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

Elizabeth Carlson
Penn State
ecc13@psu.edu

November 8, 2017

Abstract

Relative distribution – whether one is favored or disfavored by government – seems to matter to voters in many African countries. But why? I demonstrate theoretically that voters who do not have information about government revenues, and who can't determine whether their allocation represents an absolutely high share, can compare their goods to others' to generate a rough prediction of whether they would be better off under a challenger. In this model, relative distribution is a heuristic, whose relevance is conditioned by available information. I test the model among a sample of Ugandan lab participants, who, when not told how much a "leader" had to distribute, were significantly more likely to reelect the leader when their payout was greater than another player's. Those given information about the size of the pot, on the other hand, no longer responded to relative distribution. I show external validity by demonstrating that the incumbent's ethnicity, which strongly predicts whom he favors, is less important to respondents on the cross-national Afrobarometer when they report easy access to information on government revenues; other types of information do not show a similar effect. Unfortunately, though rational under low information, voting on relative distribution is not optimal under all conditions: those who are favored by a highly corrupt leader are most likely to reelect him when they are relying on relative distribution to make their choice.

One of the most common arguments in the African politics literature is that voters support leaders who favor them or their communities in the distribution of development goods. Because goods are often distributed along ethnic lines, this may manifest as ethnic voting (Ferree, 2006; Padro i Miquel, 2007; Posner, 2005). However, voters will also vote for non-coethnic politicians, as long as they expect these politicians target them or their communities with goods (Arriola et al., 2016; Ferree and Horowitz, 2010; Ichino and Nathan, 2013).[1]

The reason *why* voters prefer leaders who favor them is often left unaddressed: receiving relatively more is simply taken as a proxy for receiving more in absolute terms. However, evidence from

[1]Similar arguments have also been made about voters elsewhere in the world (Manzetti and Wilson, 2007), particularly India (Acharya et al., 2015; Chandra, 2004; Kitschelt, 2000).

multiple countries suggests that voters care whether they are receiving more or less than others, controlling for any correlation this may have with the absolute size of their allocation. Bratton et al. (2011), for example, show that vote choice is strongly correlated with perceived government discrimination in a pooled African sample, even after differences in absolute income are controlled for. Other studies find that voters reward higher absolute distributions only when these distributions are also high in relative terms. In earlier work, I show that Ugandan voters are no more likely to support "good type" politicians who wouldn't favor them then they are to support "bad type" candidates who would not distribute anything to anyone(Carlson, 2015).Wantchekon (2003) shows that Beninese voters become significantly more likely to vote for a candidate when he promises to favor them with services, and less likely to vote for him when he promises to improve these same services for everyone. Finally, de Kadt and Lieberman (2015) find that South African voters actually *punish* the incumbent for improving public services over time if their services have not improved as much as those around them.

More importantly, in countries with substantial corruption, high relative distribution is not necessarily a good proxy for high absolute distribution in the first place. A large existing literature indicates that leaders who favor their supporters with goods are less likely to engage in other types of pro-poor spending, and more likely to keep government resources for themselves (Acharya et al., 2015; Desposato, 2006; Manzetti and Wilson, 2007; Thachil and Tietelbaum, 2015). A voter who receives a large slice of a small pie distributed by a kleptocratic leader may not actually be better off than a voter who receives a smaller slice of a larger pie distributed by a leader who is less corrupt. So why do voters respond strongly to *relative* distribution, preferring leaders who favor them, when this is not necessarily the best way to ensure high *absolute* distribution?

Many existing answers to this question highlight a psychological mechanism: they emphasize voters' strong emotional responses to unequal treatment or outcomes. Horowitz's influential theory of positional psychology argues that groups who are better off under a given regime experience higher self-esteem and less status anxiety (Horowitz, 1985); this is consistent with the result of numerous studies that link happiness to relative income (e.g.Easterlin (1995)). Others argue that inequality triggers anger and resentment in a way that poverty does not (Cedarman et al., 2011;

Gurr, 2011), and link government favoritism to grievance and animosity (Ilorah, 2009; Straus, 2015). The psychological costs of inequality can be strong enough to override material concerns, leading people to take on absolute losses to improve their relative standing (Solnick and Hemenway, 1998).

If voters gain direct utility from being better off than others, this is sufficient to explain why they prioritize relative distribution. Nevertheless, in this article, I argue that a psychological mechanism is not necessary: responding to relative distribution is rational under low information. I demonstrate theoretically that voters who do not know the leader's budget, and who can't determine whether their absolute distribution is as high as it could be, can use information about relative distribution to predict whether they would be better off under a challenger. I also demonstrate that considering relative distribution does not improve outcomes for those who already have full information about government revenues. If relative distribution is important because it serves as a heuristic, therefore, the correlation between relative outcomes and vote choice should vary with information, becoming salient primarily among those who do not have information on budgets.

I test the model by experimentally manipulating the amount of information available to lab subjects sampled from Uganda, where existing literature indicates that vote choice is driven by expectations of favoritism (Carlson, 2015; Velasquez, 2013). In the game, players received a payout from a "leader", who had divided a pot of money between himself and two voters, and chose whether to reelect the leader for the same payout, or replace him with a new leader. I show that players who were not told the amount of money the leader had to distribute were substantially more likely to retain the leader when their payout was higher than the other player's. Players who were told the size of the pot, on the other hand, were no more likely to reelect the leader when they were favored than when they were disfavored. Furthermore, given a choice of information to seek out, players were more interested in absolute than relative distribution: players were twice as likely to ask the size of the pot than to ask what the other player received. These results demonstrate that even players who do not necessarily prioritize high relative distribution will still use relative distribution as a heuristic under low information.

I provide evidence that the theory explains the salience of relative distribution outside the lab, and outside Uganda, by demonstrating that ethnicity, which serves as a proxy for whether the voter

is favored by the incumbent, is significantly less predictive of vote choice among a cross-national sample of voters who report that it is easy to learn how their government spends its revenue. Critically, only information about revenues has this effect: access to other types of information does not similarly reduce the salience of ethnicity.

The theory and findings I present here resolve a puzzling discrepancy. Though government favoritism features prominently in the African politics literature as a cause for everything from vote choice to conflict, direct analysis suggests that voters' psychological responses to inequality may actually be quite muted: studies conducted in African countries find no correlation between relative income and happiness for the typical subject (Akay et al., 2012; Alem, 2013; Ravallion and Lokshin, 2005).[2] By highlighting an informational mechanism, I explain why voters might respond strongly to relative distribution even if they gain no psychological utility from being better off than others.

Unfortunately, I also show that rationality will not necessarily prevent voters from supporting kleptocrats: both the theory and the game results indicate that voters who know they are favored, but who do not know total revenues, are more likely than those with full information and those with *no* information to retain leaders who are paying out less, and stealing more, than a challenger would. Accordingly, the study provides guidance for the design of informational policy interventions. Though budgetary information improves outcomes under all conditions, information on relative distribution, provided on its own, may induce voters to reelect a leader they would otherwise correctly reject, especially in highly corrupt regimes.

Because relative distribution is highly correlated with ethnicity and reflected in the distribution of clientelist goods, this study also has implications for the broader ethnic voting and clientelism literatures. Though numerous studies have argued that ethnic voting is caused by low information (Chandra, 2004; Conroy-Krutz, 2013; Posner, 2005), this study identifies the particular information that ethnicity substitutes for, and consequently, the information voters need in order to abandon their reliance on ethnic cues. Similarly, though others have noted that low-information voters ap-

[2]In general, relative deprivation is a concern of the rich: the subjective well-being of the absolutely poor is more strongly correlated with their ability to meet their basic needs (Asadullah and Chaudhury, 2012; Graham and Pettinato, 2001; McBride, 2001).

pear to have a particularly strong preference for particularistic goods (Desposato, 2006; Shin, 2015; Wantchekon and Vermeersch, 2008), I explain where this preference may come from: only discretionary goods, which are distributed differently across candidates and voters, help low-information voters determine whether they would be better off under incumbent or challenger. In general, the theory and results here suggest that corruption and weak bureaucracy, often treated as results of ethnic voting and clientelism, may actually be the cause: where leaders are unwilling or unable to disseminate budgetary information, strategic voters have good reason to prioritize relative distribution in deciding how to vote.

1 Theory

To understand the theory intuitively, consider a voter[3] who has received a basket of goods from his elected leader: these goods might include subsidies, transfers, or local public goods like schools or roads. These goods are not sufficient to meet the voter's needs. 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[4] is a shirker who stole or wasted a large portion of a budget that would otherwise have been sufficient. Or it may mean that the incumbent is giving the voter the most he can of a 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 different distributional preferences, and she should vote against the incumbent. In the latter scenario, the voter is already experiencing the best outcome she can given existing constraints, and it would be foolish for her to vote for a challenger. In order to determine which scenario she is in, the voter would really like to know how much there was available for distribution, and thus how much more (or less) she might get under a different leader.

And, yet, is it unlikely that a voter in an African country would be able to access such infor-

[3]The model below produces exactly the same results whether we are discussing a voter or a bloc of voters, and whether the pool of resources is local, national, or sectoral. The "voter" could be a household evaluating the distribution of ration coupons; an ethnic group evaluating the distribution of school funding from the center; or an industry evaluating the distribution of subsidies.

[4]To limit confusion about which actor is doing what, I use "he" throughout to refer to a leader,"she" to refer to a voter, and "they" to refer to a group of voters.

mation. African bureaucracies tend to be weak and under-resourced, limiting the collection and dissemination of financial data. What budgetary information is released is often impenetrable to a citizenry with low levels of formal education. Additionally, precisely because budgets can be used to identify corruption or shirking, leaders have a strong incentive to hide information about their finances. We should expect *all* leaders to tell voters that they are spending efficiently and distribution is insufficient because budgets are tight. And in some cases, this may even be true: the average GDP per capita in Sub-Saharan Africa is only $3700.[5] If voters are unable to access useable information about budgets, there will remain substantial uncertainty as to whether poor outcomes occur because leaders choose to distribute little, or because leaders have little to distribute.

