Working Paper Number 75
January 2006

Aid Project Proliferation and Absorptive Capacity

By David Roodman

Abstract

Much public discussion about foreign aid has focused on whether and how to increase its quantity. But recently aid *quality* has come to the fore, by which is meant the effectiveness of the aid delivery process. This paper focuses on one process problem, the proliferation of aid projects and the associated administrative burden for recipients. It models aid delivery as a set of production activities (projects) with two inputs, the donor's aid and a recipient-side resource, and two outputs, namely, development and "throughput," which proxies for the private benefits for both donor and recipient of implementing projects, from kickbacks to career rewards for disbursing. The donor's allocation of aid across projects is taken as exogenous while the recipient's allocation of its resource is modeled and subject to a budget constraint. Unless the recipient cares purely about development, increasing aid can *reduce* development in some circumstances. Sunk costs, representing the administrative burden for the recipient of donor meetings and reports, are introduced. Using data on the distribution of projects by size and country, simulations of aid increases are run in order to examine how the project distribution evolves, how the recipient's resource allocation responds, and how this affects development if the recipient is not a pure development optimizer. With Cobb-Douglas production, a threshold is revealed beyond which marginal aid effectiveness drops sharply. It occurs when development maximization calls for the recipient to withdraw from some donor-backed projects—but the recipient does not, for the sake of throughput. Donors can push back this threshold by moving to larger projects if there are scale economies in aid projects.

Aid Project Proliferation and Absorptive Capacity[1]

Introduction

Public discussion about foreign aid has long focused on the question of "how much?" Since the early 1960s, some have called on rich countries to give 0.7 percent of their GDP in aid (Clemens and Moss 2005). The proposed International Financing Facility would double aid in the short run. The U.N. Secretary General's report, *Investing in Development*, calls for a tripling (UNDP 2005b). Recently, however, aid *quality* has come to the fore in the public policy debate (Roodman 2004; UNDP 2005a; World Bank 2005; ActionAid 2005; the Paris Declaration of 2005).

Quality can be defined as the capacity, per dollar, for aid to increase development and reduce poverty. In practice, it has much to do with efficiencies, or lack thereof, in the aid delivery process. And while one might debate whether aid quantity ought to rise in any given country, it is hard to argue that aid quality should not. Under the rubric of "aid quality" come a number of themes: untying (dropping requirements that aid be accepted in kind or spent on donor-country goods and services); selectivity (for countries deemed more deserving and propitious as recipients); harmonization (of donors' procurement, reporting, and other requirements); alignment (with the recipient's own goals and plans); coordination (among donors to prevent duplicative efforts); and proliferation (of many small aid projects). In February 2003, representatives of 40 donor agencies signed the Rome Declaration calling for greater harmonization and alignment of aid, coordination, simplification of donor practices, among other

---

[1] Forthcoming in George Mavrotas and Mark McGillivray, eds., *Development Aid: A Fresh Look* (Palgrave-Macmillan). For excellent comments, the author thanks Peter Timmer, William Cline, and participants in three seminars—at the January 2005 conference of the Global Development Network in Dakar, at the Center for Global Development in February, and at World Institute for Development Economics Research in September.

things. They reconvened in 2005 and issued the Paris Declaration, which defines a dozen indicators of a donor's aid quality and commits donors to assist in collecting them..

This paper focuses on the problem of project proliferation, which is thought to impose great administrative burdens on some recipient governments. Using available data on aid activities it first shows why one should worry that there are too many aid projects. It then defines a model in which a recipient's resources combine with aid in projects to produce development. If donor and recipient do not both maximize development, but rather such things as meeting per diems and the career benefits of disbursement for aid officials, collectively labeled "throughput," then increasing aid can reduce development. The extent to which this happens depends on the degree to which throughput and development are not complements.

The paper then moves to the continuous setting, positing a suite of projects lognormally distributed by aid size. Sunk costs are also introduced to represent the recipient-side cost of donor meetings and quarterly reports and economies of scale. A simulation is run to illustrate how the strategy that the recipient uses to allocate its resource across projects varies as total aid expands. The simulation also shows how this affects development. The relationship between total aid and the distribution of projects by size is calibrated using detailed data on aid projects covering most donor and recipient countries. When total aid increases beyond a certain point and projects proliferate in the way suggested by the data, the effective marginal product of aid, factoring in recipient behavior, declines sharply, and can even go negative. Thus the model hints at the existence of thresholds beyond which aid becomes much less effective in practice; it gives rise to a notion of absorptive capacity.

Background and motivation

According to Elliott Morss (1984), during the 1950s and 1960s most foreign assistance took the form of program aid, by which he means large infrastructure investments as well as packages of sector support, such as for agriculture, that include finance, commodities, and technical assistance. However, as concerns grew about the effectiveness of aid, legislatures demanded more evidence of results. This led by the 1970s to greater use of project aid, which "entails a more specific statement of objectives and means." The multiplication of aid projects with more narrowly and precisely defined goals and more measurable outcomes was compounded by the multiplication of donors after Western European nations and Japan recovered from World War II. Morrs identifies several troubling consequences, including lack of coordination on the donor side and lack of ownership on the recipient side. In addition is the administrative burden associated with the sheer number of projects: "[E]fforts to implement the large number of discrete, donor-financed projects, each with its own specific objectives and reporting requirements, use up far more time and effort than is appropriate."

