The Relevance of Relative Distribution: Favoritism, Information, and Vote Choice Africa

Elizabeth Carlson
Penn State
ecc13@psu.edu

October 23, 2015

Abstract

Relative distribution– how what one has received from government compares to what others have received – matters to African voters. In this article, I develop and test a model in which the salience of relative distribution is a function of Africa's particularly low-information context. Voters want to maximize their share of government provision, but without good information about available resources, cannot determine whether their share is high or low. However, they can use information about what others receive to make a second-best inference. I use an experimental game Mukono, Uganda, to show that in accordance with expectations, voters are more likely to reelect leaders who provide more to them than to another player. However, this effect only occurs when players are kept uninformed about the size of the pot being distributed. Those who know the resources available, and thus whether their share is high or low, no longer consider relative distribution in making their choice. Though relative distribution improves vote accuracy on average, it is not ideal, particularly when overall distribution is low: those who know they have been favored, but have no other information, can be induced to reelect a leader who is providing far less, and keeping far more, than a challenger would. The results imply that the salience of relative distribution to African voters is temporary, and will decline as as African governments become more transparent about budgets. They also imply that, until such transparency is achieved, providing voters in corrupt countries with information about relative distribution, while likely to change votes, will not necessarily improve accountability.

Many theories of voting in Africa highlight voters' preferences for leaders who will favor them in the distribution of scarce government resources. Getting access to resources that others are denied is known in many African countries as "eating" or "chopping" and getting a "turn to chop" is a common reason for supporting one candidate over another (Lindberg, 2003). That voters expect leaders from their own ethnic group, in particular, to favor them is arguably the most common explanation for why African voters prefer coethnic leaders (Posner, 2005; Ferree, 2006). However, others have found that African voters will support any candidate whom they expect will prioritize them or their communities in the distribution of resources, even if this leader is not a coethnic (Ichino and Nathan, 2013; Wantchekon, 2003).

An obvious explanation for this effect is that those who are favored by government will end up absolutely better off than they would have while being disfavored. However, there is evidence that African voters respond positively to favoritism, holding absolute outcomes constant. African voters' beliefs about whether their group is treated fairly by government remain a significant predictor of vote choice, controlling for perceived absolute poverty (Bratton, Bhavnani and Chen, 2011; Ferree and Horowitz, 2010). Additional evidence of the political relevance of relative outcomes is straightforward to show. Table 1 in the appendix, shows the determinants of presidential approval using data from an 18-country African public opinion survey.[1] The table shows that three measures of relative well-being[2] – the respondent's perceived personal wellbeing relative to others, the respondent's ethnic group's economic wellbeing relative to other groups, and the frequency with which the ethnic group is treated fairly by the government – are all meaningful predictors of presidential approval. Respondents' evaluations of their own absolute well-being, on the other hand, are *not* strongly related to approval ratings. We can thus begin from a premise that relative outcomes matter in and of themselves to African voters.[3]

A strong correlation between relative outcomes and political attitudes is not necessarily surprising. A large number of studies find that relative income is strongly correlated with subjective wellbeing and happiness (**??**) and that happiness predicts support for the incumbent (**?**). Consequently, studies find a correlation between relative income and support for the incumbent (Luttmer, 2004; **?**; **?**). On the other hand, the literature suggests that relative deprivation matters primarily to the rich, whose basic needs have all been met. **?** finds that while the subjective well-being of high-income Americans depends heavily on how their income compares to others in their cohort,

[1]The data is from Round Four of the Afrobarometer, which is the most recent round in which data are released pooled across countries.

[2]In this article, I will treat relative government distribution and relative outcomes as if they are interchangeable. This is because in many places in Africa, outcomes are, in fact, strongly dependent on government action. In most African countries the economy is predominantly informal and employment that provides a cash income is uncommon. Though private services exist in all sectors, the average citizen's ability to pay for these services is limited. Access to education, health care, clean water, and agricultural inputs are strongly influenced by whether government has provided or is subsidizing these goods.

[3]It is also possible, and even likely, that the correlations shown in these regressions are endogenous. Regime supporters both within and outside Africa tend to have much rosier views of reality than those who oppose the regime, and this can include biased perceptions of relative well-being (Logan et al., 2003; Carlson, 2014; Gerber and Huber, 2010). This is less of a concern for the findings I present here, however, as they are based on an experimental game in which relative distribution is exogenously manipulated and the political context is abstracted away.

the correlation between relative income and happiness for low-income Americans is almost zero. More importantly for this study, **?** finds that, in Malawi, the average respondents' relative economic well-being is not a significant determinant of happiness; a small subset of those to whom it does matter are objectively well-off. Since most people in most African countries are objectively poor, therefore, there may be another reason for the strong relationship between relative outcomes and political attitudes in Africa.

In this article, I offer an alternative explanation for the particular salience of relative distribution in African countries: in the absence of more useful information, voters can use relative distribution to make rough inferences about the availability of resources and the likelihood they would be better off under another leader with different preferences. In other words, though being favored by government may certainly have psychological benefits, voters' attention to relative distribution may also be entirely rational. African voters want to maximize what they receive from government, but have no information about the pool of resources that are available: without this information, they can't determine whether what they are receiving represents a large or small share, and thus, whether they would be likely to receive more under a new leader. Those who know what others have received, however, can calculate the total distributed resources, and determine whether their share would likely be higher or lower under a new leader who in expectation would be distributing the same total amount. In this model, relative distribution serves as a second-best type of information, and becomes irrelevant when voters have the information they need to directly calculate the share of resources they are receiving.

I test the theory with an experimental game in Mukono, Uganda. Respondents receive a payout from a "leader", who has divided a pot of money between himself, the respondent, and an additional player. The respondent then gets to choose whether to play again against the same leader, for the same payout, or try her luck with a new one. The experimental treatment is the amount of information the player has about what the other player received and the size of the pot. I show that respondents who know they have received more than the other player are more likely to retain the leader, but *only* when they are not also given information about the size of the pot. When voters know how much the leader had to distribute, and thus can precisely estimate their share,

they are no longer sensitive to whether they have been favored. I also show that while voting on relative distribution may be better than voting on no information, it is not optimal: players who know they have received more than the other player, but who have no other information, are likely to reelect leaders who are providing less and keeping more than a challenger would.

This article provides a new explanation for the relevance of relative distribution to African voters: voters prioritize high relative distribution not necessarily for its psychological value, but for its informational content. The game demonstrates that African voters can and will use relative distribution as an informational cue when they lack more direct information about available resources. This model implies that perceived favoritism will become less important as a driver of vote choice – and unequal distribution less important as a political strategy – as African governments become more transparent, and information about budgets becomes more widely available. The model also has implications for the construction of informational interventions. Specifically, the findings indicate that while providing voters with information about how their outcomes compare to others is likely to affect votes in contexts where information is low, such interventions may not improve accountability, and may in fact induce voters to support incumbents they should reject. Instead, interventions should provide information about total available resources.