Though there has not been much direct measurement of what African citizens know about budgets at the local or national level, there is evidence that very little information trickles down to the public. A large majority of Round 5 Afrobarometer respondents (77%) say that accessing information about how government spends its revenues is "difficult" or "very difficult." In Mali, most voters are unaware that their local leaders have access to earmarked funds for the provision of clean water, health clinics and schools (Gottlieb, 2016). In Uganda, the strong effects of an intervention that publicized educational disbursements from the central government imply that this information was not widely accessible, even to the teachers and administrators at the schools that were supposedly receiving the funds (Reinikka and Svensson, 2005). Meanwhile, in Malawi, a remarkable 42% of Afrobarometer respondents said they "did not know" whether the president's office was corrupt the year before an audit revealed diversion of at least $32 million - 4% of the total government budget - from a subset of government transactions in only six months. That Malawians were so uncertain about the existence of such extensive corruption implies they did not know how much should have been available for development.

However, even where a voter can't access information on budgets, she can gather reasonably reliable information about how her allocation compares to that of others. Though leaders may hide their theft of resources, they have an interest in highlighting distribution. News media routinely cover construction of roads; assignment of new teachers; and distribution of private goods like

[5] World Bank estimate for 2016 in current international dollars, purchasing power parity adjusted

scholarships, vouchers, and even vehicles or livestock. A voter who learns that others are receiving bags of grain while she goes hungry, or that another clinic is receiving new equipment that her local clinic already had, can infer that she is (dis)favored in the distribution of resources. Accordingly, African citizens report far more certainty about relative outcomes than about available resources. Over 60% of Round Five Afrobarometer respondents are willing to assert that their living conditions are either better or worse than others'; only 3% say they don't know.[6] 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 the voter can determine whether she is favored, she can also infer whether she might be better off under a new leader. A voter who observes not only that her own goods are insufficient, but that others' goods are better, knows there were resources withheld from her that might be available under a different leader with different preferences. On the other hand, if she observes that others have even worse outcomes than she does, she can infer that her share of resources is already higher than it might be and it would be risky to replace the incumbent in hopes of getting even more. If relative outcomes can help voters evaluate their leaders' performance, and are more readily observable than the total pool of resources, it is reasonable that voters would rely heavily on this information, even if they don't care about inequality for its own sake.

1.1 Formalization

Formalization of voters' decisions under different levels of information confirms that relative distribution is sufficiently correlated with absolute distribution to provide a useful heuristic when voters lack information about revenues, but insufficiently correlated to make considering relative distribution advantageous under full information. If voters care about relative distribution for strategic reasons, rather than psychological ones, they should respond to relative distribution only under partial information. The formalization also generates a non-obvious but important implication: relative distribution becomes less informative as leaders steal more, and relying on relative distribution under high corruption can lead voters to retain kleptocratic leaders they otherwise would

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

not.

To formally theorize the informational role played by relative distribution, I assume a simple game with four players: an incumbent, a challenger, and two voters (or interest groups), A and B.[7] Each candidate divides a pool of resources between himself and the two voters, according to his preferences, which are a function of his personal proclivity toward corruption (his type) and the structural conditions that determine whether A, B, or both are necessary for his electoral coalition.[8] To make the model more straightforward, I assume that both the total pool of resources and candidate preferences remain constant across time.[9]

Equations 1 - 4 are equivalent ways of representing distribution by the incumbent. Equation 1 presents total resources as the sum of the distributions to each player, where each $\theta_G xi$ represents the share of total government resources given to player x and is the player's **absolute distribution**. Equation 2 presents total resources as the sum of the share the incumbent kept for himself ($\theta_G ii$) and the remainder ($\theta_R i$), which encompasses what he distributes to A and B together. Equation 3 introduces $\theta_R xi$, which is the share of the non-stolen remainder allocated to player x and signifies that player's **relative distribution**: a is favored if $\theta_R ai > \theta_R bi$. Finally, Equation 4 presents the sum in terms of actual dollar amounts, where G, R_i, and D_{xi} represent the dollar amounts of the total budget; the remaining budget after the incumbent takes his share; and the allocation given to each player x, respectively. Throught the model and the piece a "distribution" refers to a share of one, while an "allocation" is a dollar amount.

To notate distribution by the challenger rather than the incumbent, every i in the model is simply replaced by c.

After observing their allocation D_{ai}, A must choose whether to reelect or replace the incumbent.

[7] There are, of course, many more than two voters or interest groups in most political settings, but it is also possible to divide any society into two groups - A and Not A - along whatever cleavage is salient to A.

[8] This model, and indeed the premise of the study, rests on the assumption that voters are fundamentally persuadable. Since both A and B can be bought, a candidate's decision to appeal to each group is determined by a variety of factors - electoral rules, elite bargaining, and the size of groups A and B - that are out of voters' control (Arriola, 2013; Elischer, 2013; Oyugi, 2006; Posner, 2005).

[9] In reality, leaders' preferences are almost certainly neither constant nor exogenous. To accurately model how voters' responses endogenously shape leaders' distributional decisions, however, we need to know what those responses are. I test the latter, and leave the former to future work.

Equivalent equations denoting distribution by incumbent

$$1 = \theta_G ii + \theta_G ai + \theta_G bi \tag{1}$$

$$1 = \theta_G ii + \theta_R i \tag{2}$$

$$1 = \theta_G ii + (\theta_R ai + \theta_R bi) \times \theta_R i \tag{3}$$

$$G = D_{ii} + D_{ai} + D_{bi} = D_{ii} + R_i \tag{4}$$

Because G is constant across candidates, and preferences are constant within candidates, the payout in time $t+1$ if A chooses to reelect will be the same as the payout in time t. If A chooses to replace, they will receive the distribution provided by the challenger ($\theta_G ac$). A will choose to reelect the incumbent when their absolute distribution under the incumbent is greater than the absolute distribution they expect under the challenger, or when $\theta_G ai > \mathrm{E}[\theta_G ac]$.

A cannot know for certain what their absolute distribution will be under the challenger, because the challenger has not had an opportunity to reveal his preferences. A may have some priors about what they would receive from the challenger,[10] or they may not, but they can always calculate an expectation because as a share of 1, $\theta_G ac$'s probability distribution is bounded.

However, A can never calculate an expected dollar value for G, because the distribution has no upper bound. Without G or $\mathrm{E}[G]$, A also cannot calculate expected values for any D_{xy}, since there is way to know what these allocations should sum to. The decision rule that A uses to maximize their payout thus depends heavily on what they can directly observe about G and D_{xi}.

1.2 Decision Rules

I provide three decision rules, each of which requires a different set of information. In all conditions, A can directly observe the dollar amount of their allocation (D_{ai}) and calculate an expected value for their absolute distribution under the challenger ($\mathrm{E}[\theta_G ac]$), but knowledge of G and D_{bi} vary. A summary of the outcomes under each decision rule is shown in Table 1, which includes both the case when A has no priors about distribution under either candidate, and the case in which A is

[10] I do not make any assumptions about where voters' priors come from, if they have them. Voters could gather information from campaigns or party intermediaries, or they might form expectations of the challenger based on the behavior of the incumbent. It is also plausible that voters have no priors at all, because candidate's promises or explanations are not always seen as credible (Baldwin, 2013; Keefer and Khemani, 2003; Kramon, 2016).

assigned moderately informative priors about both. The details of the simulations used to generate the statics are shown in Appendix B.

A can use **Decision Rule One** if they have **full information** about G. Since $\theta_G ai = D_{ai} \backslash G$, A will precisely calculate their absolute distribution and reelect if this is greater than the absolute distribution they expect under the challenger. A will reelect when:

$$\theta_G ai > \mathrm{E}[\theta_G ac] \tag{5}$$

Under full information, A has uncertainty only about absolute distribution under the challenger and they are able to select the candidate providing the higher absolute distribution most of the time. Table 1 shows that this rule allows A to select the candidate providing the higher absolute distribution 66% of the time without any priors; with moderately informative priors, the rate of correct votes increases to 72%.

A can use **Decision Rule Two**, when they have **no information** other than their allocation(D_{ai}). Without knowledge of G, A cannot determine their absolute distribution under the incumbent from D_{ai}. A must simply estimate $\mathrm{E}[\theta_G ai]$ according to their priors about the incumbent, and decide whether this value is greater than the absolute distribution they expect under the challenger. A will reelect when:

$$\mathrm{E}[\theta_G ai] > \mathrm{E}[\theta_G ac] \tag{6}$$

In the no information condition, there is uncertainty about absolute distribution under both incumbent and challenger, making this decision rule less efficient than Rule One, particularly when the A has no priors. Without priors, this rule performs no better than chance. With informative priors, this rule leads A to the correct outcome 65% of the time.

A can use **Decision Rule Three**, when they have **partial information** that includes their allocation and B's allocation (D_{bi}). A still does not know G, but with knowledge of D_{ai} and D_{bi}, A can calculate R_i, and their share of R_i ($\theta_R ai$). As shown in Appendix A, A's relative distribution ($\theta_R ai$) is positively correlated with their absolute distribution ($\theta_G ai$): $\theta_G ai$ is a function of, and

always increasing in, $\theta_R ai$. More importantly, A can use their share of distribution, entered into Equation 7 below, to determine whether their absolute allocation is greater than what they would expect from a challenger. A will reelect when the ratio between their current relative distribution and the relative distribution they expect under the challenger is greater than ratio of the total amounts remaining after each candidate takes a share for themselves. The greater A's relative distribution, the more likely the inequality will resolve in favor reelection. A will reelect when:

$$\frac{\theta_R ai}{\mathrm{E}[\theta_R ac]} > \frac{\mathrm{E}[\theta_R c]}{\mathrm{E}[\theta_R i]} \tag{7}$$

Under partial information, A still has uncertainty about absolute distribution under both incumbent and challenger, but uncertainty is reduced by exact knowledge of $\theta_R ai$. Accordingly, when A has no priors, Rule Three leads to the correct vote 58% of the time. With priors, Rule Three leads A to the correct vote 68% of the time. In both cases, the share of correct votes is higher than A would obtain using Rule Two under no information, but lower than they would obtain using Rule One under full information.