There is no sign that the administrative burden has lessened in the 20 years since Morrs wrote. Van de Walle and Johnston (1996) reported roughly 60 active donors and 600 ongoing projects each in Kenya and Zambia in the mid-1980s, and some 40 donors and 2,000 aid projects in Tanzania in the mid-1990s. Starting from a hypothetical project count of 600, they suggest that recipient governments typically file 2,400 quarterly reports to donors and host 1,000 "missions" from donor officials to monitor project activities. These two numbers appear to have been picked up by a speechwriter for then–World Bank President James Wolfensohn and misinterpreted as a fact about Tanzania, leading to the urban legend that Tanzania files 2,400 reports and hosts 1,000 missions each year. If anything, the figures are substantial underestimates for Tanzania.

In fact, in the spring of 2003, the Tanzanian ministry of finance took a striking step to manage the administrative burdens associated with aid. It announced that the period April-

August each year would be a "quiet time" during which only the most urgent donor missions would be received. Tanzanian officials were to use this time to prepare the central government budget. In addition, as a sort of naming-and-shaming exercise, the government began posting on the Web a cumulative list of major meetings with donors.[2]

Data suggest that the project proliferation problem extends beyond Tanzania and might be worsening. Any effort to count projects must begin with definitions because, like most important concepts, that of a "project" is complex on close examination. Is an organized effort to build ten schools ten projects or one? Where does one draw the line between large road-building projects and "programs" of support to the transportation sector? The definition used in this paper is partly principled, partly pragmatic, driven by the structure of available data. The best relevant and available data source is the Creditor Report System (CRS) database maintained by the Development Assistance Committee (DAC) in Paris. CRS Table 1 contains detailed information on individual aid commitments by bilateral and multilateral donors to fund what the CRS reporting directives refer to as "aid activities" (DAC 2002).[3] The CRS guidelines for donors reporting to the database define an aid activity only to this extent:

> An aid activity can take many forms. It could be a project or a programme, a cash transfer or delivery of goods, a training course or a research project, a debt relief operation or a contribution to an NGO (DAC 2002).

The data begin in 1973. Despite the "creditor" in its name, the database covers both grant and loan commitments, including nonconcessional ones, and even a few equity investments.

With respect to this data set, a "project" is defined here as an entry in CRS Table 1:

- that is not an equity investment;

[2] See http://www.tzdac.or.tz/Mission%20calendar.doc.
[3] CRS Table 5, which begins much more recently, has data on disbursement rather than commitments. This might seem more relevant since not all commitments are realized. But the reporting concept for this table is the financial transaction rather than the aid activity, and there can be many transactions per activity

- that fits the definition of Overseas Development Assistance, meaning that it is a grant or adequately concessional loan for a development purpose;
- whose recipient is identified as a specific country, as opposed to, say, "Africa unspecified";
- is not identified as being for administrative costs or support for nongovernmental organizations;
- is not identified as being emergency aid.

Somewhat confusingly, this definition includes Sector-Wide Action Programs and budget support, which Morss calls program aid. These are in effect very large projects, assuming commitments to them are in fact large in dollar terms. It seems appropriate to include programs since excluding them (to the extent possible with a CRS coding system not designed for the purpose) could paint a misleading picture of how much individual donors tend to proliferate their overall aid portfolios.

Table 1 shows that the number of projects in the database according to this definition nearly tripled between 1995 and 2003. Better reporting—more donors providing data, and on larger fractions of their portfolios—may account for the bulk of this increase. For this reason, comparisons based on cross-sections of the CRS database may be more meaningful than those based on time series. Table 2 therefore reports the top-ten recipients of project commitments during 2001–03. It aggregates over three years on the idea that this proxies better for ongoing activities in 2003: commitments tend to lead to disbursements and project operations over several years.[4]

[4] Examination of extracts from the World Bank Development Gateway's Accessible Information on Development Activities (AiDA) database, suggested this value for average project duration. The AiDA database draws on the CRS and other sources, and contains project start and stop dates for some entries, unlike the CRS. However, it is inferior in other respects for the purposes of this study.

Table 1. Number of reported project commitments, 1995–2003, all donors

Year	Number
1995	10,327
1996	10,626
1997	10,310
1998	14,790
1999	20,692
2000	20,847
2001	28,739
2002	25,716
2003	27,876

Table 2. Number of reported project commitments, 2001–03, top ten recipients

Country	Number
Mozambique	1,921
India	1,910
China	1,885
Russia	1,721
Ethiopia	1,677
Indonesia	1,639
Vietnam	1,609
Tanzania	1,528
Serbia & Montenegro	1,497
South Africa	1,466

Scholarly interest in proliferation also appears to be rising. Arnab Acharya, Ana Fuzzo de Lima, and Mick Moore (2006) develop indices of donors' tendency to proliferate (disperse) aid among recipients, and of the tendency of recipients' aid to be fragmented among many donors. They find that the donors that are the greatest proliferators are especially likely to aid the countries with the greatest fragmentation.

Stepenen Knack and Aminur Rahman (2004) investigate an index of fragmentation similar to that of Acharya, de Lima, and Moore as a determinant of another variable of interest, government bureaucratic quality. In their model, donors compete with each other and the government for the scarce resource of skilled nationals. Hiring a skilled professional away from the government reduces the quality of public governance, which is a public-good input to all aid projects. This creates an externality. The lower a donor's share in the recipient's aid "market," the less it internalizes this externality, and the more incentive it has to poach the best people from the government. Knack and Rahman thus predict—and appear to confirm empirically—that aid fragmentation reduces the quality of public bureaucracy.

Both of these contributions relate to the proliferation of donors, not projects. Though donor and project proliferation no doubt go hand in hand, they are distinct notions. Countries with few donors can have many projects, and vice versa. The present study therefore focuses on project proliferation and illustrates how it can limit the ability of recipient countries to absorb aid effectively.