1 Theory

To understand the theory intuitively, consider an African voter whose community has received some local public goods from the leader,[4] but these goods are not sufficient to fully meet the needs of the community. How should the voter respond at election time? The shortfall may mean the incumbent is providing the bulk of resources to other voters. It may mean he is a shirker who stole or wasted a large portion of the budget and is providing insufficient goods to everyone. Or it may mean that the incumbent is giving the voter the most he can of very limited pool of resources. In the first two scenarios, the voter would be better off under another leader who was less corrupt or who had

[4]The model assumes that the leader has discretion over distribution. This is a reasonable assumption in many African countries where administration is highly centralized and appointed ministers control all aspects of public goods. Legislatures, who control resource allocation in most established democracies, serve a largely rubber-stamp function in many African countries (van de Walle, 2001). In addition, African leaders often personally give goods such vehicles and fertilizer directly to individual voters.

different distributional preferences, and he should vote against the incumbent. In the latter case, it is unlikely that a new leader would continue to distribute such a large share of total resources to the voters' community, and he should vote to retain the incumbent. In order to determine which scenario he is in, the voter would really like to know the total amount of resources that were available for distribution, and how his share compares to the total.

And, yet, is it unlikely that the voter would be able to learn such information. Most African countries are profoundly low-information environments. African governments struggle with minimal bureaucratic capacity, which limits the collection and dissemination of budgetary data. The information that is released is often impenetrable to a citizenry with low levels of formal education. In Uganda, for example, the median citizen has received seven years' of education, the first three years of which take place in the pupil's mother tongue. Yet finance documents released online are written in English that tests at the 11th grade reading level or higher.[5]

Additionally, precisely because voters can use information about available resources to distinguish good performance from bad, leaders have good reason to deliberately obfuscate how much they have and how much they are spending. All leaders – both those who are spending resources efficiently and those who are stealing them – have an incentive to claim that public goods are insufficient because budgets are tight. And, in many cases, budgets are indeed tight. African leaders were entirely truthful when they reported that growth had stalled in the wake of the 2008 global financial crisis (**?**), or that aid had been withdrawn in reaction to the passage of laws that the Western world deemed to be human rights violations (**??**).

Though there has been little direct measurement of what African citizens know about budgets, what information there is suggests that citizens are quite uncertain about the availability of resources, and the amount their leader might be stealing. For example, Gottlieb (2012) shows that most members of a Malian sample were unaware that Mali earmarks funds for the provision of clean water, health clinics and schools at the local level. Nor does widespread "knowledge" of corruption appear to be particularly informative. In Malawi, for example, there is evidence in the literature of corruption at all levels in the Malawian government (**?**). Nevertheless, on Round Five of the

[5]Development agencies are putting in place programs such as trainings and telephone hotlines to make these data more accessible, but these programs are just beginning to take effect.

Malawi Afrobarometer survey, which took place in 2012, a remarkable 42% of Malawians said they did not know whether the president was corrupt. Following the 2013 release of a report by a British auditing firm showing clear evidence of corruption among high-level officials, the percent of respondents reporting uncertainty about the extent of corruption in the president's office dropped to 12%. However, reports like this one are rare, suggesting that most African voters remain uniformed as to total amount of resources, and the share of resources that are ultimately being distributed to them to or others.[6]

This situation, in which a principal struggles to evaluate the performance of their agent because they lack sufficient information about what level of performance is possible, is hardly uncommon, and has been solved in other cases by the use of the yardstick competition model (Shleifer, 1985; Besley and Case, 1995). In this model, the principal compares the performance of agents to one another, sanctioning those who perform the worst. Shleifer (1985) shows that regulators can determine whether firms receiving public funds are using funds efficiently by comparing the cost across firms of producing the same output. Besley and Case (1995) argue that voters compare tax rates set by leaders experiencing similar shocks, to determine whether the rate increase was necessary or whether the leader was taking excess tax revenue as rents. In an African context, Gottlieb (2012) uses a similar model to show that voters in Mali hold their local leaders to higher standards when they learn what other leaders have been able to provide to their constituents. The yardstick model usually addresses comparison across different agents under the same principal, but a similar intuition can easily be applied to a single agent acting on behalf of several different principals. By noting what the agent has provided to other principals, a principal can determine whether her agent is shirking in fulfilling his obligations to her. If the voter knows that others are receiving more than she is, for example, she also knows that there are resources available that are being withheld from her and that might be provided under a new leader with different preferences. On the other hand, if she knows that others are receiving less, she is aware that her share of resources is already higher

[6]This is not unique to Malawi or even Africa.(Ferraz and Finan, 2008) find that Brazilian voters who are given specific details about irregularities in their leaders' accounts will unelect that leader, and those not given these details will not, even though most Brazilian voters assume their leaders are at least somewhat corrupt. Similarly, Chong et al. (2011) show that providing precise information about the percent of a mayoral budget spent corruption causes voters to disengage from elections in a way that general cynicism about corruption does not.

than it might otherwise be and it would be risky to replace the incumbent in hopes of getting even more.

Fortunately, African voters have far more information about relative distribution than they do about budgets and expenditures. Relative outcomes are directly observable: what voters can't observe first-hand about the quality of life in other parts of the country, they can infer from news stories and images of failed crops, crowded classrooms, and shuttered clinics. Evidence indicates that voters use such media reports to inform themselves, if imperfectly, about relative distribution: the strength of the correlation between actual and perceived distribution increases with media exposure (Carlson, 2014). Additionally, though almost all African leaders vehemently deny that they are corrupt, many make no attempt to hide that they provide goods unequally. Leaders have an incentive to remind voters that what they get from they government is a function of how they vote: leaders have not only acknowledged that some citizens are under-served, but have actually used this as a bargaining chip, telling citizens that they would receive better goods if they were to support of the ruling regime. Accordingly, Africans voters are far more certain about how their well-being compares to others than they are about their leaders' overall provision. Over 60% of Round Five Afrobarometer respondents are willing to assert that their living conditions are either better or worse than others'; only 3% indicate they don't know.[7] Similarly, about 65% of respondents make the decisive claim that their ethnic group is either "always" or "never" treated unfairly by government; 4% report that they don't know. If information about relative distribution can help voters more accurately evaluate their candidates, and this information is more readily available, it is not surprising that voters should rely heavily on this information when casting their vote.

1.1 Formalization of the intuition

To formally theorize the informational role that could be played by relative provision, assume a simple game with four players: a leader, a challenger and two interest groups. The incumbent (which the challenger will become should he be elected), divides a pool of resources between himself and the two interest groups, according to his preferences, which are exogenous to the game. The

[7] The remainder report that their living conditions are similar to others'.

incumbent and challenger are notated I and C, respectively. The groups, notated A and B, are unitary actors.[8] The pool of resources, which is constant, is notated G.[9] The allocation of goods by the incumbent would therefore be:

$$G_i = \theta_i * G_i + \theta_a * G_i + \theta_b * G_i \tag{1}$$

where θ_x represents the share of the budget the incumbent distributes to player x. The shares are drawn from uniform distributions, where $\theta_i \in [0 - 1]$ and $\theta_a \in [0 - (1 - \theta_i)]$. The two interest groups cannot directly observe the president's preferences, but only the payouts that result from those preferences, noted D_a and D_b (which equal the product of G and θ_a and θ_b, respectively.) Equation 1 can be rewritten as follows, where D_i represents the amount of resources the incumbent kept for himself and R represents the remainder of resources distributed to A and B ($D_a + D_b$) after the incumbent takes his share.

$$G_i = D_i + D_a + D_b = D_i + R \tag{2}$$

Group A, which comprises a minimum winning coalition, must choose whether to reelect or replace the incumbent. Because G is constant across candidates, and preferences are constant within candidates, the payout in time $t + 1$ if the voters chooses to reelect will be the same as the payout in time t. If they choose to replace, they will receive a new leader whose preferences are unknown, but drawn from the same distributions as those of the previous incumbent. A will choose to reelect the incumbent when the size of his allocation under the incumbent is greater than the expected allocation under the challenger, or when $\theta_a > \mathrm{E}[\theta_a]$.[10] is When A knows G, they can

[8]There are, of course, many more than two interest groups in any political system, but it also possible to divide any society into two groups – A and Not A – along whatever cleavage is salient to A, and compare average outcomes of their chosen group to the average outcomes of the rest. The literature often divides African voters into two groups, including coethnic/non-coethnic, urban/rural, or core/opposition.