Table 1: Simulated Outcomes Under Each Decision Rule

			Correct (%)	
	Information	Reelect if:	No priors	Priors
One	Full: D_{ai} and G	$\theta_{ai} > \mathrm{E}[\theta_{ac}]$	66	72
Two	None: D_{ai}	$\mathrm{E}[\theta_{ai}] > \mathrm{E}[\theta_{ac}]$	50	65
Three	Partial: D_{ai} and D_{bi}	$\frac{\theta_R ai}{\mathrm{E}[\theta_R ac]} > \frac{\mathrm{E}[\theta_R c]}{\mathrm{E}[\theta_R i]}$	58	68

The table presents outputs from simulations described in Appendix B. A "correct vote" is one in which A selects the candidate providing the larger absolute distribution.

Critically, however, the ratio in Rule Three means that the rule becomes increasingly less helpful as the leader becomes increasingly corrupt (as $\theta_R i$ shrinks). Consider a case in which A is somewhat favored by the incumbent ($\theta_R ai = 0.6$) and expects to be somewhat disfavored under the challenger ($\mathrm{E}[\theta_R ac] = 0.4$). Given this ratio, A will be better off under the incumbent unless the challenger distributes at least 1.5 times more than the incumbent does. If the incumbent is not stealing

very much, and $\theta_R i$ is large, it is almost certain that A would be better off under the incumbent, because the challenger would never able to distribute 1.5 times more in total. If the incumbent is a kleptocrat, on the other hand, and $\theta_R i$ is small, a challenger can easily distribute 1.5 times more than the incumbent, and thus the incumbent's favoritism toward A is not a strong signal that A would be better off under the incumbent.

Indeed, as shown in Figure 1, when corruption is high, A will actually make a *worse* decision relying on Rule Three than voting at random. Figure 1 plots voting outcomes over $\theta_R i$, the share of resources being distributed to A and B combined after the incumbent takes his share. In both figures, A is favored. Figure 1a plots the likelihood A will select the candidate providing the highest absolute distribution. It demonstrates that Rule Three becomes more effective as the incumbent steals less and $\theta_R i$ grows. Where the incumbent is stealing almost everything, and $\theta_R i$ is very low, Rule Three is the least effective voting rule, producing fewer correct votes than either Rule One (full information) and Rule Two (no information). This is because, as shown in Figure 1b, it is when A votes on relative distribution that they are most likely to retain incumbents who are stealing almost everything, and distributing very little to either A or B. Altogether, the model demonstrates that while voting on relative distribution produces better outcomes *on average* than voting on no information, the reverse is true when the incumbent is capturing most state resources: those with information about relative distribution will be more likely than those with no information to reelect a leader they should not. Unfortunately, without knowledge of G, A does not know whether $\theta_R i$ is high or low, and thus whether using Rule Three will improve or worsen outcomes.

1.3 Hypotheses

The model produces a number of testable implications that should obtain if voters respond to relative distribution for its informational content. First, Rule Three should be used only by those with partial information. Those with no information, who do not know D_{bi}, can't use Rule Three, and those who know G will get better outcomes using Rule One instead.[11]

[11] Table 3 in the appendix tests alternative decision rules that combine relative distribution with other information to show that there is no case in which a strategic voter benefits from considering relative distribution under full information. Where relative distribution and absolute distribution send consistent signals, A gets the same results whether or not they add information about relative distribution. Where relative distribution and absolute distribu-

Figure 1: Vote Choice of Favored Voters, Over $\theta_R i$, by Decision Rule

(a) Cast correct vote

(b) Retained incumbent

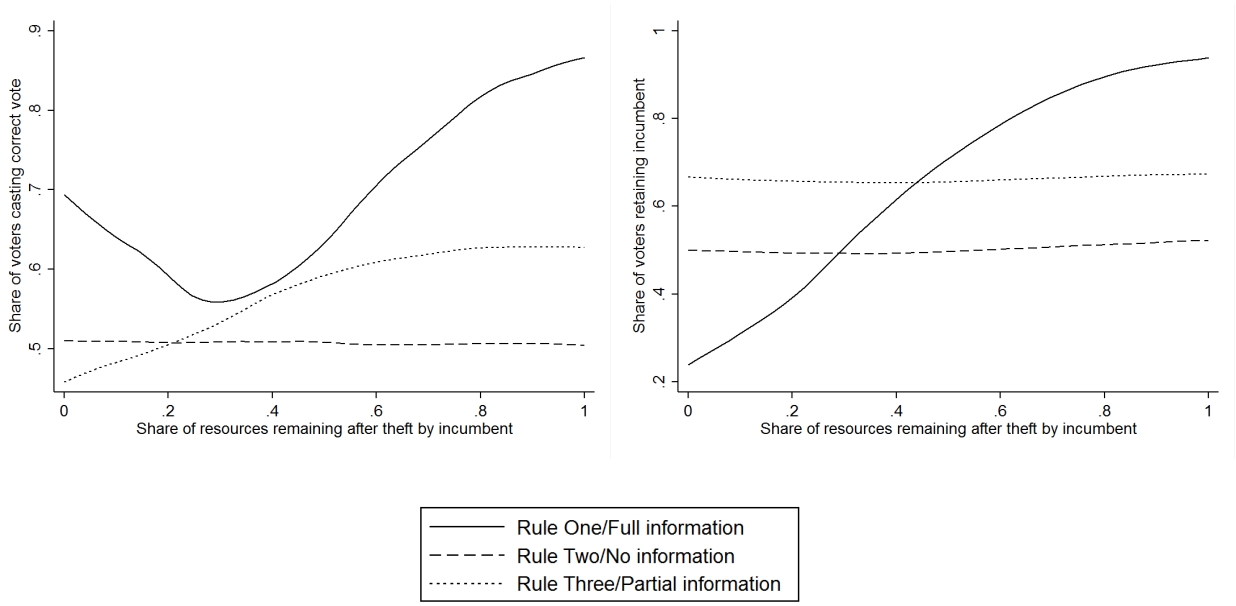

Note: Figure shows simulated likelihood that a voter who is favored by the incumbent a) casts a vote for the candidate providing the highest absolute distribution and b) votes to retain the incumbent when the voter has uninformative priors. "Favored" voters are those who are receiving more than the other voter ($\theta_R ai > \theta_R bi$). Specifications of simulation described in Appendix B.

Hypothesis One: The marginal effect of relative distribution on vote choice will be greatest under partial information.

Second, strategic voters should seek out information that helps them maximize their payout. Since A will always cast the most correct votes using Rule One, which is only possible if they know G, A should prefer to gather information on G over any other type of information. There is no reason for a payout-maximizing voter who can learn G to seek out information on D_{bi} instead.

Hypothesis Two: Voters given a choice of information will seek out G rather D_{bi}.

Finally, we should see that even where voters are strategic, reliance on relative distribution will produce poor outcomes under high corruption.

Hypothesis Three: Of voters who are favored, those with partial information will be most likely to retain leaders who are keeping more, and payout out less, than a challenger would.

Note that these hypotheses directly contrast with what we should expect if voters' response to relative distribution reflects inequality aversion or another psychological motivation. If voters gain direct utility from high relative distribution, the marginal effects of relative distribution will not be conditional on knowledge of G or on the extent of corruption: voters will reward relative distribution equally under partial and full information. Given an option, voters who value relative distribution for its own sake should also seek out information on relative distribution.

2 Behavioral Game

To determine whether voters can and will use relative distribution as heuristic to compensate for low information, I use a behavioral game in Uganda. In the game, players chose whether to reelect or replace a "leader" who had divided a pot of money between himself and two voters. The

tion provide opposite signals, considering relative distribution decreases the likelihood A will choose the candidate providing the highest absolute distribution.

experimental manipulation was how much information the player, Voter A, had about the size of pot and how it had been distributed.

I selected Uganda because it fits the scope conditions of the theory well. We have good existing evidence that vote choice in Uganda is correlated with whether voters expect to be favored by the regime (Carlson, 2015; Velasquez, 2013). Ugandan voters are also uncertain about the extent of corruption at the highest levels of the government: 14% of Afrobarometer respondents report that "everyone" in the president's office is corrupt, 14% report that no one is, and the rest either give an answer in between or say they don't know. Information about revenues is also low, with over 80% of Ugandans reporting that it is difficult or very difficult to determine how government uses revenue. It is therefore an appropriate test case for determining whether low information and uncertainty about available resources can explain voters' reliance on relative distribution.

I use a lab experiment because it provides a clean test of the effect of information. Though subjects can always be randomly assigned to hear real-world information in a survey or field experiment, it is difficult to create truthful variation in relative distribution across subjects while keeping other factors constant: every voter is either favored or disfavored on a given good, and these outcomes are likely to be correlated with other relevant variables. Perhaps more importantly, voters in the real world have pre-existing beliefs with which any provided information is likely to interact: some voters may already have precise information about budgets or distribution; some may reject information that threatens their partisan identity; and some may only react to information they find surprising. Any or all of these might cause non-compliance with the intended information treatment. In the lab, on the other hand, all information is new and players do not know the "leaders" they are playing against. If information does not change behavior in this setting, we can conclude that information itself was unhelpful, rather than that subjects already knew, disbelieved, or otherwise failed to take in the information.

A potential trade-off for a clean theoretical test is a loss of external validity. Lab games are often abstract from reality, and intentionally structured to reward strategic calculations. A finding that players act to maximize their payouts in such a game does not necessarily mean that voters' real-world behavior is similarly grounded in strategy. On the other hand, lab games can and do

pick up players' emotional responses to inequality in particular: players in dictator games routinely act against their material interest to punish distributions seen as unfair, particularly when they are one receiving less (Dawes et al., 2007; D'Exelle and Riedl, 2013; Fehr and Gachter, 2000). To increase external validity, I made several design choices that would encourage players to act on any non-strategic motivations they might have.

First, I ran the game among a sample that is reasonably representative of the population of interest. Players vary markedly across regions in their responses to unequal lab distributions (Henrich et al., 2005): evidence of an emotional response to inequality (or a lack thereof) in an American undergraduate sample would not necessarily be a good indicator of whether such attitudes are relevant to voters in any African country. Instead, I pulled a random sample from Mukono District. Mukono District is appealingly diverse: the district contains Uganda's fourth largest city, also called Mukono, along with a large number of entirely rural parishes with a very low standard of living. Though certainly not representative of all African voters (or even all Ugandan voters), the population of Mukono demonstrates much of the variation in income, education, and urbanization we would expect to find in Uganda and other Sub-Saharan African countries. The characteristics of the sample, compared to those of the samples in the Uganda and a cross-national Afrobarometer sample, are listed in Table 4 in the appendix.

