

Owning It: Accountability and Citizens' Ownership over Aid, Oil, and Taxes*

Brandon de la Cuesta[†]

Lucy Martin[‡]

Helen V. Milner[§]

Daniel L. Nielson[¶]

September 30, 2017

Abstract

Why do citizens in developing countries often choose not to punish poor use of government revenues, electorally and otherwise? Existing work suggests that low expectations drive low accountability: if citizens expect poor performance, their willingness to punish poor performance is low. Yet relatively little is known about how citizens form expectations. Drawing on recent work in cognitive psychology, we argue that citizens' feelings of ownership over collectively-owned resources is a key driver of expectations, and that low ownership will make citizens less willing to pay the (often high) costs necessary to sanction poor performance. Using a set of lab-in-the-field experiments in Ghana and Uganda, we demonstrate that ownership—defined as the extent to which citizens feel entitled to welfare gains realized by government spending—is a strong and significant predictor of citizens' willingness to punish poor performance. We then show that ownership is the mechanism behind the well-known accountability effect, in which citizens demand greater benefits from tax-based revenues than from those derived from non-tax sources such as foreign aid and oil. Using a novel measurement instrument that tracks subjects' utility over time, we provide evidence that the ownership effect works by raising citizens' expectations, causing them to suffer higher disutility when governments underperform. Finally, we show not only that ownership is an important driver of accountability pressures, but that it can be manipulated experimentally, and that increasing ownership over aid and oil can produce accountability pressures equal to those generated by taxation.

Key words: resource curse, endowment effect, field experiment, accountability, taxation

Preliminary Draft: Please do not cite or circulate without permission

*We thank Brendan Cooley, James Gilman, Delanyo Kpo, Gaétan Nandong, and Elsa Voytas for their excellent research assistance and efforts in developing and implementing the experimental protocols for this project.

[†]PhD Candidate, Department of Politics, Princeton University, Princeton NJ 08544. Email: brandon.delacuesta@princeton.edu

[‡]Assistant Professor, University North Carolina, Chapel Hill, Email: lucy.martin@unc.edu

[§]Professor of Politics and International Affairs, Princeton University, Princeton NJ 08544. Phone: 609-258-0181, Email: hmilner@princeton.edu

[¶]Professor, Brigham Young University, Provo, Utah 84602, Email: dan.nielson.byu@gmail.com

1 Introduction

In many developing countries, citizens often fail to sanction poor performance and even outright corruption by government officials.¹ More generally, there is increasing evidence that it is difficult to improve bottom-up accountability pressures: most interventions evaluated to date have either had no effect, or seem to improve accountability only under certain conditions.² Understanding when and why citizens are willing to demand accountability—through either electoral behavior or protest and other forms of activism—is therefore a critical area of research. This is especially true for countries whose budgets depend on natural resources or foreign aid, both of which are thought to undermine accountability pressures even further.³

Recent research on citizens’ willingness to take costly political action suggests that absolute levels of performance matter less than performance relative to expectations (Gottlieb 2016): citizens are most likely to punish when their expectations are high but not being met. Yet, in many cases, expectations are extremely low; citizens believe that government either cannot or will not improve, leading to few complaints when these low expectations are confirmed through corruption scandals or poor public services. However, Gottlieb (2016) also shows that improving citizens’ perception of government capacity can increase citizens’ expectations, and De Kadt and Lieberman (2016) suggest that citizens adjust expectations over time in response to existing levels of service provision. This is encouraging, as the result suggests that expectations may be malleable. Nonetheless, we still lack a strong understanding of how citizens form expectations—whether high or low—in the first place. This is especially important for understanding why citizens fail to hold governments accountable even when elections are present and citizens have at least basic information about government performance.

Previous work on expectations has focused on citizens’ understanding of government capacity, the size of the budget, and other factual aspects of governance (De Kadt and Lieberman 2016; Gottlieb 2016). This paper introduces and tests a new, psychological mechanism that can explain variation in expectations and therefore in accountability pressures regarding a core function of any government: how its budget is allocated and spent. We draw on two key findings from psychology that explain individuals’ willingness to punish how a resource is divided. First, there is substantial evidence that individuals get psychological benefits from punishing behavior that deviates from expectations; this includes punishing someone perceived as having divided resources unfairly (Fehr and Gächter 2000; Fehr and Schmidt 1999). Second, evidence suggests that individuals’ actual or perceived ownership over the resource being divided drives punishment significantly; those with formal ownership over a resource expect to receive more from it (Wu et al. 2012; Shu and Peck 2011).

However, in most psychology research, ownership over money or goods, such as coffee mugs, is clearly assigned to a particular individual. But government budgets from aid, oil, and taxes are not coffee mugs, and ownership over large and diverse sources of revenue may activate

¹Citizens can sanction government leaders by punishing incumbents during elections, or by using protests and other public forums to express dissatisfaction and impose costs on leaders in between elections. Urban protests are especially important in non-electoral regimes (Bates 1981).

²For example, both Olken (2007) and Casey, Glennerster and Miguel (2012) find that interventions to increase citizen monitoring had no impact, and while information interventions are sometimes effective (Reinikka and Svensson 2005), in many cases it has no effect on citizen behavior (see e.g Chowdhury et al (2017)).

³This paper focuses on when citizens will punish low public goods provision, high corruption, or other policies that reduce the extent to which citizens benefit from the government’s budget. This is in line with the expressed priorities of citizens in many sub-Saharan African countries, where public goods are the main policy priority.

different psychological mechanisms than ownership over small consumer goods. And beliefs about the purpose of and claim to government funds may also matter: while a rational citizen might believe that the budget “belongs” to citizens and should be used for their benefit, the same citizen may also believe in prebendalism: that budgets belong not to citizens but to the government officials who control those resources (Van de Walle 2001). It is thus not clear whether ownership extends to publicly-owned resources, or how that ownership might affect citizens’ willingness to pay the sanctioning and monitoring costs that yield good governance.

We argue here that the degree to which citizens feel psychological ownership over common-pool resources like the budget revenues from aid, oil, or taxes drives accountability demands, and that it does so by making subjects more willing to pay the costs of sanctioning. We use two separate laboratory experiments, conducted in Uganda and Ghana, to demonstrate that perceived ownership over a group fund significantly predicts subjects’ willingness to pay the costs of sanctioning a leader for how he or she allocates a budget.⁴ We then demonstrate that ownership is systematically different depending on the type of revenue at stake: while taxation induces higher ownership over the budget, windfall revenues such as aid and oil have substantially lower ownership. We demonstrate that it is these differences in ownership that drives the finding, replicated here and shown elsewhere (see e.g. Martin 2014; Paler 2013), that taxation increases accountability pressures. We show, in other words, that ownership is the causal mechanism by which taxation induces greater accountability.

However, while ownership over tax revenues is higher on average compared to aid and oil, we also find substantial variation in citizens’ feelings of ownership, and in willingness to punish, within each revenue source. This suggests that ownership is not necessarily fixed. We run additional experiments in Uganda to demonstrate that ownership is indeed malleable: treatments that assign citizens indirect ownership over portions of aid or oil revenues significantly increase both ownership and punishment. Our work thus suggests that poor governance driven by low ownership over government revenue—a key theoretical explanation for the resource curse—can be at least partially attenuated if citizens’ feelings of ownership over non-tax revenues such as aid and oil can be increased. This is especially important in sub-Saharan Africa, where a number of countries, including Uganda, have recently discovered oil and where many countries rely on aid for substantial portions of budgets. If ownership is most malleable in the years leading up to and immediately preceding the discovery of natural resources or the influx of budget support in the form of foreign aid, then efforts to foster high ownership over these non-tax revenues early on may be critical in determining whether the resource curse and its aid corollary, and the poor government performance that allegedly attends them, takes hold.

2 Background and Theory

2.1 Existing Research

When should we expect citizens to punish poor performance—in particular poor service provision or low central transfers—by leaders? For citizens to demand the ability to punish leaders, a number of conditions need to be met (Lieberman, Posner and Tsai 2014). Citizens must first receive information about government performance. They must then decide if they are satisfied with that performance; if not, they must work with other citizens to coordinate on some action

⁴We do not argue that ownership is the sole predictor of punishment: other factors could include societal fairness norms, the cost of punishment, and other individual-level factors.

to sanction officials. There must therefore exist some strategy set available to citizens who wish to sanction. This suggests that there is a set of conditions that are necessary and sufficient for citizens to successfully demand accountability from their leaders: if any single step breaks down, successful sanctions may not occur, even in the presence of high corruption or poor performance more generally.

We focus here on a particular aspect of accountability: if citizens receive information that government performance is low, what determines when they will be sufficiently upset to sanction leaders? The chain of causation that runs from dissatisfaction to political action is long and complex, and existing work suggests that it is often the case that this chain is broken in many places, particularly in the developing world (Lieberman, Posner and Tsai 2014). Yet dissatisfaction—either with the status quo or some proposed change—is the necessary if not sufficient condition for political action. In developed democracies, this fact is belied by parties’ preoccupation with the approval ratings of their leaders; in autocracies, it is belied by elites’ fear of urban uprisings and other forms of popular discontent. But enacting sanctions against leaders is costly: citizens must pay transport costs or forgo economic activity in order to vote or protest, and in some cases may face the possibility of repression or coercion. Citizens will therefore only enact sanctions when they believe that they will receive a sufficient benefit in return, or when they are sufficiently upset that political action, even when it is unsuccessful, carries with it the expressive benefits that come from voicing one’s dissatisfaction.

