vote buying. Although vote buying is much more formalized in the context of that study (the Philippines), handouts of cash and other goods are common during campaigns in Ghana and across Africa (Bleck and van de Walle 2011). Parties may view these campaign expenditures as more cost-effective ways to get votes and increase base turnout than providing and publicizing public goods provision.

7.1 Other Contexts

I justified the choice of Ghana as a case in part based on external validity. So, what can these results tell us about other countries? The results for the effect of public goods provision on

support for incumbents should be most applicable to countries with politics and governance that are similar to Ghana's. So, other countries in Africa with stable party systems and high levels of partisanship, for instance. In these contexts, the effect of public goods provision will be limited because most voters will have a preferred party they will vote for no matter what kinds of public goods are provided by the incumbent. In countries with lower levels of partisanship (i.e. where there are more swing voters), we might expect the effect of public goods provision to be larger than what I find in Ghana. In those contexts, voters can identify public goods provision as coming from a particular party, which they can then choose to vote for or against in the next election. With more swing voters, there is more of a chance for public goods to influence vote choices.

In contexts where the national government is not as responsible for local public goods provision as it is in Ghana, local public goods should have a larger effect on support for local incumbents. In Ghana, local politicians have only a moderate ability to affect how much public goods are provided. Although I do not examine public goods provided entirely by the national government in this dissertation, it seems likely that those public goods would have a larger impact on support for national incumbents than the effects I identify here.

In contexts where voters have more concentrated preferences (i.e. strongly prioritize a few items), policies in those important areas should have a larger effect on support for incumbents. This is because voters with fewer strong preferences would not be basing their vote choice on as many factors as they would in contexts where voters tend to prioritize many different policy areas. If voters care about many different things, the effect of policy in any one area would be attenuated by their concern for performance in other policy areas.

7.2 Looking Ahead in Ghana

One challenge for democratic accountability in the context I examine is that voters do not directly choose their district chief executive (DCE, similar to a mayor). Instead, DCEs are

appointed by Ghana's president. While I do find evidence that the president receives more electoral support in districts that provided more public goods, voters must also consider a variety of other policies and performance measures when choosing whether and how to vote in presidential elections. This almost certainly results in a lower importance for local government performance in voting behavior.

A recent development could have improved this situation dramatically: a proposal to elect DCEs directly had been discussed in Ghana for some time. This would have allowed voters to take local governance into consideration when voting for DCEs without it being diluted by their views on national government performance. However, when the current opposition party (the NDC) came out against a related plan to make local elections partisan, the incumbent NPP president Nana Akufo-Addo scrapped plans to bring either question to a referendum. So, for the foreseeable future, Ghanaian voters will continue to only have an indirect say in who their DCEs are.

7.3 Areas for Future Research

It is important to acknowledge the areas in which future research could improve and expand upon these findings. While I measure monetary transfers earmarked for public goods provision, future studies should collect data on the actual public goods provided to different communities. A challenge that will face researchers in this endeavor is controlling for confounding variables, as it is likely there are political factors that affect which communities receive public goods and how much support incumbents receive there.

In contexts where a panel of survey respondents is available before and after an election, researchers could take an individual level approach to these questions. Determining whether individual voters in communities that receive more public goods are more likely to support incumbents in subsequent elections would be a valuable contribution to existing work. A further improvement on the preference measurement front would be fielding a

survey that specifically asks voters to rank the importance of different government policies to avoid drawbacks with the “most important problem" question.

Another major measurement issue that is beyond the scope of this dissertation is how communities are classified as urban and rural. Every country decides on its own definition, which presents a problem for cross-national research. But some definitions, like Ghana’s, present a problem even for subnational research. Any locality in Ghana with more than 5,000 inhabitants is classified as urban, which probably ends up over-counting urban residents. The field of geography has developed several ways of classifying spaces as urban and rural (and different categories in between) that should be incorporated into political science research. The availability of georeferenced survey data would enable such an approach in a variety of contexts.

Appendix A

Local Election Data

Results of district assembly (DA) elections are not officially available from the main office of the Electoral Commission of Ghana in the capital, Accra. Researchers have generally had to go to individual districts to get local election results, which would be impractical for a whole-country study of Ghana.[1] But, I was able to acquire detailed results from the 2015 DA elections from the regional offices of the Electoral Commission of Ghana, which are located in the regional capitals (there were 10 at the time of my fieldwork).

However, none of the regional offices had detailed results from the 2010 DA elections. Instead, an Electoral Commission official in Accra gave me the list of 2010 DA election winners that was published in the government's official gazette. I then tracked down the list of winners from the 2006 DA elections at the Ghana Law School Library in Accra. I processed all of this data (described below), which will be made publicly available online.[2]

The first step in processing the 2015 DA election data was aggregating the results from individual districts into one data set. This involved combining Word and Excel files, as well

[1]There were 170 districts in Ghana at the time of the 2010 DA elections and 216 at the time of the 2015 elections.