Second, I primed players to see the game as an election. In early rounds of a game, before they master the particular strategy involved, players in lab games rely on existing models of decision-making: the particular rubric they call up is affected by the frame on the game (Binmore, 2005; Kiyonari et al., 2000; Levitt and List, 2007). I overtly analogized the game to politics, and embedded it in a survey about political attitudes, encouraging players to apply the rubrics they use in real elections. Aspects of game play provide evidence that the frame was successful. Players apparently expected leaders of their own ethnicity to favor them: players who knew only their payout were 25 points more likely to retain leaders of their own ethnicity, but did not prefer coethnics once they had direct information on relative distribution. Ethnic favoritism is not what Ugandans should expect of players in an anonymous dictator game (Habyarimana et al., 2007), but *is* what they expect of political candidates (Carlson, 2015; Conroy-Krutz, 2013). Additionally, players retained leaders

from the incumbent president's ethnic group 9 points more often than leaders from other groups; this is hard to explain unless players were analogizing these leaders to Uganda's actual president. Any pyschological responses that voters have to unequal distribution from actual politicians should thus be reflected in the way they play the game.

2.1 Game Design

Testing the impact of relative distribution requires variation in how players' payouts compare to those of other players. Since the modal dictator in a dictator game distributes the pot evenly, I controlled gameplay with a computer program rather than having players play with other people.[12] The pot in each game was drawn from a normal distribution with a mean of 3000 Ugandan shillings and a standard deviation of 500 shillings. All payouts were drawn from uniform distributions that summed to one.

Relative distribution is a key independent variable. A's precise relative distribution, however, can never be fully de-correlated from A's absolute distribution: A can never receive a lower share of R_i than of G. To help isolate the effect of relative distribution, I generated a second, dichotomous variable that captured whether A was favored ($\theta_R ai > \theta_R bi$) and set this orthogonal to whether the player's share of the total was more or less than the 25% she could expect from the challenger.[13] To allow full factorialization, no player received a payout of zero (because there would be no way for Citizen B to receive less than this) and no player received more than half the pot (because there was no way for B to receive more than this.) Player payouts thus ranged from 15% to 45% of the pot in each game. All payouts were rounded to the nearest 200 shillings, which is the smallest denomination of coin enumerators could easily access in large quantities.[14]

[12] I did not tell subjects they were playing against a computer. Instead, I told them that they were playing against other people from elsewhere in Uganda. I did this, first, because I wanted to avoid novelty effects among a sample that does not normally use computers. Second, while most people have intuitions about how other people will play a game, accurately predicting the gameplay of a computer requires a fluency with probability that players without formal education may not have. Finally, noting that the allocator was a computer would undermine the frame I very deliberately put on the game. Since this is deception, all players were debriefed after the game and offered the chance to ask questions; none did. A more thorough discussion of my use of deception is in Section D of the appendix.

[13] I pre-tested the game with a sample of fully-informed players to determine the threshold at which they were likely to reelect; from this I inferred what they expected of a challenger. Players 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.

[14] The experimental design succeeded in artificially lowering the correlation between absolute and relative distri-

Over the course of game, players received three payouts. The first payout was provided immediately after the player learned about the leader and before she chose whether to retain or replace him: this payout was not strictly necessary, as players were incentivized by the payout that resulted from their choice, but we provided it so that the promised future payout would be concrete and credible.[15] After receiving this payout, 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 request one piece of information from a list of information they were missing about the current game;[16] they could also refuse more information. Upon receiving the information, they voted once more whether to retain or replace the leader, for the appropriate payout. Players interested in maximizing their payout had an incentive to ask for information that would improve the accuracy of their vote. However, information was costless, so players who cared about something other than maximizing their payout – such as whether allocations were fair – were not disincentivized from asking for the pertinent information.

The average total payout was 2200 Ugandan shillings, which at the time of data collection was approximately US$1, or about 1/3 of a typical daily income (UBOS, 2014).

2.2 Information Treatments

The key experimental treatment in the game was the amount of information that each player was given about how much money there was in the pot and how it was divided. Each player was assigned to one of three informational conditions, which accorded with the decision rules in the theory section: 1) a "no information" condition in which players knew only their own payout; 2) a "low information" condition in which players knew their own payout and that of the other citizen, but not the total pot that the leader had to distribute; and 3) a "full information" treatment in which players knew their payout, the other player's payout, the size of the pot, and, by subtraction,

bution. In the simulations, relative and absolute distribution are correlated with an R of 0.65. In the game, the dichotomous relative and absolute distribution measures are not correlated, and each player's precise relative and absolute payouts are correlated with an R of 0.14.

[15] This pre-decision payout puts voters in the frame of losses and raises the possibility of endowment, wealth, or portfolio effects. Any of these might affect game play. However, the magnitude of the payout is not correlated with the assigned information treatment, so measures of the effect of information should be unbiased.

[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.

how much the leader kept for himself. (See treatment scripts in Appendix D.1) I limited accidental disclosure of information to those in the lower-information treatments by sending the enumerator only enough information about each game to meet the requirements of the assigned treatment, and drawing a new pot for each game so that enumerators would not be able to learn the size of the pot over time.

I also limited *mis*information by assigning each leader an ethnicity, as indicated by the city he was playing from. In Uganda, as elsewhere in Africa, ethnicity and distribution are highly intertwined. Leaving ethnicity unassigned presented a risk that players would make assumptions about it, and then infer their absolute or relative distribution based on assumed (non-)coethnicity with the leader.[17] Players' responses might also be conditioned by the leader's ethnicity: players may be more (or less) likely to punish coethnics who disfavor them, or they may hesitate to replace any leader of the president's ethnicity, especially if they think the game is a covert attempt to learn their political preferences. To avoid bias from unmeasured heterogeneous effects, I randomized the leaders' ethnicities and control for them in analyzing the results of the game.

Note that, despite my best efforts, players in the lower information conditions were not entirely uninformed. Unlike actual government budgets, payouts in the game were within the realm of daily experience: there is a limit to the amount a researcher can give away and to the amount an enumerator can carry around in cash. Though players certainly had a wide range of expectations – uninformed players asked to guess the size of the pot guessed values from 100 shillings to 60 million – the median guess was 10,000 shillings. This is more than the actual median pot size, but well with within an order of magnitude. More importantly, players' guesses as to their share of the pot were significantly correlated with their actual share of the pot ($r = 0.13$; $p = 0.01$). This means that players in the game had informative priors about G, and their share of it. As shown in Table 1, information is has less of an impact on outcomes when voters have priors, meaning that the effect of information is likely smaller than it would have been under greater uncertainty.

[17] In particular, I was concerned that players would either assume the leader was from the locally dominant ethnic group, or believe the leader was a stand-in for Uganda's actual president. As it turns out, these concerns were valid. Of a small sample of players 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. The next largest share, 20%, thought the leader was playing from President Museveni's home district.

2.3 Implementation

Players were recruited from every fifth household in two directions out from the main trading center of 50 randomly selected electoral constituencies within Mukono District. The game took place in players' homes, following an introductory survey about distribution and politics. Enumerators read the questions aloud in English or Luganda, according to the player's preference, and recorded answers on a smartphone. The game itself was played by text message. Once the player understood the game's instructions, the enumerator sent an SMS containing the player's 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.

The 569 players I include in the analysis represent only 70% of the people who were actually recruited to played the game: the study suffered from an unexpectedly high percent of missing or invalid data.[18] As a result, the study is underpowered[19] to detect significant treatment effects at standard levels: I increase power by post-stratifying on pre-treatment covariates.

Balance plots are shown in Figure 6 in the appendix. Those assigned to be favored, and those assigned to the partial information treatment, look similar in terms of age, education, employment and food insecurity. However, there are imbalances at the 10% level in gender, urban residence, and leader ethnicity. I control for all imbalanced covariates in every analysis to avoid biasing my estimates.

[18] To avoid having to pause the games due to electrical shortages, players did not play against the computer in real time. Instead, randomized games were pre-printed on paper forms, and the code of the player 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 the research assistant's residence before the data 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 whose data had far less variance than those of other enumerators, raising concerns about coaching or fabrication. Finally, I dropped any game in which the game details entered by the enumerator did not match those sent out by the research assistant.

[19] Ex ante, the study was *over*powered based on treatment effects from the simulation. Lost data combined with a smaller treatment effect due to information spillover combined to reduce power below 80% at 95% significance.

3 Experimental Results

Hypothesis One predicts that relative distribution will have the greatest impact on vote choice for those who have partial information. The results of the first portion of the experiment are consistent with this expectation.

3.1 Relative Distribution

Figure 2 presents the smoothed[20] likelihood the leader was retained, over the players' relative distribution ($\theta_R ai$), under each of the three information treatments. Next to each line are two OLS coefficients that estimate the effect of increasing the player's relative distribution from 0 to 100%.[21] The first coefficient is from an OLS that controls for absolute distribution and the imbalanced covariates of gender, urban residence and leader ethnicity. The second coefficient is from a model with all these controls, plus age, education, food insecurity, un(der) employment and enumerator fixed effects.[22]

The figure shows that the slope of the fitted line is substantially steeper for those with partial information than for those under either of the other two information conditions. The coefficients on relative distribution are at least twice as large, and only statistically significantly different from zero, under partial information. The coefficients indicate that those who are strongly favored are roughly 20 points more likely to retain the leader than those who with a very low relative distribution. Critically, voters' response to relative distribution under full information looks very similar to voters' response under no information.

Figure 3 plots the effects of relative distribution measured dichotomously: the bars represent the difference in the rate of retention between those who were favored and those who were disfavored by the leader, with 95% confidence intervals. The darker bars provide the estimated treatment effect on the treated (ATT) when players are matched on the imbalanced covariates of gender, urban

[20] I use a kernel-weighted polynomial smoother, with a bandwidth of 0.3

[21] I use linear regression to make the coefficients straightforward to interpret; a logit model generates the same results.

[22] Food insecurity is a measure of how often the respondent reports being without sufficient food that ranges from never (0 times per year) to always (365 times). Un(der)employment takes a one if the respondent reported that they were seeking employment or more hours per week.