Thus another literature is relevant, that on absorptive capacity. According to a review in by Michael Clemens and Steven Radelet (2003), a number of studies have attempted to measure absorptive capacity for aid via inclusion of linear and quadratic aid terms in growth regressions (Hadjimichael et al. 1995; Durbarry, Gemmell, and Greenaway 1998; Hansen and Tarp 2000, 2001; Hansen 2001; Lensink and White 2001; Collier and Dollar 2002; Dalgaard, Hansen, and Tarp 2002; Clemens, Radelet, and Bhavnani 2004). Clemens and Radelet report that the implied turning points in the marginal impact of aid range from 15% to 45% of recipient GDP.

Many pathways have been suggested that would cause the marginal productivity of aid to fall and even go negative as aid increases. These can be organized along the sequence of aid flow from donor to recipient central government to field implementation.[5] Donor practices—imposition of locally inappropriate project designs, internal pressure to disburse, heavy demands for meetings and reports, incompatible fiscal years and reporting requirements, and so on—are one large set of potential problems. On the recipient side, at the national level, aid discourages tax effort, distorts domestic political economy, and can cause Dutch Disease. And recipients too have administrative problems, including small and often overwhelmed line ministries, and

[5] I think Peter Timmer for this organizing principle.

corruption. Meanwhile, aid can siphon the best people away from government, as in Knack and Rahman. Many of these recipient administrative problems also pertain at the project implementation level. Finally, as the aid flow increases, monitoring and evaluation become increasingly important in order to detect the problems just listed and help ensure coordination among donors. But good monitoring and evaluation is generally lacking.

The effect modeled here has to do with the limits imposed by recipient administrative capacity, whether at the center or in the field. Implicitly, it also about the failure of donors to monitor, evaluate, and coordinate in order to limit aid project proliferation.

A microeconomic model of aid projects

The model developed here depicts the aid process as a set of production activities, one per project, with identical technologies. The key ideas in the model are:

- There is one recipient.
- Aid projects have two inputs—aid and a recipient-side resource that can be thought of as spending on capital or recurring costs or the time of officials in recipient ministries.
- The recipient has a fixed budget for its resource, which it is free to allocate among aid projects to maximize its utility.
- The donors' portfolio of projects is taken as exogenously determined.
- Aid projects have two outputs too. One is "development," which can be thought of as growth and poverty reduction. The other is "throughput" and is meant to capture the more direct benefits of aid projects to the officials involved. On the donor side, this may encompass the political dynamic of aid tying and the career benefits of being associated with disbursing projects. On the recipient side, it may include similar professional benefits, as well as the high salaries that can come with working with donors, "sitting

fees" for going to meetings, travel, per diems, and so on. On both sides, it could also reflect dynamics of corruption.

- Each aid project produces its two outputs—development and throughput—simultaneously. Throughput and development are generally complements, but not exactly so. Development happens in the recipient countries, while throughput can accrue to both donors and recipients.
- Outputs are additive. If two projects would each by themselves produce one unit of development, then together they produce two.
- In general, both donor and recipient care about both development and throughput, but can weigh them differently.
- The model allows for sunk costs on the recipient side—investments in hosting meetings, filing reports, etc., that do not directly contribute to production but consume recipient resources and are required for aid to flow.

Appendix 1 describes the general model. We confine ourselves here to two pairs of relatively intuitive examples. In the first pair, the recipient is purely developmentalist. In the second pair, it values throughput only.

Example 1. The recipient is purely developmentalist and the donor increases aid to one project. There are two projects. Production of both development and throughput is Cobb-Douglas, with constant returns to scale in the case of throughput. Let

$$
\begin{gathered}
D_i = A_i^{\alpha_D} R_i^{\rho_D} \\
B_i = A_i^{\alpha_B} R_i^{\rho_B} \\
u_R(D,B) = D \\
0 < \alpha_D, \rho_D, \alpha_B, \rho_B \le 1 \\
\alpha_B + \rho_B = 1,
\end{gathered}
$$

where A_i and R_i are in the inputs of aid and the recipient resource into project i, D_i and B_i are the outputs of development and throughput ("B" for private benefit), D and B are total development and throughput, and u_R is the recipient's utility. The recipient's problem is

$$\max_{\mathbf{R}} u_R(D, B)$$

$$\text{such that } R_1 + R_2 \le R \text{ and } R_1, R_2 \ge 0.$$

Working with the Lagrangian, the first-order conditions for the recipient's optimum work out to

$$\rho_D A_1^{\alpha_D} R_1^{\rho_D - 1} = \rho_D A_2^{\alpha_D} R_2^{\rho_D - 1} = \lambda$$
$$R_1 + R_2 = R$$

and have the solution

$$\hat{R}_i = \frac{R}{A_1^{\gamma_D} + A_2^{\gamma_D}} A_i^{\gamma_D} \tag{1}$$

where

$$\gamma_D = \frac{\alpha_D}{1 - \rho_D}.$$

Notice that γ_D is 1 exactly when there are constant returns to scale in development ($\alpha_D + \rho_D = 1$) and greater (less) than 1 when there are economies (diseconomies) of scale. Because the $A_i^{\gamma_D}$ term in (1) is the only one that varies with i, it shapes the allocation: when there are economies of scale in development, the recipient allocates its resources more than proportionally to the projects with larger aid budgets, but does the opposite if there are scale diseconomies. When there are constant returns to scale in development, the recipient allocates its resource in direct proportion to the aid funding for each project