[2]The data will be posted to the author's website (https://www.evanmorier.com/) or the Harvard Dataverse (or both). Interested researchers may also email the author (evanpmorier@gmail.com) with questions or requests.

as using optical character recognition (OCR) software on scanned images of results I was only able to acquire in hard copy. Results from some electoral areas within districts had to be manually entered due to poor OCR results. The list of 2010 winners was only available in hard copy, which I scanned, OCR'ed, and converted to CSV. I manually compared observations of obvious poor quality to the original scanned documents and corrected them.

A variety of steps were taken to facilitate the linking of 2010 winners to 2015 candidates. First, I inserted a space into any name strings where the first and last names were combined without a space (indicated by no spaces in the name string and a pattern of "<lowercase><uppercase><lowercase>". I then converted names to lowercase and removed all digits, punctuation, and extraneous spaces. I used a similar process for electoral area (EA) names, though I left punctuation and digits in strings as they are relevant for some EA names.

I then implemented an algorithm to check each candidate name from 2015 to see if it matched a 2010 winner and could therefore be labeled an incumbent. Where I could link an electoral area's name as reported in 2015 to one in 2010, candidate names were only compared within the EA to reduce the rate of false positives. For EAs without exact matches in 2010 I used a fuzzy string distance algorithm, the Jaro-Winkler distance[3] (Porter and Winkler 1997), to identify the closest matching EA name. For EAs without a fuzzy match, I compared candidate names to winners from an entire district in 2010.

I used the Jaro-Winkler (JW) distance to match candidate names. This algorithm worked well (better than others I tried) in initial tests. I chose a moderate value for the penalty factor (.15 on a scale of 0-.25), which determines the weight placed on the beginning of the string when matching. For two names to be considered a match, the JW distance between them must not exceed a given value. I tested a few different distance cutoffs, settling on .12

[3]The Jaro-Winkler distance has been shown to work particularly well when matching names (Christen 2006; Cohen, Ravikumar, and Fienberg 2003)

for comparisons within a matched EA and a more restrictive .06 for 2015 candidates in EAs that couldn't be linked to a 2010 EA.

Although this approach performed fairly well, it did not successfully match many incumbents whose first, last (and middle) names were recorded in a different order in 2010 and 2015. To address this issue, for all electoral areas where an incumbent was not identified among the 2015 candidates I rearranged each candidate's names and try matching each of the permutations. I used the closest matching permutation for each candidate; if more than one 2015 candidate in a given EA had a permutation that matched with a 2010 winner, I chose the candidate with the closest match (i.e. the lowest string distance). I found that the cutoffs that produced the lowest rate of false negatives with minimal false positives in a representative training sample were .12 for candidates in EAs that matched a 2010 EA and .075 for candidates that were compared to 2010 winners from all EAs in their district.

Appendix B

BARP Estimation and Validation

B.1 Supplementary Results from BARP Estimation

In this section, I report and discuss additional results from the BARP estimation procedure.

Figure B.1 shows the distributions of estimates for the main parameter of interest (the proportion of adult citizens in a district in 2008 who would cite infrastructure/roads as one of the three most important problems the government should address). It shows that BARP was less likely to produce extreme estimates. In particular, the number of districts with estimates between 0 and .08 is over six times as great with disaggregation. Extreme estimates are unlikely to be accurate and are instead a result of the inaccuracy inherent in constructing naive estimates from small sample sizes.

In Figure B.2, we can see how uncertainty in BARP estimates is driven by sample size. The uncertainty estimates on the y-axis are the widths of each district's 95% credible interval for the parameter estimate in 2008. The x-axis is the log of the number of respondents from that district in the sample across all waves of the Afrobarometer survey. In general, the BARP method was able to make more precise estimates for districts with more respondents interviewed in the overall sample.[1] This is due to the fact that the model has to rely more

[1]There is also a negative relationship between the number of respondents interviewed in 2008 specifically

