



Regular Article

How information about foreign aid affects public spending decisions: Evidence from a field experiment in Malawi[☆]

Brigitte Seim^a, Ryan Jablonski^{b,*}, Johan Ahlback^b

^a University of North Carolina, Chapel Hill, United States

^b London School of Economics and Political Science, United Kingdom



ARTICLE INFO

Keywords:

Foreign aid
Public spending
Fungibility
International development
Africa
Malawi

ABSTRACT

Does foreign aid shift public spending? Many worry that aid will be “fungible” in the sense that governments reallocate public funds in response to aid. If so, this could undermine development, increase the poorest’s dependency on donors, and free resources for patronage. Yet, there is little agreement about the scale or consequences of such effects. We conducted an experiment with 460 elected politicians in Malawi. We provided information about foreign aid projects in local schools to these politicians. Afterwards, politicians made real decisions about which schools to target with development goods. Politicians who received the aid information treatment were 18% less likely to target schools with existing aid. These effects increase to 22–29% when the information was plausibly novel. We find little evidence that aid information heightens targeting of political supporters or family members, or dampens support to the neediest. Instead the evidence indicates politicians allocate the development goods in line with equity concerns.

1. Introduction

When foreign aid is prevalent, do politicians make different public spending decisions than they would in the absence of aid? If so, how does aid shift public spending patterns, and which types of citizens are helped or hurt by the shift? The potential for aid to displace public spending—or be “fungible”—has been blamed for a host of development ills. In addition to undermining development goals and aid effectiveness, aid fungibility may contribute to corruption and political patronage. Additionally, by making the poorest in society dependent upon donors rather than turning to domestic authorities to meet their needs, fungibility may contribute to low government accountability and

aid dependency.

Scholars have not definitively determined how politicians’ spending decisions respond to foreign aid. The predominant view is that fungibility undermines development because politicians respond to foreign aid by shifting public spending to areas that are more politically valuable or are overlooked by donors (Easterly, 2009; Gibson et al., 2005; Bueno de Mesquita and Smith, 2009; Morrison, 2007). In this view, donors have “first dibs” on where to focus their efforts, and politicians select other beneficiaries based on alternative (and potentially less optimal) criteria. We call this a “crowding out effect” of aid on public spending.¹ Alternatively, politicians may choose to spend in line with donors, either because they view donor spending as a kind of endorsement,

[☆] This research was supported by AidData at the College of William and Mary, United States and the USAID Global Development Lab, United States through agreement AID-OAA-A-12-00096. The views expressed here do not necessarily reflect the views of AidData, USAID, or the United States Government. The project has been reviewed and approved by the Malawi National Commission on Science and Technology and by the London School of Economics Research Ethics Committee. A pre-analysis plan is available at <http://doi.org/10.17605/OSF.IO/XJ72Z>. Replication data are available at <https://doi.org/10.7910/DVN/GMQUAQ>. Invaluable research assistance was provided by Inbok Rhee, Jimmy Mkandawire, and Bright Chimatiro, and a dedicated field research team. We thank Tearfund for partnering with us on this research, and many others in the Malawi development community for data and feedback. We also appreciate generous feedback from Editor Patrick Francois, two reviewers, Daniel Nielson, Darren Hawkins, Alexandra Zeitz, Abhit Bhandari, Oliver Morrissey, Ariel BenYishay, Gabrielle Cheung, Matthew Winters, and seminar participants at Duke University, Boise State University, Brigham Young University, Nottingham University, University of Washington, and the Political Economy of International Organizations Conference. We also thank rtSurvey for their work on survey design.

* Corresponding author.

E-mail addresses: bseim@ad.unc.edu (B. Seim), r.s.jablonski@lse.ac.uk (R. Jablonski), j.m.ahlback@lse.ac.uk (J. Ahlback).

¹ In this paper, we are focused on identifying a *spatial* crowding out effect, where aid to one site “crowds out” public spending at that site and shifts it to another site. Others in the literature examine how aid “crowds out” spending from particular budget sectors.

<https://doi.org/10.1016/j.jdevec.2020.102522>

Received 3 January 2019; Received in revised form 2 March 2020; Accepted 5 June 2020

Available online 3 July 2020

0304-3878/© 2020 Elsevier B.V. All rights reserved.

or because of “flypaper” effects (Hines and Thaler, 1995). Indeed, this behavior occurs among donors themselves, with widespread clustering of foreign aid projects.² We call this a “validation effect” of aid on public spending.

While the effect of aid fungibility on public spending patterns has been tested using observational data on public spending at the national and cross-national level (e.g., Feyzioglu et al., 1998), there are no direct tests of how foreign aid or information about foreign aid affects real spending decisions by politicians at the individual level. Further, much of the literature on aid fungibility assumes that any crowding out effect would be a normatively undesirable outcome, shifting resources from programmatic and geographic areas that need them to those that do not.³ Few acknowledge that crowding out could mean politicians are effectively and efficiently shifting resources to meet citizens’ needs in response to foreign aid (Rana and Koch, 2020).

To fill these gaps in the fungibility literature, we execute an experiment among in-office elected politicians in Malawi.⁴ The effects of aid are particularly important to understand in Malawi. Malawi is among the most aid dependent countries in the world, with aid totalling 129% of central government spending (World Bank, 2019). About 10% that aid is channelled into the education sector, either through support to local government or in the form of projects executed directly by donors (Ministry of Finance, 2019). Yet, by many measures, Malawi remains among the most unequal countries in the world, especially with respect to educational spending and attainment. For instance, a child born in a rural area has only about a 6% chance of achieving a minimum level of reading mastery, compared to 21% for a child in an urban area (World Bank, 2010). According to one estimate, 10% of students in Malawi consume 68% of all education spending (UNICEF, 2015, 57).

Our experiment was designed to adjudicate between the “crowding out” and “validation” effects of aid fungibility. Specifically, we randomly assigned 460 elected local councillors (LCs) and members of parliament (MPs) to receive or not receive information about the spatial distribution of foreign aid projects at schools in their constituencies. After receiving (or not receiving) this information, the politicians made real decisions about the spatial allocation of development goods to these same schools. Schools selected by the politicians were entered into a lottery to determine which schools received the goods. About 30% of the politicians enrolled in the experiment followed up on the lottery, implying that the goods are valuable.

Our results are consistent with a spatial crowding out effect. When politicians are informed about foreign aid projects, they are 18% less likely to select schools that already received foreign aid projects. These effects are considerably smaller than those identified in most observational studies of fungibility. In assessing politicians’ pre-existing knowledge of the spatial distribution of foreign aid projects, we determined that the information we provided was novel for most politicians. Among

these politicians, for whom the information about foreign aid projects was more novel, treatment effects increase to 22–29%. This heterogeneous effect based on pre-existing knowledge suggests that the aid information treatment affects politicians’ decisions by updating their priors. In interpreting our finding, we provide qualitative evidence from interviews and open-ended survey questions that the primary mechanism underpinning the spatial crowding out effect is a fairness norm, whereby politicians seek to provide development assistance to schools that have not yet received support. While we remain agnostic regarding whether the fairness norm is socially optimal, we note it as a compelling area for future research.

We also evaluate the distributional implications of the spatial crowding out effect. We coded schools based upon their level of need and the percentage of votes received by the politician in the most recent election at a given school—most primary schools in Malawi function as polling stations during elections—as well as the attendance of his or her family members. We find no evidence that the spatial crowding out effect of information about aid benefits considerably less needy schools, increases political biases in education spending, or results in the systematic targeting of goods to schools with politicians’ family members.

The aid information treatment was randomized and delivered in the context of a broader multi-arm factorial experiment that evaluated how politicians respond to information about school need, political support, and aid in public allocation decisions. As we discuss in Section 5.2.1 below, we do not find compelling evidence that the other forms of information affected politician decision-making, suggesting that information about foreign aid may be particularly novel and relevant for politicians as they make public spending decisions. We therefore focus this article on the effects of providing information about foreign aid projects. In the SI, we further detail the additional treatment arms, consider interactions across treatments, provide tests of pre-specified hypotheses, and provide multiple comparison corrections.

The decision to examine aid fungibility via an experiment adds realism and causal leverage, but does come with trade-offs. One tradeoff is that we can only study one specific form of budgetary decision. In our experiment, the decision was how to allocate goods funded by an NGO within the education sector. While this decision might seem contrived to those familiar with more traditional budgetary processes, this is a common budgetary decision in much of the developing world; around half of local education spending in Malawi comes from NGOs or other donors and about 70% of our respondents claim to meet with donors about development issues.⁵ While we cannot make claims that this budgetary decision is wholly representative, we took pains to include a nearly comprehensive sample of politicians in our experiment. Out of 655 LCs and MPs in Malawi, our experiment involved 460. Thus, the experiment provides a realistic test of how fungibility affects politicians’ decisions within one highly aid dependent and politicized sector.

Our conclusions have important implications for policymaking. For one, the results imply that concerns about the perverse budgetary effects of aid may be overstated, and that efforts to mitigate fungibility may be misplaced. This is particularly true given the fact that we see no evidence that spatial crowding out substantially changes allocations to politically important or needy schools.

² Of the 3151 schools in our sample with foreign aid, 37% have more than one donor involved. Yet we identify no aid in 40% of schools.

³ For instance, Bueno de Mesquita and Smith (2009) assume in their model that “bilateral aid is largely fungible such that the recipient leader can spend the resources as she sees fit.” Similarly, Morrison (2007) argues that “foreign aid is a highly fungible resource and acts similarly to oil in that it provides extra resources the government can use to distribute to its key constituencies without taxation.”

⁴ This experimental design and our hypotheses were pre-registered on the Evidence in Governance and Politics website prior to analysis (<http://doi.org/10.17605/OSF.IO/XJ72Z>). We describe some minor deviations from this plan in the Supplementary Information (SI). Additionally, our research protocol was informed by an earlier pilot experiment with local councillors in 2015. This pilot project mirrored many of this project’s design features and informed our pre-specified priors (Jablonski and Seim, 2017).

⁵ In Malawi, over 37% of the government budget is funded by foreign donors (World Bank, 2019), and donors contributed funds to about 34% of all schools in Malawi during our data collection period, compared to 38% from local government. These statistics are based on an authors’ survey with teachers in 311 schools.

2. Our contribution to the literature

We build upon a rich literature exploring the fiscal effects of aid.⁶ Most conclude – as we do – that governments often respond to foreign aid by changing the composition of public spending (Dollar and Pritchett, 1998; Chatterjee et al., 2012; Marć, 2017; Feyzioglu et al., 1998; Werker et al., 2009). Yet, estimates from these studies are surprisingly disparate. Some studies document a large *negative* effect of aid on public spending (Marć, 2017; Feyzioglu et al., 1998; Werker et al., 2009; Dollar and Pritchett, 1998). Other studies find “flypaper” effects in which foreign aid *increases* public spending (van de Walle and Mu, 2007; Morrissey, 2015; Remmer, 2004). Still others find no evidence of fungibility (Pack and Pack, 1990). Estimates of the scale of fungibility similarly vary from nearly the entirety of the aid budget to more marginal effects (Werker et al., 2009; Chatterjee et al., 2012; Van de Sijpe, 2013; van de Walle and Mu, 2007).