We now examine how the production of development varies when the donor increases aid to project 1 while holding aid to the project 2 fixed. Applying the chain rule to the equation for D, the quantity of interest is

$$\frac{\partial \hat{D}}{\partial A_1} = \frac{\partial D_1}{\partial A_1} + \frac{\partial D_1}{\partial R_1}\frac{\partial \hat{R}_1}{\partial A_1} + \frac{\partial D_2}{\partial R_2}\frac{\partial \hat{R}_2}{\partial A_1}.$$

This says that the aid increase affects development in three ways: directly, through increasing the output from project 1, and indirectly, by causing the recipient to change its resource allocation to both project 1 and project 2. However, since the recipient's current resource allocation maximizes *D*, it is just at the point where a marginal shift in resources between the two projects has zero net impact on *D*—otherwise *D* would not currently be maximized. So, as in equation (7) in the Appendix, the indirect effects cancel, leaving only the first term. Thus

$$\frac{\partial \hat{D}}{\partial A_1} = \frac{\partial D_1}{\partial A_1} = \alpha_D A_1^{\alpha_D - 1} \hat{R}_1^{\rho_D}$$

Substituting for $\hat{R}_1$ with (1),

$$\frac{\partial \hat{D}}{\partial A_1} = \alpha_D A_1^{\alpha_D - 1 + \gamma_D \rho_D} \left(\frac{R}{A_1^{\gamma_D} + A_2^{\gamma_D}} \right)^{\rho_D} = \alpha_D A_1^{\gamma_D - 1} \left(\frac{R}{A_1^{\gamma_D} + A_2^{\gamma_D}} \right)^{\rho_D}.$$

(This uses the fact that $\alpha_D + \gamma_D\beta_D = \gamma_D$.) This expression is always positive.

Conclusion: if the recipient is purely developmentalist, expanding the aid pie never hurts. Adding aid expands production possibilities for the recipient, which can only cause development to increase. This is true even if the aid increase brings thousands of tiny new projects.

Example 2. The recipient is purely developmentalist and the donor moves aid from one project to another. We use the set-up from Example 1. But this time, the donor moves aid from project 2 to project 1 while keeping total aid constant. In this case,

$$\frac{\partial \hat{D}}{\partial A_1} = \frac{\partial D_1}{\partial A_1} - \frac{\partial D_2}{\partial A_2} = \alpha_D A_1^{\gamma_D - 1} \left(\frac{R}{A_1^{\gamma_D} + A_2^{\gamma_D}} \right)^{\rho_D} - \alpha_D A_2^{\gamma_D - 1} \left(\frac{R}{A_1^{\gamma_D} + A_2^{\gamma_D}} \right)^{\rho_D} = \alpha_D \left(\frac{R}{A_1^{\gamma_D} + A_2^{\gamma_D}} \right)^{\rho_D} \left(A_1^{\gamma_D - 1} - A_2^{\gamma_D - 1} \right)$$

This is positive exactly when

$$A_1^{\gamma_D - 1} > A_2^{\gamma_D - 1},$$

which is to say, when $A_1 > A_2$ and $\gamma_D > 1$ (economies of scale dominate in development) or when $A_1 < A_2$ and $\gamma_D < 1$ (diseconomies dominate).

Conclusion: When a donor moves aid from one project to another, the recipient reallocates its resources in the same direction, and if the gaining project is more economical in scale than the old one, development goes up. Otherwise, it goes down.

<u>Example 3. The recipient cares only about throughput and the donor increases aid to one project.</u> Again, we slightly modify the set-up in Example 1, this time by setting

$$u(D, B) = B.$$

The first-order conditions are nearly identical to those before:

$$\rho_B A_1^{\alpha_B} R_1^{\rho_B - 1} = \rho_B A_2^{\alpha_B} R_2^{\rho_B - 1} = \lambda$$
$$R_1 + R_2 = R.$$

The solution is the same too except that an analogously defined γ_B replaces γ_D. In this case, $\gamma_B = 1$ because we are still assuming that $\alpha_B + \rho_B = 1$, so the recipient's solution is just:

$$\hat{R}_i = \frac{R}{A} A_i \tag{2}$$

The recipient allocates its resource in direct proportion to the donor's aid to each project. If the donor increases aid to project 1 while fixing project 2's aid, then

$$\frac{\partial \hat{R}_1}{\partial A_1} = \frac{\partial}{\partial A_1}\left[R \frac{A_1}{A_1 + A_2} \right] = \frac{R A_2}{\left(A_1 + A_2\right)^2} = \frac{A_2}{A}\frac{R}{A}$$
$$\frac{\partial \hat{R}_2}{\partial A_1} = -\frac{\partial \hat{R}_1}{\partial A_1},$$

the second formula following from the fixed budget constraint. The quantity of interest is

$$\frac{\partial \hat{D}}{\partial A_1} = \frac{\partial D_1}{\partial A_1} + \frac{\partial D_1}{\partial R_1}\frac{\partial \hat{R}_1}{\partial A_1} + \frac{\partial D_2}{\partial R_2}\frac{\partial \hat{R}_2}{\partial A_1}$$
$$= \alpha_D A_1^{\alpha_D - 1} R_1^{\rho_D} + \rho_D A_1^{\alpha_D} R_1^{\rho_D - 1}\frac{A_2}{A}\frac{R}{A} - \rho_D A_2^{\alpha_D} R_2^{\rho_D - 1}\frac{A_2}{A}\frac{R}{A}$$