[9]In this model, and in the game built on the model, I am treating G as exogenous to voters' utility, and distribution to A and B as windfalls. In reality, G will be at least partially funded by public taxation, the burden of which may fall unequally on A and B (Kasara, 2007). However, Uganda, like many African governments, receives substantial funding from foreign donors, and earns revenue taxing foreign direct investment. Thus much government distribution can, in fact, be treated as a windfall from the point of view of most citizens. The same results come out of the model whether we consider D_a or D_b as pure windfalls, or as net transfers after taxes are accounted for.

[10]In reality, the challenger is known, and there may be many factors that determine how much A expects from him, including his ethnicity, party, prior record, or campaign promises. However, A's strategy is the same regardless

easily determine θ_a by dividing D_a by G.

If A does not know G, they cannot determine the θ_a that is represented by D_a or determine whether $\theta_a > \mathrm{E}[\theta_a]$. However, since the pot is fixed, A can infer that their share, whatever it is, would be higher under the challenger if the challenger would keep less than the incumbent ($\theta_i > \mathrm{E}[\theta_i]$)) or if the challenger would give less to B ($\theta_b < \mathrm{E}[\theta_b]$).) The amount that the incumbent is keeping is secret, as is the amount the challenger would keep. However, since incumbent and challenger are drawn from the same distribution, in expectation they would keep the same amount ($D_i = \mathrm{E}[D_i]$). Since D_i is the same in expectation, R is also the same in expectation. R is the sum of D_a and D_b, so a voter who knows D_a and D_b can fully calculate R and A's share of R, and compare this to the share of R A expects under the challenger. If we assume voters prior's are flat, as they would be in a information-free environment, then A is drawn from $U(0, R)$, and A's payout in expectation is $0.5 * R$. A should always choose to reelect when $D_a > D_b$, because $\theta_a > \mathrm{E}[\theta_a]$.

1.2 Model Predictions

To clarify the advantage of voting on relative distribution in a low-information context, I assess each of three decision rules, using simulated data.[11] The goal is to maximize the percent of correct votes, which means that a voter chooses to retain a leader when his payout under the leader was more than he would receive under a randomly-matched challenger, and replace him otherwise.

Under decision rule one, voters know only their own payout and thus must select an arbitrary threshold at which to reelect; the percent of correct votes varies by the selected threshold, but since the threshold is itself selected at random, this decision rule leads voters to the the correct choice as often as random chance, or 50%. Under decision rule two, in which voters have full information about the incumbent, they choose to reelect when their payout is more than 25% of the pot, which is the average payout in the simulated data. This produces a correct vote approximately 74% of the time. Finally, in the third decision rule, a voter chooses to reelect whenever his payout is

of the value of $\mathrm{E}[\theta_a]$. If A has reason to believe that $\mathrm{E}[\theta_a]$ would be particularly high(low), then their share under the incumbent would simply have to be equally high(low) in order for A to retain the incumbent.

[11] For each of 1000 draws, I generate 2000 voters, each of whom must choose between an incumbent and a challenger. The pot for all candidates is 3000 units. Each candidate's θ_i is randomly drawn from $U(0, 1)$; θ_a is drawn from $U(0, 1 - \theta_i)$ and $\theta_b = 1 - \theta_i - \theta_a$. As expected under these distributions, A and B's mean payouts are both one quarter of the pot, or 750 units.

Figure 1: Simulated Percent of Correct Votes by Decision Rule

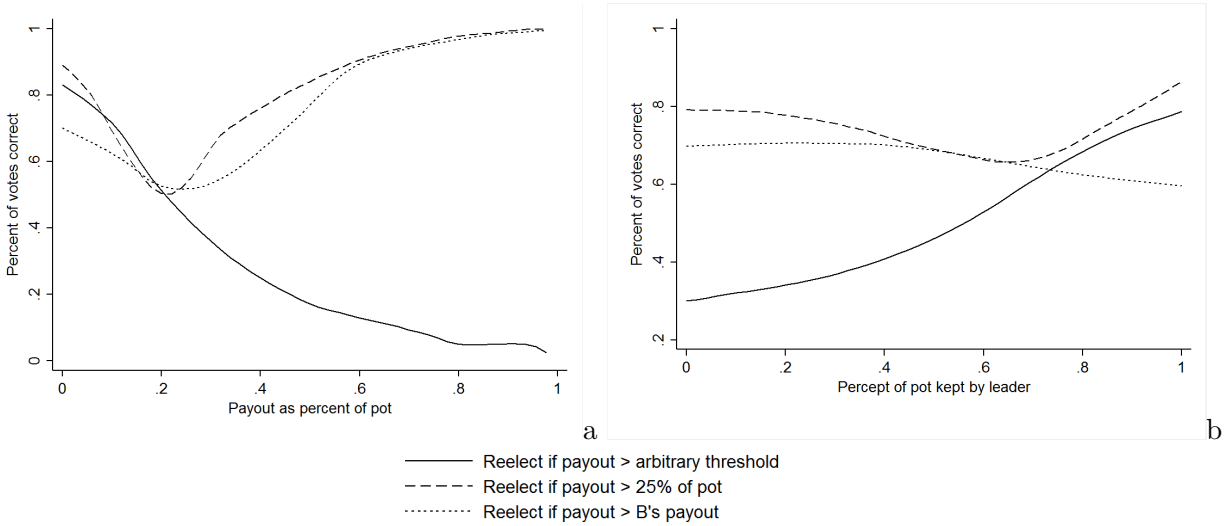

greater than B's. This decision rule produces the correct decision 66% of the time. If we assume that voters prioritize information that produces the best outcomes, this means they should use information about relative distribution if they have it, but override this information if they have information about the size of the pot.

Figure 1 shows the particular conditions under which voting on relative distribution will be the most and least effective. The figures plots the percent of correct votes cast under each decision rule, based on the leader's distributional decisions: Figure 1a plots vote accuracy against the θ_a, the percent of G paid out to the voter. Figure 1b plots the accuracy of votes against θ_i, or the percent the leader kept for himself. The figures show that, while an improvement over the arbitrary decision rule in most cases – and thus still the dominant strategy when G is not known – voting on relative distribution is actually the *least* effective decision rule when A's payout is very low. This happens because voters choosing how to vote based on relative distribution will vote for a candidate who has distributed very little to the voter, as long as he has paid even less to the other player: under both other decision rules, this type of leader is correctly rejected. Since for any given value paid to A, the percent the leader is keeping increases as B declines, voters using this decision rule are also more likely to elect leaders who have paid out very little to anyone, as shown in Figure 1b.

When payouts are low, the results of my model are very much like those of Padro i Miguel

(2007), who argues that African voters are willing to reelect kleptocratic leaders who favor them because they believe they will receive even less under an equally corrupt leader who does not favor them. However, the assumptions about ethnicity, taxation and capital mobility that underlie his model are not necessary to produce this result, nor is a context of weak institutions or instability around elections. The only necessary condition for relative distribution to matter is that voters have insufficient information to directly estimate their share of total resources.[12] Accordingly, my model is both more widely applicable than Padro i Miguel's, and suggests that the problem is more tractable than his model implies: while ethnicity and institutions are both quite sticky, my model indicates that voters will no longer rely on relative distribution as a heuristic when they have information about their share of total resources.