Figure B.1: Distribution of Estimates–BARP and Disaggregation

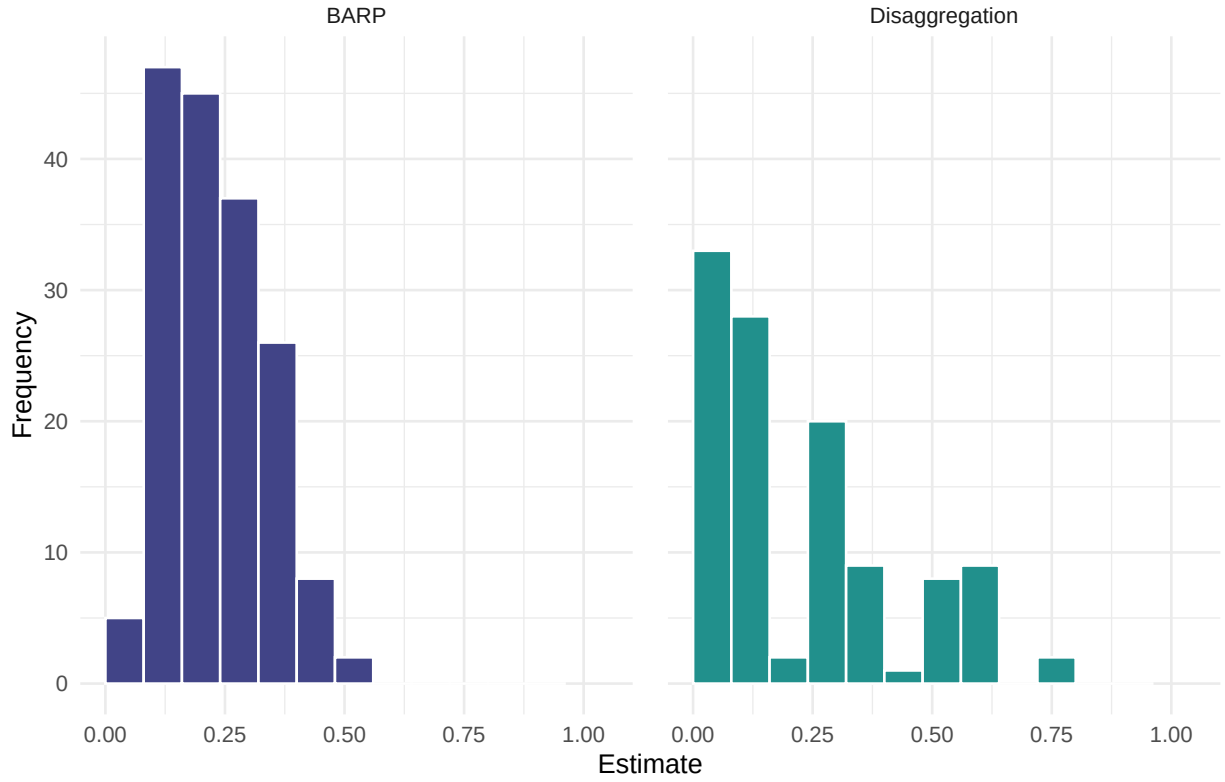

Note: Distributions of estimated proportion of adult citizens in a district in 2008 who would cite infrastructure/roads as one of the three most important problems the government should address. Disaggregation does not include estimates for districts where no respondents were interviewed by Afrobarometer in 2008 (56).

on observations from outside of the district to generate estimates for districts with fewer observations.

B.2 Validating MRP/BARP in an African Context

Multilevel regression and post-stratification (MRP) has been used extensively in the US context. Numerous studies have demonstrated its effectiveness at improving estimates of public opinion for geographic areas like states. Though it has been applied outside the US,

and the uncertainty for that estimate, but it is less clear visually because most districts have the same (small) number of respondents in a given year.