Substituting with the formula for $\hat{R}_i$ in (2),

$$\frac{\partial \hat{D}}{\partial A_1} = \alpha_D \left(\frac{R}{A}\right)^{\rho_D} A_1^{\alpha_D + \rho_D - 1} + \rho_D \left(\frac{R}{A}\right)^{\rho_D} A_1^{\alpha_D + \rho_D - 1}\frac{A_2}{A} - \rho_D \left(\frac{R}{A}\right)^{\rho_D} A_2^{\alpha_D + \rho_D - 1}\frac{A_2}{A}. \quad (3)$$

Rearrangement (the details are omitted) shows that this is greater than 0 if and only if

$$1 + \frac{\alpha_D}{\rho_D} > \frac{\dfrac{A_1}{\frac{1}{2}\left(A_1 + A_2\right)}}{\dfrac{A_1^{\alpha_D + \rho_D}}{\frac{1}{2}\left(A_1^{\alpha_D + \rho_D} + A_2^{\alpha_D + \rho_D}\right)}}.$$

It can be shown that the numerator of the grand fraction on the right is less than the denominator exactly when $A_1 > A_2$ and $\alpha_D + \rho_D > 1$, or when the opposite is true on both counts—that is, when A_1, the project gaining aid, has a more economical scale than A_2. If this holds, it guarantees the inequality since the left side is always at least 1. Then $\partial \hat{D} / \partial A_1 > 0$ and increasing aid to project 1 increases development.

But the inequality can fail in other cases—most easily when recipient-side resources are the dominant factor in the production of development (α_D is much less than ρ_D, so $\alpha_D/\rho_D + 1$ is low) and project 1, the one gaining aid, is far less economically scaled than project 2. In this case, the first right-hand term of (3), the direct effect of the aid increase on development from project 1, is dwarfed by the second and third terms, which capture the indirect effect of the recipient's resource reallocation. And these latter terms can sum to a negative value if the recipient is reallocating resources toward the project where marginal total productivity is lower.

Conclusion: If the recipient cares only about throughput, increasing aid can *reduce* development if recipient-side resources are an important ingredient in development and the project receiving the increase is of a relatively uneconomical scale.

If the recipient cares both about throughput and development rather than throughput alone, the same result should hold but the mathematics are more complex. Thus, to generalize, if the recipient is not purely developmentalist, increasing aid to uneconomically scaled projects can reduce development.

Example 4. The recipient cares only about throughput and the donor moves aid from one project to another

This time, the derivative of development with respect to the (increasing) aid to project 1 is

$$\left.\frac{\partial \hat{D}}{\partial A_1}\right|_{A_1+A_2=A} = \frac{\partial D_1}{\partial A_1} + \frac{\partial D_2}{\partial A_2}\frac{\partial A_2}{\partial A_1} + \frac{\partial D_1}{\partial R_1}\frac{\partial \hat{R}_1}{\partial A_1} + \frac{\partial D_2}{\partial R_2}\frac{\partial \hat{R}_2}{\partial A_1}$$

$$= \alpha_D A_1^{\alpha_D - 1} R_1^{\rho_D} - \alpha_D A_2^{\alpha_D - 1} R_2^{\rho_D} + \rho_D A_1^{\alpha_D} R_1^{\rho_D - 1}\frac{R}{A} - \rho_D A_2^{\alpha_D} R_2^{\rho_D - 1}\frac{R}{A}$$

$$= \alpha_D \left(\frac{R}{A}\right)^{\rho_D}\left(A_1^{\alpha_D+\rho_D-1} - A_2^{\alpha_D+\rho_D-1}\right) + \rho_D \left(\frac{R}{A}\right)^{\rho_D}\left(A_1^{\alpha_D+\rho_D-1} - A_2^{\alpha_D+\rho_D-1}\right)$$

which is positive exactly when

$$A_1^{\alpha_D+\rho_D-1} > A_2^{\alpha_D+\rho_D-1}.$$

This occurs under the same circumstances as in Example 2.

Conclusion: If the recipient cares only about throughput, when a donor moves aid from a project with lower marginal productivity to one with higher, development goes up. Otherwise, it goes down.

Note that this conclusion is underpinned by the assumption that production of throughput has constant returns to scale. If there are large enough diseconomies of scale in throughput, then it is possible for the recipient to reallocate resources in the opposite direction from the donor, so that development can go down even when the donor is shifting aid to a more economically scaled project.

The empirical distribution of aid projects

This section examines the empirical distribution of aid projects more closely, again using the CRS database. The purpose is to derive an empirical relationship between the total aid going to a country and the distribution of its projects by size. This will then serve to calibrate a simulation using a more sophisticated example of the models just presented.

The distribution of projects by size, both within and across countries, follows clear patterns. In general, the distribution in a given country or group of countries is unimodal, of course has support above zero, and skews to the right. The lognormal distribution is therefore a promising model for the distribution, as Figures 1–3 confirm. They show the distributions of project size on a logarithmic scale, for all countries together, and for China and Tanzania alone. Project sizes are in thousands of dollars. As in Table 2, tallies are for the three most recent years of data.