We follow existing voting and accountability models in assuming that this will result in threshold strategies by citizens: individuals will take action against leaders if and only if performance falls below a certain threshold. In the case of budget allocations, this could entail punishing when transfers to citizens fall below a certain point, or when the fraction of the budget leaders extract as rents becomes too high. This structure has the advantage of mapping closely onto our experiments, in which citizens explicitly form a threshold of this kind. If citizens are strictly economically rational, they should only enact sanctions on leaders when they believe that it will result in improved government policy in the future, perhaps through higher transfers or better public services; this is in line with many voting models (see e.g. Fearon 1999). However, in most cases the likelihood that an individual citizen’s participation is pivotal to who wins an election, or whether a protest is successful, is essentially zero: the citizen should expect to reap the same benefits regardless of his personal participation. This is central to the collective action problem: coordination is easier to achieve when citizens receive excludable, private benefits from taking part (Olson 2009).

A large body of research in psychology suggests that the private benefits of taking action may in part be expressive rather than economically rational: individuals are willing to engage in costly punishment behavior even when there are no economic benefits to doing so (Fehr and Gächter 2000; Fehr and Schmidt 1999). This suggests that understanding citizens’ demands for accountability requires understanding the psychological mechanisms underlying these demands. In particular, research has shown that individuals receive substantial disutility when they observe behavior that they perceive as unfair, especially towards themselves, and that this disutility is increasing in the size of the perceived norms violation (Fehr and Gächter 2000). Punishment appears to generate expressive benefits that relieve those negative emotions. Expressive benefits have been widely documented in psychology experiments, but also with regards to political behavior.⁵

This psychological mechanism is especially important in authoritarian and quasi-authoritarian countries with weak formal institutions. In such places, citizens will have a limited ability to en-

⁵See for example the literature on expressive voting (Fiorina 1976; Kan and Yang 2001; Tyran 2004)

gage in the kind of tax bargaining that is thought to drive the relatively higher accountability for tax-based revenues observed at the macro-level. They may also have limited voting rights, making it difficult to sanction officials through their voting behavior. It is in these places where the kinds of bottom-up political action emerging out of widespread discontentment—particularly protests, which tend to be located in urban areas and are highly visible signs of dissatisfaction—are likely to be most dangerous for elites. Critically, such countries continue to receive large transfers of foreign aid and many of the recent (and largest) resource discoveries are in countries of this kind.⁶ As such, understanding the role of expressive benefits in driving citizens’ thresholds, and in turn their willingness to punish, is likely to be the most valuable in places that are also the most likely to fall prey to the negative consequences of non-tax revenues such as the resource curse and its aid corollary.

When will citizens be sufficiently upset at government officials’ actions to be willing to pay the (sometimes substantial) cost of sanctioning, whether that cost be associated with voting, protesting, contacting politicians, or organizing against a government policy? Recent research suggests that absolute levels of government performance matter less than performance relative to expectations (James 2007; Gottlieb 2016). A citizen who observes poor public goods provision but never expected her local government to do better may not be willing to demand more, while a citizen who has seen that other local governments do much better may have higher expectations and be more willing to punish. This dynamic is in line with models of reference-dependent utility, in which individuals evaluate their wellbeing relative to some reference point. A similar process may apply to citizens’ demands; those with higher expectations (i.e. reference points) will be willing to demand policies or spending that accords more closely with their preferences. However, expectations and thresholds are still distinct: we argue that thresholds are determined by a number of factors, including expectations but also the costs of taking action and other personal and institutional variables.

3 An Ownership-Based Theory

What, then, determines how citizens set their expectations? We focus in this paper on one form of expectations, namely the degree to which citizens expect to benefit from the government’s budget or resources.⁷ Research in cognitive psychology has shown that “ownership” has a significant impact on individuals’ expectations and willingness to punish unfair divisions of a resource (Wu et al. 2012). Even an act as simple as randomly selecting a research subject to receive a coffee mug is strong enough to induce ownership and affect what individuals perceive as a fair price (Kahneman, Knetsch and Thaler 1990). Ownership is stronger for earned income, and it appears to become stronger the longer a good or resource has been owned (Shu and Peck 2011). Given a set division of resource, individuals are more willing to reject an offer or punish the sender when they had initial ownership over that resource (Wu et al. 2012). Existing evidence suggests that, while ownership is in part responsible for creating endowment effects or loss aversion, it is a distinct and causally prior concept (Morewedge et al. 2009; Shu and Peck 2011). Thus, we should expect ownership over a resource to induce higher expectations and, subsequently, higher punishment.

⁶For example, in the past 10 years there have been significant oil deposits found in East Africa, including Uganda, Kenya, Tanzania, and Mozambique.

⁷In addition to the ownership theory suggested here, this could depend in part on citizens’ beliefs regarding the size of the budget, or the cost of providing services, as Gottlieb shows.

Applying this to politics, research on ownership suggests that citizens should have higher expectations, and be more willing to demand accountability from leaders, when they feel ownership over the government’s budget. However, it is not obvious when this ownership should occur. Existing research in psychology has typically focused on goods for which ownership is clearly established, and assigned to a particular individual. Government budgets are distinct in that ownership is at best joint, never belonging to a particular individual.

Existing work on accountability suggests two competing views. In the first view, the structure of democracy assumes that the budget belongs to the citizens, in the sense that it should be used for their benefit. However, even this does not answer questions regarding which citizens should benefit, or how these benefits should be distributed. Indeed, the entire literature on redistribution and its accompanying debates can be understood in part as grappling with whom should benefit and by how much.

The second view proposes that it is not obvious that citizens should feel ownership over the budget at all, especially in countries where patronage and clientelism are rife. Van de Walle (2001) argues that many countries, particularly in Africa, operate via prebendalism, where resources are “owned” by the government officials who control them. In such cases, there is no assumption that citizens should benefit at all, especially if they are not clients or co-ethnics of the relevant officials (Ekeh 1975; Van de Walle 2001). Depending on which view predominates, it therefore seems likely that there is substantial variation in the degree to which citizens feel ownership over government resources, and this may drive some of the variation in demands for accountability.⁸

We argue that the source of government revenues is a key determinant of ownership and therefore of citizens’ demands for accountability. In general, governments that rely on taxation for funding are less corrupt, are more likely to democratize, and provide higher public services (Ross 2004; Timmons 2005; Baskaran and Bigsten 2013; Prichard 2015; Fisman and Gatti 2002; Brollo et al. 2013; Gadenne 2015). In contrast, funds from oil and aid are believed to be “windfall revenues” that enable corruption, undermine governance, foster repression, prolong autocratic rule, and increase conflict (Bräutigam and Knack 2004; Djankov et al. 2008; Caselli and Cunningham 2009; Morrison 2009, 2015; Smith 2008). Elites benefit from windfalls since they now have access to more funds that can be diverted to corruption and clientelism relatively free from citizen scrutiny. Citizens, on the other hand, suffer from windfalls, as money from non-tax sources is said to undermine accountability, generate resources for repression, and buy citizen quiescence toward bad governance (Mahdavy 1970; Beblawi and Luciani 1987; Chaudhry 1998; Waterbury 1998).

While there are a number of proposed mechanisms for this effect, a frequent and recurring argument holds that citizens who are taxed demand more from leaders, and are more willing to enact costly sanctions on non-accountable leaders (Paler 2013; Martin 2014).⁹ The proposed mechanism for this effect is loss aversion: citizens expect to receive their pre-tax income, and taxation pushes them below their expectation, making them more sensitive to government corruption and more willing to punish (Sandbu 2012; Paler 2013; Martin 2014).

⁸Qualitative evidence from our experiments, below, suggests support for both views: citizens who make high demands on leaders often justify their behavior by saying that government money should benefit citizens, while those with low demands often explain that government officials have many ways to use their money, and citizens should not demand very much.

⁹Other proposed mechanisms include tax bargaining, in which citizens use the threat of withholding tax revenues in order to extract concessions from the state, and informational theories, in which taxation either conveys more information about the size of the budget, or makes citizens more willing to pay for information about government spending (Schumpeter 1991; Bates and Lien 1985; Levi 1989; North and Weingast 1989; Gadenne 2015; Paler 2013; Prichard 2015).

The ownership theory provides a counterpoint to this. We do not deny that taxation induces loss in citizens, and that this is part of why citizens who are taxed are more willing to pay sanctioning costs when taxed. Our argument is that ownership is causally prior to the economic or psychological disutility that citizens feel when governments use revenues poorly. Citizens certainly get disutility when governments underperform, but *how much* disutility they feel is a function of their expectations. Aversion to loss and the well-known S-shaped utility function it produces may explain why people punish underprovision more than they reward equivalently good governance, but it can say nothing about how citizens decide what constitutes "good" or "bad" levels of provision in the first place. We argue here that the extent to which citizens feel ownership over government budgets drives performance expectations, and that these expectations determine how much of the cost of underprovision is internalized by citizens.

The loss story is thus incomplete: taxation does not simply remove money from the citizen, but directly transfers it to the state. This feature of taxation is part of what separates it from other general forms of loss, such as negative economic shocks or inflation. It also provides a clear basis through which taxed citizens can claim direct ownership over at least part of the government's budget. In contrast, aid and oil revenues lack a clear ownership mechanism: while citizens may feel ownership over money the government receives from foreign aid or oil production, there are also reasons to expect ownership to be low. Non-tax sources are never held nor earned by citizens, and they accrue (in the case of oil) or are given (as is the case with aid) directly to the government. For this reason, we expect to find more individual-level variation in ownership over non-earned revenues.

While norms could emerge that windfalls such as aid or oil should be used for the benefit of citizens (think for example of Alaska), such a norm may not emerge automatically or at all. Yet the emergence of such norms is especially important in countries such as Uganda where oil has only recently been discovered, or in Ghana, where there has been substantial public debate over the potential welfare gains from oil-based revenues. In such settings norms of ownership may be more malleable. Research in psychology suggests that relatively small cues can send strong signals about ownership in ways that affect punishment [Kahneman, Knetsch and Thaler \(1990\)](#). While these studies deal with ownership over small consumer items such as coffee cups, it is less clear that ownership over collectively owned resources such as aid and oil can be similarly manipulated.