Figure B.2: Sample Size and Degree of Uncertainty in BARP Estimates

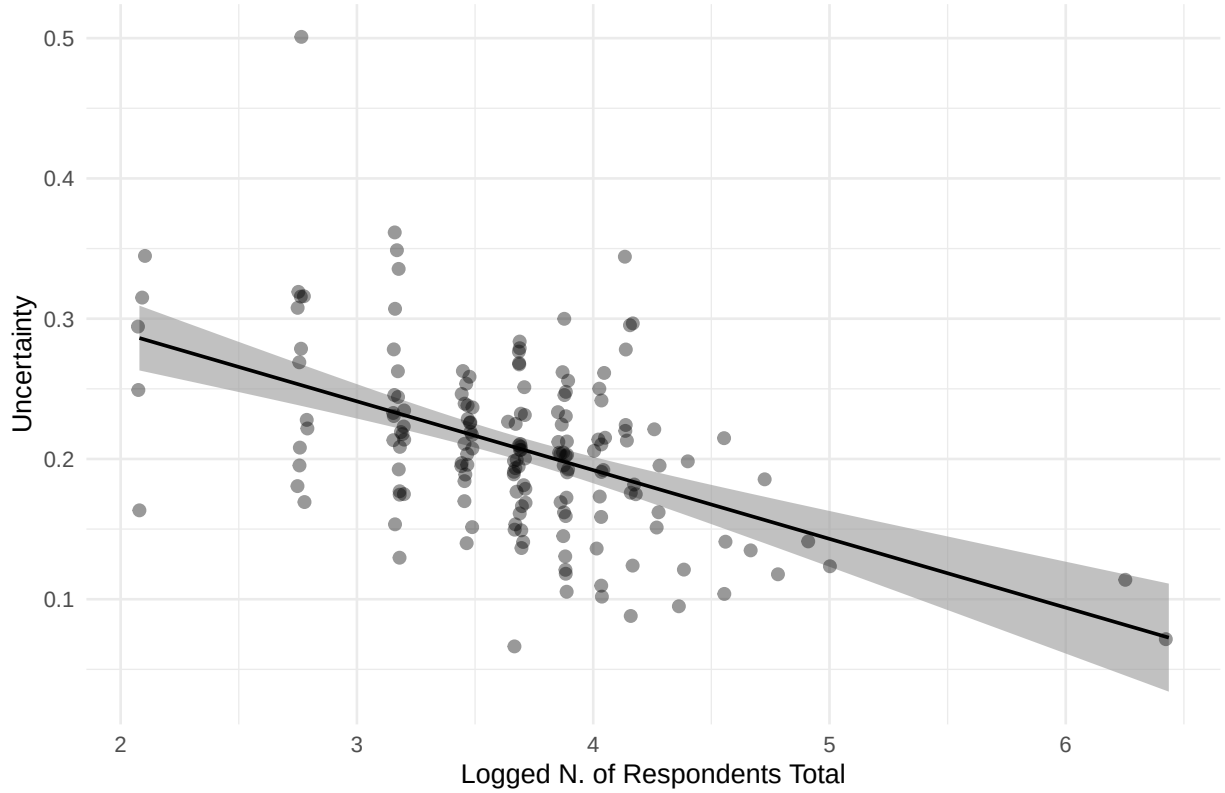

Note: Uncertainty is the width of a district's 95% credible interval as estimated using BARP. The number of respondents is pooled across all waves of Afrobarometer surveys. The parameter being estimated is the proportion of adult citizens in a district in 2008 who would cite infrastructure/roads as one of the three most important problems the government should address.

this dissertation is the first example to my knowledge of MRP or the related BARP method being applied in research on Africa.

There is no reason a priori to believe that MRP or BARP would not work or even perform worse in Africa. In any case, it is not possible to perform a completely realistic validation exercise due to the same data limitations that make MRP/BARP useful in the first place. Previous validation exercises combine many polls to create a super poll that serves as the population of a study. A random sample is then drawn to act as the test poll, MRP performed on it, and the resulting geographic estimates compared to the "true" population values from the super poll. But there is not enough public opinion data available to create

such a mega poll in Ghana.[2]

Nevertheless, it is important to at least demonstrate that the general approach works in the context of this dissertation. To do so, I use the Ghana census to estimate the proportion of internet users in each district. First, I take a regionally-representative sample from the 2010 Ghana Population and Housing Census with the same number of observations as my Afrobarometer sample. I then perform the same BARP procedure on this sample as I do on the Afrobarometer data, the only differences being that no year variable is used (the census is only from one year) and the outcome variable is the proportion of internet users in a district rather than the proportion who said infrastructure was a problem. The sample universe is the same (adult Ghanaian citizens). Finally, I compare these estimates both to the true values in each district (calculated on the full set of census microdata) and the estimates obtained by simple disaggregation (where the proportion is estimated by a simple mean of those surveyed from a district regardless of how small that sample is).

The results of this procedure are encouraging. First, the mean absolute error (MAE) of the BARP estimates (compared to the true values) was just .009, less than one percentage point. This shows that BARP estimates are quite accurate overall. By comparison, the MAE of the disaggregation estimates was .023, over 2.5 times as large. Further, in 83% of districts the BARP estimate was closer to the true value than the disaggregation estimate. And finally, when looking at the 95% credible intervals created by taking the 2.5th and 97.5th percentiles of the MCMC samples of each district's estimate, the district's true value is inside the estimated interval for 94% of districts (almost exactly 95%).

Further comparison of the BARP method to disaggregation (and the true values) is shown in Figure B.3. In the disaggregation (right) panel, points are spread more widely above and below the diagonal line (representing an estimate that is exactly the same as the true value). There are also many points for which $y = 0$ for disaggregation. By contrast, in

[2]Doing so for Africa as a whole would provide enough respondents but using MRP in a cross-national context, where each country in the sample already had nationally-representative Afrobarometer surveys taken, would not be very realistic.

the left panel (BARP estimates) the points are noticeably more tightly clustered around the diagonal line and there are none for which $y = 0$. This reflects BARP's ability to shrink estimates towards the true value.

