What Is "Country Ownership"? A Formal Exploration of the Aid Relationship

William D. Savedoff

Abstract

The aid literature and high-level accords like the Paris Declaration argue that "country ownership" is critical to the effectiveness of aid. In response, donors and recipients renamed themselves "development partners," obscuring the tendency for country ownership benefits (i.e. more successful and sustainable programs) to come at the expense of satisfying the funding countries' priorities.

This paper illustrates the tradeoff between country ownership and funders' priorities with a formal model in which aid is governed by a contract to produce a jointly desired outcome. The model generalizes the Principal-Agent approaches for studying aid which treat countries as having multiple objectives.

The new model illustrates how a recipient country's rational resource allocation choices vary with different aid contracts, whether based on lump sum payments, input-based payments, conditional payments, matching grants or outcome payments. It reveals two critical aspects of the country ownership debate. First, even when funders and recipients agree on project goals, funders can only achieve their priorities through distorting domestic allocative choices. Second, funders are likely to fully embrace country ownership only in cases where they believe alternative uses of domestic funds have integrity (as defined by the funder).

The model also shows that when funders put higher priority on achieving their goals than accommodating recipient allocation preferences, they should prefer conditional payments, matching grants, or outcome payments. Among these, the donor's preferences would depend on the relative observability of expenditures to outcomes. If instead funders embrace country ownership and seek to maximize the country's welfare, lump sum grants are better. In terms of Paris Declaration goals of sustainability, the aid contracts which are least aligned with recipient country priorities will not be sustained after aid ends unless domestic preferences are altered by a process of hysteresis.

What Is "Country Ownership"? A Formal Exploration of the Aid Relationship

William D. Savedoff
Center for Global Development

Contents

Introduction

The aid literature and high-level accords like the Paris Declaration argue that "country ownership" is critical to the effectiveness of aid. In implementing the Paris Declaration, donors and recipients renamed themselves "development partners" obscuring the tendency for country ownership benefits (i.e. more successful and sustainable programs) to come at the expense of satisfying the funding countries' priorities.

This paper will illustrate the tradeoff between country ownership and funders' priorities by using a formal model in which funders and recipients contract to produce a jointly desired development outcome. The model generalizes Principal-Agent approaches for studying development aid which treat countries as having multiple objectives rather. This contrasts with Principal-Agent models which treat the primary aid contracting problem as inducing the recipient to exert greater effort in implementation.

The new model illustrates how a recipient country's rational resource allocation choices will vary in response to different aid contracts based on lump sum payments, input-based payments, conditional payments, matching grants, and outcome payments. It reveals two critical aspects of the country ownership debate. First, even when funders and recipients agree on project goals, funders can only achieve their goal through distorting domestic allocative choices. Second, funders are likely to fully embrace country ownership only in cases where they believe alternative uses of domestic funds have integrity (as defined by the funder).

The model also shows that when funders put higher priority on achieving their goals than accommodating recipient allocation preferences, they should prefer conditional payments, matching grants, or outcome payments. Among these, the donor's preferences would depend on the relative observability of expenditures to outcomes. If instead funders embrace country ownership and seek to maximize the country's welfare, lump sum grants are better. In terms of Paris Declaration goals of sustainability, the aid contracts which are least aligned with recipient country priorities will not be sustained after aid ends unless domestic preferences are altered by a process of hysteresis.

Literature Review

Country Ownership

"Country ownership" is a term in the development assistance literature that emerged for the first time in the mid-1990s, as part of a reaction against the language of "conditionality" employed by organizations like the IMF and World Bank in the structural adjustment programs of the 1980s.[1] Instead of posing external financial assistance as a tool to be used to force governments to adopt policies designed and motivated by external agents, country

---

[1] A Google Ngram of the term " country ownership " shows no mentions in books before 1996, exponential growth until 2004 and then steady usage thereafter. (Accessed June 7, 2019).

ownership was to make countries (implicitly governments) the primary agents in choosing policies and designing programs financed by foreign aid. In this sense, they should "own" the programs. This was accompanied by changes in language regarding foreign aid from "donors and recipients" to "development partners."

The legitimation of country ownership as the primary lens for foreign assistance was established in the Paris Declaration for Aid Effectiveness (OECD 2005) and later elaborated in the Accra Agenda for Action of 2008. The principles to which development partners (both those providing funding and those receiving it) committed themselves required that each country should establish its own priorities and take the lead in designing and implementing programs. One specific outcome of these agreements was the wider application of agreements which disbursed foreign aid directly to government budgets along with technical assistance to improve the public financial management which would ultimately influence the integrity of the application of funds.

Thus, the push for country ownership may be understood primarily as an effort to shift the political relationship between those providing aid and those receiving it from a paradigm of charity and tutelage to a paradigm of partnership and mutual accountability. Nevertheless, practical experiences and research on the factors that influence the effectiveness of foreign aid also provided empirical support for the notion that programs which were country-owned did, indeed, perform better.

While the notion of country ownership received broad acceptance and continues to inform foreign aid discussions today, it has not been without criticism. Most commonly, critics argue that aid agencies represent programs as being country owned when, in fact, programs are drafted by external consultants or are accepted by governments under financial duress. Others argue that the notion is inherently undermined by the assumption that governments receiving aid are, indeed, interested in development (Booth 2012). Soon after the Paris Declaration, Buiter wrote: "When the statement 'this program is country-owned' tends to mean no more than 'this program is supported by the people who own the country', it is time to purge it from our vocabulary" (Buiter 2007). In contrast to such critiques, this paper illuminates problems with the concept of country ownership even when the government is truly representing the legitimate interests of its citizens.

Principal Agent Models

Principal-Agent models have been utilized for analyzing the relationship between foreign funders and governments that receive aid in a variety of contexts.[2] But does the principal-agent model elucidate or confuse our understanding of the foreign aid relationship? Certainly

---

[2] It is referenced in the academic literature in such works as Martens et al. 2002 and Paul 2006; in international development institutions, such as Inter-American Development Bank 1996 and World Bank 2005; and in think tanks and bilateral agencies, such as Savedoff (2010 and 2011) and Clist and Verschoor (2014).

the standard applications for which the principal-agent model was developed are only vaguely similar to the foreign aid relationship.

Some of the first and most prominent applications of this model analyzed how shareholders delegated responsibility to their managers and sharecropping arrangements between tenants and landlords. The key features of these models were a divergence between the interests of the principal and agent and the unobservability of output or effort. The principle findings involved the design of second-best contracts which exploited differences in risk aversion or other features to align the agent's interests more closely with those of the principal.[3]

Relatively few studies have used this particular framing to study foreign aid programs. Savedoff (2010 and 2011) uses such a framework to organize a discussion of aid programs that pay for results. Clist and Verschoor (2014) and Clist (2016) use such a model to assess aid contracts in terms of the incentives they create, whether paying in proportion to results or actions, and provide a comprehensive discussion of phenomena related to paying for results—such as the signal-to-noise ratio in the outcome measure, the risk of gaming, and the impact on intrinsic motivation. In both cases, the agent is assumed to produce just one output that is of interest to the principal. This approach is appropriate for aid programs where the recipient is an individual, a service provider, a firm or perhaps a community.