3.1 Testing for the Ownership Effect

The ownership theory suggests a number of testable hypotheses. Our first set of hypotheses tests for the existence of the ownership effect and makes predictions about its causal relationship between revenue source, ownership, and punishment. Previous evidence that taxation increases accountability has examined citizens' demands when taxed, relative to a case where the budget is the same but is derived from an unspecified windfall ([Martin 2014](#)). It seems reasonable that citizens will have low ownership over an unspecified fund. However, as discussed above, we know little about citizens' sense of ownership over foreign aid or natural resources. Other work has failed to find that taxation increases willingness to punish relative to oil ([Paler 2013](#)), or that it increases willingness to monitor ([de la Cuesta et al. 2017](#)). This could occur if, for example, citizens feel strong ownership over oil as a resource extracted from citizens' own land. Likewise, foreign aid is explicitly meant to benefit citizens, and so citizens could in theory feel strong ownership over it. In the experiments below, we include three versions of the "windfall" condition: an unspecified grant; foreign aid; and oil. If the theory presented above is correct, we expect the following hypotheses to hold:

Hypothesis 1 *A greater sense of ownership over the government budget should increase citizens' willingness to punish, regardless of the source of the government budget.*

Hypothesis 2 *Taxation increases willingness to punish by increasing ownership over government budgets.*

3.2 Uncovering the Mechanism Behind the Ownership Effect

While testing the above hypotheses can establish the existence of an ownership effect, our theory also makes clear predictions about the underlying mechanism that connects ownership with an increased willingness to punish poor performance. We argue that ownership affects punishment by raising citizens' expectations of what they should receive from the government's budget. In particular, citizens should experience disutility (losses) when they receive less than they expect, and this disutility will drive higher punishment. By moving citizens' expectations of the "acceptable" level of provision, ownership increases the potential losses for a given amount underprovision. As disutility rises, citizens receive a greater expressive benefit to punishment, making them more willing to pay the costs to do so. In the experiments, we implement a novel procedure for measuring citizens' utility at different parts of the experiment, allowing us to test the following hypothesis about the underlying mechanism:

Hypothesis 3 *A greater sense of ownership increases citizens' dissatisfaction with poor performance; this increases the minimum acceptable level of public goods provision below which citizens punish, increasing accountability pressures.*

Finally, we argue above that ownership may be malleable. Research on ownership in psychology has examined two methods of assigning ownership: physically giving individuals money or a good, and assigning indirect ownership through the design of the intervention. We focus on the latter option, using a treatment in the experiments below that simply tells individuals that a portion of the group fund of aid or oil money is theirs. We expect this to increase citizens' sense of ownership over the budget, and through that willingness to punish.

Hypothesis 4 *Increasing ownership over aid or oil revenues should also increase willingness to punish, eliminating differences with the levels observed under taxation.*

4 Revenue and Accountability in Ghana and Uganda

Ghana and Uganda are particularly apt locations to test the effects of different revenue sources on accountability pressures. Both countries rely on a mix of domestic taxes, foreign aid, and oil revenues, providing the opportunity to study the effects of revenue on accountability pressures. However, there are also significant differences. In Ghana oil has been a significant part of the budget since 2010, while in Uganda oil was discovered in 2011 and the government has only now started to receive oil revenues. Uganda is also much more aid-dependent than Ghana, and has a lower

tax-to-GDP ratio. This allows our experiments to test how revenue source affects accountability pressures in both aid-dependent and oil-dependent settings.

Variation in economic and political development also makes our country sample more comparable to the rest of sub-Saharan Africa. While Ghana is significantly wealthier than Uganda—Ghana’s GDP per capita is \$US 3,784 compared to Uganda’s \$1,634—both countries have development indicators that are at or near the mean for the continent (World Bank 2016). Politically, Ghana is a stable, multi-party democracy with a strong record on political freedoms. In contrast, Uganda is a semi-democracy in which the current regime has been in power since 1986, and political opposition is increasingly harassed and shut out from formal electoral competition.

Our data come from urban samples in both Uganda and Ghana. Because we expect that exposure to real-world taxation should create a strong sense of ownership over the use of tax-based revenues in the experimental context, we sample respondents from urban areas where exposure to taxation is higher than in rural areas. However, urban residents are also very aware of the non-earned revenues their governments receive from aid and oil, ensuring that both tax and non-tax sources will be salient to respondents, and, critically, that aid and oil are of sufficient importance that their use could have a significant impact on citizens’ utility. Both countries have seen intense public debate over the use of oil-based revenues, and Uganda has seen significant impacts of the misuse of aid moneys over the past five years.

There is also evidence on previous linkages between taxation and accountability in both Ghana and Uganda. Prichard (2015) discusses in detail the evolution of taxation in Ghana, including the explicit bargaining between taxpayers and citizens that has occurred over the past 20 years. This bargaining has frequently been catalyzed by popular protests among citizens, pointing to citizens’ willingness to engage in costly political action to demand accountability; this suggests that we might expect to find the types of expressive benefits from punishment that we study here. In Uganda, the national government has eliminated several forms of direct taxation—including head taxes, school fees, property taxes, and health center fees—in the run-up to elections, and many observers argue that this has led to lower accountability pressures from citizens (Persson and Rothstein 2015; Martin 2017).

5 Intervention Description

To test the hypotheses outlined in Section 2, we designed a set of laboratory experiments, conducted in Ghana and Uganda inspired by Martin (2014). Martin’s original experiments consisted of the “Tax” and “Grant” games, each played between one “Citizen” and one “Leader.” In each game the Leader had to allocate a group fund between his own salary and the Citizen; the citizen chose whether to pay a small cost to enact a fine on the Leader. The games differ only in the source of the group fund. We keep this basic structure for the experiments introduced below, and we retain many of the rules and constraints Martin used. These include the notion that taxes are exogenous and mandatory, preventing bargaining between leaders and citizens. Additionally, government budgets are constant and observable across treatments.

To test our hypotheses regarding the accountability effects of different types of revenues, we designed a set of four games, the steps on which are summarized in Table 1: the Tax, Grant, Aid, and Oil games. The steps for the Tax and Unspecified Grant versions are identical to those in Martin (2014). We then add two additional revenue source treatments: Aid and Oil. The basic

Stage	Tax Game	Windfall game		
		Unspecified Grant	Aid	Oil
1	The citizen is given a wage of 10 MU.	The citizen is given a wage of 5 MU.		
2	The citizen is taxed 5 MU. This is doubled to 10 MU and given to the leader as the group fund.	The leader is given 10 MU as the group fund.		
3	The Leader allocates 10 MU between himself and the Citizen.			
4	The Citizen observes the Leader’s decision and, based on the decision rule they specified, decides whether to pay 1 MU to have enumerators remove 4 MU from the Leader.			

Table 1: Timing of Original Tax and Grant Game. The game is identical to that of Martin (2014) except that in addition to the unspecified grant, players may be randomized into a version of the game in which the source of the grant may be either aid or oil.

steps for these games are the same as the Grant game, with one key difference. Whereas in the Grant game the source of the group fund is not specified, in the Aid and Oil games respondents are told either that the group fund is money that was given by a donor as foreign aid, or that the money comes from Ghanaian or Ugandan oil revenues. For enumeration purposes, 1 money unit (MU) was set equal to 0.50 Ghanaian Cedis or 100 Uganda Shillings (UGX).¹⁰ All enumeration employed real coins to better convey the decisions to respondents.

The source treatments were built into the game scripts used by the enumerators as well as illustrated on the game boards. During both participant training and actual gameplay enumerators stated the revenue source each time the group fund was mentioned. In order to emphasize the treatment, enumerators placed the coins representing the group fund on a tile illustrating and the source, and verbally stated the source, before moving the group fund to the leader’s tile. The game board for the Oil condition is given in Figure 3 in the appendix; game boards for the remaining conditions differ only in the image on the source tile.

Table 1 describes the citizen as making a punishment decision after the Leader allocates the group fund. For implementation purposes, Citizens were instead asked to make an ex ante decision rule; they were asked to decide which possible allocations of the group fund they would punish. This substantially increased the experiment’s power. For example, in the two-player games, enumerators would start by asking the Citizen “If the Leader kept 10 MU, and gave you 0 MU, would you pay 1 MU to punish the leader?” If the Citizen replied “yes,” the enumerator would keep asking for different allocations, increasing the share the Citizen receives in 1 MU increments.¹¹ Enumerators stopped when they received a transfer level at which the Citizen would no longer punish: this becomes the punishment threshold in the analysis below. All games were implemented using real Ghanaian or Ugandan coins to make the decisions concrete for respondents.

Because this is a single-shot game, punishment strictly decreases the citizen’s economic utility in all versions of the game. When the citizen receives no expressive benefit from punishment, the unique subgame-perfect Nash equilibrium is for the Leader to offer 0 MU to the Citizen, and for the Citizen to never sanction the leader for any transfer size. Thus, punishment is an expressive action rather than economically rational. Following Martin (2014), we refer to the Leader’s salary as the group fund to signal that Citizens have some degree of discretion over its disbursement. Note that no one gets the money that is taken away in punishment; this was repeatedly stressed to all

¹⁰At the time of data collection, exchange rates were one US dollar to 4.2 GHC or 3,500 UGX.

¹¹i.e. the next step would be to ask “If the Leader kept 9 MU, and passed you 1 MU, would you pay to punish?”

respondents.

5.1 Implementation

Slightly different sampling techniques were used for the Ghana and Uganda samples. In both countries, volunteers were recruited for three enumeration sessions per day, each consisting of 16 respondents. The Ghana experiments were conducted in Accra in June and July 2016, using subjects recruited from 8 constituencies in the Greater Accra region. Chosen constituencies fell into a “low” or “medium/high” income category.¹² In Uganda, we ran the experiments in January and February 2017 using a convenience sample from Kampala. Instead of a central enumeration site as used in Ghana, we rented 8 field sites in 2 separate areas of the city and recruited volunteers from the neighborhoods surrounding each site. In both countries, each session was randomly assigned to one of the possible treatment conditions; in Ghana assignment was blocked on the average income level in the constituency, and in Uganda it was blocked on enumeration site. Table 2 summarizes the eight treatments, and the number of total subjects assigned to each. Note that the number of observations in each cell varies due to the need to balance treatment assignment across blocks.