We make several contributions to these debates. First, to our knowledge, this is the first experimental test of how aid affects public spending decisions. Since our treatment is randomly assigned, our research design is unaffected by many of the endogeneity and measurement issues inherent in observational research on fungibility (McGillivray and Morrissey, 2000; Van de Sijpe, 2013). Similarly, existing research has failed to establish causal precedence, as donors have incentives to tailor aid conditionality provisions in anticipation of fungibility. Demand for foreign aid is related to shocks in public spending and income, making it challenging to show that one causes the other.⁷ We are also able to avoid measurement challenges in existing studies of aid and public spending.⁸

Second, most existing research on aid fungibility estimates the spillover of aid funds *across* public spending sectors. We show that aid also causes a spatial reallocation of spending *within* sectors.⁹ This is not an insignificant contribution, since many of the greatest inequities in development are spatial and existing policy solutions to address aid fungibility rely on sectoral spending restrictions. Moreover, the information asymmetries between donors and politicians make it very hard for donors to monitor the kind of local diversion of funds we identify in this study (Jablonski, 2014).

Third, existing studies have largely been unable to disentangle who benefits from the displacement of public spending, or what motivates politicians' decisions to reallocate development. As noted above, many studies conclude that fungibility implies more public spending among political supporters or in less deserving communities, or that fungibility will promote corruption. Yet, there have been few attempts to val-

⁶ Our research also builds upon methodological insights from other experimental studies of foreign aid. One analog is Findley et al. (2017), who assign information about development projects to Ugandan MPs in order to measure preferences for donor versus government projects. While their research question is quite different, they likewise show that information about aid projects is valuable and novel for politicians, and is meaningful for political decision-making. Additionally, other studies provide information about aid to citizens and find that this information can affect political attitudes, perceptions of donors, and preferences over development typologies (Dietrich et al., 2018; Baldwin and Winters, 2018; Blair and Roessler, 2018; de la Cuesta et al., 2019).

⁷ For instance, aid may target a budget shortfall. See Werker et al. (2009) for discussion and evidence of bias.

⁸ Among other problems, there is no easy way to determine how much donors intended to be spent in a particular sector in most cases, and distinguishing between on-budget and off-budget aid is not trivial (Van de Sijpe, 2013). Also, public spending data in aid dependent states is often unreliable, or potentially even strategically biased (Morrissey, 2015).

⁹ The only other studies we are aware of that look at spatial crowding out are van de Walle and Mu (2007); Wagstaff (2011). These authors estimate crowding out across road projects and health projects, respectively, in Vietnam. Their estimates are broadly consistent with the scale of effects in our study. As we do, Wagstaff (2011) concludes that these spatial crowding out effects are generally welfare improving.

idate these assumptions (Morrissey, 2015; Wagstaff, 2011). Our estimates and qualitative assessment provide further reason to doubt the generalizability of these conclusions.

3. Theory and hypotheses

When politicians make public spending decisions, they have to consider both the characteristics of the beneficiaries and the spending behavior of other development actors. This coordination problem is particularly acute in low-income developing democracies, where donors fund a significant portion of local development projects yet inconsistently coordinate with government. How do politicians adjust for spending by donors? How does this affect the characteristics of who benefits from public spending?

We illustrate the logic of a politician's distributional problem, and its potential consequences, with a simple model. Consider a politician that has to make a decision about how to allocate a fixed development budget of value $a > 0$ to one of three schools in her constituency. In making this decision, a politician has to consider both the effects of a on the economic development of the schools, $d(a) = d_1(a_1) + d_2(a_2) + d_3(a_3)$, as well as its effects on voting in the communities that rely on the school, $v(a) = v_1(a_1) + v_2(a_2) + v_3(a_3)$. We assume d and v are increasing and concave with respect to a and we define $d_i > d_{i+1}$ and $v_j > v_{j+1}$ for all i and j . We also assume a politician faces a trade-off between development and votes, so school $i = 1$ is never the same as school $j = 1$.

Absent donor spending, a politician's distributional problem is straightforward. She will choose a_i to solve the utility maximization problem in equation (1) where λ represents the weight she places on development versus votes. Thus if a politician cares more about maximizing development, she will, under most conditions, invest $a_{i=1} > 0$. If she maximizes votes, she will often invest $a_{j=1} > 0$.

$$\max_a \lambda(d_1 + d_2 + d_3) + (1 - \lambda)(v_1 + v_2 + v_3) \quad (1)$$

How might this problem change in the presence of donor spending? To explore this, suppose that, prior to the allocation of a , a donor also makes an aid investment b . Since donors are development focused, we assume they will always invest in the neediest school. This investment $b_{i=1} > 0$ will have two complementary implications for a politician's maximization problem. First, from our assumption of concavity, it follows that $\frac{\partial d_i}{\partial a_i}(b_i > 0) < \frac{\partial d_i}{\partial a_i}(b_i = 0)$. So the donor's aid will decrease the development effectiveness of a politician's investment in the neediest school, $a_{i=1}$. Second, the donor's aid will increase a politician's incentives to invest in votes rather than development. Because donor spending does not affect v and donors never allocate to maximize v , the returns to a politician from maximizing v will always be greater with aid than without aid.¹⁰ These two propositions jointly imply that donor spending will crowd out politician spending, which is our main hypothesis.¹¹

H1. When politicians learn about foreign aid spending, they will be less likely to allocate goods to schools that benefit from that aid (spatial crowding out effect).

This model also has implications for the type of schools and communities which are likely to reap benefits or bear costs as a result of crowding out. First, because donors target high d schools and increase a politician's incentives to maximize v rather than d , foreign aid will, on average, cause the benefits from a to flow to less needy schools.¹²

H2. When politicians learn about foreign aid spending, they will be less likely to spend on the neediest schools.

¹⁰ That is $\frac{\partial d(a)}{\partial b} < \frac{\partial v(a)}{\partial b}$.

¹¹ H1 is labeled HD.3 in the Pre-Analysis Plan and SI.

¹² While H2 and H3 follow from our theory, they were not included in our pre-analysis plan.

A second implication is that donor spending will increase a politician's incentives to target communities likely to vote for her in the election. This follows from our proposition that aid will increase the politician's incentives to maximize v relative to d . This is true regardless of whether politicians are able to take credit for donor investments, or whether donors take measures to prevent the political capture of development projects.

H3. When politicians learn about foreign aid spending, they will be more likely to spend on pivotal voters.

This model also illustrates how the development implications of crowding out depend upon a politician's preferences. **H1** implies that efforts by donors to target specific schools with development will often be nullified by the redirection of public spending away from these schools. Note, for instance, that school $i = 1$ will often be better off without donor funds since $b_{i=1}$ will crowd out $a_{i=1}$ even if $a_{i=1} > b_{i=1}$. However, the overall consequences for development are ambiguous and depend upon how a politician chooses to reallocate a . If a politician weights d over v in her utility, then crowding out will make it more likely that a politician reallocates from $i = 1$ to $i = 2$. In this case, we would often see a net increase in welfare for the set of schools as a whole relative to a counter-factual case in which aid did not cause crowding out. If, on the other hand, aid causes a politician to weight v over d , then aid is more likely to result in an investment in a high v but low d school, driving down welfare.

The distributional effects of crowding out in **H2** and **H3** are also potentially conditional. We assume in the discussion above that donors always target needy schools. However evidence suggests that this might not always be a reasonable assumption (Briggs, 2017; Jablonski, 2014). If instead donors fail to target based on development needs then, paradoxically, this will often incentivize a politician to place more weight on development needs. To see this, suppose that instead of targeting the neediest school, the donor instead invested $b_{i=2} > 0$. By decreasing a politician's utility from a_2 , this investment will increase rather than decrease a politician's incentives to invest $a_{i=1} > 0$.¹³ As we discuss in the conclusion, there are good reasons to believe that politicians do not view foreign aid as being particularly effectively allocated and often view their decisions as addressing inequities in donor allocation.¹⁴ This is one plausible reason why we fail to see evidence consistent with **H2** and **H3**.

Finally, we also consider alternative hypotheses and assumptions. First, there might be advantages to politicians in mimicking the spending behavior of donors due to flypaper effects (Hines and Thaler, 1995; Remmer, 2004). We refer to this alternative as a "validation effect".¹⁵ As we show below, there is no evidence of a validation effect.¹⁶ Similarly, we evaluated whether spatial crowding out is conditioned on school need and voting patterns.¹⁷ We fail to find evidence of these conditional effects as well.

4. Research context

Understanding how donor choices regarding project placement affect public spending is particularly important in Malawi. Malawi is among the most aid dependent countries in the world, with aid representing over 37% of the government's budget and an even larger proportion of overall development allocations (World Bank, 2019). In addition to providing budget support to local government, between 2011

and 2016, donors directly funded projects in approximately 34% of primary schools, which is roughly comparable to the percent of schools (38%) who received projects funded by the local government.¹⁸

4.1. Local government in Malawi

Within Malawi, our experiment takes place at the local government level. The *de jure* decision-making body within the local government is the District Council.¹⁹ Councils have an average budget of approximately US\$5 million, 11% of which is dedicated to education.²⁰

Elected local councillors (LCs) are the voting members on the councils. They are elected in single member constituencies (wards) every five years.²¹ Local elections were held in May 2014, and 462 LCs were elected. Out of these 462 LCs, 335 participated in our experiment.

Also in 2014, 197 members of parliament (MPs) were elected in single member constituencies.²² Out of these 197 MPs, 125 participated in our experiment. MPs are not voting members of local councils though they are influential in allocation decisions at the local level. For example, one LC discussed a time when he mobilized the community to make bricks for a community hall, and then the MP "convinced the Council to change the project and use the bricks to construct [a] girls' hostel at another school...and the bricks were moved and used on [that other] project."

4.2. Local government and development decision-making in Malawi

Regardless of the Council dynamics, every interviewed MP, LC, and District Commissioner agreed that a primary role of the Council and all elected politicians is to allocate development projects in the area, both from government coffers and in partnership with donors. For example, one LC said his primary responsibility is to, "[take] the development from the District [Council] to the ward when there is money from local development funds and other donors from different organizations." One MP stated a primary role of his position is to "attend full Council meetings [that] concentrate on projects which can develop the district."

The potential for politicization in development decisions is widely acknowledged by the citizens of Malawi. We probed the perceptions of Malawian citizens in a series of focus group discussions. One participant said, "Most politicians choose development to where they get more votes, in order to punish those who didn't vote for him." In a more positive spin on this same phenomenon, a participant in another focus group discussion said, "Most politicians want to appreciate the people who voted for him."

4.3. Relationship between the government and donors and NGOs

The experiment we conduct mimics the way elected politicians make decisions about NGO-funded projects. In the interviews we conducted, almost all politicians mentioned working with NGOs. In fact, when asked to cite an example of a development project the elected politician brought to his or her constituency, most mentioned a project that was implemented (and funded) in partnership with an NGO, rather than one

¹³ Politician utility from a_2 equals $\lambda d_2(a) + (1 - \lambda)v_2(a)$. This quantity is strictly decreasing in b_2 .

¹⁴ See also Section B.3 in the SI for an empirical assessment of donor spending.

¹⁵ Referred to as HD.1 in our Pre-Analysis Plan and SI.

¹⁶ We also predicted that this validation effect would be particularly strong when politicians interacted frequently with donors (HD.2 in the pre-analysis plan). We fail to find evidence that this is the case.

¹⁷ Referred to as HD.4 and HD.5 in the pre-analysis plan.

¹⁸ These statistics are based on a survey of teachers in 311 schools across Malawi. Project-level data we collected from a subset of donors suggests that 57% of primary schools had at least one donor-funded project in 2011–2016. The discrepancy in these figures is likely due to information gaps among teachers regarding project funding sources.

¹⁹ Within urban areas, these are called "town councils" or "city councils."