This particular form of principal-agency is not going to apply as well to the case where one government provides foreign aid to another. First, when aid is given from one government to another, neither the principal nor the agent are individuals; rather they are institutions whose preference structures are not well-ordered or stable. Second, the provider of aid generally has some interest in improving the recipient's condition; unlike, say, shareholders who may treat the Chief Executive Officer as a means to increase their own income. Third, countries receiving foreign aid from another country generally have other sources of funding, whether from their own domestic tax base, sale of natural resources, or as assistance from other funders. Furthermore, the countries involved in aid relationships each have multiple objectives. In these ways, the foreign aid relationship between governments differs from the original uses of the principal-agent model in important ways (Paul 2006). Thus, while the principal-agent may still offer insights for the foreign aid relationship, the findings must be qualified and interpreted differently when addressing relationships between countries.

Most papers which address the aid relationship between governments tend to emphasize the lattermost characteristic: multiple objectives. The other features are highlighted much less, i.e. governments tend to be treated as unitary decision making units; the funder is assumed to have altruistic motivations for providing aid; and domestic sources of funding are typically ignored. Nevertheless, the existence of multiple objectives fundamentally alters the structure

---

[3] For presentations of the essential features of principal-agent models and reviews of their applications, see Stiglitz 1989; Sappington 1991; Laffont and Martimort 2009.

and implications of the aid relationship, generating an important divergence between the preference structure of the principal and the agent. For example, Murshed & Sen (1995) analyze a situation in which a country wants foreign aid and the ability to increase military spending; while the funder wishes to constrain military spending. Azam & Laffont (2003) analyze a case where local elites have some interest in helping the local poor, but not to the extent preferred by the foreign aid agency; thus, they investigate conditions in which the local elites' preference for diverting aid to their own benefit can be attenuated. Over (1981) analyzes the intertemporal paradox in which funders and recipients agree on an investment project only to reveal (time and again) that after the investment is made, the recipient is not motivated to finance the recurrent costs necessary to generate returns from those investments. In each of these studies, the interests of the funder and the recipient overlap. However, recipients have at least one other goal which leads them to face an opportunity cost and tradeoffs that differ from the allocational preferences of the funder.

This paper explores the foreign aid relationship within this approach which treats principals and agents as having overlapping goals. It is a principal-agent model in the formal sense of establishing a logic of sequential contract negotiation which satisfies both parties to differing degrees as a consequence of strategic behavior. It requires overlapping interests in at least one common goal, it assumes recipients have other sources of funding so that they are not completely beholden to the funder, and it assumes the recipient has multiple objectives.

The Basic Model

This paper develops a model of aid in which the funder and recipient share a development objective (E_R) which can be thought of as number of children who get an education. The funder maximizes its utility by allocating resources between the production of goods and services in its own country (X_F) and grants to produce the development objective (E_R) in the recipient's country subject to its own budget constraint. The recipient maximizes its utility by allocating resources between the development objective (E_R) and "Anything Else" (A_R) subject to its budget constraint.

$$\text{Funder's Utility} = U_F(E_R, X_F)$$

$$\text{Recipient's Utility} = U_R(E_R, A_R)$$

The interpretation of the results will depend critically on the character of "Anything Else." In particular, this paper will consider the implications of three distinct uses of domestic resources other than spending on E:

- *Development Goods.* In addition to spending on the agreed development goal E, a country might be interested in spending on other goods or services which are important to development such as public health, roads, or household access to energy.

- *Non-Development Goods.* The government might be interested in spending money on things that are not considered development goods but which it still considers

important – sports arenas, military equipment (Murshed and Sen 1995), or mausoleums. This might also include spending that serves to buttress political support for the current regime that isn't in the public interest.

- *Corruption for personal gain.* Funds going to personal use by public officials.

The funder's perspective on funds that are diverted from spending on E to A are likely to differ if it considers them legitimate development goods rather than diversion for corrupt purposes. The non-development goods represent an intermediate category which may not be valued by the funder per se unless it views the realization of the country's own priorities as a positive outcome. These differences will be explored further in the discussion below.

The relationship between the funder and recipient is modelled as a sequential process in which the funder starts by offering a contract. The recipient accepts the contract only if it can reach a higher utility than before under the terms of that contract. After the agreement is signed, the recipient chooses an allocation between E_R and A_R to maximize its utility under the new budget and cost constraints set by the aid agreement.

For simplicity subscripts R and F will be eliminated where it is otherwise clear. For ease of modelling, utility will be specified with Cobb-Douglas functions as:[4]

- Funder Utility: $U(E,X) = E^{\beta} X^{(1-\beta)}$
- Recipient Utility: $U(E,A) = E^{\alpha} A^{(1-\alpha)}$

A linear function will be used to describe the costs of producing E and A, where P_e and P_a are the unit costs of producing E and A, respectively:

- Total Cost of production $= P_e E + P_a A$

Choosing these functional forms will have several implications which conform reasonably to a number of stylized facts.

First, the amount of aid provided by the funder is only a function of its total income and not the amount of E produced by the recipient. This is a direct consequence of choosing the Cobb-Douglas utility function for which the budget share allocated to any good is equal to its exponent. So, for example, the Funder's aid budget (G) will simply equal $\beta *$ (Funder's Total Income).

This implies that funders are not persuaded to donate more or less based on how much E is produced. This is not an unreasonable stylized fact given that donor countries often debate

---

[4] The Cobb-Douglas specification is appropriate for any context in which there is a smooth and concave tradeoff in preferences between the two products. The one way in which the Cobb-Douglas function unnaturally constrains the preference structure is that it is homothetic and this leads the marginal rate of substitution between the two products to be invariant with income. The paper will qualify the results which are purely a result of this structure and not a consequence of the more generalizable features of the function.

the size of their total aid budget independently of what is achieved. The frequently discussed target of allocating 0.7% of GDP to aid is a prime example of this dynamic. It is less likely to hold for individual countries, since the donor has an opportunity to shift funding from one country to another. However, it is still defensible on the empirical grounds that bilateral funding between any pair of countries is highly correlated from one year to the next.

Second, because the total grant amount will be the same in all cases, any differences in the production of E will be due to substitution effects induced by the contract modality rather than the amount of aid itself. This, again, is a direct consequence of the choice of a Cobb-Douglas functional form for which the marginal rate of substitution is invariant to pure income effects. This feature is reasonable for most countries for whom aid is a relatively small share of their total income; but may be consequential in cases where aid has more significant effects on national public budgets.

Finally, the production function remains the same in all the cases, implicitly assuming no economies of scale or efficiency improvements as a result of the aid modality. With economies of scale or improvements in efficiency, the tradeoffs facing recipient countries would change accordingly and affect the relative costs of producing A and E.

The rest of this paper will proceed by solving the model for a base case (with no aid) followed by six cases that illustrate the impact of different aid modalities. The difference between these cases is determined by the variables that affect the budget available to the recipient with one exception; in the input-based aid case, the aid agreement operates through subsidizing the costs of production. The cases can be described in the following formal terms.

In the *No Aid Case*, the recipient country's budget is simply its domestic resources (Y).

- Case 0: Budget = Y

In the *Lump Sum Grant Case*, the recipient country's budget is its domestic resources (Y) plus a grant (G) from the donor country. The amount of the grant (G) is that determined by the donor's own utility maximizing resource allocation decision as outlined earlier.