Figure 2: Retention Rate over Relative Distribution, by Information Treatment

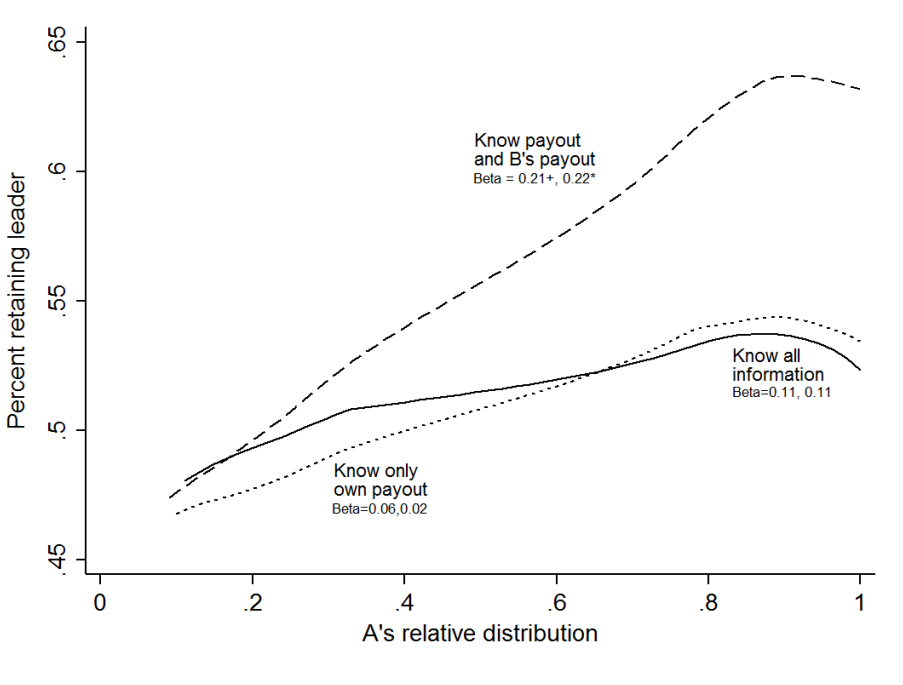

† $p < 0.10$ * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note:The figure plots the rate of retention over the player's share of distribution, under each information treatment. Coefficients are from two models. The first is the effect of relative distribution, controlling for absolute share of distribution, gender, urban residence and the leader's ethnicity. The second is from a model that adds controls for education, age, food insecurity, un(der)employment, and enumerator fixed effects.

Figure 3: Effect of Being Favored on Likelihood of Retaining Leader, by Information

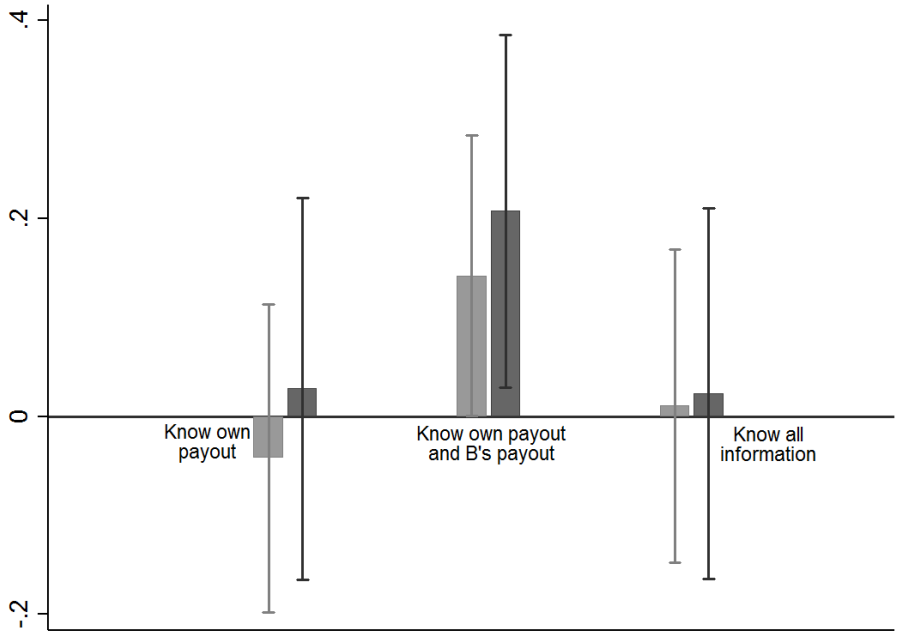

Note: The figure plots the difference in the rate of retention between those who are favored and disfavored by the leader. An effect size of zero indicates that those who are favored and disfavored retain the leader at the same rate. Darker bars present ATT, post-stratifying on gender, urban residence, and the ethnicity of the leader in the game. The lighter bars show the estimated ATT with additional stratification on age, food insecurity, education, enumerator ID, and location.

residence and the ethnicity of the leader. The lighter bars provide the treatment effects when the sample is also post-stratified on age, education, food insecurity, unemployment, and enumerator. The figure again shows that the effect of being favored is largest for those with partial information. In the partial information condition, the effect of relative distribution is quite large: those who are favored are 14-16 percentage more likely to retain the leader; both effects are again statistically significant at the 5% level. The effect of relative distribution under both no information and under full information are close to zero or even slightly negative.

The reader will note that the standard errors are large: despite substantial differences in the estimated effects of relative distribution across information treatments, these effects are not statistically significantly different from one another. I cannot rule out that being favored has the same effect under all information treatments.[23] In the following sections, I build the argument by

[23]It is also possible that absolute treatment effects are masking the contradictory effects of absolute and relative

providing additional evidence consistent with the two other implications of the model.

3.2 Information-seeking

Hypothesis Two proposes that, given a choice, voters will seek out information about the size of the pot, rather than relative distribution, because the former will allow them to use a more effective decision rule. The results of the second part of the game, in which players were allowed to seek out information, provide evidence in support of this hypotheses.

Figure 4 provides the percent of players who didn't have information about the other players' payout or the size of the pot, respectively, who asked for that particular information, with 95% confidence intervals on the means. Twice as many players requested information about the size of the pot (29%) than requested information about the other players' payout (14%), a difference that is significant at the 1% level.

Figure 4: Percent of players requesting given information

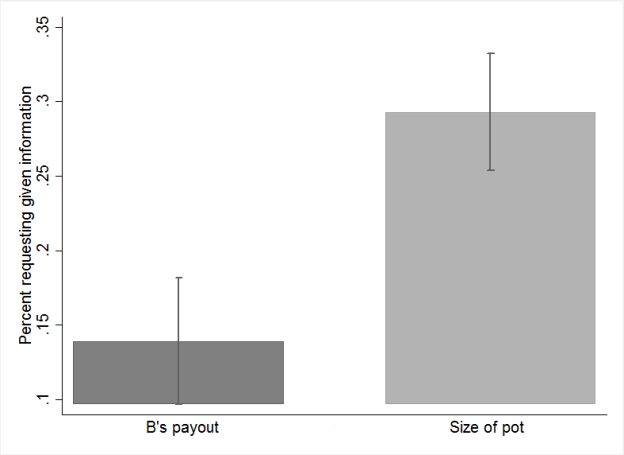

Note: The figure presents the percent of players missing each type of information who requested it, with 95% confidence intervals.

These results are consistent with a model in which relative distribution is a second-best type of information that becomes valuable to voters when they can't access information about the total

distribution for some respondents, which reduced the overall effect of relative distribution such that, with error, the effect appeared to be zero. Table 5 in the appendix shows that this is not the case: in a regression that assesses the effects of relative and absolute distribution simultaneously, the coefficient on relative distribution for players with full information is highly insignificant and very small relative to that of absolute distribution. Players in the game simply did not respond to relative distribution under full information.

pool of resources. Combined with the results above, they also indicate that voters do not need to prioritize relative distribution in order to respond to it under low information.

3.3 Retention of Kleptocrats

The final hypothesis predicted that voters who are favored by a highly kleptocratic leader will be more likely to reelect him when they have partial information than when they have either full or no information. Figure 5 shows the percent of all players who chose to retain a leader who was keeping more than the average amount of 50% of the pot.

Figure 5: Percent of Favored Players Retaining Highly Corrupt Leader, by Information

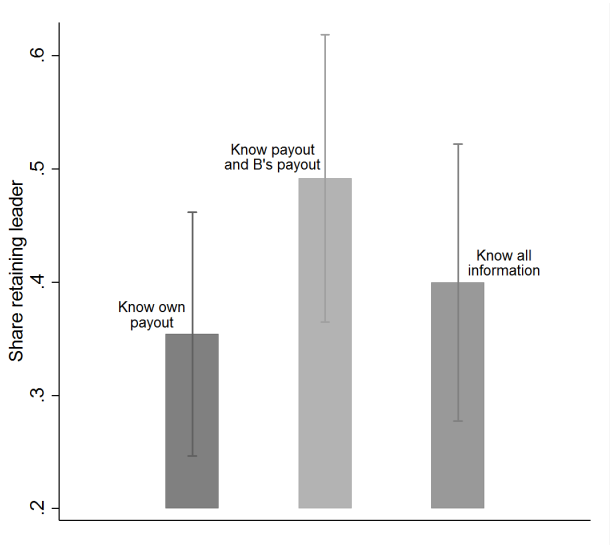

Note: Figure presents share of players choosing to retain a leader keeping more than 50% of the pot for himself, with 95% confidence intervals.

The results are consistent with the expectations of the model: those who knew they were favored but didn't know the size of the pot were the most inclined of any group to retain the leader. As above, the error bars are large and retention rates are not statistically significantly different across information treatments. However, the size of the effects is substantively meaningful: favored voters in the partial information conditions were 14 percentage points more likely to select a leader who was stealing than players in the no information condition, and 9 points more likely than those in the full information condition. These results are consistent with the prediction of the model, that

giving players information about relative outcomes in an otherwise low information context may increase the likelihood they will choose to reelect a leader they should replace. Both full information and *no* information produce better outcomes than information on relative distribution where theft is high.

4 External Validity

Though consistent with the model I propose, the results of the lab game have two primary shortcomings: I can't always detect a statistically significant difference across information treatments, and real voters in Uganda or elsewhere may not respond to actual politicians the way Ugandan subjects respond in the lab. In this section, I provide evidence that the effect of information about government revenues is both statistically significant and relevant in an observational, regional sample.