Figure 1. Distribution of project commitments by size, all countries, 2001–03

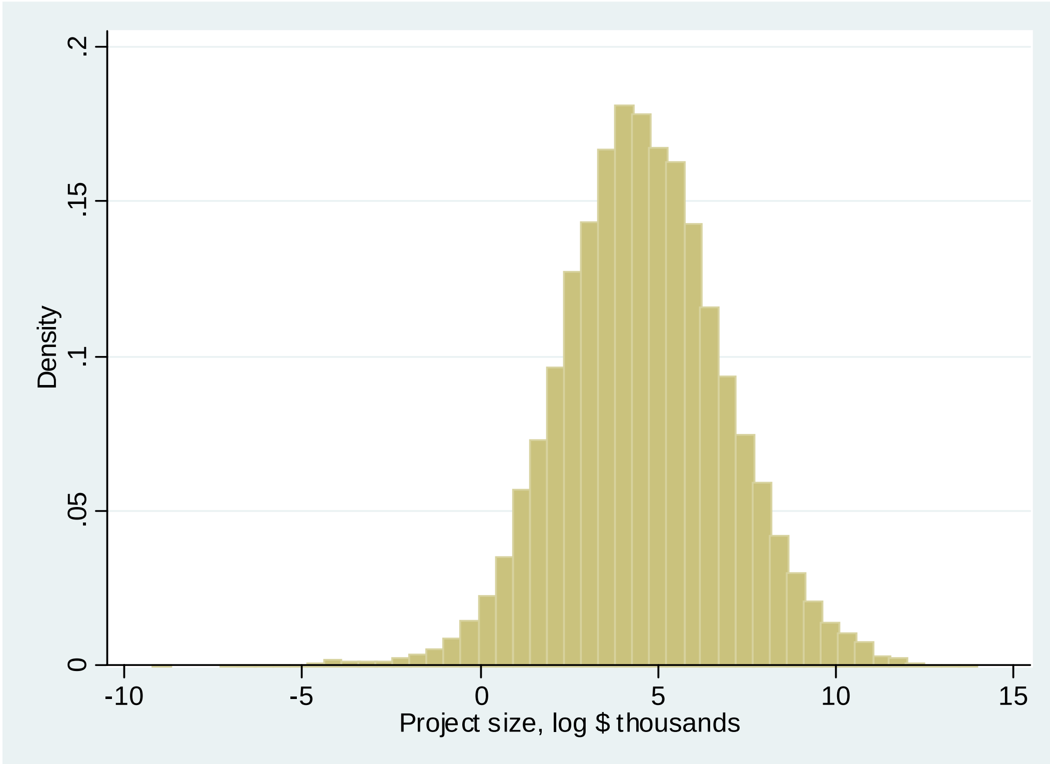

Figure 2. Distribution of projects by size, Tanzania, 2001–03

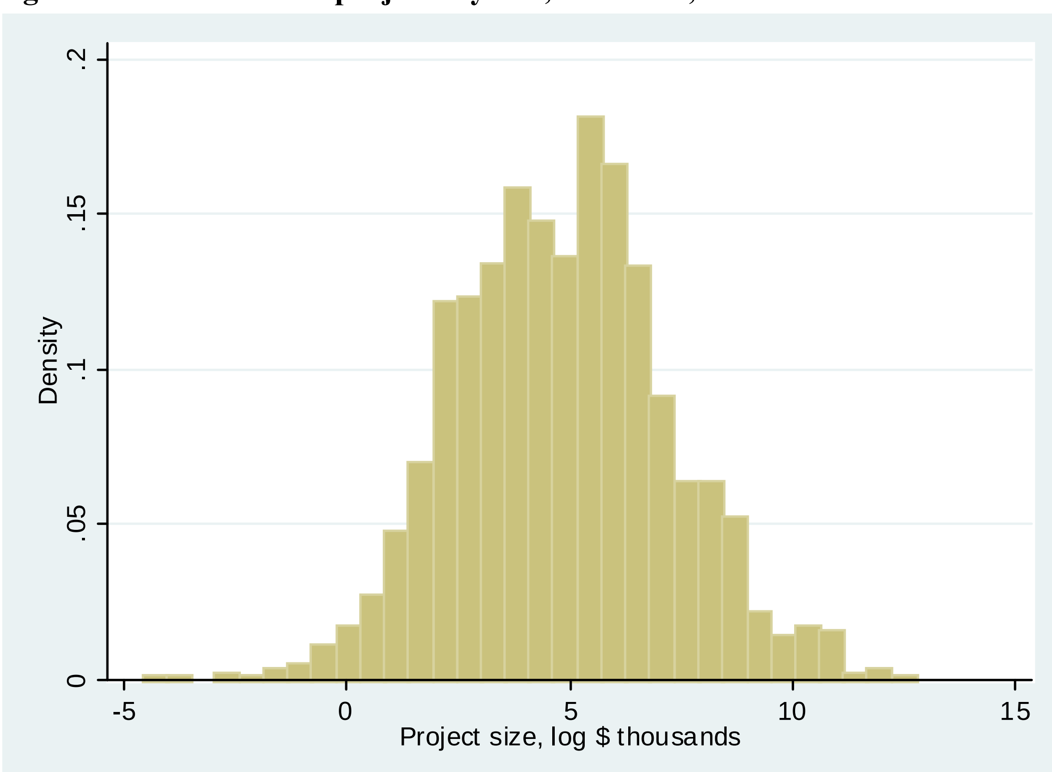

Figure 3. Distribution of projects by size, China, 2001–03

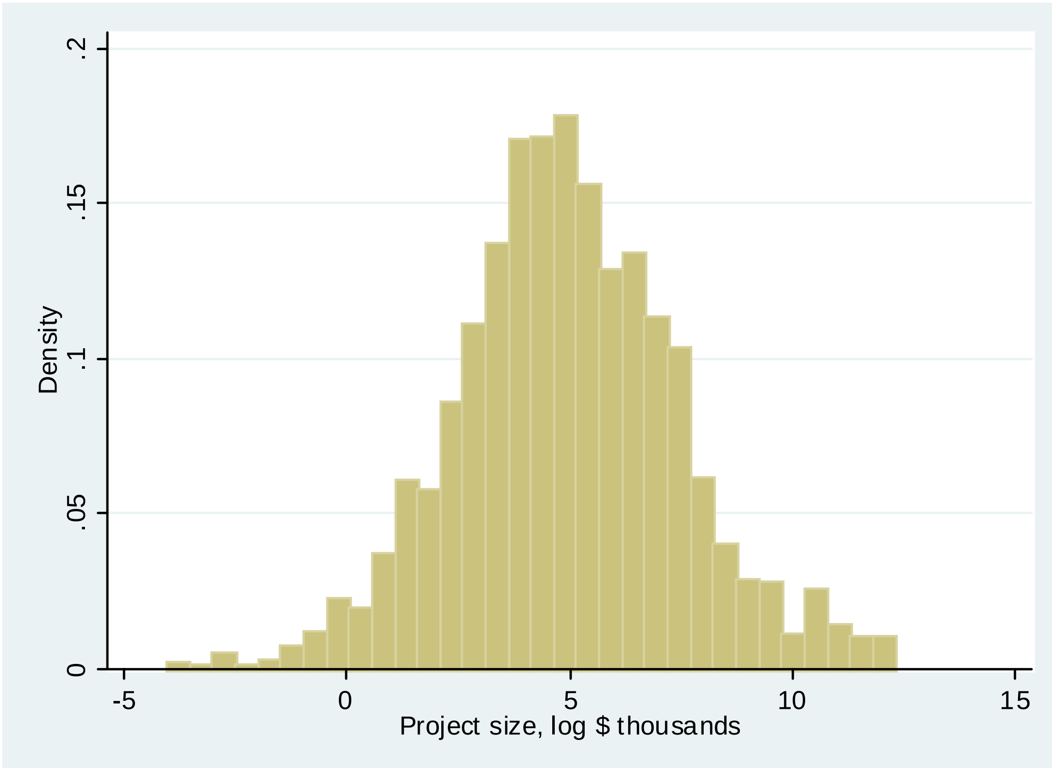

Given the quality of the lognormal fit, the distribution of projects by size in a given country is well characterized by the mean and standard deviation of the corresponding distribution in log space, μ and σ. Moreover, these turn out to be strongly correlated with total project commitments. As a result, one can make a good prediction of the distribution of projects by size in a country given only the total amount, A, of aid committed to those projects.

To formalize this idea and rigorously estimate the best-fit lines, I fit the following parameterized distribution to the full 2001–03 data set using maximum likelihood estimation:

$$\begin{aligned} \mu &= \mu_0 + c_\mu \ln A \\ \sigma &= \sigma_0 + c_\sigma \ln A \\ h(x) &= \frac{1}{\sqrt{2\pi}\sigma x} e^{-\frac{1}{2\sigma^2}(\ln x - \mu)^2} \end{aligned} \qquad (4)$$

where x is the size of a project, A is total project funding commitments in thousands of dollars for 2001–03, μ_0, σ_0, c_μ, and c_σ are constants, and $h(x)$ is the (lognormal) probability density for projects of size x.

I run the estimation two ways. In the first, there is a single lognormal project distribution for each recipient, as in Figures 1–3. The variables μ, σ, and A all refer to the full set of projects in a given recipient country. In the second, the sparser distributions of each donor-recipient pair are modeled, to allow for donor-level heterogeneity. Here, μ, σ, and A, refer to the distribution of projects of a given donor in a given recipient country. (See Table 3.) Almost all parameters are strongly different from 0. The parameters c_μ and c_σ are significantly higher in the second regression, meaning that for a given donor in a given recipient country, the project distribution shifts more quickly for a given aid increase than it does when aggregating across all donors.