Figure B.3: Accuracy of Estimation Methods in Census Simulation

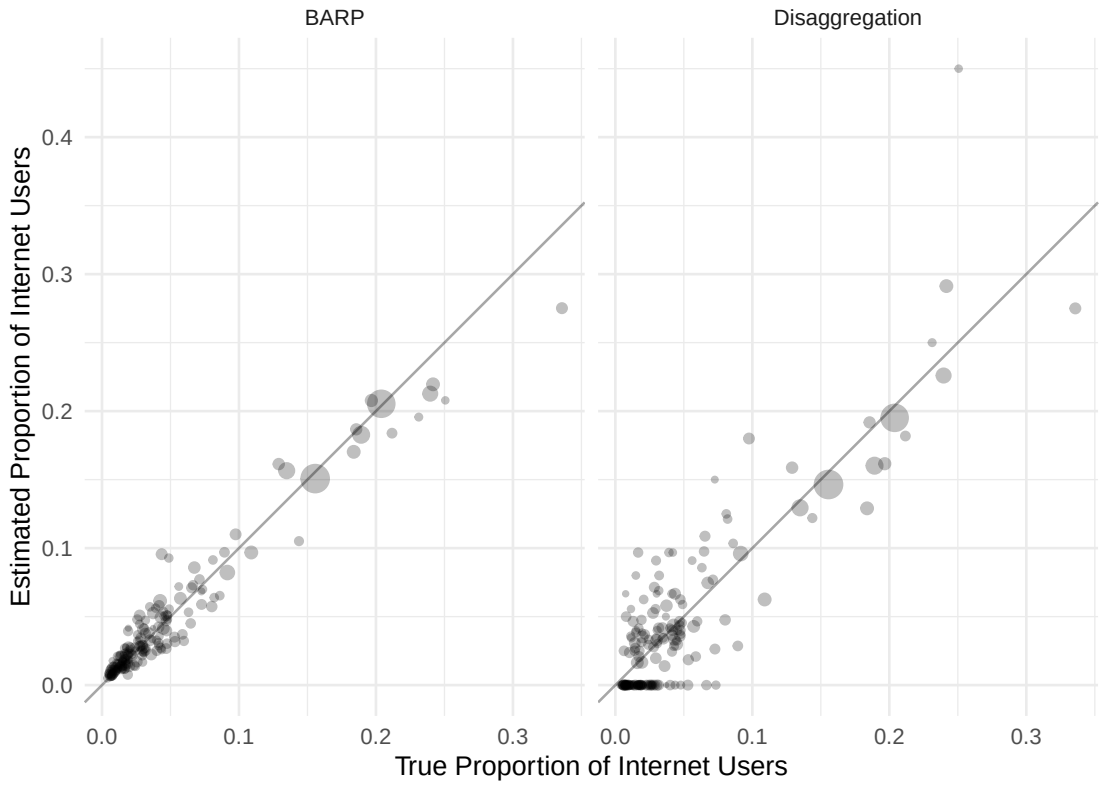

Note: Point size based on district population. Points on the diagonal line represent an estimate that is exactly the same as the true value. Points above (below) the line show an over- (under-) estimate.

There are some differences between this exercise and the analysis performed on Afrobarometer data. Here, I use demographics to predict a non-opinion variable (whether the person uses the internet), whereas I use BARP to predict an opinion question (an unavoidable difference, as no opinion questions are asked on the census). In addition, the internet user variable does not vary much across districts (SD = .06. vs. .14 for the infrastructure opinion variable across districts), so the BARP (and disaggregated) estimates may be more accurate than they would otherwise be. The standard deviation of the district-level internet user variable in the sample is .06, whereas a similarly-calculated standard deviation for the

Afrobarometer infrastructure variable is .14. This exercise also only uses one point in time while I predict for distinct years using the Afrobarometer data.

But there are also some differences that make this exercise a potential understatement of the power of BARP. Sampling from the population (census) was random, stratified on region, whereas Afrobarometer uses cluster sampling (one cluster is chosen and then eight households are selected in that area). This means that the disaggregated estimates from my simulation are more representative of districts than they would have been under clustered sampling. Further, there were individuals from all 170 districts in the sample, whereas this is never the case in a given Afrobarometer sample. So, a disaggregated estimate wouldn't even be possible for many districts. Finally, because the internet use rate was so low in most districts, the disaggregated estimate ended up being zero for 65 out of 170 districts. This floor effect means that the error rate of disaggregated estimates is lower than it would have otherwise been.

In sum, this exercise show that MRP/BARP can work on African data in general. Combined with existing evidence on MRP/BARP, this exercise shows that MRP/BARP can be successfully used to increase the accuracy of small-area estimation from non-representative samples in African contexts.

Bibliography

Achen, Christopher H., and Larry M. Bartels. 2004. "Blind Retrospection: Electoral Responses to Drought, Flu, and Shark Attacks."

———. 2016. *Democracy for Realists: Why Elections do not Produce Responsive Government.* Princeton, NJ: Princeton University Press.

Afrobarometer. 2017. "Ghanaians favour election of metropolitan, municipal, and district chief executives (MMDCEs)." `https://web.archive.org/web/20200929173009/https://afrobarometer.org/press/ghanaians-favour-election-metropolitan-municipal-and-district-chief-executives-mmdces`.

———. 2019. "Infrastructure/Roads Ghanaians' Most Important Problem, New Afrobarometer Study Shows." `https://web.archive.org/web/20191121121719/https://afrobarometer.org/press/infrastructureroads-ghanaians-most-important-problem-new-afrobarometer-study-shows`.

Altonji, Joseph G, Todd E Elder, and Christopher R Taber. 2005. "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools." *Journal of Political Economy* 113 (1): 151–184.

Arriola, Leonardo R., Donghyun Danny Choi, and Matthew Gichohi. 2017. "Political Endorsements and Cross-Ethnic Voting in Africa." *Working Paper.*

Ashworth, Scott. 2012. "Electoral Accountability: Recent Theoretical and Empirical Work." *Annual Review of Political Science* 15:183–201.

Ashworth, Scott, Ethan Bueno de Mesquita, and Amanda Friedenberg. 2018. "Learning About Voter Rationality." *American Journal of Political Science* 62 (1): 37–54.

Ayee, Joseph. 2013. "The Political Economy of the Creation of Districts in Ghana." *Journal of Asian and African Studies* 48 (5): 623–645.