- Case 1: Budget = Y + G

In the *Input-Reimbursement Case*, the recipient country's budget is its domestic resources (Y) but its costs of producing E are reduced by the amount of the input-based grant (g).

- Case 2: Budget = Y and cost of production = $(P_e E - g) + P_a A$

In the *Conditional Payment Case*, the recipient country's budget is its domestic resources (Y) plus a payment (M) which the recipient receives if and only if it spends at least a certain amount (Σ) on E.

- Case 3: Budget = Y + M, if and only if $P_e E \geq \Sigma$

In the *Matching Grant Case*, the recipient country's budget is its domestic resources (Y) plus a payment which is a proportion (m) of the recipient's spending on E.

- Case 4: Budget = $Y + m\,P_e\,E$

In the *Unit Payment Case*, the recipient country's budget is its domestic resources (Y) plus a payment (r) for each unit of E produced by the recipient.

- Case 5: Budget = $Y + rE$

In the *Unit Payment for Additional Production Case*, the recipient country's budget is its domestic resources (Y) plus a payment (r) for each unit of E produced by the recipient *in excess of the initial level* of E (denoted, E_0).

- Case 6: Budget = $Y + r(E - E_0)$

Table 1. Summary of model and cases

Recipient's utility function	$U(E,A) = E^{\alpha}\,A^{(1-\alpha)}$
Funder's utility function	$U(E,X) = E^{\beta}X^{(1-\beta)}$ implies that $G = \beta *$ (Funder Income)
Total cost of production	$P_e\,E + P_a\,A$

	Cases
Case 0: No Aid	Budget = Y
Case 1: Lump Sum Grant	Budget = Y + (G)rant
Case 2: Input reimbursement	Budget = Y and Total Cost to recipient = $(P_e\,E - g) + P_a\,A$
Case 3: Conditional payment	Budget = Y + M
Case 4: Matching grant	Budget = $Y + m\,P_e\,E$
Case 5: Unit payment	Budget = Y + rE
Case 6: Unit payment for additional production	Budget = $Y + r(E - E_0)$

Results

This section solves the base case of the model along with the 6 aid cases in terms of the resulting impact on resource allocations and production of A and E. It then summarizes the results and calculates the funder's and recipient's utilities.

The No Aid Case

In this model, the recipient simply maximizes welfare by allocating its own resources between Education and Anything Else (A).

The recipient chooses to allocate its resources (Y) between producing E and A so as to maximize its utility, subject to the constraint of the production function. Using the Lagrangean function this would be:

$$\text{Max } \mathcal{L} = U(E,A) - \lambda[Y - (P_e E + P_a A)]$$

To maximize utility, the country will allocate resources to producing E up to the point where the ratio of its marginal utility relative to its price is equal to the same ratio for A. The solution for this first order condition is: $\alpha/(1-\alpha) = (P_e * E)/(P_a * A)$

Because of the Cobb-Douglas specification of the utility function, the optimal level of E in this base case (E_0^*) can be characterized in terms of the utility function preference parameter (α), total spending (i.e., the budget Y), and the cost of producing each unit of E (P_e):

$$E_0^* = \alpha Y / P_e$$

This can be depicted in a standard diagram in which maximum utility is reached when E and A are chosen at the point where the utility curve is tangent to the budget line. That is, at the optimal allocation, the slopes of the two curves are identical and represent the tradeoff in producing E and A—which is the ratio of the prices in the production function ($-P_a/P_e$).

Figure 1. No aid case
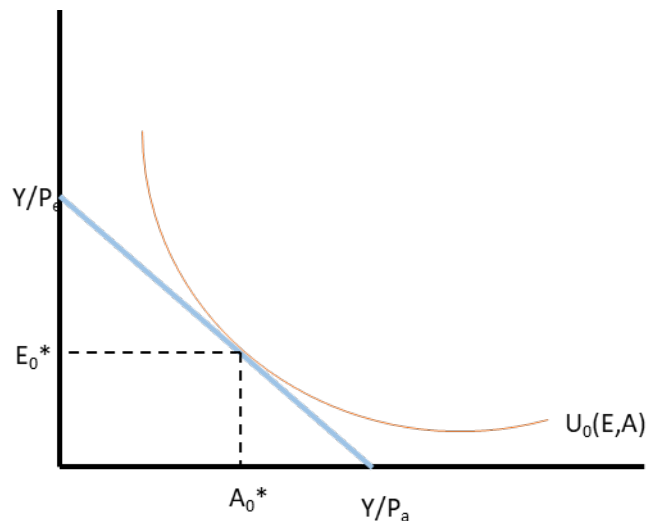
The Lump Sum Grant Case

In the case of a lump sum grant, the country's budget increases by the amount of the grant (G).

Even if the agreement sets targets for E, after signing the agreement and receiving the funds, the recipient would choose to allocate resources between E and A that would maximize its utility subject to the new, relaxed, budget constraint.

$$\text{Max } \mathcal{L} = U(E,A) - \lambda[(Y + G) - (P_e E + P_a A)]$$

Solving for the recipient's optimal allocation as before but with the new budget constraint yields a first order condition of: $\alpha/(1-\alpha) = (P_e * E)/(P_a * A)$. This is the same as the No Aid Case because the Cobb-Douglas functional form is invariant to the country's income level and because the slopes of the utility function and budget constraint $(-P_a/P_e)$ are also the same because costs of production have not changed.[5] However, the optimal amount of E will now be larger because of the additional resources as:

$$E_1* = (\alpha Y + G)/ P_e$$

Essentially, putting all of the grant into producing E would raise the recipient's utility but not by as much as if the recipient were to allocate a share of the grant (αG) to E and the remainder to A. Thus, the resulting allocation would look like this:

Figure 2. Lump sum grant case
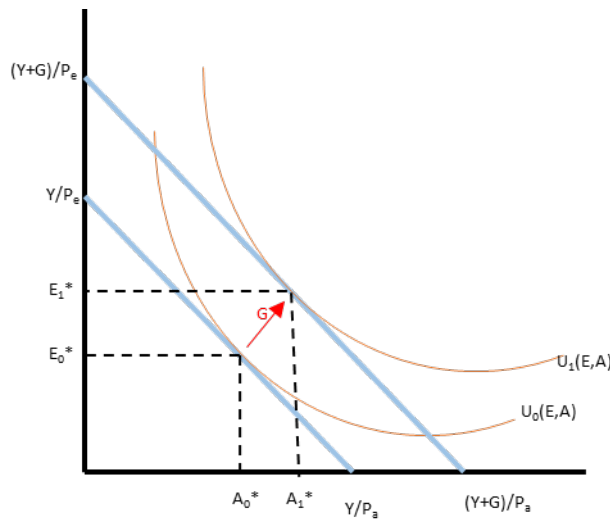
If the funder only wants to maximize the recipient's welfare, then this is not a bad outcome. However, if the funder wanted spending on E to increase by G it will be disappointed by the allocation of a significant share of the funds $(1-\alpha)G$ going to produce A.

The Input-Reimbursement Case

If the funder offers to reimburse the inputs used for producing E for an amount (g) and which is capped at G, then it appears that the recipient is forced to spend on E and not A.

---

[5] An alternative functional form for utility would lead to a different outcome, but in this case it would still only be affected by the income effect and not the substitution effect.