Table 3. Maximum likelihood estimates of model parameters

	By recipient	By donor and recipient
c_μ	0.2707	0.3642
	(0.0065)**	(0.0036)**
μ_0	0.0546	0.0729
	(0.0819)	(0.0325)*
c_σ	0.1058	0.1284
	(0.0043)**	(0.0025)**
σ_0	0.9580	0.9859
	(0.0541)**	(0.0223)**
N	79,426	79,426

Standard errors in parentheses. Sample excludes Israel and recipients with population below 1 million.
* significant at 5%; ** significant at 1%

The regression modeling a single distribution for each recipient is the basis for the fit lines in Figures 4–6. These show the empirical average and standard deviation of log project size by recipient for the 2001–03 data, as well as the number of projects. Data points are labeled with 3-letter ISO country codes. The best-fit curve in the final graph, which plots the number of projects against total aid is computed using the fact that the average of a lognormal distribution is $e^{\mu+\sigma^2/2}$. (Aitchison and Brown 1963, p. 8), so that

$$N = \frac{A}{e^{\mu + \frac{\sigma^2}{2}}}.$$

According to the upper-left cell of Table 3, a factor-of-10 increase in total project aid to a country lifts the average log project size by

$$0.2707 \times \ln 10 = 0.6233$$

which is equivalent to multiplying the geometric average of project size (which in a lognormal distribution is also the median, e^{μ}) by a factor of

$$e^{0.6233} = 1.865.$$

Intuitively speaking, a 10-fold increase in total aid to a country is associated in the data with an 86.5% rise in representative project size. Since average project size rises much more slowly than total aid, the *number* of projects goes up. Along the same lines, regression 2 suggests that if a single donor increases its aid for a country 10-fold, its own typical project size goes up 131%. Of course these statistics tell us little about the true direction of causality. But they are useful for simulations.