2 Behavioral Game

I test the model with a behavioral game, in which players choose whether to reelect a "leader", or play against a new one, after learning at least some information about how the leader had divided a pot of money between himself and two citizens. The experimental manipulation is how much information the player, Citizen A, has about the size of pot or how it has been distributed. If the theory is correct, a player's whose payout is larger (smaller) than Citizen B's should be more (less) likely to retain the leader, but *only* when the player does not know the total pool of resources that were available for distribution.

A lab experimental design allows a clean test of the role of information because it improves compliance with treatment. Though it is always possible to randomly assign people to hear real-world information, one can't randomize whether the information they hear is new to them. The real-world information that would be relevant here – the relative distribution of government resources – is something that many Ugandans already have access to. More importantly, those who already have this information are systematically different from those who have not. Additionally, voters

[12] In fact, the reason that Padro i Miguel must include electoral instability in his models is because he assumes full information. Unless political transitions are costly, it is difficult to explain why political competition would not improve outcomes when voters have full information about the net share of resources they are receiving under all leaders.

tend to ignore or justify information that challenges their beliefs about their leaders; this happens in Uganda as well as any number of other countries (Logan et al., 2003; Carlson, 2014; Gerber and Huber, 2010). The likelihood that a respondent will take in information about the incumbent's performance also not random, but is correlated their pre-existing attitudes.

What this means is that real-world information provided in the context of a field or survey experiment may not have an effect, even if it that same information had been critical at some point in the past to shaping the attitudes respondents currently hold. In a behavioral game, on the other hand, no respondents come with pre-existing knowledge of the game, nor are anonymous players likely to trigger information-resistance in the way that real political actors would. In an semi-autocratic county like Uganda, the abstraction away from real politics also makes is more likely we will be able to measure respondents' true vote preferences; voters in many African countries are afraid to voice opposition to the regime and would be likely to say they support the incumbent regardless of the provided information (Weghorst, 2011). If information has the effect I posit, therefore, we should be able to observe it in the lab experimental context; if we do not find an effect, that is strong evidence that the theory itself is wrong.

2.1 Game Design

Testing the model requires that all respondents receive either more or less than the other citizen. Since the modal dictator in a dictator game distributes the pot evenly, respondents played against a computer rather than another person.[13] The pot in each game was drawn from a normal distribution with a mean of 3000 Ugandan shillings. The leader's share was drawn from a uniform distribution ranging from zero to the full size of the pot, with the remainder split between the respondent and the other citizen. I randomly assigned all respondents to receive a payout that was either more

[13]I did not tell respondents they were playing against a computer, but rather that they were playing against other people. I did this for the very simple reason that most Ugandans are not familiar with computer simulations, or even computers themselves – 85% of respondents on Round Five of the Ugandan Afrobarometer report that they never use one. More importantly, the theory requires that players have an expectation of what an unknown challenger would distribute, to which they can compare what they are currently receiving. Though most people have relatively reliable intuitions about how other people will play a game, accurately predicting a computer's game-play requires fluency with probabilities that respondents with little formal education may not have. Since this is deception, all respondents were debriefed after the game and offered the chance to ask questions or withdraw their data; none did. A more thorough discussion of my use of deception, the potential for harm, and the texts of the instructions and debriefing scripts, are in section X of the appendix.

or less than what Citizen B received. To isolate the effects of high relative distribution from a high absolute payout, I then separately randomized whether the respondent's payout was more or less than the 25% of the pot he could expect under the challenger.[14] No respondent received a payout of zero (because there would be no way for the other player to receive less than this) and no respondent received more than half the pot (because there is no way for the other player to receive more than this.) Respondent payouts thus ranged from 15% to 45% of the pot in each game. All payouts were rounded to the nearest 100 shillings, which is the smallest denomination of coin in common circulation.

Respondents received three payouts. After receiving the first payout,[15] the player decided whether to retain the leader and receive the same payout a second time, or take her chances on a new leader who might pay more or less. After this second payout, players were given the opportunity to ask for one additional piece of information about the leader they were currently playing against,[16] and then choose once again whether to reelect or replace, for the appropriate payout. The average total payout was 2200 Ugandan shillings, which at the time of data collection was approximately US$1, or about half the daily income of the average Ugandan; this amount is enough to buy a meal or several litres of clean water from a public tap. This payout should thus be enough incentive to override respondents' concerns about social desirability or politeness when choosing whether to reelect or replace their leader.

2.2 Controlling access to information

The key experimental treatment in the game was the amount of information that players have about how much money there was in the pot and how it was divided. There were three informational treatments: 1) a "no information" condition in which voters knew only their own payout; 2) a

[14] What players actually expect a typical payout to be is an empirical question, and I did not assume *a priori* that players would hold expectations that matched the model. However, in pre-testing the game with respondents who had full information, the 25% threshold worked well: Respondents who received more than 25% of the pot chose to retain their leader 85% of the time, while those who received less than 25% of the pot chose to retain only 22% of the time.

[15] This payout was not necessary for purposes of the game; respondents were incentivized by the payout that *followed* their decision. We provided this payout to make the later, promised payouts credible and concrete.

[16] For those who chose to reelect in round one, this was the first leader they were assigned. For those who chose to replace, this was the second leader.

"relative distribution" treatment in which respondents knew their own payout and that of the other player, but not the total pot that the leader had to distribute; 3) a "full information" treatment in which respondents knew their payout, the other player's payout, the size of the pot, and, by subtraction, how much the leader kept. I limited spillover of information into the lower-information treatments in two ways. First, I drew a new pot for each game so that players (and enumerators) would not be able to learn the exact size of the pot across multiple rounds. Second, enumerators had only enough information about each game to meet the requirements of the assigned treatment, so that they could not accidentally disclose information that respondents were not supposed to have.

I also limited *mis*information by assigning each leader an ethnicity, as cued by the location the leader was playing from.[17] Existing research indicates that Ugandans have strong expectations of how leaders of different ethnicities are expected to distribute goods (Conroy-Krutz, 2013; Carlson, 2015), though it's not clear that this expectation extends to anonymous dictators in dictator games (Habyarimana et al., 2007). Nevertheless, leaving ethnicity unassigned presented a risk that respondents make assumptions about relative distribution that I was neither measuring nor controlling. I was concerned about two things in particular. First, I was concerned that respondents would assume the leader was from the locally dominant ethnic group, and thus would favor any player also from the area. Second, I was concerned that respondents would believe the "leader" was a stand-in for Uganda's actual president, Yoweri Museveni, who would favor those of his own ethnicity.[18] It is also possible that players' responses are conditioned by the ethnicity of the leader: respondents may be more (or less) likely to punish coethnics who disfavor them, or they may be hesitant to replace any leader of the president's ethnicity, regardless of distribution, especially if they think the game is some sort of political test.[19] To avoid these unmeasured effects, I randomly assigned each leader an ethnicity and control for ethnicity in analyzing the results of the game.

[17] Most Ugandan ethnic groups have a clear home territory.

[18] As it turns out, these concerns were valid. Of a small sample of respondents not given information about leader ethnicity, and asked to guess, the plurality (47%) thought the leader was from either their own city or the neighboring one, which would mean they were likely to share the ethnicity of the leader. The next largest share, 20%, thought the leader was playing from Museveni's home district.