If my model is correct, and voters respond more strongly to relative distribution when they don't have information about budgets, we should see that they also respond more strongly to candidate ethnicity. In many African countries, the ethnic match between leader and voter is highly correlated with whether the voter is favored in the distribution of critical government goods (Ejdemyr et al., 2016; Franck and Rainer, 2012; Hodler and Raschky, 2014; Kramon and Posner, 2010). More importantly, ethnicity often predicts whether voters *believe* they are or will be favored (Carlson, 2015; Nathan, 2016). Indeed, anticipated favoritism is generally cited as the reason that voters in African countries prefer coethnic politicians (Ichino and Nathan, 2013; Posner, 2005). Table 6 in the appendix shows that, in the Afrobarometer data used below, shared ethnicity with the incumbent significantly reduces an index of lived poverty,[24] improves current and expected subjective well-being, and decreases perceptions of unfair treatment by government.[25]

The analysis includes survey data from 38,994 subjects in 25 African countries included in

[24]The indicator is an index of reported shortages in food, water, fuel, cash and healthcare.

[25]There is a direct measure of relative subjective economic well-being on the Afrobarometer, which is not correlated with ethnicity. However, it is unclear to whom respondents are comparing themselves when they answer. If they are comparing themselves to those about whom they have the most information – other members of their community, who are also likely to be coethnics – this measure provides no information about relative distribution along ethnic lines.

Round 5 of the Afrobarometer public opinion survey.[26] I select this round because respondents were asked how easy it is for them to learn how government uses the revenues it collects. I collapse the variable into a 1 if the respondent says it is "easy" or "very easy" to acquire information on revenues, and a zero if the respondent says it is "difficult" or "very difficult". If my model is correct, we should see that ethnicity, which is my proxy for perceived favoritism, is less predictive of vote choice when voters can access budgetary information. In the models in Table 2, I regress reported support for the incumbent on the respondent's ethnicity and reported access to financial information; the key independent variable is the interaction between these two indicators, which I expect to be negative.

Those who report easy access to financial information are bound to be systematically different from those who do not: they are likely to be more educated, and to live in urban areas where information is easier to come by. They are also likely to live in more democratic countries with greater bureaucratic capacity. To account for confounders, I control for a variety of respondent-level characteristics that might be correlated with both reported information and vote choice, including education, media exposure, political interest and evaluations of the government performance. I add country fixed effects to account for country-level differences in transparency and the political salience of ethnicity.[27]

As an additional check on the proposed mechanism, I add placebo tests by interacting ethnicity with education and two other indicators that capture the overall information and bureaucratic environment. The first is a measure of how easily the respondent can determine what she owes in taxes; this captures the overall availability of financial information and the respondent's skill at acquiring it. The second variable is the ease with which respondents can enroll their children in public school, which captures government capacity and respondents' ability to navigate the bureaucracy. If the theory is correct, it will specifically be information about revenues, rather than respondents' sophistication, general transparency or bureaucratic quality, that conditions the relationship between ethnicity and vote choice.

[26]There are 34 countries included in Round 5, but surveys in nine of these countries do not ask the respondent's ethnicity, so those countries are dropped.

[27]I get substantively similar results with a random effects or GEE model with effects at the level of the country or the village.

Table 2: Effect of Relative Distribution on Support for Incumbent, Conditional on Information

	(1)	(2)
Coethnic	0.163*	0.126
	(0.061)	(0.100)
Education	-0.018**	-0.018*
	(0.006)	(0.007)
Easy to learn gov't revenue	0.078***	0.082***
	(0.020)	(0.021)
Easy to learn taxes owed	0.006	0.001
	(0.006)	(0.007)
Easy to enroll in school	0.027**	0.025**
	(0.008)	(0.009)
Coethnic × revenue easy	-0.066**	-0.081*
	(0.023)	(0.031)
Coethnic × education		-0.001
		(0.012)
Coethnic × taxes easy		0.015
		(0.017)
Coethnic × school easy		0.004
		(0.018)
N	17745	17745

Standard errors in parentheses

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

Multi-level OLS regression with country fixed effects and robust standard errors. Dependent variable is reported support for the incumbent party or president. Unreported controls included in both models include personal economic well-being, perceptions of the country's economy, lived poverty, rural residence, gender, news exposure, computer access, strength of ethnic identity, and a belief that the survey is being run by the government.

The results are consistent with the theory: voters who have more information about government revenues are significantly less influenced by the ethnicity of the incumbent. Model One shows that among those who report that accessing financial information is difficult or very difficult, those who are coethnics of the incumbent are 16 percentage points more likely to support him than are non-coethnics. Among those who say it is easy or very easy to access budgetary information, on the other hand, coethnics are only 10 points more likely to support the incumbent, a 40% decrease.

Model Two confirms that the effect is particular to information about government revenue: no other interaction is substantively or statistically significant, and controlling for these additional measures increases the magnitude of the interaction effect. Among those who have information about revenues, the marginal effect of coethnicity in Model Two is only 4.5 percentage points. If, as the literature suggests, voters rely on ethnicity as a signal about relative distribution, these results provide evidence that relative distribution is less useful as a heuristic when voters have direct information about government revenue. The results of the cross-country analysis provide evidence that the results of the experiment, which show that relative distribution is less relevant to the choices of those who have information about the total pool of resources, are valid outside of the lab, and outside of Uganda.

5 Conclusion

In this article, I argue that when voters in African countries reward leaders who favor them, they do so not (or not only) because of the psychological effects of inequality, but because relative distribution provides critical information. Voters who do not have information about available resources, and thus cannot determine whether the goods they have received from government represent a high or low share, can observe what others have received and roughly infer whether they would be better off under a different leader. Using an experimental game, I provide evidence that Ugandan voters who didn't know how much an elected "leader" had to distribute were more likely to retain him when their payout was higher than another player's. When players had full information, on the other hand, relative distribution had little effect on their choice. Furthermore, given an option to seek out information, players were far more likely to request information on the total available resources than on the equality of distribution across players. These results suggest that players did not particularly prioritize equality across players, but that such a preference was not necessary in order for them to respond to relative distribution under low information.

Though almost always better than voting on no information, voters who rely on relative distribution can make sub-optimal choices, especially where corruption is high: voters who know only that they are favored by their leader are the most likely to support a leader who is paying out

less than a challenger would. Leaders who wish to retain office while keeping maximum rents can therefore *induce* this type of voting by obscuring information about total revenues and instead reminding supporters that they are favored.

This is consistent with the pattern of goods provision, and information, we actually see in many African countries, suggesting these countries are in an equilibrium that is beneficial to leaders and that will remain until an information shock disrupts voting patterns. This problematic equilibrium provides theoretical justification for policy interventions that provide financial information to voters. However, my findings also indicate that such interventions are not without potential for harm. The shift in voting that often follows an informational intervention may not always reflect increased accountability. If voters are given information about relative distribution they will use it, but, in the absence of additional information about the total revenues, this can cause voters to support a kleptocratic leader they would otherwise have rejected. The risk of harm will be greatest in countries with high levels of corruption. Full information about budgets always provides better outcomes that information about relative outcomes, so budgetary information should be the priority.[28]

My results suggest that providing information about budgets will also have other downstream effects. If, as many have argued, ethnicity matters in Africa because it is cue about who is or will be favored by the government, then ethnic voting should decline as relative distribution itself becomes less salient. The results of the Afrobarometer analysis provide evidence that this is indeed the case.

Additionally, information about budgets may help disrupt clientelism. My theory implies that low-information voters reward leaders who favor them with clientelist goods not necessarily because they prefer these goods, but because there is variation across leaders in how these goods are distributed. Even if voters value public or programmatic goods, these goods can't help voters determine whether they are better off under incumbent or challenger, and vote choice will appear to be uncorrelated with these types of goods. Dissemination of budgetary information, which would reduce the need for voters to rely on targeted distribution as a heuristic, might provide the opportunity to voters to reveal a preference for programmatic goods.

[28]For an example of such an intervention on the ground, see Chong et al. (2011) or Paler (2013).

There are, of course, limits to the extent my theory can travel beyond the particular context in which I test it. My model cannot explain why (or whether) voters reward relative distribution where information levels are high: in these countries, relative distribution is not necessary as a heuristic. It also unnecessary to explain voting where corruption is known to be very low: as theft approaches zero, relative and absolute distribution become perfectly collinear and there is no trade-off between relative and absolute distribution to puzzle over. These scope conditions mean that my model is almost certainly inappropriate to explain voting in most advanced democracies. But they do not preclude its applicability to other developing countries outside of Africa. The things that tend to distinguish Africa from other regions, such as ethnic salience, are unnecessary to the model. The condition that *is* necessary - lack of information - is common characteristic of countries in the developing world. Thus, the theory and tests may also provide insight into voter choice in the newer democracies of Latin America, Asia and Europe.

References

Acharya, A., Roemer, J. E., and Somanathan, R. (2015). Caste, Corruption and Political Competition in India. *Research in Economics*, 69(3):336–352.

Akay, A., Martinsson, P., and Medhin, H. (2012). Does Positional Concern Matter in Poor Societies? Evidence from a Survey Experiment in Rural Ethiopia. *World Development*, 40(2):428 – 435.

Alem, Y. (2013). Relative Standing and Life-Satisfaction: Does Unobserved Heterogeneity Matter? Working Papers in Economics 579, University of Gothenburg, Department of Economics.

Arriola, L., Choi, D. D., and Gichohi, M. (2016). Political endorsements and cross-ethnic voting. Unpublished manuscript.

Arriola, L. R. (2013). Capital and Opposition in Africa: Coalition Building in Multiethnic Societies. *World Politics*, 65(2):233–272.

Asadullah, M. N. and Chaudhury, N. (2012). Subjective well-being and relative poverty in rural Bangladesh. *Journal of Economic Psychology*, 33(5):940 – 950.

Baldwin, K. (2013). Why vote with the chief? political connections and public goods provision in zambia. *American Journal of Political Science*, 57(4):794–809.

Binmore, K. (2005). Economic man - or straw man? *Behavioral and Brain Sciences*, 28(06):795–815.

Bratton, M., Bhavnani, R., and Chen, T.-H. (2011). Voting Intentions in Africa: Ethnic, Economic, or Partisan. Afrobarometer Working Paper 127.