However, as we'll see, the fungibility of funding leads to an outcome identical to the lump sum grant approach.

Under this agreement, the recipient maximizes its utility subject to a budget constraint set by its domestic resources (Y) and with a grant that subsidizes spending on E:

$$\text{Max } \mathcal{L} = U(E,A) - \lambda[Y - ((P_e E - g) + P_a A)]$$

By offering to reimburse inputs, the funder envisions a process where the country chooses to produce A_0^* and applies all of the g funds to producing E. However, from the country's perspective, this would lead to a suboptimal allocation. By spending all of the g funds on E and reallocating some of its domestic resources to producing more A, the country is able to maximize its welfare in line with the tradeoff determined by the relative costs of production $(-P_a/P_e)$.

Among international aid agencies this phenomenon is referred to as *fungibility* because even when grant funds are applied directly to one purpose (E), it does not hinder the recipient from taking advantage of the additional funds to free up resources that can be applied to other purposes.

The recipient can ensure that it spends g on inputs for E and thereby obtain the maximum grant. However, as a result of reallocating some of its domestic budget from producing E to producing A, it is able to maximize utility at E_2^* and A_2^*. In accounting terms, it spent g external funds and $(\alpha Y - (1-\alpha)g)$ domestic funds to produce E_2^*. Then it spent $((1-\alpha)Y + (1-\alpha)g)$ of domestic revenues to produce A_2^* without using external funds. Despite the different sources of funds, the resulting output is identical to the lump sum grant case.

Figure 3. Input reimbursement case
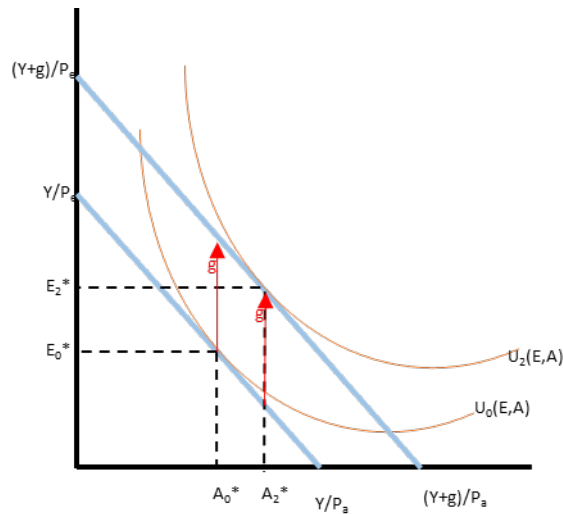
The Conditional Payment Case

To avoid the fungibility that occurs under the *Input Reimbursement Case* (and for other reasons as well), aid agencies have sometimes made agreements that stipulate the grant will be paid when the recipient spends a minimum amount (Σ) on E. In this case, the recipient solves the following problem:

$$\text{Max } \mathcal{L} = U(E,A) - \lambda[(Y + M - (P_e E + P_a A)] - \varphi (E - \Sigma)$$

The conditional payment (M) is additional to domestic income (Y). A final term is added to account for the constraint imposed by requiring the minimum level of spending on E. Assuming that the agreement is binding (i.e. $\Sigma > P_e E_0$), the recipient spends exactly Σ on E to qualify for the grant and spends the remainder of its budget on A.

Under certain conditions, these conditional payments solve the fungibility problem because the recipient cannot reallocate domestic spending to A without losing the grant. So, unlike the *Input Reimbursement Case*, the recipient does spend the additional funds on E. However, the relative prices of E and A do not change. Consequently, the recipient's utility (U_3) after spending the grant on E is lower than it would be if it could reallocate some of the resources to A (U_2). Nevertheless, the recipient's utility is still higher than it would be in the *No Aid Case*.

For the Conditional Payment agreement to work in this way, two very important assumptions must hold. First, the production function must hold in the sense that spending continues to produce E at the rate P_e. If this direct link were not maintained, the donor country could still ensure that Σ is spent on E, but it could not be sure that the desired level of E would be produced. The second critical assumption is that the donor can accurately observe domestic spending. If instead the recipient can manipulate financial records, its ability to obtain the grant while reallocating funds to A would lead this agreement to function more like the *Input Reimbursement Case* with the fungibility problem intact.

Figure 4. Conditional payment case
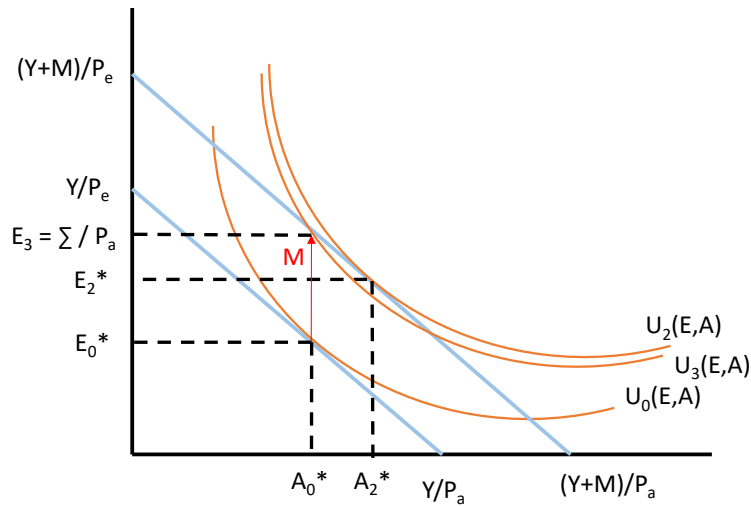
The Matching Grant Case

Another approach makes aid conditional on domestic spending in a different way. Rather than setting a minimum spending level, the donor can offer to pay a proportion (m) of the recipient's spending on E. In this case, the recipient solves the following problem:

$$\text{Max } \mathcal{L} = U(E,A) - \lambda[(Y + m\, P_e\, E - (P_e\, E + P_a\, A)\,]$$

Solving for the recipient's optimal allocation yields a first order condition of:

$$\alpha/(1-\alpha) = (1-m)\,(P_e * E)/(P_a * A)$$

In this case, the budget constraint's slope changes. By spending everything on E, the recipient would maximize the aid it receives and have a budget of (1+m) Y. By contrast, spending everything on A would mean no aid is received and the budget would simply be domestic resources (Y). By shifting the tradeoff between spending on E and A, the donor induces the recipient to spend more than otherwise on E. However, since it still values A, there is a limit at which it will sacrifice additional aid to ensure a certain amount of A.

The optimal amount of E produced is:

$$E_4^* = \alpha\, Y / (1-m)\, P_e$$

Unlike the Conditional Payment Case, the matching grant achieves its aim by altering the tradeoff between A and E. Consequently, the recipient has no desire to alter the allocation further. Given the specific nature of the Cobb-Douglas Utility function in this exercise, the donor is able to induce the recipient to put all of the aid toward E and the resulting allocation is identical to that of the *Conditional Payment Case*. The recipient's utility is lower than the *Lump Sum Grant Case* but is higher than the *No Aid Case*.

The assumptions required for this result are identical to the ones required for the *Conditional Payment Case*. Domestic and grant money has to be spent to produce E at the same rate as before and the donor has to be able to accurately observe spending on E.