[19] Enumerators reported that many respondents were suspicious of the game. One respondents even refused to take her payout, thinking it might be some sort of bribe.

Approximately one third of respondents were assigned a leader of their own ethnicity. One quarter were assigned a leader of President Museveni's ethnicity.[20]

The results presented below show that I was able to effectively control information about relative distribution: those who were not given direct information about relative distribution were not able to infer it. However, I was not able to fully control respondents' information about the size of the pot. It is reasonable to assume that voters have no information about the size of the pot when "the pot" is a national budget in the billions of dollars. In the game setting, however, there are limits to how much a researcher would be giving away, as well as on how much enumerators could realistically be carrying. Accordingly, respondents had reasonably accurate priors about the size of the pot. As part of a manipulation check, I asked respondents who had not been told the size of the pot to guess how much money there was. The median guess was 10,000 shillings, which is larger than the true median pot size, but well within one order of magnitude.[21] More importantly, there is a significant correlation between respondents' guesses about the pot and reality: those who received a payout that was above the 25% threshold in actuality were 9% more likely to guess a pot that placed their payout above the threshold. This means that players in the no information treatment vote correctly more frequently than the 50% we would expect if they truly had no information. The effect of providing direct information about the pot may be smaller that it would have been under greater uncertainty.

2.3 Details of Implementation

The game took place in respondents' homes, following an introductory survey that gathered pre-treatment covariates including demographics, local public goods and political attitudes. Enumerators read the questions in either English or Luganda, according to the respondent's preference[22] and recorded answers on a smartphone. Enumerators carried pre-printed lists of respondent codes, in which treatment assignments were embedded, and entered these into the phone at the beginning

[20]The results indicate that respondents don't particularly care about coethnicity – they are equally likely to retain coethnic and non-coethnic leaders – but they *are* concerned about appearing anti-incumbent – leaders from Museveni's ethnic group were 9% more likely to be retained that those of other groups.

[21]For reasons that are not immediately obvious, women guessed significantly larger pots than men. I control for gender in the robustness checks.

[22]Respondents who were not fluent in either English or Luganda were not recruited.

of the survey to prompt the appropriate instructions. The game itself was played by text message. Once respondents understood the game's instructions, the enumerator sent an SMS containing the respondents' code to a research assistant, who responded with the appropriate payout information from a randomly-drawn game. As a quality control measure, enumerators recorded the game-play data they received via text into the survey form.

3 Sample and balance

The experiment took place in May and June of 2014 in Mukono District, which is directly to the east of Uganda's capital city, Kampala. Uganda is appealing as a case for this study because political behavior in Uganda is well-researched compared to many African countries and thus we have good existing evidence to inform our expectations. Mukono District in particular is appealing for the diversity of environments within a relative small geographic area: the district contains Uganda's fourth largest city, also called Mukono, along with a large number of entirely rural parishes with a low standard of living. The Mukono sample, though somewhat more urban than a typical Ugandan sample, provides a reasonably good approximation of the larger Ugandan population. The 565 respondents were selected from fifty electoral constituencies in Mukono. Unfortunately, this sample represents only 70% of the people who actually played the game: the study suffered from an unexpectedly high percent of missing or invalid data.[23] As a result, the study is underpowered. In the analysis below, I will present standard errors and p-values, but will not rely on them to determine whether effects are meaningful.

Balance tables are shown in Table 2 in the appendix. They demonstrate that randomization was for the most part effective. Both payout and information treatments are equally likely to be assigned to respondents across age, education, poverty, and the ethnicity of the assigned leader.

[23]To avoid having to pause ongoing games due to unexpected electrical shortages, respondents did not play against the computer in real time. Instead, randomized games were pre-printed on paper forms, and the code of the respondent who played each game was recorded by hand by the research assistant who was sending out information about the games. A bag containing several hundred of these paper records was stolen from a research assistant before they could be entered. Since enumerators received some information about the game, and entered this into the survey form, we were able to retroactively match some, but not all, of these records. I also dropped all data collected by an enumerator who was later caught fabricating data. Finally, I dropped any game in which the game details entered by the enumerator did not match those sent out by the research assistant.

There are slight imbalances in gender and urban residence: when I control for these imbalances, the result becomes stronger.

4 Experimental Results

Figure One shows the estimated effect of relative distribution on vote choice under the three different information treatments. Specifically, the figure plots the added likelihood that the voter will choose to retain a leader when he has received more than the other player. The darker bars are the raw effects with 95% confidence intervals. The lighter bars show the estimated effects controlling for additional characteristics of the player and the game. I control for gender and urban residence, as well as the assigned ethnicity of the leader. Since the choice to replace the leader is plausibly connected to risk aversion, I also control for factors that may affect risk tolerance, including age, poverty, and education; this latter variable may also affect voters' ability to appropriately estimate the risks and probabilities in the game. Finally, I add enumerator fixed effects. The full models from which the estimates are drawn are shown in Table 3 in the appendix.

If relative distribution matters in the way that I propose, we should see that players who have information about relative provision use this information to make their choice, and become more likely to retain their leader. However, if they also have information about the size of the pot, relative distribution should no longer matter. Figure 1 shows that this is indeed the case. Those who are not informed about relative distribution are equally likely to retain leaders who have favored or disfavored them. Those who have information about relative distribution, and *only* information about relative distribution, respond to this information: they choose to retain leaders 10% - 13% more often when those leaders have favored them. Those who have full information, on the other hand, act as if they have no information about relative distribution. As predicted, those who can directly calculate their share of the pot do not care about relative distribution, with an estimated effect size of almost exactly zero.

Figure 2: Impact of High Relative Payout on Retention of Leader, by Information Treatment

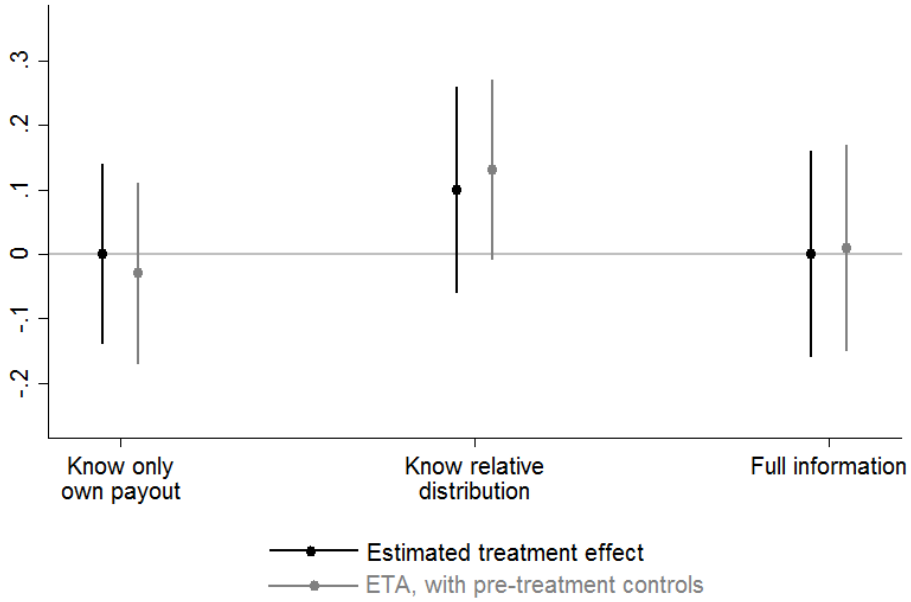

5 The Usefulness of Information

The results of the game provide evidence that players respond to relative distribution when that is the only information they have; when they have other information that lets them directly calculate whether their share is above or below what they would expect from a challenger, they don't reward the leader for favoring them. This supports my theory that relative distribution is an informational cue, that becomes salient in the absence of more helpful information. In this section, I show that, in support of the hypothesis, players find relative distribution informative, but not as informative as the size of the pot.