²⁰ These statistics are based on 2011–2012 budgets, the most recent data available. An exchange rate of MK700 = US\$1 was used.

²¹ There are 462 wards in Malawi. On average, they are about 180 square kilometers in size and have about 15 primary schools.

²² There are 197 constituencies in Malawi. On average, they are about 430 square kilometers in size and have about 30 primary schools.

implemented directly by the government. As one MP said, “Sometimes the NGO goes to the District Commissioner and gets [my] phone number and the NGO calls me directly and talk about development.” A LC similarly stated, “Almost [all] NGOs, when they want to introduce a project in my area, these NGOs they do approach me first.”

Even though these projects are funded by NGOs, the allocation of these projects is nonetheless often politicized. Elected politicians will claim credit for projects funded by NGOs in discussions with constituents and other stakeholders. In turn, their constituents give them credit for projects initiated in the ward under their tenure, regardless of the ultimate funding source. For example, in a survey we conducted among 164 head teachers at primary schools across Malawi just prior to the experiment, 27% could identify a particular project completed at their school that they attributed to the LC. Out of these, the majority (71%) were projects that could have been funded by either a NGO or government funds (or by the LC personally). Only four percent were identified as government-funded projects, and 24% were identified as projects funded by non-governmental organizations. This demonstrates that elected politicians are given credit for development projects funded and executed by non-governmental organizations within their constituencies and, as such, are incentivized to carefully consider the allocation of these projects.

Yet, while politicians value their connections with NGOs, they are also often frustrated by the frequency with which donors—particularly larger international donors—fail to consider local development priorities. As one District Commissioner said, “I would not say the relationship [with NGOs and donors] is productive. They come to fund their own projects, not projects that the Council wants. What we would have loved is for them to come and look at the Council’s District Development Plan. From the Plan, look at the needs and priorities. The problem with NGOs is that they are accountable to their donors, not the partners they work with.”

Local government officials widely agree that a primary issue in the relationship with donors is that donors ask for only limited input from government officials, and often late in the decision process. As one District Commissioner said, “They don’t consult; they come with already framed projects. They come when they have already made a decision. In actual sense the [Council] is there just to endorse what they have already planned.” In general, our interviews clearly portrayed a pattern of donors consulting local government officials on a limited basis, and local government officials perceiving that donor development projects are not allocated optimally as a result.

5. Research design

5.1. Overview

This research is based on a field experiment conducted among 125 in-office MPs and 335 in-office LCs in Malawi, or 63% and 73% of each theoretical population, respectively. In addition to the experiment, we conducted 32 semi-structured interviews with LCs, MPs, District Commissioners, and Area Development Committees, as well as four focus group discussions with Malawian citizens. These interviews and focus group discussions asked questions about decision-making, transparency, accountability, and relationships across government stakeholders and donors. Finally, we also conducted a survey among 2000 citizens and teachers across 311 schools in 60 of the 462 wards in Malawi. The survey asked questions about local school conditions and perceptions of government and donor performance.

5.2. Experiment design

In order to evaluate how politicians make decisions about public spending and the allocation of goods, we conducted an experiment among LCs and MPs. The experiment was conducted in Malawi between March and June 2016. In partnership with a UK-based NGO (Tearfund),

we offered participants the opportunity to choose schools in their ward or constituency to be eligible to receive school supplies. In face-to-face interactions with trained Malawian RAs, each politician was presented with a map that included three schools from their constituency. The three schools that appeared on the map were randomly selected from a comprehensive list of primary schools in the politician’s constituency. The politician was then asked to determine which of the three schools should receive an education good. Specifically, the survey asked “When you are ready, please tell me which school you would like to choose to receive a set of [school supply]. Please take your time in making this decision.” The maps, examples of which are shown in Fig. 1 below, were shown to the politician on portable tablets, and could be studied by him or her in detail before each allocation decision was made.²³

Each politician repeated this process three times, so they consecutively selected three schools out of nine to receive development goods. Each decision involved the allocation of a different kind of good—either solar lamps, teacher supply kits, or English dictionaries. The order of goods being allocated was randomly assigned. The goods being allocated in the experiment were chosen in consultation with teachers and civil society members, and are goods that are both highly desired and needed in most communities.

Significantly, these were not hypothetical decisions. Following the experiment, the three schools chosen by each politicians were entered into a public lottery. Approximately 20% of the selected schools were chosen in this lottery to receive goods. The details of the lottery were provided to each politician before they made the allocation decision, making the decision costly and meaningful. Our discussions with project stakeholders, as well as repeated follow-up requests by the participating LCs and MPs, indicated that the allocated goods were highly valued by both politicians and schools.

To evaluate how information about foreign aid projects influenced the politicians’ allocation decisions, we randomly assigned an aid information treatment at the map level. The aid information treatment provided information about development projects supported by major international donors in Malawi in the past five years (since 2011) at each school. Specifically, the information detailed how many donor-supported projects had been carried out at each individual school, and the type of donor support provided at that school. This aid information treatment was randomly assigned at the map level, and within respondent-level blocks. An example map for the treatment group is shown in Fig. 1, Panel A. It contains a side panel with information about foreign aid projects. An example control map is shown in Fig. 1, Panel B. It contains no information about foreign aid projects.

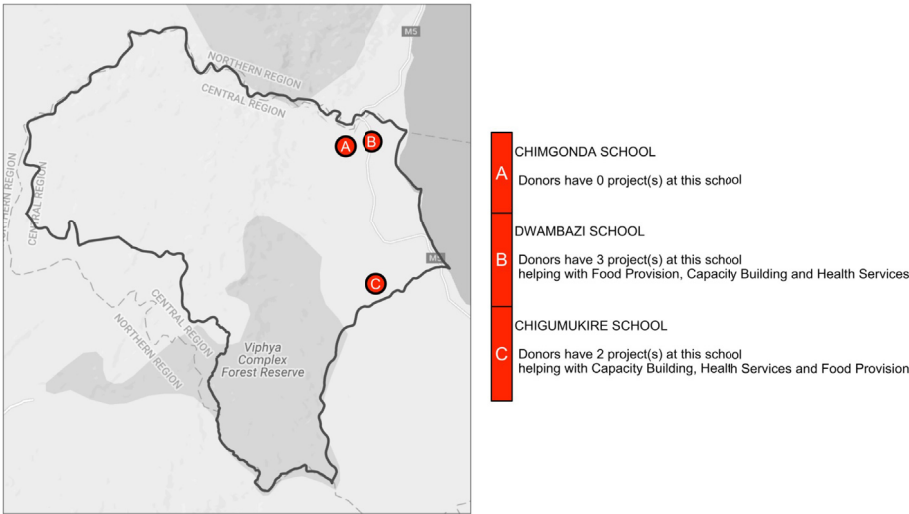
Our expectation is that the aid information treatment will cause politicians to update their priors about the number and type of foreign aid projects in the mapped schools. As a consequence, by comparing politicians’ decisions in response to treatment versus control maps, we can isolate the causal effect of learning about foreign aid projects on the allocation decisions of politicians. Our primary estimand is therefore an intention to treat (ITT) effect in which we assume that politicians update their priors about foreign aid projects in the expected direction in response to information.

We provide evidence below to support this assumption for most respondents. However, since not all respondents will update in the same way, these ITT estimates may not effectively characterize the effect of learning about foreign aid projects. In order to derive more realistic estimates of learning, we also took steps to measure politicians’ pre-existing knowledge of foreign aid projects in their constituency. Following the decisions, we provided all politicians with a random list of three school in their constituency which were not used as part of the experiment. We then asked the politician a series of questions about how many for-

²³ This experimental design is similar to those used in the choice experiment literature to model consumer behavior. For reviews in health and ecological economics, see Clark et al. (2014) and Hoyos (2010).

A. Treatment Map

Schools in Your Ward



B. Control Map

Schools in Your Ward



Fig. 1. Treatment and control maps.

eign aid projects were in each school, and which donors were involved. Politicians with low pre-existing knowledge about foreign aid projects should be more likely to update their beliefs in response to the aid information treatment, and thus—if our results are a result of information updating—we should see stronger effects among this subset of respondents.

A related compliance concern is that some politicians may misinterpret the information on the map. To ameliorate this concern, prior to the experiment, we also provided a training map to all politicians in order to assist them in understanding the information provided via

the treatment. This map depicted schools outside Malawi and provided hypothetical information about school uniforms. Politicians were only permitted to continue the survey once they had demonstrated that they could correctly answer questions about the map. This training exercise appears to have been effective and we were able to identify no difference in the response to the aid information treatment among those who initially had difficulty in understanding the training map.

5.2.1. Other treatments in broader experiment

The aid information treatment was randomized and delivered in the

context of a broader multi-arm factorial experiment.²⁴ Due to the factorial design, all treatments are orthogonal to each other, enabling independent analysis of each treatment separately.

In the broader experiment, we also evaluated the effect of providing information about school need and political support in a full factorial design. Therefore, each map displayed either individual information treatments, a combination of several information treatments, or no information treatments at all. The school need information treatment took the form of Government of Malawi data on school characteristics (see Section 5.4.2). The political support information treatment took the form of vote share of the MP/LC at the nearest polling station in the previous election (see Section 5.4.3). As specified in our pre-analysis plan, we expected information about political support and need to increase allocation to politically pivotal and needy schools.

Prior to providing any of these information treatments, we randomly assigned two transparency treatments. Specifically, we told a random sample of politicians that his/her allocation decisions would either be announced on local radio or in a report to donors.²⁵ We predicted that sharing decisions with donors might encourage politicians to align their decisions with donor preferences. We test this claim in SI Section B.2. We find no evidence that politicians respond to aid information differently when their decisions are more transparent.

The randomisation proceeded so that each sampled politician was first assigned into one of four transparency treatment combinations. Each politician then received three maps, each containing one out of eight potential information treatment combinations. Altogether, therefore, there were 32 different possible combinations of transparency and information treatments. In Section D.1 of the SI, we provide further descriptions of the different treatment arms, examples of maps with different combinations of the information treatments (Figs. D.1–D.8), and a table delineating the number of maps that received different combinations of the transparency and information treatments (Table D.1).

5.3. Sampling and administering of survey

Out of 655 sitting MPs and LCs, we sampled 82%.²⁶ Out of these, we had a total response rate of 85.2%: 94.9% for LCs and 66.8% for MPs. The lower response rate among MPs was primarily because many were unavailable due to travel or legislative commitments. None of the subjects that we did reach refused to participate. Balance and attrition tests included in the SI indicate that the final sample is reasonably representative of politicians in Malawi.²⁷ There is no significant difference in attrition across treatment conditions. The sampled politicians are also well distributed geographically across Malawi, as shown in Fig. 2.

In total, the sampled MPs were provided with a total of 370 school maps.²⁸ Out of these, 179 (48.4%) included information about foreign aid projects. The 335 sampled LCs were provided with a total of 882

maps. Out of these, 442 (50.11%) included information about foreign aid projects.

The survey was carried out by a team of trained Malawian research assistants. Interviews were typically conducted in the home constituency of the politicians or in the capital city of Lilongwe. All interviews were conducted in English, which is the language of official business in Malawi.²⁹ We provide an example survey in the SI Section D.

5.4. Data

To carry out the experimental design, we collected school-level data on aid, need, and political support at each school, as well as assessed the politicians' prior knowledge on these topics. We also collected a broad range of other data on individuals, schools, and ward/constituencies.