Figure 5. Matching grant case
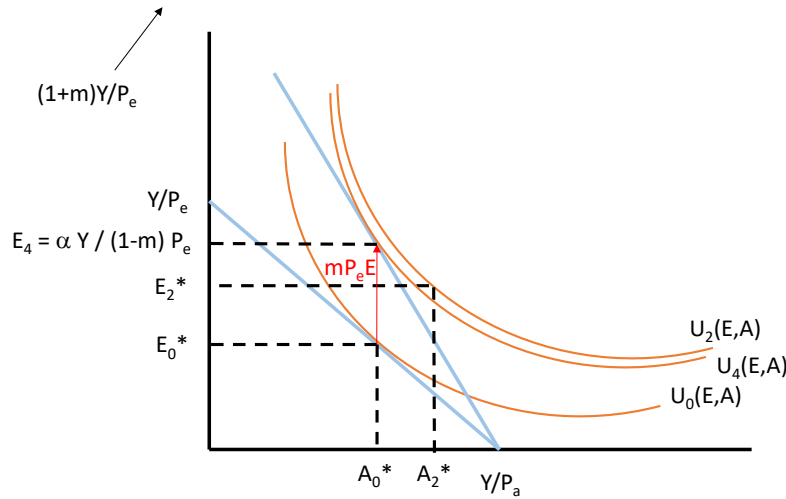
The Unit Payment Case

So far, we have considered conventional aid modalities which represent more than 90 percent of overseas development assistance. However, in the last two decades, agencies have been experimenting with modalities which disburse funds in proportion to "results" which are sometimes defined in terms of outcomes as defined here.[6]

The funder might offer a contract in which it agrees to pay a certain amount (r) for each and every unit of E produced by the country. This contract creates an opportunity for the country to expand its available resources beyond Y by an amount (r * E) that depends entirely on its own choices about how much E to produce. However, in addition to providing the country with more resources, it also changes the tradeoff between E and A because now, instead of costing the country P_e for each unit of E produced, it only has to provide $(P_e - r)$ for each unit of E. In formal terms, this shifts the cost ratio between E and A, leading the country to choose to produce more E relative to A than under the *Lump Sum Grant Case* and the same as the *Matching Grant Case*.

The recipient solves the following problem:

---

[6] Aid programs that pay for results vary considerably with regard to the character of the results and indicators they use. At one extreme, indicators may be chosen which are indistinguishable from the inputs and processes financed by conventional programs. At the other extreme, indicators will reflect a true outcome. One of the few examples of the latter is the Brazilian Amazon Fund to which Norway pays $5 per averted ton of carbon emissions due to reducing deforestation. See Perakis and Savedoff 2015 for a discussion of the range of results payment programs by recipient and type of result.

$$\text{Max } \mathcal{L} = U(E,A) - \lambda[(Y - ((P_e - r)E + P_a A)]$$

At the optimal allocation of resources, the first order condition is:

$$\alpha/(1-\alpha) = [(P_e - r) * E)]/(P_a * A)$$

and the slope of the budget line is $(Y - P_a A)/(P_e - r)$.

Solving for the optimal production of E yields:

$$E_5* = \alpha Y/(P_e - r)$$

If the funder knows how the recipient will respond, it will not choose r randomly. Rather, the funder will chooses its optimal offer r* to maximize E_5* subject to its own constraint for the total grant. Solving the funder's problem yields an optimal payment per unit of E of

$$r* = [G/(\alpha Y + G)] P_e$$

For technical reasons related to the Cobb-Douglas function, [7] the recipient will respond to an offer of r* per unit of E by putting all of the grant money toward the production of E, reaching a level E_5* and A_0*. The recipient will be willing to accept this offer because it still yields higher utility than the *No Aid Case*. Rather than forcing the recipient to spend all grant funds on E (which is generally impossible due both to sovereign protections and administrative infeasibility), the funder's offer *induces* the recipient to allocate the additional resources to E by shifting the relative cost of the two outcomes.

The key assumptions for the Unit Payment Case to induce more E are not as restrictive as in the *Conditional Payment* and *Matching Grant Cases*. The donor is not concerned over the rate at which spending produces E because aid is paid in proportion to the amount of E, not to the amount *spent* on E. However, the Unit Payment Case does require that E be observable. It is an empirical question whether observing an outcome like education (e.g. student test-scores) is easier or harder than verifying spending (e.g. performance audits that document actual expenditures).

---

[7] The Cobb-Douglas function is homothetic because the slope of the curves along any ray from the origin are the same. It is this feature that leads to the result that neatly allocates all grant resources to E. For non-homothetic functional forms, the resulting allocation could yield somewhat less of an allocation to E or more depending on the particular functional form and parameters.

Figure 6: Unit payment case
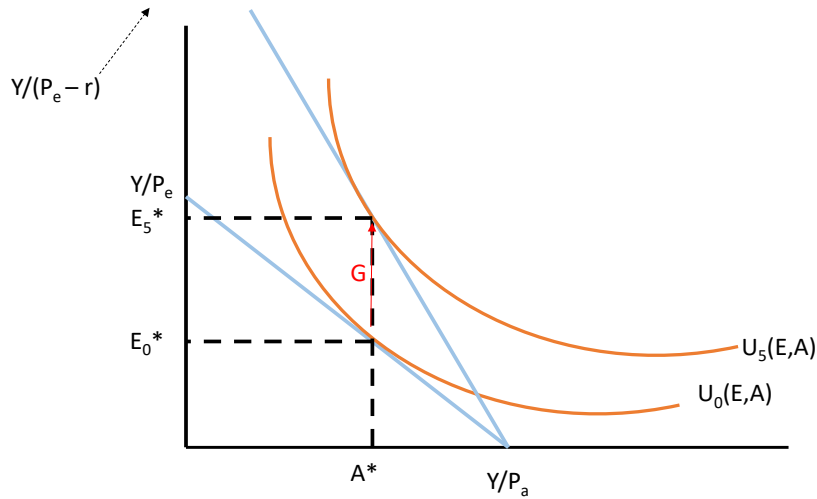
Unit Payment for Additional Production Case

The example above establishes a payment for each and every unit of E produced by the recipient. However, the funder could also offer a contract in which it agrees to pay for each unit of E produced in excess of some initial level. As an example, DFID's RBA program with Ethiopia offered about £150 for each student who completed lower secondary school *above a baseline* established by student completion in a prior year (see Perakis and Savedoff 2015 for a brief description and further references).