As noted, there was a second portion of the game: respondents were allowed to ask for one more piece of information about the leader they were playing, to use in making a choice that would trigger a payout. This incentivized respondents to ask for information that would help them make a better decision. They were not forced to ask for information: "none" was an option. Information was costless, so those who chose to receive no additional information presumably felt that additional information would not be at all relevant to their choice. We should expect that respondents find

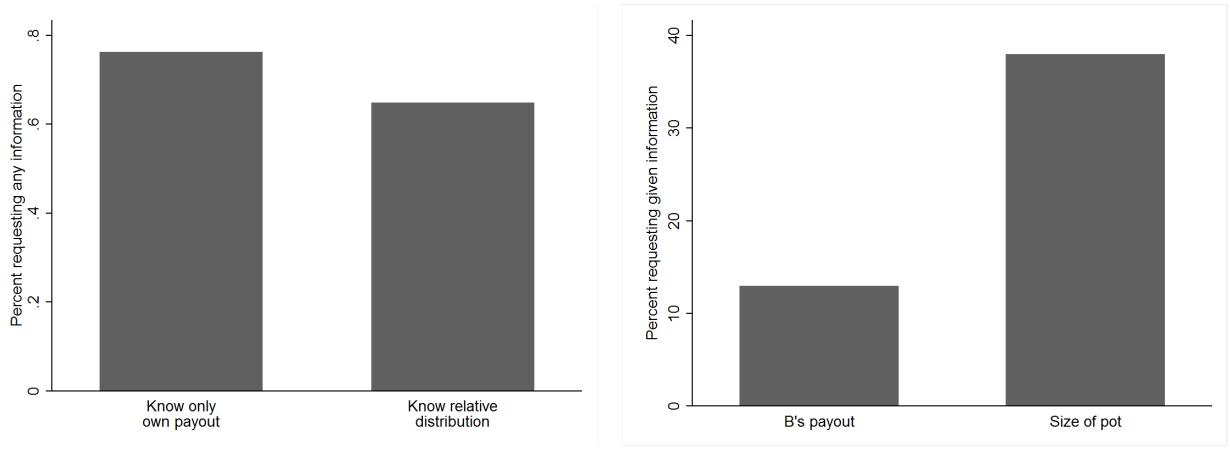

(a) Percent requesting any information

(b) Percent requesting given information

information about both the size of the pot and relative distribution useful, but when given a choice, should prefer to know information about the size of the pot.

Figure 3a shows the percent of respondents in the no information and relative distribution treatments who requested additional information (those in the full-information treatment are not included, since there was no additional information for them to learn.) The results show that respondents find relative distribution informative: those who knew what the other player received were less likely to ask for any more information, meaning they felt they had sufficient information to make a good choice. However, as predicted, relative distribution is not as useful to respondents as the size of the pot. Figure 3b shows the percent of respondents who didn't have information about the other players' payout or the size of the pot, respectively, who asked for this information. Almost three times as many people requested information about the size of the pot than requested information about the other players' payout. This is consistent with a theory in which relative distribution is a second-best type of information.

6 Vote Accuracy and Welfare Impacts

A final implication of the theory, as shown by the simulation, is that those who have information about relative distribution cast more correct votes than those who have information only about their own payout, unless the payout is very low or the percent the leader has kept is very high. I

Figure 4: Percent retaining leader when payout less than 25% of pot

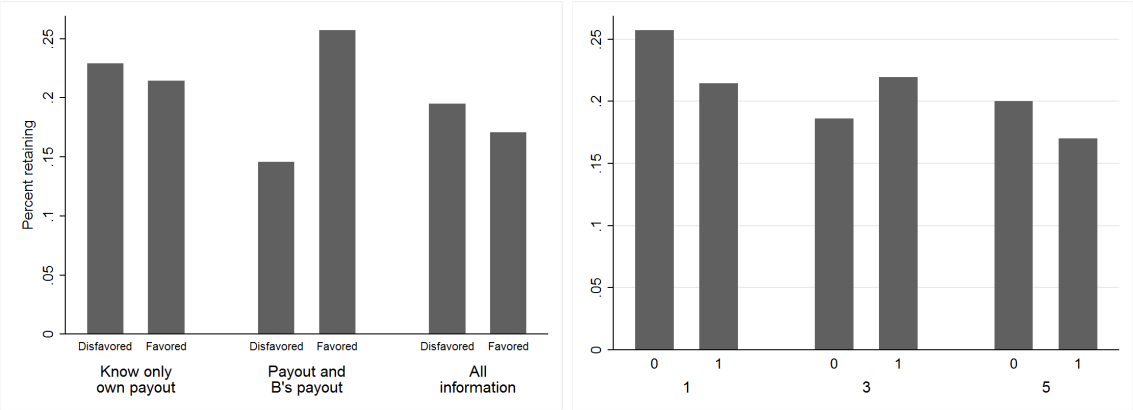

can't directly test this here: in the game, I artificially constrained relative and absolute distribution to be uncorrelated, meaning that those voting on relative distribution will necessarily vote correctly only 50% of the time. However, I *can* demonstrate that, in keeping with the model, voters who have only information about relative distribution can be induced to support leaders providing very low payouts, as long as that leader is favoring them.

The first figure in Figure 4 shows the percent of players who chose to retain the leader when their payout was less than 25% of the pot – in other words, those who retained a leader they should not have – across all three information treatments, based on whether or not they were favored. The results show that voting on relative distribution will not necessarily improve the accuracy of votes when payouts are low. Those who know they are favored but don't know the size of the pot are the most inclined of any group to retain the leader despite their low payout. Since B's payout and the percent the leader kept are directly correlated, this also means that those who know they are favored are more likely to retain leaders who are keeping a larger share of the pot: this result is shown in the second half of Figure 4. To the extent that these findings extend beyond the lab, they suggest that providing voters with information about relative outcomes in an otherwise low information context, rather than improving accountability, can *increase* the likelihood that voters will choose to reelect a leader who is paying less and keeping more than a challenger would.

6.1 Implications for Information Interventions

Both the simulation and the experimental results show that information about relative distribution can pose problems for accountability when payouts and overall distribution are very low. In the original model, and the game, low payouts were not the norm: leaders' distributive preferences were exogenous to the game, and set to be above this threshold half of time. However, it is quite possible, and even likely, that actual leaders' distributive preferences are *not* exogenous, but instead strategic attempts to earn a majority of votes with the minimum expenditure. Figure 5 shows the vote share a leader would receive under each decision rule if Group A received a given share of resources and Group B received nothing; in other words, the X axis plots both A's payout and the total distribution. Those with full information, who can calculate their share, are sensitive to what they are receiving, and the leader will only be reelected if he provides A with at least 25% of the pot. Those who vote on relative distribution, on the other hand, will always choose to reelect. As long as Group B receives nothing, Group A will vote to retain the incumbent if they receive any amount greater than zero, even if the amount of distribution is negligibly small. If we assume that leaders distribute the minimum that will enable them to earn a majority of votes, leaders of populations that vote on relative provision, will have little incentive to distribute any goods to anyone.