5.4.1. Aid information

To collect information on foreign aid used for the aid information treatment, we focused the data collection on the main foreign donors active in the primary education sector in Malawi, and the projects these donors had carried out in individual primary schools in the past five years (since 2011). Following consultations with local stakeholders and practitioners active within the aid sector in Malawi, we identified the main donors whose project activities included the primary education sector. When approaching each of these donors, we asked them to provide detailed data on their project activities since 2011, including the type of intervention and the name and location of the recipient school. Donors were also asked to cross-validate our list of active donors in the sector, and to suggest organizations that were not on the list.³⁰ As we discuss below, the politicians in our experiment had little or no knowledge about most of these foreign aid projects, and were not involved in their allocation.³¹

In total, 3151 primary schools received 4566 foreign aid projects from this set of donors between 2011 and 2016. This constitutes 57% of the 5438 primary schools in Malawi for which we had location data. The number of foreign aid projects in each school varied from 0 to 4. Fig. 3 displays the total sample of primary schools in Malawi with no projects supported by donors (in grey) versus those with at least one project (in blue).

The total number of primary schools that were included in the school maps presented to the 125 sampled MPs was 1109. Of these, 683 (62.03%) contained at least one foreign aid project. The average number of projects per school was 0.95, ranging from 0 to 4. For the 335 sampled LCs, the total number of primary schools presented in the maps was 2,646, of which 1545 (58.39%) contained at least one foreign aid project. The average number of projects was 0.88, again ranging from 0 to 4.

²⁹ RAs were trained to clarify terms in the respondent's local language.

³⁰ The organizations from which data on aid projects were obtained include Department for International Development (DFID), Deutsche Gesellschaft für Internationale Zusammenarbeit (GIZ), German Development Cooperation (KfW), Norwegian Embassy, Save the Children, United States Agency for International Development (USAID), United Nations Children's Fund (UNICEF), United Nations Development Programme (UNDP), Volunteer Service Overseas (VSO), World Food Programme (WFP), and the World Bank. Organizations that were identified as active in the education sector, but that failed to respond to our queries, include Japan International Cooperation Agency (JICA), OXFAM, United Nations Population Fund (UNFPA), and World Vision.

³¹ Like most of the Malawian aid portfolio for these donors, these education projects were almost entirely off-budget and implemented by donors or non-governmental implementing partners. Government ministries were consulted on some projects. However, we could find no evidence that council authorities or parliamentary representatives in benefiting constituencies had influence or insight into the process of allocating these projects. In the SI, we plot the plot the characteristics of schools associated with donor spending and find no significant association with the political characteristics of communities.

²⁴ Due to space constraints, we have focused this manuscript on the results surrounding the aid information treatment, though we present the full set of pre-specified analyses for all information treatments in the SI.

²⁵ A few months after the study, a report was delivered to donors and a radio script was broadcast on Zodiak radio about the decisions of the politicians randomly assigned to these transparency treatments.

²⁶ The main reasons for excluding some politicians from the sample were the unavailability of electoral data or data on school need, or because the number of schools in their ward or constituency was not sufficient for carrying out the experiment.

²⁷ Since we exclude wards with very few primary schools, our sample is biased towards wards with a greater number of schools relative to enrollment.

²⁸ As a rule, each politician was provided with three school maps and one test map, each containing 3 schools. However, in a few rare cases this was not possible, due to limited numbers of schools within sampled constituencies. Therefore a very small number of the politicians participating in the survey only received two or even one school map.

A. Constituencies

B. Wards

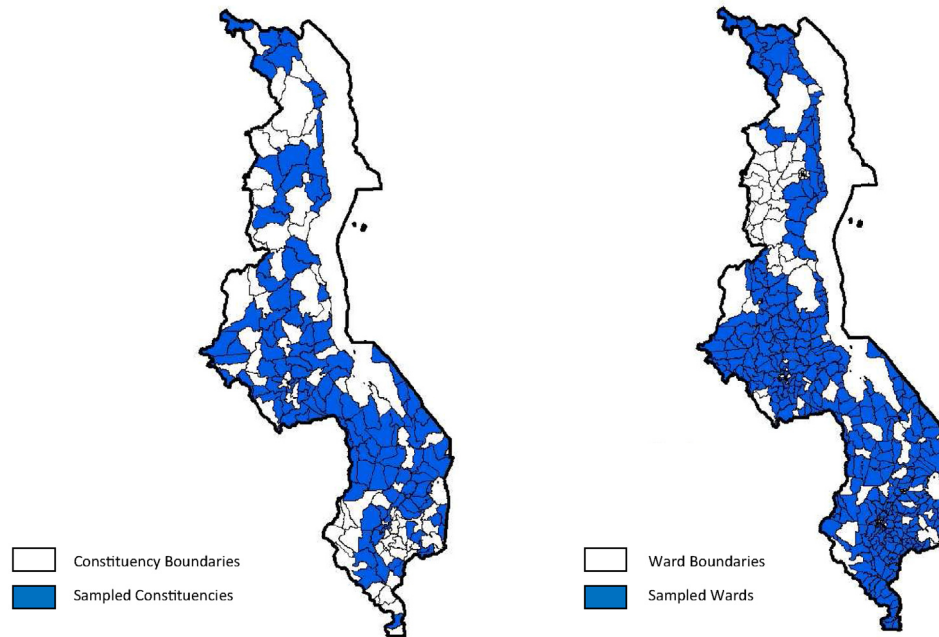


Fig. 2. Sampled constituencies and wards.

We create two variables to measure aid at the school level: *Aid Project Count* equalling the $\log(+1)$ of the number of projects in each school, and *Aid Good Types* equalling the $\log(+1)$ of the number of project types. Aid project types include capacity building, construction, health services, food provision, community support, gender issues, and teacher training. Some aid projects encapsulate several project types. Since politicians might care both about the number of foreign aid projects and the scale of donors' involvement in a school, and since both types of information were provided in the aid information treatment, we predicted that both variables would have similar effects on spending outcomes. Seventy-three (73%) of our treatment maps contained variation across schools in the number of foreign aid projects.

5.4.2. School need information

In order to test H2, we require information about the level of need at a school. For this we rely on official school-level statistics from the Education Management Information System (EMIS) at the Malawi Ministry of Education Science and Technology. These data are from 2014 and encompass over 99% of all schools in Malawi. They are collected approximately biannually by district education offices through the support of local headmasters. These data have been collected and refined over multiple years and independent assessment exercises on these data suggest a high level of reliability (Bernbaum and Moses, 2011).

Though not an exhaustive assessment of school need, these data allow us to measure three highly visible characteristics of need. First, we measure structural overcrowding using the ratio of students per classroom. Structural overcrowding is among the more severe problems facing schools in Malawi: on average, primary school classrooms have 138 students each, though some have more than 300. Second, we measure teacher overcrowding using the number of students per teacher. Due to chronic problems of low or unpaid salaries, teachers in Malawi are often heavily over-committed and underpaid. Primary school teachers are expected to teach 75 students on average, though some have more than 200. (The global average is 23 students per teacher; World Bank, 2019.) Third, we measure the quality of existing classrooms by looking at the ratio of temporary classrooms to permanent classrooms.

The quality of temporary classrooms vary in Malawi, but they are often of extremely poor quality—sometimes a lean-to or a borrowed residence.

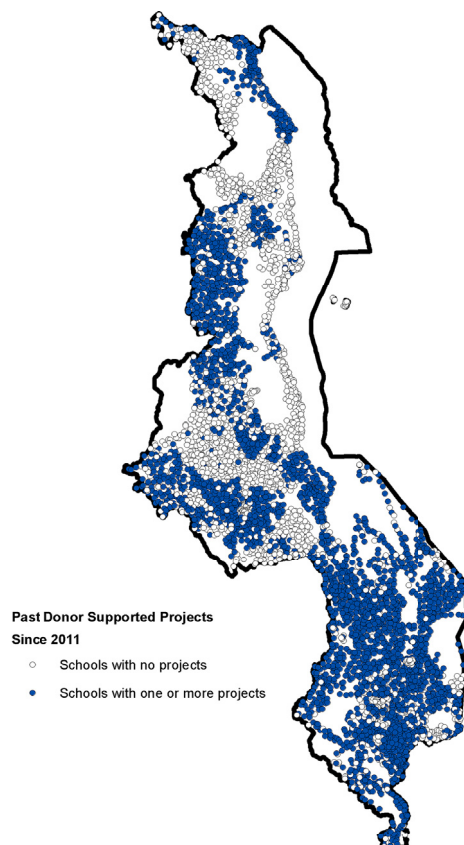
These measures generally align with the priorities of teachers themselves. In our survey of teachers, we asked head teachers to name, in order of priority, the important needs of the school. The highest priority issues by far (named by over 60% of head teachers and citizens) were overcrowding in classrooms or teacher houses. Teachers also frequently mentioned needing more staff, various facility improvements including electricity, and learning materials. Additionally, in our interviews with politicians about their development decisions in the education sector, they most frequently mentioned enrollment levels, the number of classrooms, and the number of teachers houses. That said, there are some need-based characteristics that these data do not capture: for instance, several politicians also mentioned that they use measures of school quality and achievement, such as the passing rate, or that they simply examine the “look of the infrastructure,” or “just see the nature of the school”.

In our analysis, we analyze the effects of each of these components of need separately. In addition, as specified in our pre-analysis plan, we create an overall index, *School Need*, which is equal to the sum of the z-scores of the three measures of school need.³²

5.4.3. Political support information

In order to measure the political characteristics of communities, we collected polling station level data from the Malawi Electoral Commission on the votes received by all candidates for LC and MP seats. A large proportion (68%) of the schools in our sample were also polling stations, allowing us to directly measure political support in those communities. For those schools in our sample (32%) which were not used

³² $SchoolNeed = \frac{x - \mu_1}{\sigma_1} + \frac{x - \mu_2}{\sigma_2} + \frac{x - \mu_3}{\sigma_3}$ where μ_i and σ_i indicate the within ward/constituency means and standard deviations of students per teacher, students per classroom, and proportion of temporary classrooms for all available primary schools in Malawi.



Note: The map displays 5,438 primary schools in Malawi, differentiating between the schools that received no foreign aid projects between 2011–2016 versus those that received at least one in that period. The data are based on interviews with major donors in Malawi.

Fig. 3. Primary schools and foreign aid projects.

as polling stations, we measure political support by using the geographically nearest polling station to the school.

One challenge we face is how to determine whether an allocation decision was politically motivated. There is considerable debate about which citizens are most likely to be targeted by pork barrel or clientelistic allocations, and studies suggest that such strategies are contingent on the social and institutional environments faced by politicians (e.g., Stokes et al., 2013). Despite this diversity in strategy, most studies in multi-ethnic developing democracies like Malawi conclude that pork barrel allocations will be targeted to core supporters of politicians due to politicians' greater ability to organize voting and turnout in communities where they have pre-existing social or ethnic ties, clientelistic networks, or information (Jablonski, 2014; Stokes et al., 2013; Keefer and Vlaicu, 2008; Burgess et al., 2015; Kramon and Posner, 2013).