Baldwin, Kate, and Matthew Winters. 2017. "Do Aid Flows that Bypass National Governments Alter Citizen and Elite Political Engagement?" In *International Political Economy Society Conference.*

Banful, Afua Branoah. 2011. "Do Formula-based Intergovernmental Transfer Mechanisms Eliminate Politically Motivated Targeting? Evidence from Ghana." *Journal of Development Economics* 96 (2): 380–390.

Barro, Robert J. 1973. "The Control of Politicians: An Economic Model." *Public Choice:* 19–42.

Bates, Robert. 1983. "Modernization, Ethnic Competition, and the Rationality of Politics of Contemporary Africa." In *State versus Ethnic Claims: African Policy Dilemmas*, edited by Donald Rothchild and Victor Olorunsola, 152–171. Boulder, CO: Westview.

Bechtel, Michael M, and Jens Hainmueller. 2011. "How Lasting Is Voter Gratitude? An Analysis of the Short-and Long-Term Electoral Returns to Beneficial Policy." *American Journal of Political Science* 55 (4): 852–868.

Besley, Timothy. 2006. *Principled Agents?: The Political Economy of Good Government.* Oxford: Oxford University Press.

Besley, Timothy, and Robin Burgess. 2002. "The Political Economy of Government Responsiveness: Theory and Evidence from India." *The Quarterly Journal of Economics* 117 (4): 1415–1451.

Bhavnani, Rikhil, and Mark Copelovitch. 2018. "The Political Impact of Monetary Shocks: Evidence from India's 2016 Demonetization."

Bisbee, James. 2019. "BARP: Improving Mister P Using Bayesian Additive Regression Trees." *American Political Science Review* 113 (4): 1060–1065.

Blattman, Christopher, Mathilde Emeriau, and Nathan Fiala. 2016. "Do Anti-poverty Programs Sway Voters? Experimental Evidence from Uganda."

Bleck, Jaimie, and Nicolas van de Walle. 2011. "Parties and Issues in Francophone West Africa: Towards a Theory of Non-Mobilization." *Democratization* 18 (5): 1125–1145.

Bratton, Michael, Ravi Bhavnani, and Tse-Hsin Chen. 2012. "Voting Intentions in Africa: Ethnic, Economic or Partisan?" *Commonwealth and Comparative Politics* 50 (1): 27–52.

Bratton, Michael, and Mwangi S. Kimenyi. 2008. "Voting in Kenya: Putting Ethnicity in Perspective." *Journal of Eastern African Studies* 2 (2): 272–289.

Bratton, Michael, Robert Mattes, and Emmanuel Gyimah-Boadi. 2005. *Public Opinion, Democracy, and Market Reform in Africa.* Cambridge: Cambridge University Press.

Bratton, Michael, and Nicholas Van de Walle. 1997. *Democratic Experiments in Africa: Regime Transitions in Comparative Perspective.* Cambridge: Cambridge University Press.

Briggs, Ryan C. 2018. "Receiving Foreign Aid Can Reduce Support for Incumbent Presidents." *Political Research Quarterly.*

Broniecki, Philipp, Lucas Leemann, and Reto Wüest. n.d. "Improved Multilevel Regression with Post-Stratification Through Machine Learning (autoMrP)."

Bursztyn, Leonardo. 2016. "Poverty and the Political Economy of Public Education Spending: Evidence from Brazil." *Journal of the European Economic Association* 14 (5): 1101–1128.

Canes-Wrone, Brandice, David W Brady, and John F Cogan. 2002. "Out of Step, Out of Office: Electoral Accountability and House Members' Voting." *American Political Science Review* 96 (1): 127–140.

Carlson, Elizabeth. 2015. "Ethnic Voting and Accountability in Africa: A Choice Experiment in Uganda." *World Politics* 67 (2): 353–385.

Cattaneo, Matias D., Nicolás Idrobo, and Rocío Titiunik. 2018. "A Practical Introduction to Regression Discontinuity Designs: Volume II." In *Cambridge Elements: Quantitative and Computational Methods for Social Science.* Cambridge University Press.

———. 2019. "A Practical Introduction to Regression Discontinuity Designs: Foundations." In *Cambridge Elements: Quantitative and Computational Methods for Social Science.* Cambridge University Press.

Caughey, Devin, and Christopher Warshaw. 2019. "Public Opinion in Subnational Politics." *The Journal of Politics* 81 (1): 352–63.

Chipman, Hugh A, Edward I George, and Robert E McCulloch. 2010. "Bart: Bayesian Additive Regression Trees." *The Annals of Applied Statistics* 4 (1): 266–298.

Christen, Peter. 2006. "A Comparison of Personal Name Matching: Techniques and Practical Issues." In *Sixth IEEE International Conference on Data Mining-Workshops (ICDMW'06)*, 290–294. IEEE.

Cohen, William W, Pradeep Ravikumar, and Stephen E Fienberg. 2003. "A Comparison of String Distance Metrics for Name-Matching Tasks." In *American Association for Artificial Intelligence*, 2003:73–78.

Cole, Shawn, Andrew Healy, and Eric Werker. 2012. "Do Voters Demand Responsive Governments? Evidence from Indian Disaster Relief." *Journal of Development Economics* 97:167–181.