The results have implications for the design of the ever-increasing number of informational interventions intended to improve accountability in the developing world. There is a shift to include data on relative outcomes in existing interventions. (**?**), for example, argue based on feedback from Indian respondents that government performance "score cards" should include information about absolute and relative outcomes. Similarly, (Gottlieb, 2012) notes that interventions that provide information on relative outcomes shift attitudes and vote shares more than information that does not allow respondents to compare across units. What my findings highlight is that the shift in voting that occurs when voters learn about relative distribution may not actually improve accountability, especially when overall distribution is low. When voters have only this information they will use it, but when payouts are low, doing so produces worse outcomes than voting on no information at all. Thus the priority for those developing informational interventions should be on ensuring

Figure 5: Percent Retaining Leader if B Receives Nothing, By Decision Rule

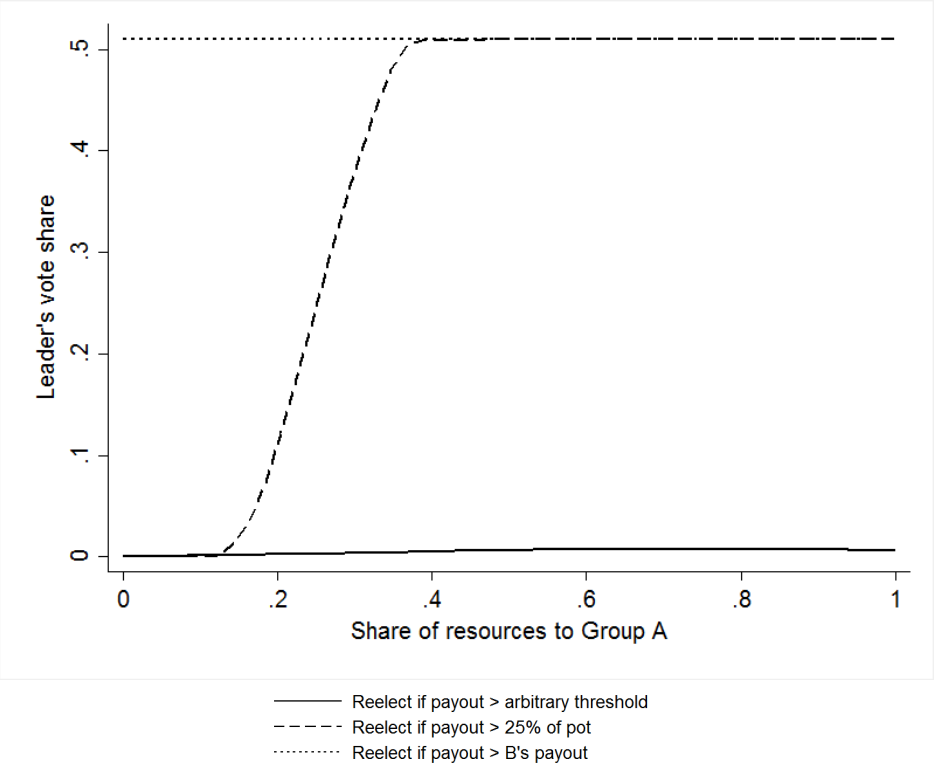

that respondents receive information about the total amount of resources that are available to be distributed, as that is the only type of information that consistently improves accountability.

7 Conclusion

In this study, I propose and test a model to explain the salience of relative distribution to voters in African countries. I argue that relative distribution serves an informational function. I use an experimental game to show that when voters only have information about relative distribution, they are much more likely to retain leaders who favor them in distributing a pot of resources. However, when they also have information about the total amount of resources, relative distribution no longer influences their vote. This indicates that, at least in a lab setting, receiving high relative distribution not valuable to voters beyond its use as a clue to where their leader falls in a distribution of distributive preferences. When voters use relative distribution as their heuristic, it can produce sub-optimal outcomes; in particular, it can induce respondents to reelect leaders who are giving them less than they would receive under a challenger. Fortunately, to the extent that this type of voting occurs in a low information environment, it can be resolved by the distribution of information about available budgets and resources.

References

Besley, Timothy and Anne Case. 1995. "Incumbent Behavior: Vote-Seeking, Tax-Setting and Yardstick Competition." *The American Economic Review* 85, no.1 :25–45.

Bratton, Michael, Ravi Bhavnani and Tse-Hsin Chen. 2011. "Voting Intentions in Africa: Ethnic, Economic, or Partisan." Afrobarometer Working Paper 127.

Carlson, Elizabeth. 2014. "Partisan Bias and Perceptions of Provision in Uganda." Presented to the Conference on Subnational Research in Comparative Politics, Brown University, 2013.

Carlson, Elizabeth. 2015. "Ethnic Voting and Accountability: A Choice Experiment in Uganda." *World Politics* 67, no.2 .

Chong, Alberto, Ana de la O, Dean Karlan and Leonard Wantchekon. 2011. "Looking Beyond the Incumbent: The Effects of Exposing Corruption on Electoral Outcomes." NBER Working Paper 17676.

Conroy-Krutz, Jeffrey. 2013. "Information and Ethnic Politics in Africa." *British Journal of Political Science* 43, no.2 :345–373.

Ferraz, C and F. Finan. 2008. "Exposing Corrupt Politicians: The Effect of Brazil's Publicly Released Audits on Electoral Returns." *The Quarterly Journal of Economics* 123, no.2 :703–745.

Ferree, Karen. 2006. "Explaining South Africa's Racial Census." *Journal of Politics* 68:803–815.

Ferree, Karen and Jeremy Horowitz. 2010. "Ties that Bind? The Rise and Decline of Ethno-Regional Partisanship in Malawi, 1994-2009." *Democratization* 13, no.3 :534–563.

Gerber, Alan S. and Gregory A. Huber. 2010. "Partisanship, Political Control and Economic Assessments." *American Journal of Political Science* 54, no.1 :153–173.

Gottlieb, Jessica. 2012. "Does Raising Voter Expectations Improve Accountability? A Field Experiment in Mali." Working Paper.

Habyarimana, James, Macartan Humphreys, Daniel Posner and Jeremy Weinstein. 2007. "Why Does Ethnic Diveristy Undermine Public Goods Provision?" *American Political Science Review* 101, no.4 :709–725.

Ichino, Nahomi and Noah Nathan. 2013. "Crossing the Line: Local Ethnic Geography and Voting in Ghana." *American Political Science Review* 107, no.2 :344–361.

Kasara, Kamuli. 2007. "Tax Me if You Can: Ethnic Geography, Democracy, and the Taxation of Agriculture in Africa." *American Political Science Review* 101, no.1 :159–172.

Lindberg, Staffen. 2003. ""It's Our Time to Chop": Do Elections in Africa Feed Neopatrimonialism rather than Counteract It?" *Democratization* 10, no.2 :121–40.

Logan, Carolyn J, Nansozi Muwanga, Robert Sentamu and Michael Bratton. 2003. "Insiders and Outsiders: Varying Perceptions of Democracy and Governance in Uganda." Afrobarometer Working Paper 27.

Luttmer, Erzo. 2004. "Neighbors as Negatives: Relative Earnings and Well-Being." NBER Working Paper 10667.

Padro i Miguel, Gerard. 2007. "The Control of Politicians in Divided Societies: The Politics of Fear." *Review of Economic Studies* 74, no.4 :1259–1274.