Sessions were similarly structured in each country sample. At the beginning of each session, subjects were given a short group training, lasting approximately 10 minutes, which laid out the basic rules of the assigned game. After group training, enumerators then administered a short on-on-one training with each subject, explaining a sample round of the game and probing subjects on their comprehension of the key game steps, particularly the allocation decision. After one-on-one training was complete, subjects were then sent back to the group training room and called up one at a time to complete five single-shot rounds of the game.¹³ At the start of rounds 2-5, respondents were told what had happened in the previous round, but were not told the decisions of any other respondents. Subjects were instructed not to speak about the game between rounds and were monitored at all times by project staff to ensure this rule was followed. At the end of the five rounds, respondents completed an outtake survey. They were then paid a show-up fee, plus their earnings from all five rounds.

Within each game session, we randomly assigned subjects to the role of citizen or leader at a ratio of 3 citizens per leader. In the first round, each citizen was randomly assigned to a play with a leader. During the game, each citizen received the transfer decided by the leader to whom he or she was assigned. Because each leader played with multiple citizens, one was selected to serve as the leader’s pair, and his or her punishment threshold determined whether or not the leader was punished. In each subsequent round, the subjects’ roles remained the same, but citizen-leader pairs were re-randomized. Our randomization algorithm took an arbitrary n subjects, k leaders, and l rounds as arguments and returned a series of pairings that satisfy the above criteria. To stress that each round was a single-shot game, in between rounds enumerators stressed to respondents that the pairings would be different than in the previous round.

¹²We avoid high-income enclaves of the city because of low recruitment rates during piloting, due both to the difficulty of finding high-income individuals at home during the day and also because the compensation offered was far less attractive to high-income individuals.

¹³Respondents in Uganda also played a practice round, that was not payout-relevant, before the five rounds began.

Country	Tax	Aid	Oil	Grant
Uganda	143	142	135	143
Ghana	228	215	217	204

Table 2: Number of respondent observations in each of the eight treatment conditions. The randomization algorithm used gave optimal balance within each block given design constraints. Cell sizes are the final number of subject-round observations included in the dataset.

6 Data

The resulting dataset used in the analysis below consists of 1680 respondents across the two countries.¹⁴ Table 2 reports distribution of subjects across our two-dimensional design. Small discrepancies in cell size are driven by missingness due to attrition in recruiting or subjects leaving the session early. In the Ghana data there also exists some mild imbalance between cells due to the need to ensure equal numbers of unique treatments across blocks; the analysis explicitly controls for these sources of imbalance. There was also a set of cross-cutting treatments delivered alongside the source conditions in Ghana; these are discussed in greater detail in the appendix. We control for these additional effects where appropriate.

A potential concern with studies of this type is that the use of convenience sampling produces subject pools that are highly skewed towards young, unemployed men who are available during the day and for whom the compensation is especially attractive. For the purposes of our experiment, it is possible that convenience sampling would also have the unfortunate effect of producing a sample in which the number of subjects who had experience paying direct taxes—a key group of interest—was small. A highly skewed sample thus reduces the external validity of findings, an especially important concern for this study since we are interested in accountability pressures broadly, not within a narrow and unrepresentative sample.

To address this possibility, subjects in Ghana were recruited via random-walk sampling from a central polling station. Each session of 16 was recruited from a single polling station and then provided transport to the field office, located in Adabraca, Accra. While recent representative statistics on Metropolitan Accra are difficult to obtain, a comparison with a large-scale study of 5,484 respondents from 1,250 households conducted from 2008 to 2010 (Fink, Weeks and Hill 2012) suggests that our sample performs favorably in terms of representativeness given that we did not sample from high-income areas of the city. As we will see later, the results for this more representative sample mirror those of the convenience sample in Uganda, reducing concerns about potential selection and the poor external validity it might produce.

Table 3 reports the means of several socioeconomic characteristics of interest (Column 1) alongside the Fink et al estimates where available (Column 2). On balance, our sampling strategy yielded a sample that is approximately gender-balanced (52.5% women) and considerably wealthier than a pure convenience sample. While our sample is relatively less educated, we nearly match the Fink figures for age, employment, and ethnicity. Approximately 60% of our sample was employed, with 11% formally employed by a firm and 12.6% as traders. Critically for our purposes, over half were primary earners, and a full 31% paid some form of direct tax in the previous 6 months.

¹⁴This number represents only subjects who played the game as Citizens. Those who played as Leaders are not included as they did not set a threshold.

Statistic	Sample Mean	Fink et al. Estimate
Age	32.238	29
Female*	52.5	62.4
Employed	62.2	60.2
No Schooling	5.5	17.6
Completed Primary	33.6	21.9
Completed Secondary	23.1	52.0
Ga	56.7	42.2
Akan	30.7	31.0
Ewe	6.5	12.4
Household Income	604.579	NA

Table 3: Summary Statistics for Experimental Sample in Ghana. The Fink et al sample is highly imbalanced along the gender dimension because the sample was recruited based on participation in an earlier survey of at least one adult woman in the household. As such, these estimates should be taken as rough estimates rather than as definitive values for a representative sample.

The average per-month household income was 604 GHC, slightly higher than the inflation-adjusted average of 479 GHC reported for the Greater Accra region by Ghana’s national statistics bureau in 2008. Expanding the sample of educated, high-income respondents with experience paying taxes was our primary motivation for a more rigorous sampling strategy than is often used in experimental games. Nonetheless, to the extent that our sample in Ghana more closely resembles the broader population, this also increases the external validity of our findings.

7 Outcomes of Interest

Our key outcome of interest is the “punishment threshold” for each citizen, defined as the smallest transfer made by the Leader at which the Citizen does not punish the Leader. For example, if a Citizen reports that she would punish the Leader if he passed back 3 MU or less of the 10 MU group fund, but not 4 MU, her punishment threshold for that round is 4 MU. For clarity, our analysis scales the punishment thresholds in both Uganda and Ghana such that they are directly comparable to each other, with 1 MU representing 100 Ugandan shillings and 0.5 Ghanaian cedis respectively.

To measure ownership, we use an item similar to that of Pierce et al (2001) in their study of ownership effects. At the end of the game, each subject was given an outtake survey in which the following question was asked:

How much do you agree or disagree with the following statement: “the group fund belonged to me.”

Subjects who weakly or strongly disagreed were given a value of zero and those who weakly or strongly agreed were coded as a one. The ownership measure captures the degree to which citizens felt they had a claim to the group fund. As per Section 2, we expect ownership to play a strong role in driving expectations of provision which is represented in the game by citizens’ punishment thresholds. While we dichotomize this measure for the ease of exposition, all results presented below hold for its original ordinal construction ranging from zero (strong disagreement) to three (strong agreement).

Finally, Hypothesis 3 requires precise measurement of subjects’ utility. To do so, we use a “utility ladder” to capture participants’ subjective utility at four distinct points in the game. The ladder is a 21-point scale ranging from 0 to 20. At the beginning of the round, subjects are anchored at a value of ten by the enumerator. We set the citizen’s reference point mid-way up the ladder when the wage for the round is received. We then ask the citizen to update their ladder when the leader is given the group fund; after the Citizen observes the Leader’s allocation; and at the end of the round (e.g. after any punishment has occurred and final payouts realized). The utility ladder is included in Figure 4 in the appendix. The ladder was implemented in Uganda only. Below is a brief summary of the ladder implementation protocol:

1. Subjects are given wage and their reference point set at 10
2. After the Leader is given control of group fund and tax is paid (tax condition only), the subject is asked the following (**Ladder1**):

Now that the Leader has the group fund, where are you on the ladder?

3. Citizen observes the Leader’s chosen transfer and is asked the following (**Ladder2**):

Now that the Leader has the group fund, where are you on the ladder?

The ladder measure of interest for the purposes of testing Hypothesis 3 is **Ladder2**, which gives subjects’ utility after they have observed the Leader’s transfer. Given that the utility ladder is a novel measurement instrument, there may be some concerns about construct validity. We show in Section 9 that the ladder responds to Leader transfers exactly as expected, with higher Leader transfers having a large, positive, and statistically significant effect on subjects’ ladder positions.¹⁵

8 Establishing the Ownership Effect: Evidence from Ghana

Because the ladder was implemented in Uganda only, we utilize the Ghana study to test Hypotheses 1 and 2 and the Uganda study to test Hypotheses 3 and 4. Because of space considerations we report the Ugandan results for Hypotheses 1 and 2 in the appendix. While a difference in protocol across the countries prevents us from testing Hypothesis 3 in its entirety in Uganda, we replicate all remaining tests in Uganda and find nearly identical results.¹⁶ We use replication primarily to establish that the protocols were implemented correctly in both countries. Nonetheless, this is one of the few instances in political science of implementation of a nearly identical lab-in-the field protocol across countries of which we are aware. Replication of the major results across two countries provides strong evidence that the ownership effect is a general one and not driven by the particular conditions of either Uganda or Ghana.

¹⁵The ladder is, in fact, sufficiently sensitive that it allows us to recover the S-shaped utility curve postulated by prospect theory. Figure XX in the appendix plots subjects’ change in utility from **Ladder1** to **Ladder2** as a function of the difference between subjects’ expected transfer (their threshold) and the realized transfer (the transfer amount chosen by the Leader). The figure reveals exactly what should be observed if the ladder were working as intended: subjects suffer substantial disutility from lower-than-expected transfers but only modest gains for higher-than-expected transfers, resulting in a steep downward curve on the left of the reference point. In Table XX in the appendix, we also show that it captures the expressive benefits to punishment postulated by existing work in psychology that demonstrates that punishment can relieve negative emotions, increasing subjective well-being.

¹⁶These results are available in Section XX in the appendix.