Building on the literature and our field research, our pre-specified expectation is that electorally motivated politicians would prefer to target allocations in communities where they received a large proportion of votes in the last election. We call this variable *Incumbent Percent*, which equals the percentage of votes received by the incumbent politician in the nearest polling station to a school. Additionally, we coded whether a politician's family member attended a particular school, since these schools might be particularly likely to benefit from networks of patronage. We coded this by asking the politicians to indicate which

schools their children or their family's children attended.³³ In the SI we re-estimate our results using alternative measures of political support.³⁴

5.4.4. Respondent priors

To assess the politicians' prior knowledge of donors and foreign aid projects, we conducted surveys of all respondents, testing their ability to describe characteristics of schools in their ward or constituency. This was done by presenting the politician with a map of their ward/constituency containing three schools and asking which schools on the map received projects sponsored by large donors in the past five years. In addition, they were also asked to list the name(s) of any major donor(s) that had supported a project in one of the schools shown on the map. For the first question regarding which schools had received

³³ This question was asked after the assignment of treatment, raising potential concerns about post-treatment bias. However we see no indication that politicians assigned to the aid information treatment were more or less likely to respond to this question ($p = 0.67$). Nor do we find that politicians are less likely to name schools within treatment maps ($p = 0.88$).

³⁴ We lack sufficient data to test whether politicians are targeting co-ethnic voters. However, we expect little co-ethnic targeting in this context. Ethnicity typically does not vary extensively within MP or LC constituencies, and competing candidates for elected office often share ethnicity. Seventy-eight percent (78%) of LCs indicated to us that their ward consists primarily of one ethnic group.

Table 1
Sample statistics.

Variable	All	Treatment	Control	LCs	MPs
Aid Project Count	0.532 (0.478)	0.542 (0.481)	0.526 (0.476)	0.521 (0.477)	0.558 (0.479)
Aid Good Types	0.723 (0.675)	0.741 (0.677)	0.713 (0.674)	0.699 (0.668)	0.783 (0.688)
School Need Index	−0.028 (1.813)	0.038 (1.828)	−0.067 (1.803)	−0.015 (−1.806)	−0.059 (−1.831)
Incumbent Percent	0.469 (0.218)	0.47 (0.218)	0.469 (0.218)	0.492 (0.215)	0.416 (0.215)
Family Attends School	0.062 (0.242)	0.063 (0.243)	0.062 (0.241)	0.077 (0.266)	0.028 (0.165)
Log Enrollment	6.131 (1.524)	6.145 (1.521)	6.123 (1.527)	6.12 (1.544)	6.158 (1.475)
Pop Density at School	9.712 (19.6)	9.184 (13.611)	10.03 (22.446)	9.786 (16.697)	9.519 (25.719)
Knowledge of Donors	0.124 (0.224)	0.12 (0.223)	0.126 (0.224)	0.122 (0.223)	0.127 (0.227)

Note: Table shows sample means for each variable and sub-group with standard deviations in parentheses.

most aid, approximately 24% of MPs and 21% of LCs were able to provide a correct answer: a rate which is only slightly better than random chance.³⁵ When it came to the latter question about which donors had provided support, less than 4% of the respondents were able to provide a correct answer. This confirms that the information we provided was novel: politicians started the interaction with little knowledge of the distribution of foreign aid projects in their constituency, and almost no knowledge of which major donors had provided projects.³⁶ This relatively low level of knowledge is perhaps reflective of the fact that most of the respondents had only been in office for 32 years at the time of the survey, though also reflects the low level of transparency around many donor initiatives in Malawi.³⁷ We plot responses to these knowledge questions in the SI.

We briefly summarize the main analysis variables in Table 1 and include more detailed summary, attrition and sample selection statistics in the SI.

6. Estimation

We are interested in the odds that a school is selected in each of a respondent's three choice sets (maps), and seek to estimate how these odds differ conditional on the characteristics of the school and the treatment assignment. We estimate these odds using a conditional logit (fixed effects) estimator conditioned on each choice set. Formally, let Y_{nsi} indicate whether politician n chooses school i in map choice set s . Let z_{is} be the variables specific to a school i , such as whether previous donor projects have been carried out there. We can represent the probability of selecting a given school in a set s conditional on z_{is} using the conditional logit specification in equation (2).³⁸

$$P(Y_{nsi} = 1 | z_{is}) = \frac{e^{\beta z_{is}}}{\sum_{j=1}^J e^{\beta z_{js}}} \text{ for } j = 1, 2, 3 \quad (2)$$

³⁵ Respondents had the choice of selecting school A, B, or C, all schools, or no schools. Additionally they could say they didn't know.

³⁶ These questions were asked post-treatment so one might worry that respondents were strategic in their response to this question. However, we find that responses to this question are indistinguishable between treatment and control ($p = 0.80$).

³⁷ This low level of knowledge about foreign aid projects is also consistent with Baldwin and Winters (2018) who show that less than 4% of citizens in Uganda can identify the donor of a local development project.

³⁸ The conditional logit specification has the disadvantage of assuming independence of irrelevant alternatives and having limited flexibility in modelling heterogeneity across respondents. In the SI we also show consistent results using a mixed logit specification, which extends the conditional logit probability by allowing β to vary across respondents.

We are primarily interested in evaluating how the effects of z_{is} vary with the treatment assignment. Let $t_s \in [0, 1]$ be our randomly assigned treatment of information at the map level. Our treatment equals one if map s has been assigned to a treatment group and zero if it is in a control group. To estimate the conditional effects of z_{is} , we interact t_s with z_{is} as in equation (3):

$$P(Y_{nsi} = 1) = \phi(\beta_1 z_i + \beta_2 t_s z_i + \gamma X_{is} + e_{nsi}) \quad (3)$$

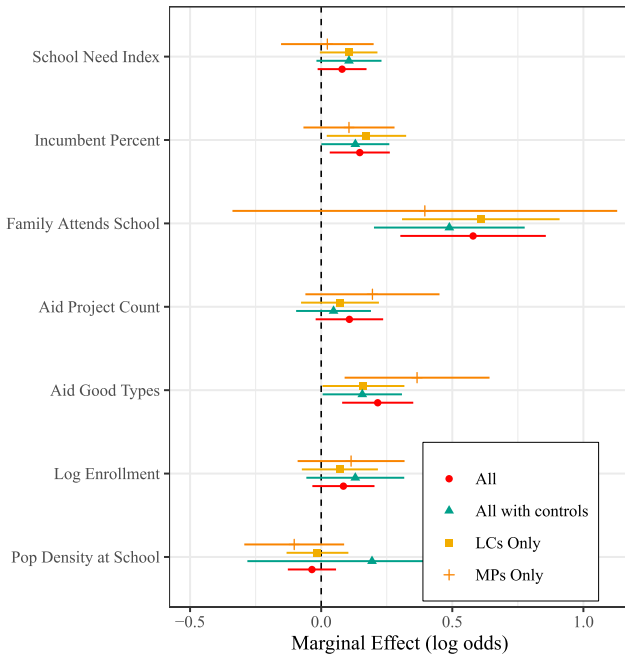
Where ϕ is the conditional logit estimator in equation (2). X_i is a vector of control variables which are specific to a school, or an interaction of respondent and school-specific variables. Note that since this is a within choice set estimator, the coefficient on t_s is not estimated.

We include estimates both with and without control variables for all our models. Our pre-specified control variables, which vary at the school level, include *Log Permanent Classrooms*, *Log Temporary Classrooms*, *Log Teacher Houses Permanent*, *Log Teacher Houses Temporary*, *Opposition Percent Votes (for MP and LC)*, *Log Enrollment*, *Number of Aid Projects*, *Family Attends School*, *Incumbent Percent at Polling Station*, and *School Need Index*. Summary statistics and coding details for these variables are provided in the SI.³⁹ Our primary interest is in β_2 which tells us the difference in the effects of z_i in the treatment group relative to the control group. We cluster our errors at the respondent level.

We include two measures of z_i . First, we use *Aid Project Count*, which is the log(+1) of the number of projects in each school. Second, we use *Aid Good Types* which is the log(+1) of the number of project types. Aid project types include capacity building, construction, health services, food provision, community support, gender issues, and teacher training. Some aid projects encapsulate several project types. This second measure is intended to capture the scale of the projects, and we expect it to have similar effects as the first measure on respondent incentives.

We also anticipate that the effect of providing information about aid projects may vary with the pre-existing knowledge of each politician about the school. As discussed, we collected information about the pre-existing knowledge a politician had about foreign aid projects at a random selection of schools in the politician's ward or constituency not used in the experiment. We expect this to be a reasonable proxy for the amount of information about foreign aid projects held by the politician prior to receiving the treatment. Let k_n be the level of information about foreign aid projects and donors held by politician n about these three schools. We can then estimate how the effect of t_s varies with k_n using equation (4):

³⁹ Missing data in control variables are imputed as specified in the pre-analysis plan using the mean value for the lowest level of aggregation available (map, ward or district).



Note: This figure shows the coefficients of separate conditional logistic regressions of school selection on baseline variables. The sample is limited to maps that do not contain treatment information related to the school characteristic. 95% confidence intervals are shown in the horizontal lines. Standard errors are clustered on politician. Continuous variables are normalized for comparison purposes. Control variables include *Log Permanent Classrooms*, *Log Temporary Classrooms*, *Log Teacher Houses Permanent*, *Log Teacher Houses Temporary*, *Opposition Percent Votes* (for MP and LC), *Log Enrollment*, *Number of Aid Projects*, *Family Attends School*, *Incumbent Percent at Polling Station*, and *School Need Index*.

Fig. 4. The effects of school characteristics on school selection.

$$P(Y_{nsi} = 1) = \phi(\beta_1 z_i + \beta_2 t_s z_i + \beta_3 z_i k_n + \beta_4 t_s z_i k_n + \gamma X_{is} + e_{nsi}) \quad (4)$$

We estimate other heterogeneous effects in a similar fashion.

7. Results

7.1. Odds of school selection in control group

We begin by estimating the odds in the control group that a politician selects a school in a map based upon the characteristics of that school. The results in Fig. 4 are broadly consistent with other qualitative and quantitative evidence on the distribution of public resources in Malawi and elsewhere. They are also consistent with our pre-specified priors. Both need and political support appear to play a role in politicians' allocation decisions. A one standard deviation increase in a school's need index increases the odds that a school is selected by 1.08. A one standard deviation increase in a school's percentage of votes for the politician increases the odds of selection by 1.15. Politicians are also much more likely to select schools where they have family members attending, even controlling for other factors. Such a school is 1.8 times more likely to be selected.

Also worth noting is that schools in the control group are more likely to be selected when they have received more projects or types of projects from donors. Given the weak pre-existing knowledge surrounding foreign aid projects, we do not interpret this to mean that politicians in the control group target schools with more aid spending; rather, we see this as evidence that donors and politicians often have

similar preferences.⁴⁰

7.2. Average aid information treatment effects

We next consider how the odds of school selection vary between treatment and control groups. We first evaluate in Table 2 whether the aid information treatment causes politicians to be more or less likely to select schools with existing foreign aid projects. On average, receiving information about foreign aid projects decreases the odds of a school with one foreign aid project being selected by 0.26 ($p = 0.055$). (On average, schools have 0.9 aid projects.) We also see an insignificant and smaller effect size among MPs compared to LCs.⁴¹

We next evaluate whether the odds of school selection vary depending upon how many categories of goods have been delivered by donors to a school. The estimates in Table 3 suggest that when politicians learn from the aid information treatment that there are three categories of goods being delivered by donors at a school (the average is 2.6), the odds of that school being selected decrease by 0.42 ($p = 0.02$) on average.

Since we conduct multiple tests of our main hypotheses, it is possible that the uncorrected p-values overstate the evidence against the null hypotheses. In Tables 2 and 3 we include p-values accounting for the false discovery rate using the procedure specified in Benjamini and Hochberg (1995). We provide estimates controlling both for the two tests of H1 and for tests of all pre-registered hypotheses.⁴² In the SI, we provide further details about this procedure and show estimates of type one error rates under the sharp global null hypotheses that all treatments in the study as a whole had zero effect.