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

Shleifer, Andrei. 1985. "A theory of yardstick competition." *Rand Journal of Economics* 16, no.3 :319–327.

van de Walle, Nicolas. 2001. *African Economies and the Politics of Permanent Crisis, 1979-1999*. New York: Cambridge University Press.

Wantchekon, Leonard. 2003. "Clientelism and Voting Behavior:A Field Experiment in Benin." *World Politics* 55, no.3 :399–422.

Weghorst, Keith. 2011. "Political Attitudes and Response Bias in Semi-Democratic Regimes: A Survey Experiment Comparing the List Experiment and Randomized Response in Tanzania." Working Paper.

A Relative outcomes and presidential approval ratings

Table 1: Determinants of Presidental Approval

Variable	**Coefficient** (Std. Err.)
Equation 1 : presapp	
Sociotropic evaluation	0.290** (0.020)
Pocketbook evaluation	0.026 (0.021)
Personally better off than others?	0.106** (0.021)
Ethnic group better off than others?	0.090** (0.021)
Ethnic group treated unfairly?	-0.342** (0.021)
Coethnic of incumbent	0.286** (0.050)
Education	-0.063** (0.011)
Urban resident	-0.157** (0.043)
Believes gov't running survey	0.104** (0.040)
Intercept	0.669** (0.242)
Equation 2 : lnsig2u	
Intercept	-0.378 (0.369)
N	16784
Log-likelihood	-8556.534
$\chi^2_{(9)}$	934.352

Significance levels : † : 10% ∗ : 5% ∗∗ : 1%

B Ethics and Debriefing

I could not test the hypothesis by having real players play a dictator game in real time because uneven distributions across players were critical for the design: most players in a dictator game divide the pot evenly. Other researchers running laboratory games in African countries have developed techniques for controlling game play without computers, but I find these solutions unsatisfactory. One option was to recruit respondents to play leaders in the pre-study period, and keep only those games in which the leader divides the pot in ways that are useful for the study design. During the study, respondents would then be randomly assigned to receive one of these retained games. Another option is to assign a confederate to play the leader and other citizen in every game, and have them "play" in real time according to the results of the computer randomization. However, both of these strategies, while not technically deceptive, still require misleading respondents about the nature of the game and manipulating the game without the respondents full consent to be manipulated. These techniques are at best only marginally more ethical, and, in the absence of a debriefing statement, arguably *less* ethical than outright deception combined with debriefing. These techniques also raise ethical questions about whether the individual(s) playing the leader and other citizen, including project confederates, must be allowed to keep the appropriate amount every time respondents play a game against them.

There are two potential risks of harm from the deception. One is respondents' sense of frustration or embarrassment to have been manipulated. Fortunately, only three respondents had a negative reaction to learning they had been playing against a researcher: all three thought researchers should have been able to give them more money in the game. The other risk of harm is that respondents would disbelieve any future researchers who sampled from the same population, even if these researchers were being entirely truthful. Mukono is a popular place for research, so this is not a trivial concern. However, as it turns out, respondents did not particularly believe us in the first place. Enumerators reporte that many respondents though they were government agents. One respondent even refused to take her winnings in the game because she thought it was some of bribe. The debriefing statement seems to have reassured a number of respondents that we were, in fact, researchers, and that the game was actually part of a study.

B.1 Game instruction and debriefing scripts

The instructions of the game, delivered in the respondent's preferred language, read:

> Now I am going to ask you take part in an exercise. It will give you the chance to earn some money. In this exercise, you will be playing a citizen. There will be two other players. One of them will be playing a leader. The third player will be playing a citizen, just like you are. Both of the other players are playing from other locations. I will receive information about them via SMS. I will show you and tell you what the SMS's say. The person who is playing the leader has a pot of money. He gets to decide how much to keep for himself, and then he will divide the rest between you and the other citizen. He knows where you are playing from, but nothing else about you. He will tell us how much to give you and we will give it to you. Then you will decide whether or not to reelect the leader or replace him with someone else. If you reelect him, you will get the same amount of money a second time. If you vote to replace the person playing the leader, we will randomly assign you a new leader and he will decide how much he wants to give you. This new leader may have more or less money than the first leader did, and he may decide to give you a larger or smaller share. There is a risk that this new leader will give you less than the first leader did, but also a chance that he will give you more. So, let us pick your first leader. I will describe him, and give you the money that he wants you to have. Then you will decide whether you want to retain him and receive the same amount of money again, or whether you want to replace him with a new leader who may give you a different amount of money.

The debriefing script, delivered in the respondent's preferred language, read:

> Now that the game is finished, I need to let you know that the person playing the "leader" was actually a researcher who selected your payout and that of the other player with the help of a computer program. Any money that was left over from the pot rather than being given to you will not kept by the researcher personally, but will be paid out to another respondent who we will interview later in the project. Do you

have any questions?”

C Balance

Table 2: Balance Tables

	Female	Age	Education	Urban	Coethnic leader	Munyankole leader
Below 25%	0.48	37	8	0.36	0.26	0.23
Above 25%	0.50	38	8	0.36	0.34	0.23
Less than other player	0.46	0.36	8	0.36	0.31	0.25
More than other player	0.52	0.38	8	0.36	0.31	0.21
Payout only	0.47	37	9	0.36	0.33	0.26
Payout + B's payout	0.52	37	8	0.28	0.25	0.22
Payout + pot	0.50	38	8	0.44	0.30	0.24
All information	0.48	37	8	0.37	0.32	0.19

D Robustness

Table 3: Effect of Relative Distribution on Choice, Controlling for Covariates

	(1) retain	(2) retain	(3) retain	(4) retain	(5) retain	(6) retain
Received more than B	-0.136	0.598	0.0243	-0.267	0.582	-0.387
	(0.308)	(0.346)	(0.349)	(0.529)	(0.606)	(0.778)
Age	0.0129	0.00290	-0.00429	-0.000601	-0.0191	0.0103
	(0.0129)	(0.0153)	(0.0147)	(0.0247)	(0.0310)	(0.0313)
Female	0.344	0.217	0.423	0.776	0.401	1.610*
	(0.305)	(0.338)	(0.356)	(0.544)	(0.609)	(0.820)
Urban	-0.447	-0.803	-0.318	0.587	-0.247	-0.462
	(0.322)	(0.414)	(0.383)	(0.581)	(0.740)	(0.897)
Education	-0.0612	-0.0807*	0.0149	-0.157*	-0.0703	-0.106
	(0.0338)	(0.0387)	(0.0448)	(0.0644)	(0.0704)	(0.0975)
Poverty	0.00217	0.000318	-0.00364	0.000818	0.00163	-0.0123
	(0.00240)	(0.00269)	(0.00248)	(0.00357)	(0.00366)	(0.0107)
Same ethnicity as leader	0.828*	0.233	-0.0159	0.253	-0.304	0.637
	(0.349)	(0.403)	(0.408)	(0.620)	(0.722)	(0.978)
Leader of president's ethnicity	0.646	0.595	0.170	0.229	-1.173	1.947
	(0.381)	(0.433)	(0.458)	(0.638)	(1.146)	(1.039)
Constant	-0.348	0.389	0.0190	-0.747	-0.402	-2.236
	(0.687)	(0.787)	(0.863)	(1.264)	(1.516)	(1.725)
lnsig2u						
_cons	-1.562	-1.718	-0.217	-2.239	-1.381	0.315
	(0.980)	(1.225)	(0.622)	(4.080)	(2.458)	(1.037)
N	207	170	167	104	82	81

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$