Croke, Kevin. 2017. "The Impact of Mass Bed Net Distribution Programs on Politics: Evidence from Tanzania."

Cruz, Cesi, Philip Keefer, Julien Labonne, and Francesco Trebbi. 2018. "Making Policies Matter: Voter Responses to Campaign Promises." *National Bureau of Economic Research Working Paper.*

Cruz, Cesi, and Christina J. Schneider. 2014. "Foreign Aid and Undeserved Credit-Claiming." *American Journal of Political Science* 61 (2): 396–408.

De Kadt, Daniel, and Evan S Lieberman. 2017. "Nuanced Accountability: Voter Responses to Service Delivery in Southern Africa." *Forthcoming, British Journal of Political Science.*

De La O, Ana L. 2013. "Do Conditional Cash Transfers Affect Electoral Behavior? Evidence from a Randomized Experiment in Mexico." *Journal of Political Philosophy* 57 (1): 1–14.

Debrah, Emmanuel. 2014. "The Politics of Decentralization in Ghana's Fourth Republic." *African Studies Review:* 49–69.

Driscoll, Barry. 2015. "The Perverse Effects of Political Competition: Building Capacity for Patronage in Ghana." PhD diss., University of Wisconsin–Madison.

Duch, Raymond M, and Randolph T Stevenson. 2008. *The Economic Vote: How Political and Economic Institutions Condition Election Results.* New York: Cambridge University Press.

Dunning, Thad, Guy Grossman, Macartan Humphreys, Susan Hyde, and Craig McIntosh. 2018. *Metaketa I: The Limits of Electoral Accountability.* New York: Cambridge University Press. Forthcoming.

Dunning, Thad, and Lauren Harrison. 2010. "Cross-Cutting Cleavages and Ethnic Voting: An Experimental Study of Cousinage in Mali." *American Political Science Review* 104 (1): 21–39.

Ferejohn, John. 1986. "Incumbent Performance and Electoral Control." *Public Choice* 50 (1): 5–25.

Ferraz, Claudio, and Frederico Finan. 2011. "Electoral Accountability and Corruption: Evidence from the Audits of Local Governments." *The American Economic Review* 101 (4): 1274–1311.

Fiorina, Morris. 1981. *Retrospective Voting in American National Elections.* New Haven, CT: Yale University Press.

Gottlieb, Jessica. 2016. "Greater Expectations: A Field Experiment to Improve Accountability in Mali." *American Journal of Political Science* 60 (1): 143–157.

Grossman, Guy, Macartan Humphreys, and Gabriella Sacramone-Lutz. 2014. "" I wld like u WMP to extend electricity 2 our village": On Information Technology and Interest Articulation." *American Political Science Review* 108 (3): 688–705.

Gyimah-Boadi, Emmanuel. 1999. "Institutionalizing Credible Elections in Ghana." In *The Self-Restraining State: Power and Accountability in New Democracies,* edited by Andreas Schedler, Larry Diamond, and Marc F. Plattner, 105–121. Boulder, CO: Lynne Rienner Publishers.

Harding, Robin. 2015. "Attribution and Accountability: Voting for Roads in Ghana." *World Politics* 67 (4): 656–689.

Harding, Robin, and David Stasavage. 2013. "What Democracy Does (and Doesn't Do) for Basic Services: School Fees, School Inputs, and African Elections." *The Journal of Politics* 76 (1): 229–245.

Healy, Andrew, and Neil Malhotra. 2013. "Retrospective Voting Reconsidered." *Annual Review of Political Science* 16:285–306.

Hicken, Allen. 2011. "Clientelism." *Annual Review of Political Science* 14:289–310.

Horowitz, Donald. 1985. *Ethnic Groups in Conflict.* Berkeley: University of California Press.

Ichino, Nahomi, and Noah L Nathan. 2013. "Crossing the Line: Local Ethnic Geography and Voting in Ghana." *American Political Science Review* 107 (2): 344–361.

Ichino, Nahomi, Martin Williams, and Erik Wibbels. 2018. "The Political Geography of Government Projects: Evidence from +/− 40,000 Projects in Ghana." *AidData Working Paper #58.*

Imai, Kosuke, Gary King, and Carlos Velasco Rivera. 2017. "Do Nonpartisan Programmatic Policies Generate Partisan Electoral Effects? Evidence from Two Large Scale Randomized Evaluations."

Jeffries, Richard, and Clare Thomas. 1993. "The Ghanaian Elections of 1992." *African Affairs* 92 (368): 331–366.

Kapelner, Adam, and Justin Bleich. 2013. "bartMachine: Machine Learning with Bayesian Additive Regression Trees." *arXiv e-prints* 1312.2171, arXiv:1312.2171. arXiv: 1312.2171. https://arxiv.org/abs/1312.2171.

———. 2014. "Prediction with Missing Data via Bayesian Additive Regression Trees." *arXiv e-prints* 1306.0618, arXiv:1306.0618. arXiv: 1306.0618. https://arxiv.org/abs/1306.0618.

Keele, Luke, and Randolph T. Stevenson. 2020. "Causal Interaction and Effect Modification: Same Model, Different Concepts." *Political Science Research and Methods:* 1–9.