These effects are large and substantively important. Our estimates suggest that, in a world where politicians are fully informed, each additional aid project from a major donor in a school would displace about 18% of public discretionary projects in that school, or 22–29% among those politicians for whom this information is more likely to be novel (see discussion below). Given that approximately 57% of schools in Malawi benefit from some major donor's foreign aid, this represents a substantial potential redistribution of resources.

7.3. Comparing across information treatments

In Table 4, we compare the effects of providing information about aid to the effects of the other two information treatment conditions in the experiment (information about need and political support). As discussed above in Section 5.2.1, these other two information treatments provided information about the level of need at the school and the percentage of votes the respondent received in the school's community.

Table 4 shows that the effects of these other treatments are smaller and more consistent with the null hypotheses, though there is a small significant positive effect of providing information about school need. One possible reason for the weaker effects of the other information treatments is that respondents find it easier to learn about the needs and political characteristics of schools, whereas information about donor spending is less accessible. Consistent with this explanation, respondents were considerably less likely to be able to answer questions about foreign aid projects in a school compared to questions

⁴⁰ In SI Section B we compare donor and politician spending. There is a small positive correlation between donor and politician spending in the control group, though only a couple school characteristics significantly predict both donor and politician allocation decisions.

⁴¹ This may be due to the fact that LCs value this information more. We find, for instance, that 81% of LCs claim they find the information useful compared to 64% of MPs. However, these differences between offices should be interpreted with caution due to the sample size for MPs.

⁴² We had five pre-registered hypotheses testing the main effects of the aid information treatment. In the pre-analysis plan, these are labeled as HD.1, HD.2, HD.3, HD.4 and HD.5.

Table 2

The effects of information about past aid.

	All (1)	All with Controls (2)	LCs (3)	MPs (4)
Aid Treatment* Aid Project Count	−0.424* (0.235)	−0.393* (0.241)	−0.558** (0.279)	−0.114 (0.440)
Aid Project Count	0.317** (0.143)	0.178 (0.152)	0.250 (0.169)	0.486* (0.271)
Observations	3738	3728	2634	1104
R ²	0.001	0.019	0.002	0.004
Benjamini-Hochberg Adj. (H1)	p = 0.055	p = 0.086	p = 0.031	p = 0.154
Benjamini-Hochberg Adj. (all pre-registered)	p = 0.137	p = 0.215	p = 0.076	p = 0.385

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

Standard errors are clustered on politician. Control variables include *Log Permanent Classrooms*, *Log Temporary Classrooms*, *Log Teacher Houses Permanent*, *Log Teacher Houses Temporary*, *Opposition Percent Votes*, *Log Enrollment*, *Number of Aid Projects*, *Family Attends School*, *Incumbent Percent at Polling Station*, and *School Need Index*. Note that the constituent term for Aid Treatment is co-linear with the map-level fixed effect, and does not have a meaningful coefficient estimate. Benjamini-Hochberg p-values correct for multiple comparisons across the two tests of H1 (row one) and across all tests of our five pre-registered hypotheses about the aid information treatment (row two).

Table 3

The effects of information about past aid and aid categories.

	All (1)	All with Controls (2)	LCs (3)	MPs (4)
Aid Treatment* Aid Good Types	−0.398** (0.176)	−0.368** (0.180)	−0.521*** (0.210)	−0.174 (0.328)
Aid Good Types	0.322*** (0.108)	0.223** (0.115)	0.241** (0.127)	0.531*** (0.210)
Observations	3738	3728	2634	1104
R ²	0.003	0.019	0.002	0.008
Benjamini-Hochberg Adj. (H1)	p = 0.038	p = 0.073	p = 0.018	p = 0.154
Benjamini-Hochberg Adj. (all pre-registered)	p = 0.096	p = 0.182	p = 0.045	p = 0.385

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

Standard errors are clustered on politician. Control variables include *Log Permanent Classrooms*, *Log Temporary Classrooms*, *Log Teacher Houses Permanent*, *Log Teacher Houses Temporary*, *Opposition Percent Votes*, *Log Enrollment*, *Number of Aid Projects*, *Family Attends School*, *Incumbent Percent at Polling Station*, and *School Need Index*. Note that the constituent term for Aid Treatment is co-linear with the map-level fixed effect, and does not have a meaningful coefficient estimate. Benjamini-Hochberg p-values correct for multiple comparisons across the two tests of H1 (row one) and across all tests of our five pre-registered hypotheses about the aid information treatment (row two).

about school need and political support.⁴³ In SI Table B.1, we evaluate interaction effects across information treatment conditions, and in SI Tables B.2–B.4, we use several specifications to assess interaction effects across information and transparency treatments. We do not find evidence of significant interactions between the aid information treatment and other treatments in the experiment.⁴⁴

7.4. Information updating

Our theory implies that politicians should be responding to the aid information treatment because they have updated their priors about the locations of foreign aid projects in their constituency. We consider three tests of this claim. First, we estimate our treatment effects conditional on how well-informed politicians are about foreign aid projects in their constituencies prior to receiving the aid information treatment. If our treatment effects are driven by information updating, we would expect our effects to hold primarily among the subset of politicians with less pre-existing knowledge. Second, we estimate effects conditional on how frequently politicians interact with donors. The frequency of a politician's interaction with donors is another reasonable proxy for how

knowledgeable a politician is about foreign aid projects, and we would expect our effects to hold primarily among those with less frequent interactions. Third, we test whether our aid information treatment has a stronger effect among those subjects who indicated that they learned something from the information or found the information useful.

The results in Fig. 5 are generally consistent with an information updating mechanism and, as discussed below, inconsistent with social desirability bias. We only see significant spatial crowding out when politicians are less knowledgeable about foreign aid projects prior to receiving the aid information treatment, when politicians interact with donors less frequently, or when politicians indicate that the information is useful. Moreover, consistent with information updating, our treatment effect estimates approach zero among those politicians who already have considerable knowledge about foreign aid projects in their constituency or who do not claim to have learned.⁴⁵

These conditional treatment effects are potentially more credible estimates of how aid crowds out public spending, as these politicians are those most likely to have updated their priors in response to the treatment. Among those politicians who lack pre-existing knowledge of foreign aid projects in their constituencies or who indicate learning from the information treatment, the odds of selecting a school with one project decrease by 0.40 ($p = 0.04$) and 0.59 ($p = 0.10$) respectively; or 22% and 29% relative to control group means.

⁴⁵ In most cases, we cannot identify a statistically significant difference between these subgroups.

⁴³ See SI Section C.3.

⁴⁴ As highlighted in Muralidharan et al. (2019), if our assumption of zero interaction across treatments is not correct, the estimates from a non-saturated model are “a composite treatment effect that includes a weighted-average of the interactions with other treatments.”

Table 4
Comparison with other information treatments.

	All Treatments (1)	All Treatments (2)	Need Treatments (3)	Voting Treatments (4)
Aid Project Count*Aid Treatment	−0.390* (0.236)			
Aid Project Count	0.318** (0.144)			
Aid Good Types*Aid Treatment		−0.361** (0.177)		
Aid Good Types		0.319*** (0.109)		
School Need Index*Need Treatment	0.062* (0.037)	0.062* (0.037)	0.060* (0.037)	
School Need Index	0.043 (0.026)	0.043 (0.026)	0.044* (0.026)	
Incumbent Percent*Voting Treatment	0.133 (0.413)	0.125 (0.413)		0.132 (0.411)
Incumbent Percent	0.704** (0.295)	0.700** (0.295)		0.684** (0.293)
Observations	3738	3738	3738	3738
R ²	0.010	0.011	0.005	0.004

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

Standard errors are clustered on politician.

We also see evidence that the aid information treatment caused politicians to claim that they learned something at the conclusion of the experiment. Politicians who were assigned to one or more aid information treatments were 0.19 ($p = 0.01$) times more likely to claim that they “learned anything new” from the experimental interaction.⁴⁶

7.5. How spatial crowding out affects allocation decisions

We next consider how the aid information treatment changes the way in which politicians chose to distribute funds. Many worry that aid fungibility will increase the tendency for politicians to spend money on corruption or patronage, or that politicians will spend more on richer areas, making the poor increasingly dependent upon donors for their welfare. To test these claims, we interact the aid information treatment with school-level variables intended to measure corruption, patronage, and need. By comparing the effects of these variables on spending between treatment and control groups, we can determine whether the aid information treatment caused politicians to make significantly different allocation decisions.

In Table 5 and Fig. 6 we see that allocation decisions in treatment and control groups are similar. Treatment information appears to result in slightly more goods being allocated to larger and less overcrowded schools, as well as to schools where politicians’ family members attend. However these differences are small and not statistically significant. We do see evidence that politicians in the treatment group are more likely to select schools with high enrollment. This effect was not anticipated, though we think it worth further investigation. One potential reason for this effect is that politicians may wish to target areas where they can influence a larger number of voters.

As we discussed in the theoretical section, small effects are consistent with politicians trying to make effective development decisions for their constituencies rather than using their increased discretion to target pivotal voters or family members. That said, there may be small differences which cannot be identified with our sample. This is particularly true for identifying treatment effects on the selection of family members’ schools since this represents only a small proportion of schools in our sample (5%). However, overall these results suggest that arguments that crowding out will promote aid dependency among the very poorest, or facilitate patronage spending or corruption, are overstated. In the

next section we discuss qualitative evidence that politicians reallocate funds out of concern for the welfare of their constituents.

In the SI, we also consider other ways to operationalize the targeting of pivotal voters: we find no significant evidence that the aid information treatment causes politicians to target communities with a higher victory margin, fewer opposition voters, or higher turnout for the politician. We also include tests for non-linearities in these effects, as we might expect non-linear effects if politicians were targeting areas with more indifferent or “swing” voters. In each case we fail to reject the null hypothesis of no effect.