8.1 Hypothesis 1: Ownership over the government budget should increase citizens’ willingness to punish, regardless of the budget’s source.

The first and most important test of the ownership effect is simply whether strength of ownership *matters* in determining subjects’ willingness to punish.¹⁷ To do so, we estimate the following OLS model:

$$Y_i = \alpha + \beta \text{Ownership}_i + \delta \text{Source}_i + \gamma \mathbf{X}_i + \epsilon_i$$

where Y_{ij} is subject i ’s punishment threshold in round j and **Ownership** is the independent variable of interest, a binary indicator for whether subjects felt ownership over the group fund. **Source** controls for the effect of each revenue source on punishment thresholds, ensuring that β identifies the independent effect of ownership on punishment thresholds. The vector \mathbf{X}_i contains enumerator fixed-effects as well as dummies for additional cross-cutting treatments not analyzed here.¹⁸ Results are reported in Table 4.

The results show that high ownership is a substantively strong and statistically significant predictor of subject thresholds. Subjects who report weak or strong ownership ask for larger transfers from Leaders, with strong ownership corresponding to an increase in transfer thresholds of 0.90 MU ($p \approx 0$). The independent effect of ownership on punishment behavior is stronger even than the effect of taxation: the effect size is approximately 50% larger than the next-largest substantive effect estimated in the experiment.¹⁹

8.2 Hypothesis 2: Taxation increases willingness to punish by increasing ownership over government budgets; the ownership effect is the causal mechanism by which taxation increases accountability

Having established that ownership increases willingness to punish, we now show that this effect drives the major finding of the experimental literature on taxation, namely that taxation increases accountability pressures (see e.g. Martin 2014; Paler 2013). Our hypothesis is that ownership drives the relationship between taxation and accountability: taxation increases accountability by increasing ownership, which, in turn, leads citizens to demand more from government. Establishing this causal chain thus requires providing evidence for three distinct empirical claims: that taxation increases willingness to punish in a general sense; that taxation increases ownership; and that ownership drives punishment behavior.

¹⁷The ownership question was added three days after data collection began, resulting in the loss of 131 subjects from our sample. Of the total, 51 were assigned to the Oil condition, 28 to Grant, 29 to Aid, and 23 to Tax.

¹⁸Two cross-cutting treatments, one manipulating the punishment probability and one introducing a valence prime during the group training, are discussed in the pre-analysis plan and are the subject of ongoing analysis. Because the number of sessions in each block (36) was not a multiple of the number of unique treatments (24), there was minor imbalance in the valence dimension. We thus include the valence term in our estimating equation to control for the effect of this imbalance.

¹⁹One potential concern in estimating the effect of ownership is that ownership may vary with pre-treatment covariates. To account for this possibility, we estimated models with a range of pre-treatment covariates. The coefficient on **Ownership** is stable across these models, nearly identical to that presented in Table 4, and remains strongly significant. Models with pre-treatment covariates included are available in Table XX in the appendix.

<i>DV: Subject Threshold</i>	
Ownership	0.900*** (0.195)
Enumerator FE	Yes
Add'l Treatment FE	Yes
Observations	2297
Adjusted R ²	0.197
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01

Table 4: Impact of Ownership on Punishment Behavior. The reference category for the estimates on the Ownership coefficient are those that report weak or strong disagreement with the statement that the group fund “belonged to them.” The positive and highly significant result demonstrates that ownership over the government budget causes citizens to demand 0.9 MUs more from the Leader than low-ownership types. This effect is not only statistically strong but substantively meaningful; it is approximately twice as large as that of taxation. Model includes source and enumerator fixed-effects as well as dummies for Uncertain Punishment and Valence treatments included but omitted for presentation purposes. Subject-clustered standard errors in parentheses.

The results presented in Table 4 demonstrate the third portion of this causal chain. In this section, we confirm the remaining two. First, we demonstrate taxation’s positive effect on willingness to punish. We then provide evidence of that taxation causes citizens to feel greater ownership over the government budget. This causal chain gives rise to a straightforward prediction: that low-ownership types in the Windfall condition should demand much less than those in the Tax condition, while high-ownership types should demand much more. We conclude the section by presenting evidence of precisely this pattern.

8.2.1 Empirical Implication 1: Taxation Should Increase Willingness to Punish

Hypothesis 2 implies that taxation has some effect on punishment, and as such we must first estimate the overall effect of taxation. To do so, we estimate the following model:

$$Y_{ij} = \alpha + \beta \text{Source}_i + \gamma \text{AdditionalTreatments}_i + \delta \text{Enumerator}_i + \epsilon_i$$

where Y_i is subject i ’s average transfer threshold in round j —that is, the lowest Leader transfer at which the Citizen would not choose to punish. The effect of the source conditions (Aid, Oil, Tax or Grant) is given by the β coefficient vector, while the effect of Uncertain Punishment game and Valence treatments (both implemented in Ghana only) is given by γ . The term δ is the coefficient on a vector of dummies for the enumerator who administered the game to the subject.²⁰

²⁰In our pre-registered design, we noted that we would use randomization inference with OLS as a robustness check. However, due to the replacement of an enumerator partway through the project and the use of substitute enumerators on several occasions due to illness, there was imbalance in the enumerator-game distribution. Given that, as in Martin (2014), there are large enumerator effects with respect to subjects’ thresholds, appropriate randomization inference would require covariate adjustment via regression, a step that, in our view, removes its principal advantage of being a non-parametric estimation method. All results are robust to using a subject-average dataset in which subject i ’s average threshold across all rounds is the dependent variable. Standard errors are clustered at the individual-level to account for within-subject correlation of the error term.

<i>DV: Subject Threshold</i>	
Avg. Δ in Tax Condition	
Windfall	0.416*** (0.164)
Aid	0.416** (0.193)
Oil	0.425** (0.199)
Grant	0.407** (0.196)
Enumerator FE	Yes
Add'l Treatment FE	Yes
Observations	2721
Adjusted R ²	0.144

Note: *p<0.1; **p<0.05; ***p<0.01

Table 5: Average Treatment Effects Relative to Baseline Condition. Positive differences represent greater transfers in the Tax condition relative to the reference group. The large and significant positive effect is consistent with Empirical Implication 1, and demonstrates that taxation increases willingness to punish for both named (Oil, Aid) and unnamed (Grant) non-tax sources. All reported p-values are one-sided (greater) to account for directional predictions. Subject-clustered standard errors in parentheses. Reported n represents total number of subject-round observations used to estimate effects. Enumerator, valence and uncertainty treatment fixed-effects included but not reported for presentation purposes.

We report the results of this model in Table 5. The first comparison of interest is Tax vs non-Tax, in which all subjects who received the Oil, Aid or Grant conditions were pooled into a single non-Tax category. Those in the Tax condition demanded, on average, 0.45 MU more from the Leader than those in the non-Tax conditions ($p = 0.001$). Moreover, as rows 2-4 demonstrates, this result is not due to a single large effect for a particular source but rather consistently higher thresholds in the Tax condition relative to all non-Tax sources.²¹

Given that the mean threshold across all non-tax conditions is 4.34 MU, these treatment effects increase thresholds by approximately 10% relative to the baseline. Overall, the data thus clearly bears the prediction of significant positive differences in thresholds for those in the Tax condition relative to those in non-Tax conditions, providing clear evidence that taxation produces stronger willingness to punish.

The results in Table 5 also reveal that, in general, subjects behaved similarly for all non-tax revenue sources. Why do subjects not respond differentially for our two named non-tax sources (Oil and Aid) compared to the Grant condition? Our theory would predict that thresholds for Oil and Aid would increase relative to the Grant condition only if subjects felt greater ownership over these sources than for the generic, unnamed group fund in the Grant condition. As Figure 1 demonstrates, there is little difference in ownership across non-tax sources. This result is consistent with qualitative work that suggests citizens have very weak intrinsic ownership over non-tax revenues, owing largely to low priors about their ability to control or observe expenditures from these sources and to a

²¹Because our theory makes clear directional predictions, we report one-sided p-values unless otherwise noted. However, all results reported below are robust to the use of two-sided p-values.

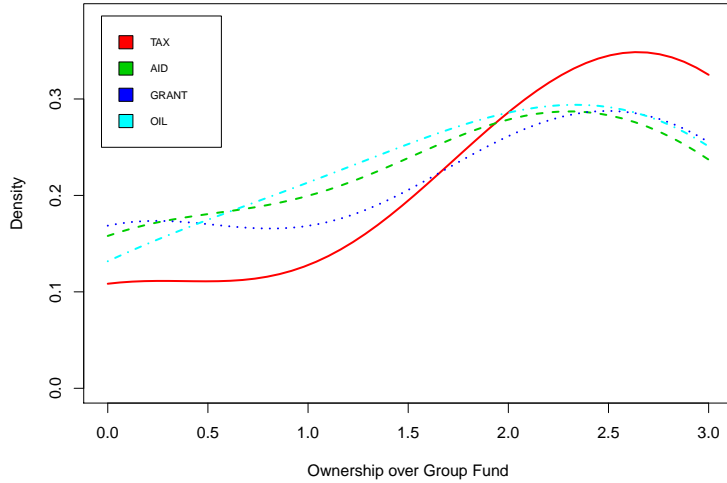


Figure 1: Variation in Strength of Ownership. Figure shows the kernel-smooth densities of ownership measure for each treatment condition. Figure produced using `sm` package in R. The ANCOVA test for distributional equivalence is highly significant ($p \approx 0$). The higher density in the right tail for the ownership condition represents a general trend toward higher ownership rather than a few especially low ownership values in the non-tax conditions or a few especially large values in the Tax condition.

widespread belief that corruption is particularly acute for these sources.