Carlson, E. (2015). Ethnic Voting and Accountability in Africa: A Choice Experiment in Uganda. *World Politics*, 67(2):253–285.

Cedarman, L.-E., Weidmann, N. B., and Gleditsch, K. S. (2011). Horizontal inequalities and ethnonationalist civil war: A global comparison. *American Political Science Review*, 105:478–495.

Chandra, K. (2004). *Why Ethnic Parties Suceed*. Cambridge University Press, New York.

Chong, A., de la O, A., Karlan, D., and Wantchekon, L. (2011). Looking Beyond the Incumbent: The Effects of Exposing Corruption on Electoral Outcomes. NBER Working Paper 17676.

Cilliers, J., Dube, O., and Siddiqi, B. (2015). The white-man effect: How foreigner presence affects behavior in experiments. *Journal of Economic Behavior and Organization*, 118:397 – 414.

Conroy-Krutz, J. (2013). Information and Ethnic Politics in Africa. *British Journal of Political Science*, 43(2):345–373.

Dawes, C. T., Fowler, J. H., Johnson, T., McElreath, R., and Smirnov, O. (2007). Egalitarian motives in humans. *Nature*, 446:794–796.

de Kadt, D. and Lieberman, E. S. (2015). Do citizens reward good service? voter responses to basic service provision in south africa. Afrobarometer Working Paper no. 161.

Desposato, S. W. (2006). Parties for rent? ambition, ideology, and party switching in brazil's chamber of deputies. *American Journal of Political Science*, 50(1):62–80.

D'Exelle, B. and Riedl, A. (2013). Resource allocations and disapproval voting in unequal groups. *Journal of Conflict Resolution*, 57(4):627–652.

Easterlin, R. (1995). Will raising the incomes of all increase the happiness of all? *Journal of Economic Behavior and Organization*, 27:35–47.

Ejdemyr, S., Kramon, E., and Robinson, A. L. (2016). Segregation, ethnic favoritism and the strategic targeting of public goods. Unpublished manuscript.

Elischer, S. (2013). *Political Parties in Africa: Ethnicity and Party Formation*. Cambridge University Press, Cambridge.

Fehr, E. and Gachter, S. (2000). Fairness and retaliation: The economics of reciprocity. *Journal of Economic Perspectives*, 14(3):159–181.

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

Ferree, K. and Horowitz, J. (2010). Ties that Bind? The Rise and Decline of Ethno-Regional Partisanship in Malawi, 1994-2009. *Democratization*, 13(3):534–563.

Franck, R. and Rainer, I. (2012). Does the Leader's Ethnicity Matter? Ethnic Favoritism, Education, and Health in Sub-Saharan Africa. *American Political Science Review*, 106:294–325.

Gottlieb, J. (2016). Greater expectations: A field experiment to improve accountability in mali. *American Journal of Political Science*, 60(1):143 – 157.

Graham, C. and Pettinato, S. (2001). Frustrated Achievers: Winners, Losers and Subjective Well-Being in New Market Economies. Brookings Institution Center on Social and Economic Dynamics Working Paper No. 21.

Gurr, T. R. (2011). *Why Men Rebel*. Paradigm Publishers, New York.

Habyarimana, J., Humphreys, M., Posner, D., and Weinstein, J. (2007). Why Does Ethnic Diveristy Undermine Public Goods Provision? *American Political Science Review*, 101(4):709–725.

Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., McElreath, R., Alvard, M., Barr, A., Ensminger, J., Henrich, N. S., Hill, K., Gil-White, F., Gurven, M., Marlowe, F. W., Patton, J. Q., and Tracer, D. (2005). "Economic man" in cross-cultural perspective: Behavioral experiments in 15 small-scale societies. *Behavioral and Brain Sciences*, 28(06):795–815.

Hodler, R. and Raschky, P. A. (2014). Regional favoritism. *The Quarterly Journal of Economics*.

Horowitz, D. (1985). *Ethnic Groups in Conflict*. University of California Press, Berkeley.

Houser, D. and Xiao, E. (2009). Inequality-Seeking Punishment. Working Papers 1009, George Mason University, Interdisciplinary Center for Economic Science.

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

Ilorah, R. (2009). Ethnic bias, favouritism and development in africa. *Development Southern Africa*, 26(5):695–707.

Keefer, P. and Khemani, S. (2003). Democracy, Public Expenditures,and the Poor. World Bank Policy Research Working Paper 3164.

Kitschelt, H. (2000). Linkages between citizens and politicians in democratic polities. *Comparative Political Studies*, 33(6-7):845–879.

Kiyonari, T., Tanida, S., and Yamagishi, T. (2000). Social exchange and reciprocity: confusion or a heuristic? *Evolution and human behavior*, 21(6):411–427.

Kramon, E. (2016). Electoral handouts as information: Explaining unmonitored vote buying. *World Politics*, 68(3):454–498.

Kramon, E. and Posner, D. (2010). Who is Favored? How the Outcomes One Studies Affect the Answers One Gets. Working Paper.

Levitt, S. D. and List, J. A. (2007). What do laboratory experiments measuring social preferences reveal about the real world? *Journal of Economic Perspectives*, 21(2):153–174.

Manzetti, L. and Wilson, C. J. (2007). Why do corrupt governments maintain public support? *Comparative Political Studies*, 40(8):949–970.

McBride, M. (2001). Relative-income effects on subjective well-being in the cross-section. *Journal of Economic Behavior & Organization*, 45(3):251 – 278.

Nathan, N. L. (2016). Local ethnic geography, expectations of favoritism, and voting in urban ghana. *Comparative Political Studies*, 49(14):1896–1929.

Oyugi, W. O. (2006). Coalition politics and coalition governments in africa. *Journal of Contemporary African Studies*, 24(1):53–79.

Padro i Miquel, G. (2007). The Control of Politicians in Divided Societies: The Politics of Fear. *Review of Economic Studies*, 74(4):1259–1274.

Paler, L. (2013). Keeping the public purse: An experiment in windfalls, taxes, and the incentives to restrain government. *American Political Science Review*, 107:706–725.

Posner, D. (2005). *Institutions and Ethnic Politics in Africa*. Cambridge University Press, Cambridge.

Ravallion, M. and Lokshin, M. (2005). Who Cares about Relative Deprivation? World Bank Policy Research Working Papers.

Reinikka, R. and Svensson, J. (2005). Fighting corruption to improve schooling: Evidence from a newspaper campaign in uganda. *Journal of the European Economic Association*, 3(2-3):259–267.

Shin, J. H. (2015). Voter Demands for Patronage: Evidence from Indonesia. *Journal of East Asian Studies*, 15:127–151.

Solnick, S. J. and Hemenway, D. (1998). Is more always better?: A survey on positional concerns. *Journal of Economic Behavior & Organization*, 37(3):373 – 383.

Straus, S. (2015). *Making and Unmaking Nations: The Origins and Dynamics of Genocide in Contemporary Africa*. Cornell University Press, Ithaca.

Thachil, T. and Tietelbaum, E. (2015). Ethnic Parties and Public Spending: New Theory and Evidence from India. *Comparative Political Studies*, pages 1–32.

UBOS (2014). Uganda national household survey 2012/2013. Technical report, Uganda Bureau of Statistics, Kampala.

Velasquez, M. (2013). From subject to citizen: Evidence of geographic variation in ethnic political mobilization. Unpublished manuscript.

Wantchekon, L. (2003). Clientelism and Voting Behavior:A Field Experiment in Benin. *World Politics*, 55(3):399–422.

Wantchekon, L. and Vermeersch, C. (2008). Information, social networks, and the demand for public goods: Experimental evidence from benin. In *Generating Genuine Demand for Accountability: Public Opinion and State Responsiveness*. World Bank Publications, Washington, D.C.

A Proof for Decision Rule Three

A's share of G is a function of their share of R: $\theta_G ai = \theta_R ai * \theta_R i$. Therefore if A reelects when $\mathrm{E}[\theta_G ai] > \mathrm{E}[\theta_G ac]$, then A will also reelect when

$$\theta_R ai * \mathrm{E}[\theta_R i] > \mathrm{E}[\theta_R ac] * \mathrm{E}[\theta_R c]$$

or when

$$\frac{\theta_R ai}{\mathrm{E}[\theta_R ac} > \frac{\mathrm{E}[\theta_R c]}{\mathrm{E}[\theta_R i]}$$

B Simulation details

In the simulations, I generate 50,000 versions of A, each of whom is assigned an incumbent and challenger. The total pool of resources is set at 1. The amount each leader keeps in rents is drawn from $U(0,1]$; R_x is the remainder. A is then assigned a payout drawn from $U(0,R_x]$, and B is assigned whatever remains. On average, the leader keeps 50% of the pot and A and B are each assigned 25%. A then decides whether to retain the incumbent, according to the three decision rules.

To simulate a world in which A has no priors about the incumbent, I generate E[$\theta_G ai$] and E[$\theta_G ac$] from uniform distributions. To simulate a world in which A has moderately informative priors, I first draw each $\theta_G xy$, and then draw E[$\theta_G xy$] from a normal distribution centered on $\theta_G xy$, with a standard deviation of 0.4; expected and actual $\theta_G xy$ are correlated with an R of 0.45.

C Simulated Outcomes of Alternative Models

Table 3: Alternative Decision Rules Under Full Information

	Information	Reelect if:	Correct (%) No priors	Correct (%) Priors
Alt One	D_{bi} and G	$\theta_{ai} > \mathrm{E}[\theta_{ac}]$ or $\frac{\theta_R ai}{\mathrm{E}[\theta_R ac]} > \frac{\mathrm{E}[\theta_R c]}{\mathrm{E}[\theta_R i]}$	62	70
Alt Two	D_{bi} and G	$\theta_G ai + \theta_R ai > \mathrm{E}[\theta_G ac + \theta_R ac]$	65	70

D Ethics and Debriefing

The benefit of deception is that I was able to avoid distortion of players' behavior in the game, and thus accurately capture how players respond to information. Introducing a computer had the potential to create confusion both because computers are unfamiliar to most Ugandans – 85% of respondents on Round 5 of the Uganda Afrobarometer say they "never" use a computer – and because predicting the payout from a computer "challenger" is not necessarily intuitive. Confusion, in turn,[29] increases the likelihood that players will to turn to the enumerators for help, providing an opportunity for enumerators - whether intentionally or inadvertently - to guide players to particular choices. Additionally, as shown by Houser and Xiao (2009), players tend to respond more strategically and less emotionally to computer allocators than to human ones; using a computer allocator increases the abstraction of the game and makes the results even less externally valid to actual elections. All of these substantially reduce the study's potential to improve understanding of the role of information in voters' choices.