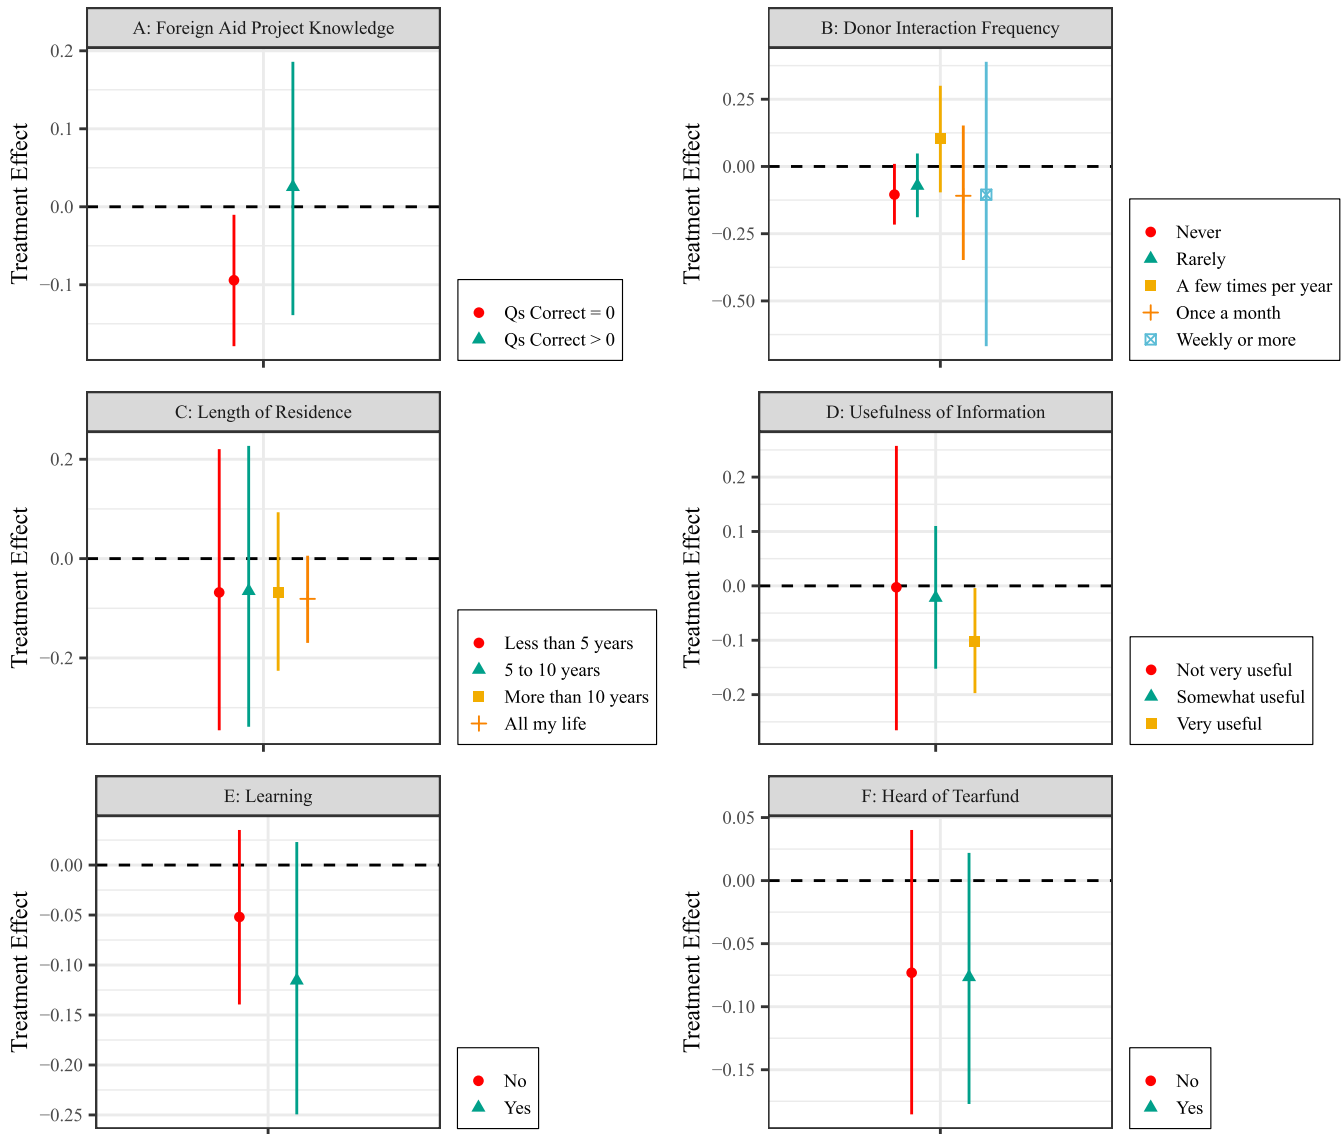
8. Possible mechanisms

8.1. Fairness norms

As we note in our theory section, we were uncertain *a priori* whether or not the influence of donors would result in politicians complementing foreign aid projects at the same sites (validation effect) or substituting for foreign aid projects at sites with fewer projects (spatial crowding out effect). To gather qualitative evidence regarding these mechanisms, we asked the politicians in our study for a brief explanation regarding the choices they made. These explanations suggest causal mechanisms to explain the spatial crowding out effect we observe.

First, only 1% of the sampled politicians mentioned any explanation related to the validation hypothesis, whereas 17% mentioned an explanation suggestive of a crowding out effect. Further, another 16% mentioned a reason related more generally to a fairness norm—the idea of distributing development projects to areas that have not received any. As one LC said when explaining his decision, “This school is located far away from other schools and there is no support from donors so if this school is supported it can be good.” Another said, “The school [does] not receive any support from donors and this can be the first one and help the school.” An MP noted in explanation, “The road is far away...and no access and some development organizations are reluctant to support those schools due to [the] road [being] impassable.” These types of explanations were not more common among those who received the aid information treatment, which indicates that this norm is present even in the absence of information that primes it. We view this qualitative evidence of a fairness norm as suggestive evidence that the spatial crowding out effect might generalize beyond foreign aid fungibility and to the fungibility or substitutability of development more generally.

⁴⁶ Treatment effects among the sample of those who claimed to learn and did not learn are −0.67 ($p = 0.10$) and −0.30 ($p = 0.26$) respectively.



Note: This figure shows the average effect of the aid information treatment on the log odds of selecting a school with an aid project. In Panel A are the results conditional on the politician's donor knowledge score. In Panel B are the results conditional on the intensity of donor interaction. In Panel C are results conditional on a politician's indication of the usefulness of the information. In Panel D are results conditional on a politician's length of residence in a constituency. In Panel E are results conditional on whether a politician indicated that she learned something from our interaction. These estimates are based on a triple interaction of Aid Treatment, Aid Project Count and the conditioning variable. Standard errors are clustered on politician. Vertical lines indicate the 95% confidence intervals.

Fig. 5. Heterogeneous treatment effects by respondent knowledge and background.

As we noted in Section 3, this fairness norm offers a plausible explanation for why spatial crowding out might not benefit less needy or more politically pivotal communities. If politicians weight public welfare highly in their decision-making, and see donors as allocating aid in a way that is not optimal, then our model implies that foreign aid will incentivize politicians to spend more in highly needy communities than they would have in the absence of aid.

8.2. Social desirability bias

One alternative explanation for the findings is that the experiment participants were making decisions in line with what they believe donors, the research team, or their peers would want them to do, rather than what they believe is best for their constituents. While we acknowl-

edge that we cannot entirely eliminate social desirability bias, we consider several pieces of evidence contrary to this alternative explanation.

First, politicians viewed the decision to allocate aid through our experiment as a meaningful decision that had real consequences. An estimated 30% of the sampled politicians contacted us following data collection to confirm when the lottery would be held and the goods delivered. Thus, even if the politicians were considering donor or research team preferences, it is unlikely their concerns for these preferences would override their concerns for the preferences of voters.

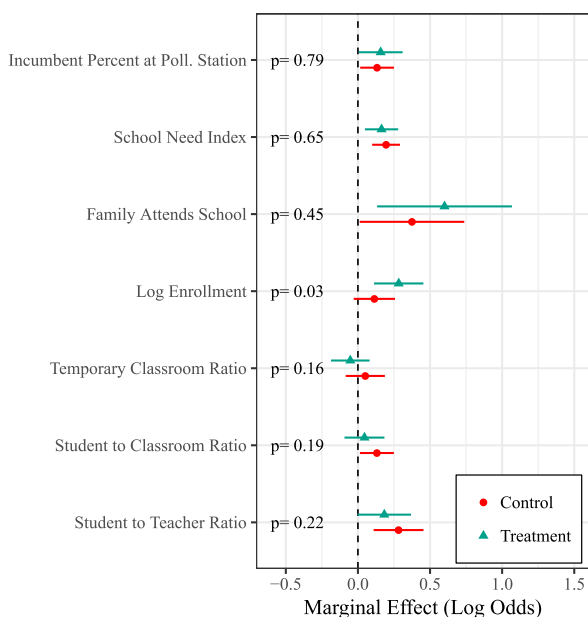
Second, it seems plausible that the direction of any social desirability bias would be the opposite of the results we find. Politicians would presumably expect that decisions that validate donor choices, or at least are orthogonal to donor choices, would be more desirable to donors than decisions that deliberately shift away from donor allocation pat-

Table 5
The effect of treatment by school characteristic.

	(1)	(2)	(3)	(4)
Aid Treatment* Incumbent Percent	0.030 (0.093)			
Aid Treatment* School Need Index		−0.023 (0.071)		
Aid Treatment* Family Attends School			0.261 (0.291)	
Aid Treatment* Log Enrollment				0.164** (0.084)
Aid Treatment* Aid Project Count	−0.177* (0.114)	−0.182* (0.114)	−0.184* (0.114)	−0.225** (0.116)
Aid Project Count	0.039 (0.072)	0.019 (0.072)	0.042 (0.072)	0.054 (0.072)
Incumbent Percent	0.129** (0.061)			
School Need Index		0.185*** (0.051)		
Family Attends School	0.491*** (0.145)	0.465*** (0.145)	0.385** (0.185)	0.494*** (0.145)
Log Enrollment				0.143* (0.070)
Observations	3718	3728	3728	3728
R ²	0.019	0.023	0.019	0.020

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

This table shows estimates from a conditional logistic regression of school selection on an interaction of treatment and school characteristics. Standard errors are clustered on politician. Control variables include *Log Permanent Classrooms*, *Log Temporary Classrooms*, *Log Teacher Houses Permanent*, *Log Teacher Houses Temporary*, *Opposition Percent Votes*, *Log Enrollment*, *Number of Aid Projects*, *Family Attends School*, *Incumbent Percent at Polling Station*, and *School Need Index*.



Note: This figure shows the coefficients of separate conditional logit regressions of school selection on baseline variables by treatment status (derived from Table 5) with 95% confidence intervals. Standard errors are clustered on politician. Continuous variables are normalized for comparison purposes. On the left side we include p-values for whether the observed difference between treatment and control is inconsistent with the null hypothesis.

Fig. 6. The effects of school characteristics on school selection by treatment status.

terns. As discussed above, 174 politicians outright said that they were choosing a school *because* it had not been supported by donors or other development projects.

Third, we examined the open-ended explanations provided by the respondent for indications of donor desirability bias. Only five explanations specifically mention Tearfund and only six mention “you” (as in the researcher), so qualitative evidence for social desirability bias is

weak. Similarly, we fail to see any significant differences in treatment effects among those respondents who had heard of Tearfund ($p = 0.91$) or interacted with Tearfund ($p = 0.63$).

Fourth, the subgroup analysis based on the frequency of interaction with donors also diminishes concerns about social desirability bias. In the presence of social desirability bias, we would expect that politicians who interact more with donors to be especially concerned about the repercussions of their choices. Yet, as discussed above, politicians who interact with donors are *less* likely to experience aid information treatment effects.

Finally, to the extent social desirability concerns *are* affecting the real spending decisions made in our experiment, such concerns are likely shifting other real spending decisions as well. In other words, perhaps this kind of “social desirability bias” is not social desirability bias in a limited research sense at all, but instead is social or peer pressure shifting behavior in the real world as well as in our research context.

8.3. Priming

Another alternative source of bias is that politicians are making responses due to priming effects. That is, the information provided may not inform politicians so much as prime them to consider foreign aid in their spending decisions. While this is difficult to rule out entirely, the heterogeneous effects in Fig. 5 undermine this explanation. If our results were primarily driven by priming, we would see little evidence that pre-existing knowledge about foreign aid projects was associated with larger treatment effects or that respondents report learning from the information in the experiment.

9. Conclusion

Scholars and policy-makers involved in international development cooperation have long expressed concerns that foreign aid earmarked for specific development objectives are directly or indirectly used to fund other — possibly less productive — government expenditures. The oft-cited statement by Paul Rosenstein-Rodan — “When the World Bank thinks it’s financing an electric power station, it is really financing a brothel.” — dates as far back as 1947, and such concerns have been

noted in several World Bank World Development Reports. Scholars and policymakers have linked aid fungibility to many negative outcomes that may undermine overall development objectives, including corruption, patronage, and a shifting of resources from areas that qualitatively need them to those that do not. Nevertheless, past research relying primarily on observational data has struggled to reach an agreement on the scale of aid fungibility, or its exact consequences for development.