8.2.2 Empirical Implication 2: Taxation Increases Ownership over Government Budgets

We now establish that taxation increases ownership, as predicted by Hypothesis 2. Pooling all non-tax conditions together, we estimate the following logistic regression:

$$Pr(\text{Ownership}_i = 1) = \text{logit}^{-1}(\alpha + \beta \text{Tax}_i + \gamma \text{AdditionalTreatments}_i + \delta \text{Enumerator}_i + \epsilon_i)$$

We find that subjects in the Tax condition are 16.5% more likely to report feeling ownership over the group fund ($p \approx 0$).²² The raw ordering ownership across conditions is also as expected: average ownership is highest in the Tax condition, with 77.4% of respondents feeling ownership over the group fund, followed by Oil at 69.5% and Aid at 62.9% and Grant at 62.8%. In order to demonstrate that this ordering is being driven by a general distributional shift and not a few outliers, Figure 1 plots the kernel-smoothed densities of the ownership variable on its original scale. As the figure demonstrates, higher average ownership in the tax condition is driven by greater density in the weak and strong ownership bins.²³

²²This result holds for a more basic model that does not account for imbalance in the cross-cutting treatments or enumerator fixed-effects. The effect is slightly smaller at 13.9% but still highly significant ($p = 0.0015$).

²³Tellingly, only in the Tax condition is the average response above a weak feeling of ownership over the group fund. This helps to explain why we observed similar overall effects of taxation regardless of the particular source (e.g. aid, oil) of the non-tax revenue.

	<i>DV: Subject Threshold</i>		
	Baseline Results	High Ownership	Low Ownership
Tax Treatment Effect	0.416** (0.164)	0.182 (0.175)	0.916*** (0.193)
Enumerator FE	Yes	Yes	Yes
Add'tl Treatment FE	Yes	Yes	Yes
Observations	2721	1349	1820
Adjusted R ²	0.144	0.162	0.196
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01		

Table 6: Effect of Taxation Relative to Windfall Samples. Results demonstrate that the effect shown in Column 1 is being driven entirely by low-ownership types in the Windfall condition (Column 3), and that high-ownership types in the Windfall condition (Column 2) do not behave differently than those in the tax condition. Enumerator fixed-effects as well as dummies for Uncertain Punishment and Valence treatments included but omitted for presentation purposes. Standard errors clustered at the subject level.

8.2.3 Empirical Implication 3: The Effect of Taxation is Driven by Low-Ownership Types in the Windfall Condition

Our theory posits that ownership should be a major cause of the differential punishment behavior observed in Table 5. Put simply, subjects who have higher ownership over non-tax sources should behave more like taxed citizens, while those with lower ownership should exhibit even larger treatment effects than those found in Table 5 above. To test this, we estimate the same equation used to test Empirical Implication 1, but subset the Windfall sample into high- and low-ownership types, identifying taxation’s effect on punishment when compared against each of the two Windfall samples. We report the results of these two models alongside the overall effect in Table 6 below.

The results in Columns 2 and 3 strongly support our claim that it taxation’s effect on ownership, not taxation *per se*, that is driving the differential willingness to punish, demonstrating that the average effect given in Column 1 is driven entirely by a large difference in willingness to punish between taxed citizens and low-ownership types in the Windfall condition. Together, these results support the major claim of this paper: that taxation increases accountability by engendering stronger ownership over the resulting government budget.

9 Unpacking the Ownership Effect: Evidence from Uganda

Section 8 provides strong evidence in favor of the paper’s two major hypotheses: that ownership drives willingness to punish and that taxation increases willingness to punish by increasing ownership. Using data from a nearly identical protocol implemented in Uganda, we provide evidence in favor of the mechanism proposed in Hypothesis 3: that ownership increases citizens’ dissatisfaction with lower-than-expected provision, increasing the expressive benefits to punishment and making them more willing to pay the costs to do so. We also show, consistent with Hypothesis 4, that while taxation produces ownership effects, feelings of ownership are malleable and subjects can be made to feel high ownership, even over non-tax revenues. We demonstrate that ownership can be

<i>DV: Ladder2 Position</i>	
Ownership	-1.37*** (0.307)
Transfer	1.822*** (0.043)
Enumerator FE	Yes
Observations	2749
Adjusted R ²	0.144
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01

Table 7: Effect of Ownership on Subject Utility. Subject-clustered standard errors in parentheses. Negative (positive) coefficients represent an average decline (increase) in ladder values. The coefficient on **Ownership** is consistent with Hypothesis 3 and demonstrates that ownership causes a large and significant increase in dissatisfaction. The large, positive effect on **Transfer** validates the utility ladder as a measurement instrument, since subjects should feel strictly better off as transfers increase.

manipulated experimentally, and that increasing it can yield accountability pressures for non-tax revenues equal to those generated by taxation.

9.1 Hypothesis 3: For a given transfer, Citizens with high ownership will have lower utility

To establish that ownership impacts subjects’ response to the transfer chosen by the Leader, we utilize the second of our three ladder measurements. **Ladder2** was taken after the subjects’ observed the transfer made by the Leader, and thus represents the subjects’ dissatisfaction with the Leader’s choice in light of the minimum acceptable amount chosen by the subject beforehand. Using this measure, we estimate the following model:

$$Ladder2_{ij} = \alpha + \beta Ownership_i + \delta Transfer_i + \gamma Source_i + \epsilon_i$$

where **Ladder2_{ij}** is Citizen *i*’s second ladder value in round *j*, **Transfer** is the amount of provision chosen by the Leader, and **Ownership** is the subjects’ level of ownership over the government budget. The inclusion of the **Transfer** variable accounts for the (large) effect that the size of the transfer the Citizen receives has on his or her ladder value. As a result, the coefficient on **Ownership** gives the expected change in subjects’ utility given an equivalent transfer from the Leader but different ownership values. Dummies for the treatment condition (**Source**) are included to account for the higher ownership values induced by taxation.

The results from the model are reported in Table 7. As expected, the coefficient on **Ownership** is large and significant, indicating that subjects who felt ownership over the group fund were, on average, 1.4 rungs lower on **Ladder2** than those who did not.

9.2 Hypothesis 4: Any treatment that increases ownership over windfall revenues should also drive willingness to punish upwards, towards the levels observed under taxation.

The results thus far suggest that citizens are more willing to engage in costly sanctioning behavior when tax revenues are at stake, and that this effect is driven by a stronger sense of ownership over tax revenues. The results of the previous suggestion demonstrate that the mechanism behind ownership is rooted in loss: as ownership drives up expectations, citizen dissatisfaction with poor performance increases. This, in turn, increases their expressive benefit to punishment, making them more willing to pay the costs necessary to sanction politicians.

An important question raised by these results is whether ownership is a sticky quantity driven by cumulative experience and expectations about different resources, or whether it can be altered in the short-term. This is an important question from both a theoretical and policy perspective: if ownership is malleable, interventions designed to increase ownership over a source, particularly in the form of information campaigns, may be successful in creating high levels of ownership for non-tax sources.

To test this, we conducted two additional experiments in Uganda identical in all respects to the design presented above but with an additional “ownership” condition for oil and aid revenues. The first four treatments were identical to the four certain punishment games reported above (Tax, Grant, Aid, and Oil). The additional “ownership” treatments were a variant on the oil and aid games with certain punishment. In the “Oil Ownership” and “Aid Ownership” conditions, citizens are told that a specific portion of oil or aid revenues actually belongs to them. In the script for each round, the enumerator tells the citizen that 500 Ugandan Shillings (UGX) of the 1,000 UGX group fund “represents the share of the oil/aid money that belongs to you, as the citizen.” This closely mimics the script for the tax game, with one key difference: the 500 UGX that “belongs” to the citizen in the oil and aid ownership conditions is never actually held by them, and as such the citizen makes no direct contribution to the group fund.

Given the strong evidence that ownership effects are driving the difference between willingness to punish in the tax and non-tax treatments, we expect that the oil and aid ownership treatments will lead to an increase in punishment thresholds relative to the “basic oil” and “basic aid” games. It should be noted that Uganda is an excellent case to run these experiments; the country recently discovered oil, and qualitative interviews suggest that, while citizens vary in the degree to which they view oil money as theirs, views on oil may still be malleable. Additionally, Uganda is highly aid-dependent, with one third of its budget attributed to budget-support aid and significant off-budget aid coming through project aid equivalent to roughly 40 percent of the budget.

Table 8 reports the results of a model identical to that estimated by Equation 8.2.2. The coefficient on `BaseOil` gives the baseline effect to which the ownership treatment should be compared: the average effect of taxation on thresholds when the reference group is those in the Oil condition. As reported previously, this effect is highly significant; taxation clearly increases thresholds. As the `OilOwnership` coefficient makes clear, the ownership treatment raises thresholds dramatically, cutting the effect of taxation by a third and make average thresholds in the Oil Ownership condition indistinguishable from those in the Tax condition ($p = 0.26$). This reduction is precisely as Hypothesis 4 would predict: giving (indirect) ownership over the non-tax revenue yields thresholds statistically equivalent to those under taxation. Results for the aid ownership condition are stronger still, but analysis is still in process.

<i>Control Group</i>	<i>Tax Effect</i>	<i>Standard Error</i>	<i>p-value</i>	<i>n</i>
Base Oil	0.493***	0.175	0.005	1668
Oil Ownership	0.212	0.187	0.26	1716

Table 8: Ownership Treatment Effects on Subject Thresholds. Both rows report the results of an OLS model with enumerator fixed-effects where the dependent variable is the threshold below which the respondent would punish the leader. One-sided p-values were used to correspond to the directional nature of the tests according to our hypotheses. The coefficient on `BaseOil` represents the effect of taxation when the control group is those in the Oil condition, while that of `OilOwnership` is the effect relative to those in the Oil Ownership treatment. The small and statistically insignificant effect for `OilOwnership` is consistent with Hypothesis 4, suggesting that giving indirect ownership over the group fund is sufficient to produce significantly stronger accountability effects.

10 Discussion and Conclusion

What drives citizens’ willingness to hold leaders accountable for government spending behavior? And why does taxation appear to increase accountability pressures? We have argued here that the answer to both questions is rooted in what we have termed the ownership effect. Drawing on recent research in cognitive psychology, we proposed an ownership-based framework to understand when and under what conditions citizens will hold leaders accountable for poor performance.