Formally, the recipient's costs are now subsidized so the total cost of production becomes:

$$(P_e\,E - r\,(E - E_0)) + P_a\,A$$

The budget constraint faced by the recipient becomes:

$$Y + r\,(E - E_0)$$

and the recipient solves the following function:

$$\text{Max } \mathcal{L} = U(E,A) - \lambda[Y - ((P_e\,E - r\,(E - E_0)) + P_a\,A)].$$

The first order condition for maximization is:

$$\alpha/(1-\alpha) = [(P_e - r) * E)]/(P_a * A)$$

The result is that utility is maximized when the recipient's preference curve is tangent to the budget slope which is now "kinked" at the baseline level E_0 and which, above that baseline level, has a slope of:

$$(Y - rE_0 - P_aA) / (P_e - r)$$

15

The recipient's maximum utility is reached by producing E_6^* which is equal to:

$$E_6^* = \alpha \ (Y - r \ E_0) \ / \ (P_e - r)$$

As in the previous case, a utility maximizing funder would not choose the payment randomly. Rather it would choose the marginal payment r^* to maximize E_6^*. The derivation is provided in an appendix, but requires that r^* be the solution of this quadratic equation:

$$(1 - \alpha) \ E_0 \ r^{*2} + (\alpha \ Y + G - P_e \ E_0) \ r^* - P_e \ G = 0$$

By paying for each additional unit of education under the same aid cap (G), the funder is able to offer a more generous unit payment and therefore shift the relative prices even further. The recipient responds to this price shift by producing more E than in the previous case, even reallocating resources away from A in order to obtain the full amount of G. The recipient is still willing to accept this agreement because its utility is higher than the *No Aid Case*.

Figure 7. Unit payment for additional production case
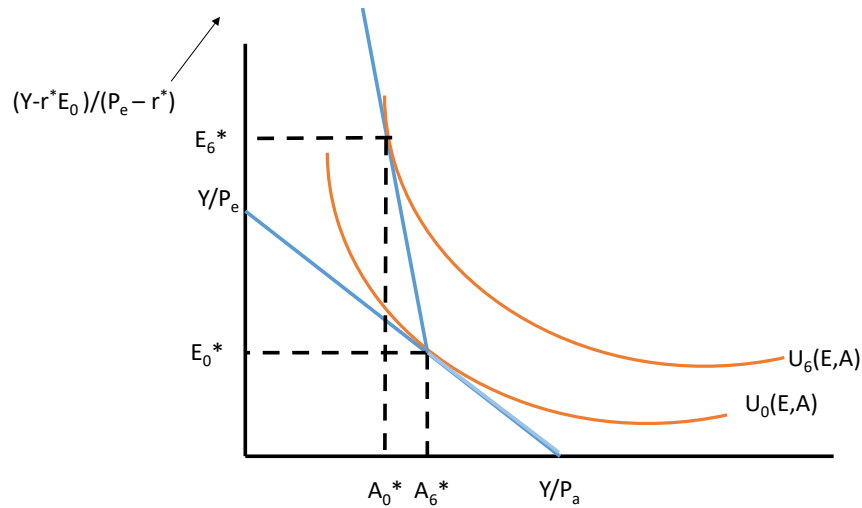
In sum, the recipient views all six of the aid agreements as preferable to the *No Aid Case* and among these, the lump sum grant allows it to get the most utility. The funder will prefer the outcome which generates the most E for the fixed grant amount G, which is the unit payment for additional production. From the funder's perspective, the next best options are the conditional payment, matching grant, and unit payment agreements. The differences among the latter agreements will depend critically on the relative observability of inputs, expenditures and outcomes, as well as the degree to which the recipient can influence the efficiency with which it produces E and A. The effectiveness of such agreements would also depend on the funders' internal politics and broader goals such as sustainability and the weights it gives to the realization of the recipient's own priorities. The next section looks at the welfare results for the different cases in more detail and continues with a discussion of some of these real world considerations.

Discussion

To facilitate a discussion of the previous cases, four parameters are given numerical values and the optimal levels of E and A are calculated along with the resulting utility for the recipient and donor. The results are summarized in Table 2.

Table 2. Summary of cases

	Maximum Utility	Optimal payment per unit	Donor Goal for ΔE	Optimal E	Optimal A	Total Budget Envelope	Total Spend	Recipient Utility	Donor Utility*
Case 0	50			50.0	50.0	1000	1000	50.0	50.0
Case 1	55		10	55.0	55.0	1100	1100	55.0	52.4
Case 2	55		10	55.0	55.0	1100	1100	55.0	52.4
Case 3	55	100	10	60.0	50.0	1100	1100	54.8	54.8
Case 4	55	0.167	10	60.0	50.0	1100	1100	54.8	54.8
Case 5	55	1.667	10	60.0	50.0	1100	1100	54.8	54.8
Case 6	55	4.633	10	71.6	38.4	1100	1100	52.4	59.8

*Assumes that donor utility is 50 in the No Aid scenario and calculates utility relative to that baseline for indicative purposes.

In all cases, the recipient is better off with aid than without. However, the recipient's utility is highest in the cases for which it receives a lump sum grant or input-reimbursement (55). The conditional payment, matching grant, and unit payment agreements yield less utility (54.8); while the unit payment for additional production yields the smallest utility gain relative to no aid at all (from 50 to 52.4). From the funder's perspective, the unit payment for additional production yields the largest utility gain. The funder prefers the lump sum payment and input-reimbursement agreements to the No Aid case, but achieves lower utility than if it offered the other four types of aid. The funder's utility is maximized in the unit payment for additional production case.

The tradeoffs between funder and recipient utility make sense. Recipients gain the highest utility by allocating some share of the additional resources to E and A, both of which it values. By contrast, any funds the recipient allocates to producing A provide the funder with no utility. Though both funder and recipient benefit from the production of E and end up with higher utility, there is a tradeoff between how much net utility is produced and who enjoys it.

As a result, we can make the following observations:

- In a sequential game where the funder's utility function has a fixed share of income going to E, the funder would best achieve its goals by offering a "unit payment for additional output" (Case 6) contract. The recipient would always accept this contract and E would grow at the expense of A.

17

- If it were impossible to set a baseline for the unit payment for additional output agreement, the funder would best achieve its goals with either the conditional payment, matching grant, or unit payment. Setting the parameters of the conditional payment and matching grant agreements would require knowing just how much spending is required to achieve the goal and the funder would have to be able to observe spending with certainty. The only requirement for the unit payment approach is to be able to accurately measure the outcome (E). All three of these cases achieve a significant shift in E without leakage to A.

- With uncertainty over production functions, the lump sum grant and input reimbursement agreements will overestimate potential production of E due to reallocation by the recipient to A. They will also overestimate the potential production of E when expectations of technical efficiency are systematically biased upward.

- In the unit payment contracts (Cases 5 & 6), the funder does not need to know the recipient's commitment to E or precisely estimate utility function parameters. It can offer per-unit subsidies of E and iterate based on the recipient's response. Any failure to produce E, simply saves the funder money which it can re-apply in later iterations.

- If the assumption of homothetic utility functions is relaxed, it will affect the cardinal but not the ordinal ranking of the contract offers.

If we relax the assumptions leading the funder to set a fixed foreign aid grant independent of the amount of E, then the funder can achieve even greater amounts of E. In particular, it could maximize the amount of E by increasing the unit payment from $4.63 to $5.86 which would increase the total grant from $100 to $207, while achieving 85 units of E instead of 71 (see Appendix for details).

Having explored the direct implications of the model as if it were true, what does it imply about the big questions of development aid when considered in a broader context? In particular:

- Should funders avoid or embrace agreements which distort the recipient's choices?
- What happens when aid stops?
- Why do funders prefer paying for input instead of results?

Should Funders Avoid or Embrace Agreements Which Distort the Recipient's Choices?