In this article, we examine how information about foreign aid projects affects public spending and distributional decisions at the individual politician level. In particular, we look at how politicians take into account existing foreign aid projects when allocating development goods within the education sector. We hypothesize that they may either seek to shift funds away from areas that have already been targeted by international donors (a spatial crowding out effect) or that they choose to align their allocation decisions with those of international donors (validation effect). These effects may also activate other distributional biases, including corruption and patronage. To evaluate these claims, we implemented a field experiment in which politicians made real and meaningful decisions regarding the provision of different goods to primary schools in their constituencies. In the process of making these decisions, we provided them with randomly assigned information about foreign aid projects in their constituencies, allowing us to adjudicate between the spatial crowding out and validation effects.

The results align with the spatial crowding out effect. When politicians received information about an existing foreign aid project in a given school, they were 18% less likely to target that school with development goods. Among those politicians for whom the information about foreign aid projects was more novel, these effects were considerably stronger (22–29%), indicating that they were due to information updating, rather than social desirability bias. This is, to our knowledge, the first experimental evidence to confirm the aid fungibility hypothesis.

Our experiment involved 70% of all sitting elected politicians across Malawi. Thus we expect the findings presented here are generalizable to other Malawian elected politicians making decisions about how to allocate donor-funded, government-allocated development projects. However, one limitation of the experiment design is that it does not examine all types of distributional decisions the elected politicians make. It is possible that elected politicians allocating development projects funded with tax revenue instead of donor funds would be more or less sensitive to a crowding out effect. Politicians may feel more constrained in such an environment; or alternatively logrolling may actually increase incentives to reallocate public funds. Similarly, politicians making spending decisions over a larger pool of funds rather than small-scale specific development goods, or over projects that can be directed to a specific household instead of an entire school, may be less constrained by fairness norms and more emboldened to engage in corruption, patronage, or targeting. Finally, it is possible that the effects we observe would vary depending on the identity of prominent donors in the constituency. These are all plausible mechanisms we are unable to explore but present as promising avenues for future research.

A key strength of our design is that we are able to attribute decisions at an individual level, and can rule out the possibility that our results are driven by higher order political actors or the nature of the institutionalized decision process. Moreover—as in many low-revenue contexts—local politicians in Malawi have little ability to use tax revenue to fund discretionary projects. Thus, by involving an NGO partner, we were able to make the spending decision both meaningful and realistic. We leave a more complete examination of group vs. individual decisions for future research.

The findings of this study have some important implications for aid effectiveness. For one, the existence of spatial crowding out means that evaluating the specific impact of aid programs is more difficult than often appreciated. Spatial spillover from aid recipient sites to government recipient sites will often violate the identifying assumption of non-interference and bias treatment estimates (Baird et al., 2014). Our findings suggest this bias will often be negative, implying that researchers

may have often under-estimated the true impact of aid programs.

Our findings also imply that the difference between budgetary aid and project aid may be less meaningful than is usually appreciated. Donors often use project aid instead of budget aid to achieve greater oversight in the end-use of development funds and to circumvent political bias and corruption. On one hand, the crowding out effects we identify suggest that circumventing political influence over development spending is much harder than often appreciated. On the other hand, we find no evidence that crowding out has any negative effects for development outcomes, suggesting that donor ring fencing efforts are sometimes misplaced.

These results also offer good news for aid effectiveness. The scale of spatial crowding out of public funds that we identify in our experiment is smaller than most estimates from cross-national observational studies, and several times smaller than the least optimistic estimates. Moreover, the results also suggest that the possible negative outcomes of aid fungibility may be overstated. We find no evidence that the spatial crowding out effect benefits schools with politicians' family members in attendance, or that funds are systematically targeted to more politically important or economically less needy schools. Instead, qualitative evidence collected in the study suggests that many participants in the experiment made their allocation decisions in line with a fairness norm, whereby they sought to provide development assistance primarily to those schools that had not yet received support.

Author statement

Brigitte Seim: Supervision, Conceptualization, Methodology, Investigation, Writing, Project administration, Funding acquisition. **Ryan Jablonski:** Supervision, Conceptualization, Methodology, Formal analysis, Investigation, Writing, Funding acquisition. **Johan Ahlback:** Investigation, Data curation, Writing, Project administration.

Appendix. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jdeveco.2020.102522>.

References

- Baird, S., Bohren, J.A., McIntosh, C., Ozler, B., 2014. Designing Experiments to Measure Spillover Effects. World Bank Policy Research Working Paper. .
- Baldwin, K., Winters, M.S., 2018. Bypass Aid and Perceptions of Local Government Performance and Legitimacy. AidData Working Paper 56.
- Benjamini, Y., Hochberg, Y., 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. Roy. Stat. Soc. Ser. B (Methodol.)* 57, 289–300.
- Bernbaum, M., Moses, K., 2011. EQUIP2 Lessons Learned in Education: A Guide to Education Project Design, Evaluation, and Implementation Based on Experiences from EQUIP2 Projects in Malawi, Uganda, and Zambia Education Management Information Systems. <https://www.fhi360.org/sites/default/files/media/documents/EQUIP2%20L1%20E-MIS%20AAR.pdf>.
- Blair, R.A., Roessler, P., 2018. The Effects of Chinese Aid on State Legitimacy in Africa: Cross-National and Sub-national Evidence from Surveys, Survey Experiments, and Behavioral Games. AidData Working Paper 59.
- Briggs, R.C., 2017. Does foreign aid target the poorest? *Int. Organ.* 71, 187–206.
- Bueno de Mesquita, B., Smith, A., 2009. A political economy of aid. *Int. Organ.* 63, 309–340.
- Burgess, R., Jedwab, R., Miguel, E., Morjaria, A., Padró i Miquel, G., 2015. The value of democracy: evidence from road building in Kenya. *Am. Econ. Rev.* 105, 1817–1851.
- Chatterjee, S., Giuliano, P., Kaya, I., 2012. Where has all the money gone? foreign aid and the composition of government spending. *B E J. Macroecon.* 12.
- Clark, M.D., Determann, D., Petrou, S., Moro, D., de Bekker-Grob, E.W., 2014. Discrete choice experiments in health economics: a review of the literature. *Pharmacoeconomics* 32, 883–902.
- de la Cuesta, B., Milner, H.V., Nielson, D.L., Knack, S.F., 2019. Oil and aid revenue produce equal demands for accountability as taxes in Ghana and Uganda. *Proc. Natl. Acad. Sci.* 116, 17717–17722.
- Dietrich, S., Mahmud, M., Winters, M.S., 2018. Foreign aid, foreign policy, and domestic government legitimacy: experimental evidence from Bangladesh. *J. Polit.* 80, 133–148.

- Dollar, D., Pritchett, L., 1998. Assessing Aid: what Works, what Doesn't, and Why. World Bank Policy Research Working Paper URL: <http://documents.worldbank.org/curated/en/612481468764422935/Assessing-aid-what-works-what-doesnt-and-why>.
- Easterly, W., 2009. Can the west save africa? *J. Econ. Lit.* 47, 373–447.
- Feyzioglu, T., Swaroop, V., Zhu, M., 1998. A panel data analysis of the fungibility of foreign aid. *World Bank Econ. Rev.* 12, 29–58.
- Findley, M.G., Harris, A.S., Milner, H.V., Nielson, D.L., 2017. Who controls foreign aid? elite versus public perceptions of donor influence in aid-dependent Uganda. *Int. Organ.* 71, 633–663.
- Gibson, C.C.K.A., Ostrom, E., Shivakumar, S., 2005. *The Samaritan's Dilemma: the Political Economy of Development Aid*. Oxford University Press, Oxford.
- Hines, J.R., Thaler, R.H., 1995. The flypaper effect. *J. Econ. Perspect.* 9, 217–226.
- Hoyos, D., 2010. The state of the art of environmental valuation with discrete choice experiments. *Ecol. Econ.* 69, 1595–1603.
- Jablonski, R.S., 2014. How aid targets votes: the impact of electoral incentives on foreign aid distribution. *World Polit.* 66, 293–330.
- Jablonski, R., Seim, B., 2017. How Transparency Affects Distributional Politics: A Field Experiment Among Elected Incumbents in Malawi. AidData Working Paper #45, Williamsburg, VA: AidData at William & Mary.
- Keefer, P., Vlaicu, R., 2008. Democracy, credibility, and clientelism. *J. Law Econ. Organ.* 24, 371–406.
- Kramon, E., Posner, D.N., 2013. Who benefits from distributive Politics? How the outcome one studies affects the answer one gets. *Perspect. Polit.* 11, 461–474.
- Marć, Ł., 2017. The impact of aid on total government expenditures: new evidence on fungibility. *Rev. Dev. Econ.* 21, 627–663.
- McGillivray, M., Morrissey, O., 2000. Aid fungibility in assessing aid: red herring or true concern? *J. Int. Dev.* 12, 413–428.
- Ministry of Finance, 2019. Malawi Aid Management Platform. URL: <https://malawiaid.finance.gov.mw/portal/>.
- Morrison, K.M., 2007. Natural resources, aid, and democratization: a best-case scenario. *Publ. Choice* 131, 365–386.
- Morrissey, O., 2015. Aid and government fiscal behavior: assessing recent evidence. *World Dev.* 69, 98–105.
- Muralidharan, K., Romero, M., Wüthrich, K., 2019. Factorial Designs, Model Selection, and (Incorrect) Inference in Randomized Experiments. National Bureau of Economic Research Working Paper.
- Pack, H., Pack, J.R., 1990. Is foreign aid fungible? the case of Indonesia. *Econ. J.* 100, 188–194.
- Rana, Z., Koch, D.J., 2020. Is it time to 'decolonise' the fungibility debate? *Third World Q.* 41, 42–57.
- Remmer, K.L., 2004. Does foreign aid promote the expansion of government? *Am. J. Polit. Sci.* 48, 77–92.
- Van de Sijpe, N., 2013. Is Foreign Aid Fungible? Evidence from the Education and Health Sectors. The World Bank.
- Stokes, S.C., Dunning, T., Nazareno, M., Brusco, V., 2013. *Brokers, Voters, Clientelism: the Puzzle of Distributive Politics*. Cambridge University Press, Cambridge.
- UNICEF, 2015. *The Investment Case for Education and Equity*. URL: https://www.unicef.org/publications/index_78727.html.
- Wagstaff, A., 2011. Fungibility and the impact of development assistance: evidence from vietnam's health sector. *J. Dev. Econ.* 94, 62–73.
- van de Walle, D., Mu, R., 2007. Fungibility and the flypaper effect of project aid: micro-evidence for vietnam. *J. Dev. Econ.* 2, 667–685.
- Werker, E., Ahmed, F.Z., Cohen, C., 2009. How is foreign aid spent? evidence from a natural experiment. *Am. Econ. J. Macroecon.* 1, 225–244.
- World Bank, 2010. *The Education System in Malawi*. URL: https://poledakar.iiep.unesco.org/sites/default/files/fields/publication_files/resen_malawi_2010_0.pdf.
- World Bank, 2019. *World Development Indicators*. URL: <http://data.worldbank.org/data-catalog/world-development-indicators>. (Accessed 0 August 2019).