We set out to test two sets of hypotheses that flow directly from a model of sanctioning rooted in the degree to which citizens feel ownership over collectively-owned resources. First, we argued that ownership should drive willingness to sanction and that ownership should explain why taxation produces greater accountability pressures. Using a lab-in-the-field experiment in Ghana, we showed strong support for both hypotheses: ownership has a substantively large and statistically significant effect on subjects’ willingness to punish leaders for their spending behavior, and taxation increases citizens willingness to punish by increasing ownership.

Our second set of hypotheses addressed the natural question that follow from these results: by what mechanism does ownership increase willingness to punish, and can ownership be increased through experimental interventions? To answer these questions, we replicated and extended our experimental protocol in Uganda, a proto-typical African state where oil has been recently discovered and concerns about the effect of the influx of non-tax revenues are high. Using a novel instrument to measure subjective utility at several points during the intervention, we demonstrated that ownership drives reference points, raising citizens’ expectations about the amount of provision they would receive from leaders.

We then used an additional experimental intervention to test our final hypothesis: that ownership is malleable, and can be increased through intervention. Using a subtle variation on our existing protocol designed to give subjects direct ownership over oil and aid revenues, we found that increasing ownership can induce untaxed subjects to demand accountability equal to that of taxed subjects. This result is also consistent with an empirical implication of our second hypothesis: that citizens with high ownership over aid and oil revenues demand accountability equal to those in the Tax condition. Together, these results thus constitute compelling evidence in favor of the ownership effect.

One might also ask how general the ownership effect is, and what scope conditions might

exist. The replication of the protocol and the empirical findings in two countries is evidence, at the very least, that this effect holds for much of sub-Saharan Africa, especially in light of the fact that Uganda and Ghana are frequent cases precisely because they represent two ends of the developmental and governance spectrum in the sub-continent. Given that the ownership effect is ultimately a psychological mechanism, we see little reason why it should not hold in most other places. Of course, while we expect ownership to be a strong predictor of punishment behavior in all cases, the extent to which it mediates the relationship between revenue source and accountability pressures will vary with the degree of intrinsic ownership citizens feel over non-tax revenues. Where citizens have high ownership over non-tax sources, we would not expect ownership to account for observed differences in accountability pressures—although we would also expect, all else equal, that accountability for non-tax sources should be high in such places.

The finding that differential accountability across sources is driven by differential levels of ownership also has important policy implications. As developing countries, particularly in sub-Saharan Africa, discover and monetize large natural resource endowments, a critical concern of donors and international institutions is how to avoid the resource curse. Our results suggest that accountability pressures for these sources are low precisely because citizens feel little ownership over these sources. As such, interventions that can enhance the level of ownership felt by citizens for all government revenues but especially for non-tax sources may therefore help to improve democratic accountability in weakly institutionalized countries.

Our results also raise several avenues for future research. Greatest among them, in our view, is investigating whether indirect taxation generates the same degree of ownership over government budgets as direct taxation. There are compelling reasons to think that the lower visibility of indirect taxation—both in terms of magnitude and their contribution to total government revenue—may less ownership and thus weaken accountability. This possibility is especially worrying because in most developing countries very few people pay income taxes, or direct taxes of any sort.

Instead, developing countries, particularly in sub-Saharan Africa, are relying increasingly on indirect taxes such as the value-added tax (VAT) for their tax-based revenues. Driven in part by pressure from donor countries and organizations, indirect taxes are seen as an attractive alternative to direct taxation in low-capacity, high-corruption environments because they are relatively easier to collect and are thought to be less vulnerable to rent-seeking. Yet it is not clear whether indirect taxation produces sufficiently high ownership to yield the accountability pressures we observe from taxation. If indirect taxation yields weaker accountability pressures, its value as a solution in low-capacity environments may be considerably lower than is currently believed.

References

- Baskaran, Thushyanthan and Arne Bigsten. 2013. “Fiscal capacity and the quality of government in sub-Saharan Africa.” *World Development* 45:92–107.
- Bates, Robert. 1981. “States and markets in tropical Africa: The political basis of agricultural policy.” *Berkeley: University of California Press, series on social choice and political economy* .
- Bates, Robert and Da-hsiang Donald Lien. 1985. “Taxation , Development , Representative Government.” *Politics and Society* 14(1):53–70.
- Beblawi, Hazem and Giacomo Luciani. 1987. *The Rentier State*. Croom Helm.

- Bräutigam, Deborah A and Stephen Knack. 2004. "Foreign Aid, Institutions, and Governance in Sub-Saharan Africa." *Economic Development and Cultural Change* 52(2):255–285.
- Brollo, Fernanda, Tommaso Nannicini, Roberto Perotti and Guido Tabellini. 2013. "The Political Resource Curse." *American Economic Review* 103(5):1759–96.
- Caselli, Francesco and Tom Cunningham. 2009. "Leader behaviour and the natural resource curse." *Oxford Economic Papers* 61(4):628–650.
- Casey, Katherine E., Rachel Glennerster and Edward Miguel. 2012. "Reshaping Institutions: Evidence on Aid Impacts Using a Pre-analysis Plan." *Quarterly Journal of Economics* 127(4):1755–1812.
- Chaudhry, Kiren Aziz. 1998. *The Price of Wealth: Economies and Institutions in the Middle East*. Cornell University Press.
- De Kadt, Daniel and Evan S Lieberman. 2016. Nuanced Accountability: Voter Responses to Service Delivery in Southern Africa. Technical report Mimeo: Massachusetts Institute of Technology.
- de la Cuesta, Brandon, Helen V Milner, Daniel Nielson and Steve Knack. 2017. "Taxation Without Representation? Experimental evidence from Ghana and Uganda on citizen action toward oil, aid, and taxes."
- Djankov, Simeon, Rafael La Porta, Florencio Lopez-de Silanes and Andrei Shleifer. 2008. "The law and economics of self-dealing." *Journal of Financial Economics* 88(3):430–465.
- Ekeh, Peter P. 1975. "Colonialism and the two publics in Africa: A theoretical statement." *Comparative studies in society and history* 17(01):91–112.
- Fearon, James D. 1999. Electoral Accountability and the Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance. In *Democracy, Accountability, and Representation*, ed. Adam Przeworski, Susan C Stokes and Bernard Manin. Cambridge University Press.
- Fehr, Ernst and Klaus M Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation." *The Quarterly Journal of Economics* 114(3):817–868.
- Fehr, Ernst and Simon Gächter. 2000. "Cooperation and Punishment in Public Goods Experiments." *American Economic Review* 90(4):980–994.
- Fiorina, Morris P. 1976. "The voting decision: instrumental and expressive aspects." *The Journal of Politics* 38(2):390–413.
- Fisman, Raymond and Roberta Gatti. 2002. "Decentralization and Corruption: Evidence from US Federal Transfer Programs." *Public Choice* 113(1-2):25–35.
- Gadenne, Lucie. 2015. "Tax Me, But Spend Wisely: The Political Economy of Taxes, Evidence from Brazilian Local Governments." <https://sites.google.com/site/lgadenne/home/research> .
- Gottlieb, Jessica. 2016. "Greater Expectations: A Field Experiment to Improve Accountability in Mali." *American Journal of Political Science* 60(1):143–157.
- James, Oliver. 2007. "Evaluating the expectations disconfirmation and expectations anchoring approaches to citizen satisfaction with local public services." *Journal of Public Administration Research and Theory* 19(1):107–123.

- Kahneman, Daniel, Jack L. Knetsch and Richard H. Thaler. 1990. "Experimental Tests of the Endowment Effect and the Coase Theorem." *The Journal of Political Economy* 98(6):1325–1348.
- Kan, Kamhon and Cheng-Chen Yang. 2001. "On expressive voting: Evidence from the 1988 US presidential election." *Public Choice* 108(3):295–312.
- Levi, Margaret. 1989. *Of Rule and Revenue*. University of California Press.
- Lieberman, Evan S, Daniel N Posner and Lily L Tsai. 2014. "Does information lead to more active citizenship? Evidence from an education intervention in rural Kenya." *World Development* 60:69–83.
- Mahdavy, H. 1970. The Patterns and Problems of Economic Development in Rentier States: the Case of Iran. In *Studies in the Economic History of the Middle East*, ed. M.A. Cook. pp. 428–677.
- Martin, Lucy. 2014. "Taxation , Loss Aversion , and Accountability: Theory and Experimental Evidence for Taxation's Effect on Citizen Behavior."
- Martin, Lucy. 2017. "Strategic Taxation." *Working Paper* .
- Morewedge, Carey K., Lisa L. Shu, Daniel T. Gilbert and Timothy D. Wilson. 2009. "Bad riddance or good rubbish? Ownership and not loss aversion causes the endowment effect." *Journal of Experimental Social Psychology* 45(4):947–951.
URL: <http://dx.doi.org/10.1016/j.jesp.2009.05.014>
- Morrison, Kevin M. 2009. "Oil, Nontax Revenue, and the Redistributive Foundations of Regime Stability." *International Organization* 63(01):107.
- Morrison, Kevin M. 2015. "Natural Resources and Development." *Emerging Trends in the Social and Behavioral Sciences* .
URL: <http://search.ebscohost.com/login.aspx?direct=true&db=ecn&AN=0972994&site=ehost-live>
- North, Douglas C. and Barry R. Weingast. 1989. "Constitutions and Commitment: The Evolution of Institutions Governing Public Choice in Seventeenth-Century England." *The Journal of Economic History* 49(04):803–832.
- Olken, Benjamin A. 2007. "Monitoring Corruption: Evidence From a Field Experiment in Indonesia." *Journal of Political Economy* 115(2):200–249.
- Olson, Mancur. 2009. *The logic of collective action*. Vol. 124 Harvard University Press.
- Paler, Laura. 2013. "Keeping the Public Purse: An Experiment in Windfalls, Taxes, and the Incentives to Restrain Government." *American Political Science Review* 107(04):706–725.
URL: http://www.journals.cambridge.org/abstract_S0003055413000415
- Persson, Anna and Bo Rothstein. 2015. "It's My Money: Why Big Government May Be Good Government." *Comparative Politics* 47(2):231–249.
- Prichard, Wilson. 2015. *Taxation, Responsiveness and Accountability in Sub-Saharan Africa: The Dynamics of Tax Bargaining*. Cambridge University Press.