There are two potential risks of harm from the deception. One is players' sense of frustration or embarrassment at having been manipulated. This effect is likely to be temporary. In addition, though always possible *ex ante*, it also does not seem to have occurred in this case: only three players had a negative reaction to learning they had actually been playing against a research assistant with a computer, and all three were irritated, not because of the deception, but because they thought a foreign researcher should have been able to give them a more generous payout.[30] The other risk of harm is contamination of the subject pool: players might disbelieve any future researchers sampling the same population, even if these researchers are being entirely truthful. Mukono is a popular place for research, so this is not a trivial concern. However, this risk is predicated on an assumption that players *will* trust researchers as long as they are never deceived; in Uganda, this is not the

[29]Note that the study already introduced some technology: enumerators recorded responses on smartphone and gameplay took place via SMS. However, Ugandans are far more familiar with mobile technology than they are with computers: at the time of study, the Afrobarometer indicates that 65% of Ugandans owned a mobile phone, and almost 70% used a mobile phone at least a few times per week.

[30]This reaction suggests that alerting players ahead of time that the payouts were controlled directly or indirectly by a researcher might also have caused some players to entirely reject the game's electoral mechanism for increasing their payout, choosing instead to attempt to bargain directly with the enumerator. This is consistent with some recent evidence that African lab subjects act strategically, not to increase their payout in the game, but to demonstrate real-world neediness to presumably-wealthy foreign observers (Cilliers et al., 2015).

case. Ugandans, like citizens in other countries where political dissent may be punished, are often highly suspicious of the motivations of interviewers who arrive at their homes, especially those who are asking political questions.[31] The debriefing statement that followed the deception, rather than heightening player's distress, seems to have reassured a number of players that we were in fact researchers and not political operatives. Thus, in my opinion, deception ensured the usefulness of the study without adding substantially to its costs.

The relevant question, of course, is whether I could have controlled game play, and accurately measured the effect of the information treatment, without involving computers or using deception at all. I believe I could not have done so. Other researchers running laboratory games in African countries have developed techniques that are intended to accomplish this, but I find these techniques unsatisfactory. One option is to recruit subjects 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, players would be randomly assigned to receive one of these retained, real games. Another option is to assign Ugandan confederates to play the leader and other citizen in every game, and have them play in real time as dictated by a list of pre-randomized games. However, both of these strategies, while not deceptive in the narrowest sense of the term, still require misleading players about the nature of gameplay and manipulating the game without players' fully informed consent to be manipulated. In the absence of a debriefing acknowledging this manipulation, these techniques are arguably less ethical than outright deception combined with debriefing. Thus I chose the latter approach. The scripts for recruitment, the game, and the debriefing follow:

D.1 Game instruction and debriefing scripts

The recruitment script, delivered in the player's preferred language, read:

> *To whomever answers door:* Hello *(or other appropriate greeting).* My name is .... I am working for a university researcher who is conducting a study. I am looking for someone in this household who is *(gender)* and between the ages of *(age)* and *(age).* Is there

[31] In this study in particular, one player refused to take her winnings because she believed it might be a bribe

someone here like that?

To appropriate player: Hello *(or other appropriate greeting).* My name is ... I am working for a university researcher who is conducting a study on the conditions in this community, and how well politicians are serving you here. The survey will take about 45 minutes. Are you interested in talking to me about this today? *If yes:* Good. Let me just check that you are the person we are looking for: Can you tell me how old you are? What language would you like to use for the survey?

If person is of correct age and gender, and language is one enumerator speaks: Excellent. Let me tell you more about this. I will ask you some questions about your community, and about your political opinions. Then you will have the opportunity to earn a small amount of money by playing a game that will help us learn what you want from your leaders. The information you tell me will be sent by SMS to a professor in the United States. Then it will be erased from my phone. I will have no record of your participation or your answers. I will not record your name at any time. About 2000 other people, from other parts of Uganda, will also be interviewed.[32]

If you agree to the survey, but then find out there are questions you don't want to answer, you don't have to answer them. You will still be allowed to play the game and earn money. Do you have any questions? Are you still interested in participating?

The instructions of the game, delivered in the player'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

[32]This number includes subjects enrolled in other parts of the same study.

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.

Sample game scripts read:

No information: The payout is 700 shillings. The leader is playing from Soroti. Would you like to reelect this leader or get a new leader?

Low information: The payout is 700 shillings. The leader is playing from Soroti. The other player received 1300. Would you like to reelect this leader or get a new leader?

High information: The payout is 700 shillings. The leader is playing from Soroti. The total pot is 2600. The other player received 1300. Would you like to reelect this leader or get a new leader?

The debriefing script, delivered in the player'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 player who we will interview later in the project. Do you have any questions about this?"

E Sample characteristics

Table 4: Sample vs. Round 5 Afrobarometer

	Sample	Afrobarometer	
		(Uganda)	(Full)
Mean age	37	35	37
% female	49	50	50
Median education	8	7	7
% Urban	33	14	40
Food secure	67	44	50

Compares characteristics of study sample to characteristics of the sample on the Round 5 Uganda Afrobarometer. "Food secure" respondents report that they are "never" without sufficient food for their household.

F Balance

Figure 6: Covariate balance

(a) Assigned to be favored

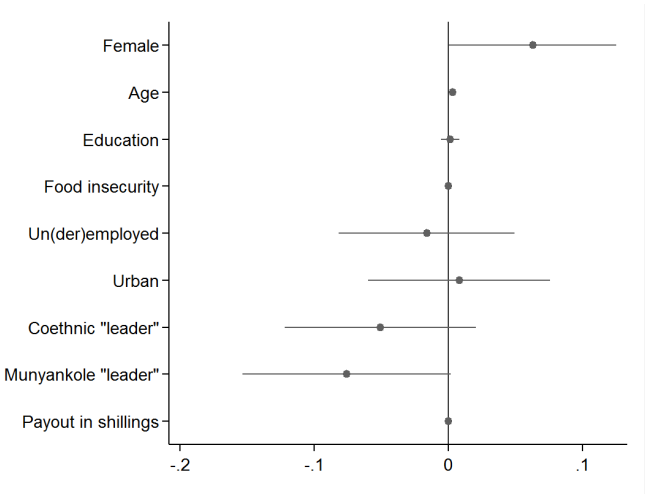

(b) Assigned to partial information treatment

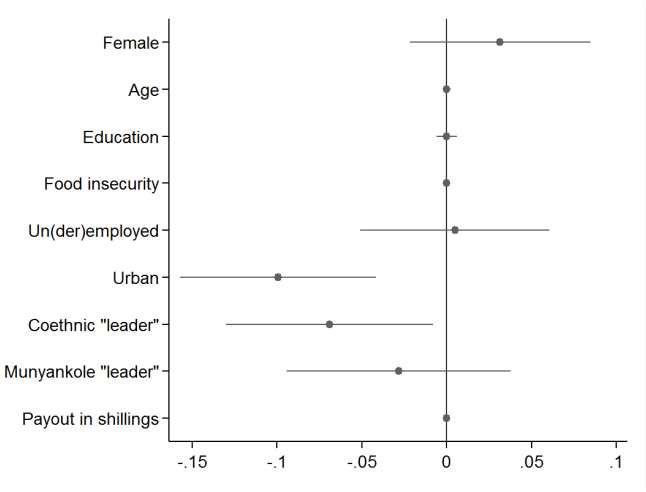

Standardized coefficients of linear model regressing treatment assignment on subject characteristics, with 90% confidence intervals.

Table 5: Effect of Relative and Absolute Distribution on Retention Under Full Information

	(1)	(2)
Share of distribution ($\theta_R ai$)	0.112	
	(0.126)	
Total share of pot ($\theta_G ai$)	2.438***	
	(0.273)	
Favored		-0.007
		(0.062)
More than 25% of pot		0.549***
		(0.065)
N	167	167

Standard errors in parentheses

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

OLS Model. Dependent variable takes a one if player retained the leader. The sample includes only players with full information about both relative and absolute distribution. Controls not shown include the assigned ethnicity of the leader as well as respondent gender, age, education, un(der)employment, food security and urban location.

Table 6: Effect of Ethnicity on Reported Economic Outcomes, Afrobarometer

	(1) Well-being vs. last year	(2) Well-being in one year	(3) Lived poverty index	(4) Treated unfairly by gov't
Coethnic w/ incumbent	0.112***	0.035*	-0.050***	-0.304***
	(0.013)	(0.015)	(0.010)	(0.012)
Education	0.026***	0.025***	-0.066***	0.007*
	(0.003)	(0.004)	(0.002)	(0.003)
Rural	-0.022+	0.001	0.173***	0.022+
	(0.012)	(0.014)	(0.010)	(0.012)
Female	-0.007	0.037**	-0.055***	-0.054***
	(0.011)	(0.012)	(0.009)	(0.010)
News exposure	0.016***	0.027***	-0.031***	-0.008*
	(0.004)	(0.005)	(0.003)	(0.004)
Computer access	0.131***	0.096***	-0.203***	0.014
	(0.025)	(0.027)	(0.019)	(0.023)
Strong national identity	-0.005	-0.031***	-0.010**	-0.110***
	(0.005)	(0.006)	(0.004)	(0.005)
Muslim	0.067***	0.138***	-0.065***	-0.052***
	(0.017)	(0.019)	(0.013)	(0.015)
Gov't interviewer	-0.016	0.020	0.009	-0.006
	(0.012)	(0.014)	(0.009)	(0.011)
Bystanders present	-0.013	-0.003	0.101***	0.043***
	(0.013)	(0.014)	(0.010)	(0.011)
Constant	2.862***	3.557***	0.367***	1.149***
	(0.055)	(0.066)	(0.053)	(0.057)
N	32450	29958	32296	31518

Standard errors in parentheses

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Models 1-3 are regressions; Model 4 uses logit. All models include country fixed effects. Data is from Round 5 of the Afrobarometer public opinion survey.