Whether or not funders are willing to distort domestic recipient choices depends critically on the character of "Anything Else." If the recipient's alternative use of funds is for goods and services that the funder deems positive for development—e.g., roads, energy, healthcare—then the funder may be more open to providing lump sum payments and accepting the subsequent reallocation of resources. If, however, the alternative use of funds is for non-

development activities (e.g., fancy government buildings) or for personal gain (i.e. corruption), then the funder will most likely want to impede the subsequent resource allocation. However, given the fungibility of resources, this cannot be achieved by subsidizing the funder's preferred output; rather, it can only be achieved by accurately measuring spending, outputs or outcomes and linking funding to them.

The funder's willingness to embrace the recipient's allocative choices also depends on its primary motivation. If the funder is motivated by charitable or redistributive impulses, then the recipient's utility might be more important and, again, the funder might be favorable to the lump sum grant. Additionally, if the funder sees the recipient country's development more broadly as a goal, it may prefer lump sum grants as a way to empower the recipient to follow its own development path with additional resources. If, on the other hand, the motivation for aid is expressly to support specific merit goods, like education, then the funder will be less willing to accept the subsequent resource allocation and, as before, the link of funding to spending, outputs or outcomes gives it the strongest chance of success.

What Happens after Aid Stops?

When doing analysis with comparative statics, as in this model, the situation reverts to its original status once the policy change is lifted. The lump-sum grant and input cases don't lead to any change in allocations, so the only impact of stopping aid is to reduce the budget and proportionally reduce production of both goods (A and E). In the case of unit payment contracts, once the funder stops offering aid, the relative costs of the two goods return to their initial ratio and the recipient returns to an allocation based on the no aid case.

In the real world, however, the unit payment cases could change a number of things that might alter the *ex post* situation. First, the relative cost of the two goods (A and E) may change as a result of spending time producing a higher level of E. For example, economies of scale in the production of E could reduce the relative cost of E, in which case the recipient will choose to produce more E even without the aid. Producing more E might require fixed investments (e.g. building schools, designing textbooks, and improving management) that could reduce future marginal costs and thereby also reduce the relative cost of E. Another way that relative costs might permanently change is if the production of more E creates rigidities in the production process, such as expanding employment in a civil service system where it is difficult to dismiss staff. In any of these situations, the production of more E ends up reducing the relative cost of E and sustains the higher rate of production without aid.

Another way that the situation will change is if the higher production of E leads to a permanent shift in the welfare function of the recipient. If, in experiencing more E, the population comes to expect and demand those higher levels, the government may be obliged to provide it. In this case, the utility function might shift by altering the relative preference for E (i.e., a change in the α parameter) or by establishing a minimum required level of E (i.e. $U = (E-\hat{E})^{\alpha}A^{(1-\alpha)}$). Thus, the situation after aid stops, may be very different from the situation *ex ante* as a result of changing allocations during the term of the aid agreement.

Why Do Funders Prefer Paying for Input Instead of Results?

The model presented here suggests that recipients would negotiate for lump sum or input payment agreements, while funders would prefer unit payment agreements. Nevertheless, paying for outputs or outcomes is quite rare in foreign aid—despite several decades of discussion and promotion of the idea. This would be the outcome if recipients had the stronger negotiating position, but that seems unlikely. Instead, it seems that funders have other reasons to prefer paying for inputs over paying for results.

It is plausible that funders choose to pay for inputs rather than results because it provides a narrative that is consistent with the aim (1) to assist poorer countries, (2) to be *seen* supporting poorer countries, (3) to claim that funding is increasing merit goods, and (4) to avoid exposing the divergence between the funder's and the recipient's priorities. Such a narrative helps the funder justify aid to domestic taxpayers, domestic auditors, and the international community.

Paying for inputs satisfies the first two of these aims by assuring that funds are disbursed (regardless of what they achieve). It satisfies the third aim by providing receipts and documentation about spending on the merit good. Finally, it satisfies the fourth aim because the aid agreement can proclaim full agreement in prioritizing the merit good when it is signed; while allowing the recipient, after the fact, to reallocate domestic funds to achieve its preferred allocation. To do this, aid agencies have to continually interpret the failure to generate the desired levels of output as a consequence of some obstacles or hindrances rather than admitting that they might be a rational choice by the recipient which systematically contradicts the aspirational commitments established in the aid agreements.

Conditional payments and matching grants are intermediate between the input reimbursement and unit payment cases. By linking aid to a country's domestic spending on a merit good, it satisfies the four aims above. However, unless the spending truly generates outcomes at the expected rate, these approaches can only guarantee spending not outcomes. Furthermore, they rely critically on the ability to observe spending. Most aid programs rely on financial audits of government accounts which cannot actually prove the funds were spent as reported. That would require performance audits which are not common.

Unit payments are likely to be effective at either increasing production of the merit good or, if the recipient does not increase production, to save the funder money that would go to other purposes. Despite this fundamental advantage, unit payments do not necessarily satisfy the four aims above.

When recipients do increase output and receive payments for results, the funder can satisfy the first two criteria (assisting poorer countries and being *seen* to support poorer countries); however, in a number of cases, fewer results will occur and payments will not occur. Therefore, the funder is likely to have a share of agreements which do not disburse. Even though non-payment is a "success"—in the sense of not providing aid to countries that are failing to achieve increased production of the merit good—the lower disbursements will be

viewed by taxpayers, evaluators, and the international community as lack of commitment to providing aid.[8]

The third aim—to show that funding is dedicated to the increase of a merit good—is undercut by those who argue the *ex post* payment is merely a windfall for the recipient, who would have produced more of the merit good without the incentive.

Finally, paying for results fails to satisfy the fourth aim because it exposes rather than obscures the divergence of goals between funder and recipient. If the recipient claims that producing more E is a priority and this position is matched by an offer of unit payments from a funder, then failing to produce more E could be evidence of obstacles or inconsistency. If failure is due to obstacles, then the funder is criticized for not having provided the technical assistance necessary to overcome those obstacles (as if the aid agency has greater omnipotence than the recipient in resolving obstacles which may be local and political in nature). If the failure cannot be ascribed to obstacles, or the obstacles are clearly within the recipient's scope of control, then failure demonstrates that the statement of priority—no matter how well-intentioned at the time the agreement is signed—is not a real commitment in the sense that the recipient will act upon it.

Qualifications

The main qualification to the results of this paper are that the insights are derived not from empirical cases but from a model. Therefore, the appropriate interpretation is not that the results are fundamentally true but rather should be interpreted in the sense of "if-then." *If* funders and recipients have multiple and overlapping objectives *then* there is an intrinsic tension between the notion of country ownership and divergent preferences over allocations between competing uses of funds.

A further qualification is that countries are not monolithic. The aid relationship may be better characterized as a coalition of like-minded groups or institutions across countries than as a negotiation between coherent uniform positions taken by one country and another. In this case, the model is quite inappropriate and the outcomes of different agreements must be analyzed with a more complex characterization of actors and interests.

The relationships sustained by foreign aid also tend to be longstanding, in which case the relationship bears less relationship to a comparative static exercise and more to a repeat game. The likelihood that two countries will continue to sign subsequent aid agreements creates the possibility of conditioning the amount of future aid on performance under earlier grants even if, within any particular agreement, the payments reimburse the purchase of inputs. However, conditioning future agreements on performance still requires that the

---

[8] An example of this phenomenon can be seen in the nonpayment success of the agreement between Norway and Indonesia to pay for reduced deforestation, see Seymour et al. 2015.

outcomes are measured and that the funder can credibly commit to the follow through—conditions that are not commonly observed.