- Reinikka, Ritva and Jakob Svensson. 2005. "Fighting corruption to improve schooling: Evidence from a newspaper campaign in Uganda." *Journal of the European Economic Association* 3(2-3):259–267.
- Ross, Michael L. 2004. "Does Taxation Lead to Representation?" *British Journal of Political Science* 34(02):229–249.
- Sandbu, Martin E. 2012. Direct distribution of natural resource revenues as a policy for peacebuilding. In *High-Value Natural Resource and Peacebuilding*, ed. P. Lujala and S.A. Rustad. Vol. 1 pp. 275–290.
- Schumpeter, Joseph A. 1991. The crisis of the tax state. In *The Economics and Sociology of Capitalism*, ed. Richard Swedberg. Princeton: Princeton University Press pp. 99–141.
- Shu, Suzanne B and Joann Peck. 2011. "Psychological ownership and affective reaction: Emotional attachment process variables and the endowment effect." *Journal of Consumer Psychology* 21(4):439–452.
- Smith, Alistair. 2008. "The Perils of Unearned Income." *The Journal of Politics* 70(03):780–793.
- Timmons, Jeffrey F. 2005. "The Fiscal Contract: States, Taxes, and Public Services." *World Politics* 57(4):530–67.
- Tyran, Jean-Robert. 2004. "Voting when money and morals conflict: an experimental test of expressive voting." *Journal of Public Economics* 88(7):1645–1664.
- Van de Walle, Nicolas. 2001. *African economies and the politics of permanent crisis, 1979-1999*. Cambridge University Press.
- Waterbury, John. 1998. The State and Economic Transition in the Middle East and North Africa. In *Prospects for Middle Eastern and North African Economies: From Boom to Bust and Back?*, ed. Nemat Shafik. Macmillian pp. 159–177.
- World Bank. 2016. *World Bank Development Indicators*. Washington, D.C.: The World Bank.
URL: <http://data.worldbank.org/data-catalog/world-development-indicators>
- Wu, Yin, Jie Hu, Eric van Dijk, Marijke C Leliveld and Xiaolin Zhou. 2012. "Brain activity in fairness consideration during asset distribution: does the initial ownership play a role?" *PloS one* 7(6):e39627.

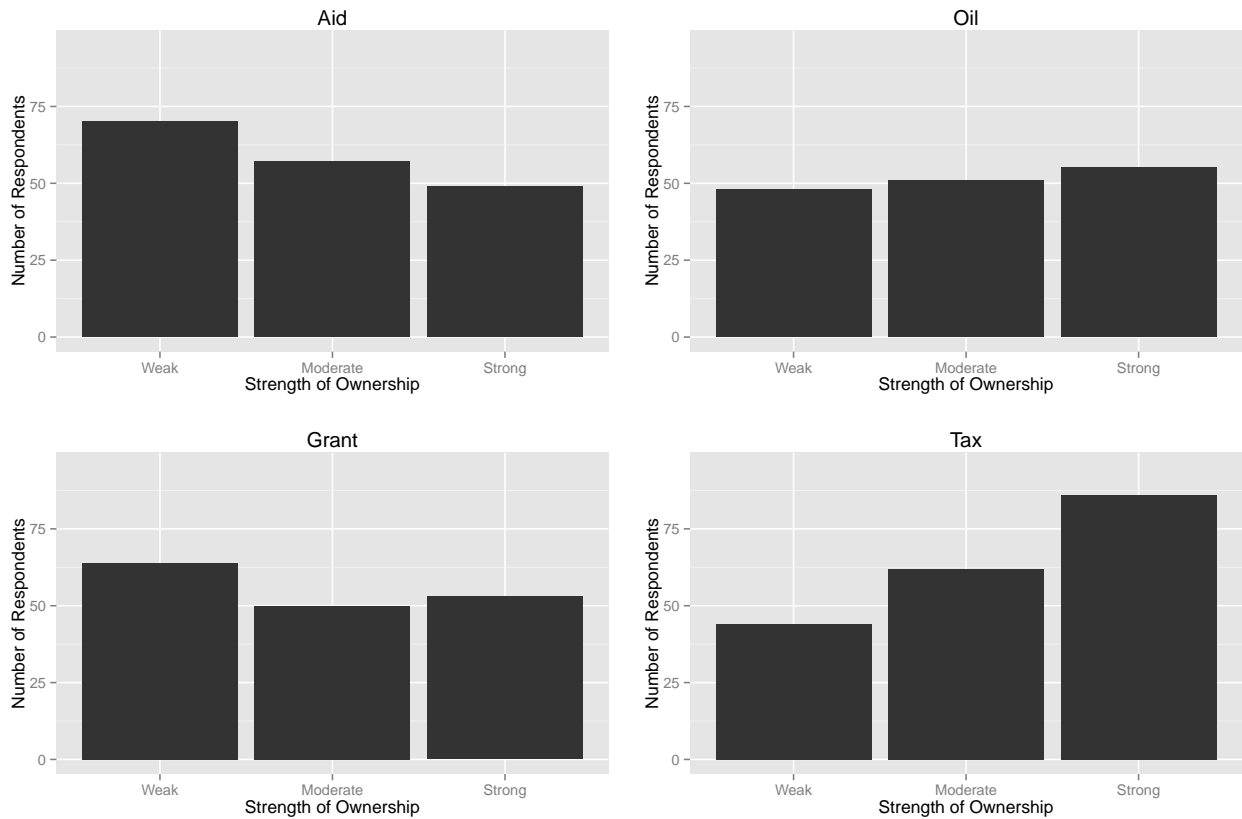


Figure 2: Distribution of Ownership Measure by Source Condition. The “weak” category contains respondents who weakly or strongly disagreed with the statement that “the group fund belonged to me”. The “moderate” and “strong” categories contains respondents who agreed or strongly agreed respectively. The Tax and Aid conditions are the most clearly distinguishable, with the former being weighted towards strong ownership and the latter weak ownership. Oil occupies an intermediate position slight imbalance towards strong ownership.

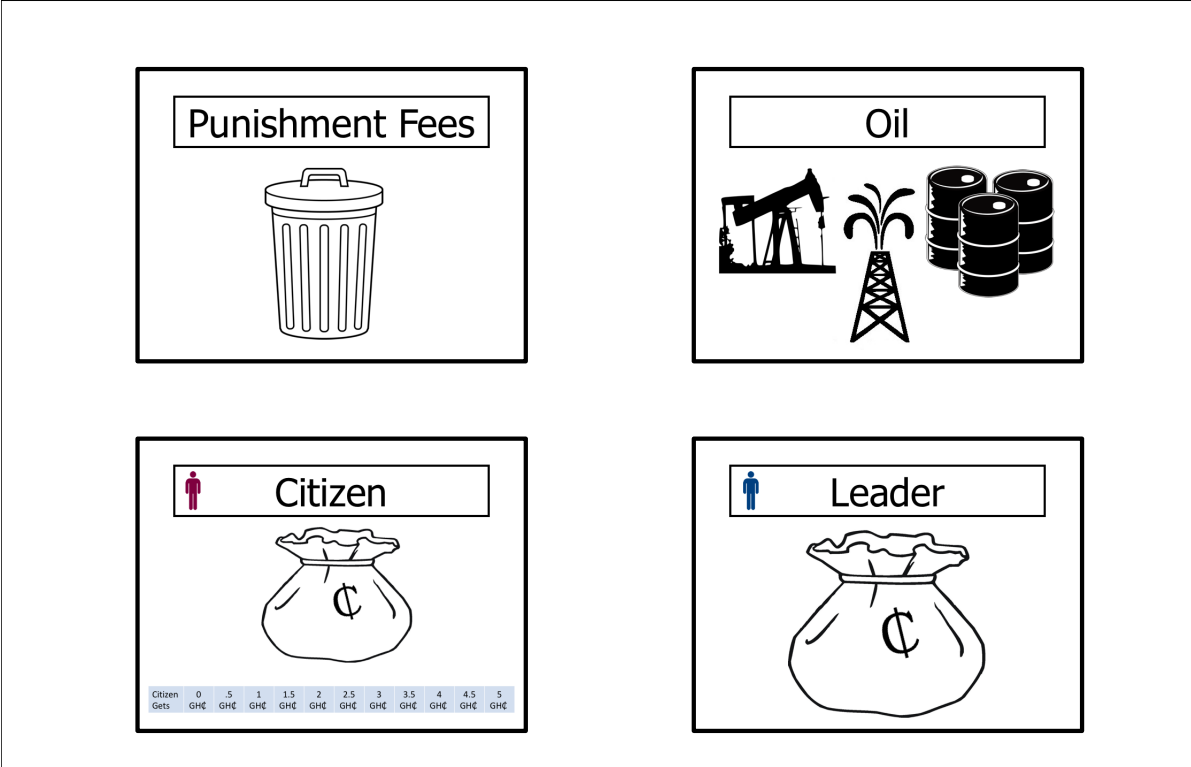


Figure 3: Example of Game Board, Oil Condition. Game boards for the Aid, Tax and Grant conditions differ only with respect to the image on the source tile. The Grant condition source tile has no image and is simply titled “Group Fund”. To emphasize the source of the group fund, the enumerator places the money on the source tile before transferring it to the Leader. At the time the Citizen sets her transfer threshold, the revenue is on the Leader’s tile.



Figure 4: Utility Ladder