Key, Valdimer Orlando, et al. 1966. *The Responsible Electorate.* Harvard University Press Cambridge, MA.

Kitschelt, Herbert, and Steven I Wilkinson. 2007. *Patrons, Clients and Policies: Patterns of Democratic Accountability and Political Competition.* New York: Cambridge University Press.

Kuenzi, Michelle, and Gina MS Lambright. 2011. "Who Votes in Africa? an Examination of Electoral Participation in 10 African Countries." *Party Politics* 17 (6): 767–799.

Labonne, Julien. 2013. "The Local Electoral Impacts of Conditional Cash Transfers: Evidence from a Field Experiment." *Journal of Development Economics* 104:73–88.

Lax, Jeffrey R, and Justin H Phillips. 2009. "How Should We Estimate Public Opinion in the States?" *American Journal of Political Science* 53 (1): 107–121.

Lee, David S., and Thomas Lemieux. 2010. "Regression Discontinuity Designs in Economics." *Journal of Economic Literature* 48 (June): 281–355.

Lipset, Seymour Martin, and Political Man. 1960. *The Social Bases of Politics.* Baltimore: The Johns Hopkins University Press.

McCrary, Justin. 2008. "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test." *Journal of Econometrics* 142 (2): 698–714.

Nathan, Noah L. 2016. "Local Ethnic Geography, Expectations of Favoritism, and Voting in Urban Ghana." *Comparative Political Studies* 49 (14): 1896–1929.

———. 2019a. "Electoral Consequences of Colonial Invention: Brokers, Chiefs, and Distribution in Northern Ghana." *World Politics* 71, no. 3 (July): 417–456.

———. 2019b. *Electoral Politics and Africa's Urban Transition: Class and Ethnicity in Ghana.* Cambridge University Press.

Ndegwa, Stephen N. 2002. "Decentralization in Africa: A Stocktaking Survey." World Bank Africa Region Working Paper Series (40).

Nunn, Nathan, and Leonard Wantchekon. 2011. "The Slave Trade and the Origins of Mistrust in Africa." *American Economic Review* 101 (7): 3221–3252.

Ofosu, George. 2021. "When Do Partisans Cross the Party Line?" `https://www.georgeofosu.com/files/ofosu_crossing_partylines1116112020.pdf`.

Pande, Rohini. 2011. "Can Informed Voters Enforce Better Governance? Experiments in Low-Income Democracies." *Annual Review of Economics* 3 (1): 215–237.

Pop-Eleches, Cristian, and Grigore Pop-Eleches. 2012. "Targeted Government Spending and Political Preferences." *Quarterly Journal of Political Science* 7 (3): 285–320.

Porter, Edward H, and William E Winkler. 1997. "Approximate String Comparison and Its Effect on an Advanced Record Linkage System." In *US Bureau of the Census, Research Report.* Citeseer.

Posner, Daniel N. 2005. *Institutions and Ethnic Politics in Africa.* New York: Cambridge University Press.

Przeworski, Adam, Susan C Stokes, and Bernard Manin. 1999. *Democracy, Accountability, and Representation.* New York: Cambridge University Press.

Riedl, Rachel Beatty. 2014. *Authoritarian Origins of Democratic Party Systems in Africa.* New York: Cambridge University Press.

Stokes, Susan C, Thad Dunning, Marcelo Nazareno, and Valeria Brusco. 2013. *Brokers, Voters, and Clientelism: The Puzzle of Distributive Politics.* New York: Cambridge University Press.

Tsai, Lily, Minh Trinh, and Shiyao Liu. 2017. "The Role of Retributive Justice in Citizen Evaluations of Government: The Case of China."

van de Walle, Nicolas. 2007. "Meet the Old Boss, Same as the New Boss? The Evolution of Political Clientelism in Africa." In *Patrons, Clients, and Policies: Patterns of Democratic Accountability and Political Competition,* edited by Herbert Kitschelt and Steven I. Wilkinson, 50–67. New York: Cambridge University Press.

Vries, Catherine E, and Nathalie Giger. 2014. "Holding Governments Accountable? Individual Heterogeneity in Performance Voting." *European journal of political research* 53 (2): 345–362.

Weghorst, Keith R., and Staffan I. Lindberg. 2013. "What Drives the Swing Voter in Africa?" *American Journal of Political Science* 57 (3): 717–734.

Williams, Martin. 2017. "The Political Economy of Unfinished Development Projects: Corruption, Clientelism, or Collective Choice?" *American Political Science Review* 111 (4): 705–723.

Wlezien, Christopher. 2005. "On the Salience of Political Issues: The Problem with 'Most Important Problem'." *Electoral Studies* 24 (4): 555–579.

World Bank. 2018. "Implementation Completion and Results Report IDA 48780 on a Credit in the Amount of SDR 112.1 Million (US $175 Million Equivalent) to the Republic of Ghana for the Ghana Local Government Capacity Support Project," no. ICR00004516.

Young, Crawford. 1976. *The Politics of Cultural Pluralism.* Madison: University of Wisconsin Press.