Finally, few aid agreements aim at a single goal in the way that this model is structured. Funders may provide grants to increase a particular merit good; however, even in such cases, the agreements typically address a wide range of ancillary purposes. For example, even within programs seeking to increase the level of educational attainment, agencies and governments will commit to a range of activities related to improving how the public education system functions. When the aid relationship, in addition, involves programs in other sectors and broader goals regarding such things as public financial management, adherence to democratic norms, and actions in other spheres of geopolitical interest, then the idea that programs are designed with one specific tradeoff in mind becomes untenable.

These qualifications are necessary to avoid treating the model's results as fundamentally true, rather than seeing the model as a useful way of illustrating the implicit tension between country ownership and the divergent aims of funders and recipients. In this spirit, "Essentially, all models are wrong, but some are useful. However, the approximate nature of the model must always be borne in mind...."[9]

Conclusions

In the last two decades, those working in foreign aid have come to acknowledge the agency of those receiving aid and the primacy of their role in achieving goals and promoting development. While written policies and speeches affirm this fact and aim to implement it, in practice, the fundamental structure of foreign aid has not changed and "country ownership" is only achieved when the funder and the recipient's goals truly align or when the funder is willing to let go of its own priorities and endorse the recipient's priorities. In other cases, funders can only achieve their goals by altering the opportunity costs of domestic allocation decisions, which still provides benefits to the recipient but fewer benefits than if the aid had been provided as a pure grant or input payment.

Whether or not funders are willing to accept the recipient country's allocational choices depends critically on the alternative uses of funds. If funders see recipient's allocating funds to goods that they deem positive for development, they are more likely to be comfortable with lump sum or input payment grants. When reallocations go to legitimate purposes that are not considered essential to development or diverted to personal gain, the funder's tolerance is tested. Yet, the only instrument that can truly avoid such reallocations are those that induce a change in behavior by conditioning aid on production of the funder's preferred good.

In practice, funders have preferred to paper over the divergence of interests rather than pay for outputs. Funders are probably reluctant to disburse aid in proportion to outputs or outcomes because the mechanism is not readily reconciled with the prevailing narrative of

---

[9] *Box, G. E. P.; Draper, N. R. 1987. Empirical Model-Building and Response Surfaces, John Wiley & Sons,* p. 424.

foreign aid which holds funders accountable for disbursement and excuses recipients when projects fail by pleading "lack of capacity."

If funders successfully induced recipient countries to allocate more resources toward the funder's preferred outputs, it is unclear whether such a change would be sustainable. Higher output of the merit good is likely to be sustained after the end of foreign aid only if the act of expanding production leads to scale economies, institutional inertia, or entitlements that permanently alter the opportunity costs faced by the recipient country.

The model presented here illustrates how countries that receive aid always "own" their programs in the sense that they ultimately determine what domestic government actions take place and the overall allocation of public resources. This places funders in a bind. They can either raise the costs of aid by establishing elaborate procedures to provide evidence that funding was spent on the funder's priorities or offer aid agreements that alter the opportunity costs faced by governments to induce the funder's preferred outcome. No amount of policy dialogue or negotiation will alter this basic dilemma.

References

Buiter, Williem H. 2007. "Country ownership: a term whose time has gone." *Development in Practice*. 17 (4/5): 647-652.View Article

Laffont, Jean-Jacques, and David Martimort. 2009. *The theory of incentives: the principal-agent model*. Princeton: Princeton University press.

Martens, Bertin, Uwe Mummert, Peter Murrell, and Paul Seabright. 2002. *The institutional economics of foreign aid*. Cambridge University Press.

Murshed, S. Mansoob, and Somnath Sen. 1995. "Aid conditionality and military expenditure reduction in developing countries: models of asymmetric information." *The Economic Journal* 105(429): 498-509.

OECD (Organization of Economic Cooperation and Development). 2008. "The Paris Declaration for aid Effectiveness and the Accra Agenda for Action." Paris: Organization for Economic Co-operation and Development. http://www.oecd.org/document/18/0,3746,en_2649_3236398_35401554_1_1_1_1,00.html

Over, Mead. 1981. "On the Care and Feeding of a Gift Horse: The Recurrent Cost Problem and Optimal Reduction of Recurrent Inputs." Research Memorandum No. 79. Williamstown, MA: Center for Development Economics, Williams College.

Paul, E., 2006. "A survey of the theoretical economic literature on foreign aid." *Asian‐Pacific Economic Literature*, 20(1):1-17.

Perakis, Rita and William D. Savedoff. 2015. "Does Results-Based Aid Change Anything? Pecuniary Interests, Attention, Accountability and Discretion in Four Case Studies," CGD Policy Paper 052, Washington, DC: Center for Global Development. February.

Sappington, David EM. 1991. "Incentives in principal-agent relationships." *Journal of economic Perspectives* 5(2): 45-66.

Savedoff, William D. 2011. "Incentive Proliferation? Making Sense of a New Wave of Development Programs." A CGD Essay. Center for Global Development: Washington, DC.

Savedoff, William D. 2010. "Basic Economics of Results-Based Financing in Health." RBF Health. Washington, DC: World Bank. https://www.rbfhealth.org/sites/rbf/files/RBF%20Economics_0.pdf , last accessed July 23, 2019.

Seymour, Frances, Nancy Birdsall and William Savedoff. 2015. "The Indonesia-Norway REDD+ Agreement: A Glass Half-Full." CGD Policy Paper 56. Washington, DC: Center for Global Development.

Stiglitz, Joseph E. 1989. "Principal and agent." In *Allocation, information and markets*, pp. 241-253. London: Palgrave Macmillan.

Appendix 1. Alternative Model 4 in Which Funder Responds to Higher Effectiveness by Increasing G

This appendix shows what would happen if instead of optimizing over a fixed grant amount, the funder had a utility function that led it to optimize the amount of E that will be produced by allowing the grant to be larger. The equations for the budget line, 1st order condition and optimal amount of E are all identical to Case 6.

- Budget line slope = $(Y - rE_0 - P_aA)/(P_e - r)$
- 1st order condition: $\alpha/(1-\alpha) = [(P_e - r) * E)]/(P_a * A)$
- $E_6^* = \alpha \ (Y - r E_0) /(P_e - r)$

However, the funder solves a different maximization problem by choosing G* that maximizes E_6^*. That is, the funder first solves $P_e E^{*2} - 2 P_e E_0 E^* + \alpha^2 Y$ when $E_0 = 0$ and then solves for r* as

$$r^* = (P_e E^* - \alpha Y)/(E - \alpha E_0)$$

which yields $G^* = r^*E^*$.

In the numerical example, this would require raising the marginal price from 4.63 to 5.86 and thereby doubling the grant from \$100 to \$207, but achieving 85 units of E instead of 71.

If funder solves for G that maximizes E under Case 6 …		
	A	10
	B	-1,000
	C	12,500
	E* (Case 6)	85.36
Solving for r*	r*	5.86
Solving for G*	G*	207.11

In sum, the funder would be able to achieve about 14 more units of education by raising r from 4.63 to 5.86; but the higher unit payment and larger amount of education requires more than doubling the grant from \$100 to \$207.