Figure 4. Average log project size vs. total project funding, by recipient, 2001–03

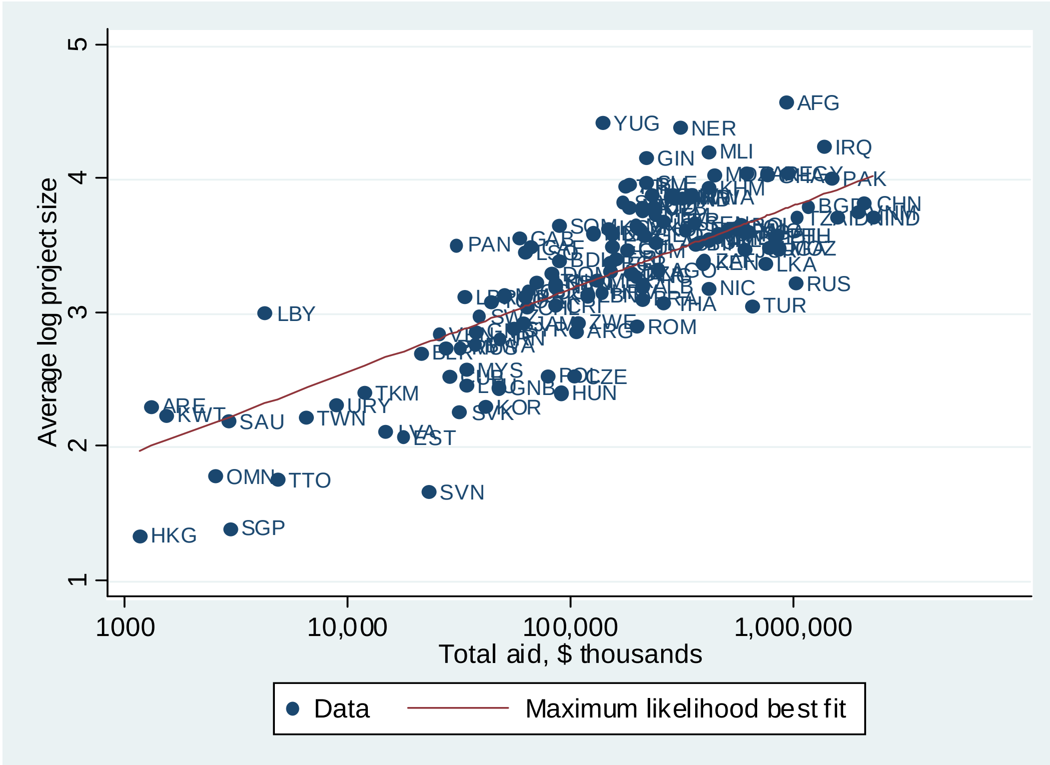

Figure 5. Standard deviation of log project size vs. total project funding, by recipient, 2001–03

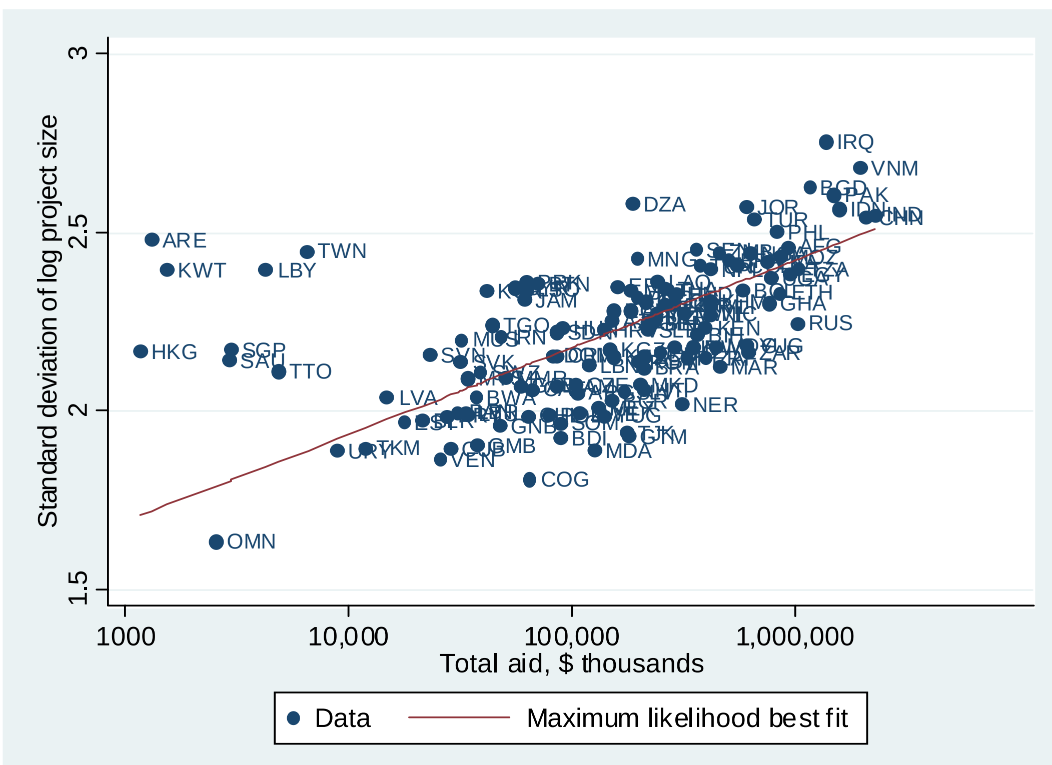

Figure 6. Number of projects vs. total project funding, by recipient, 2001–03

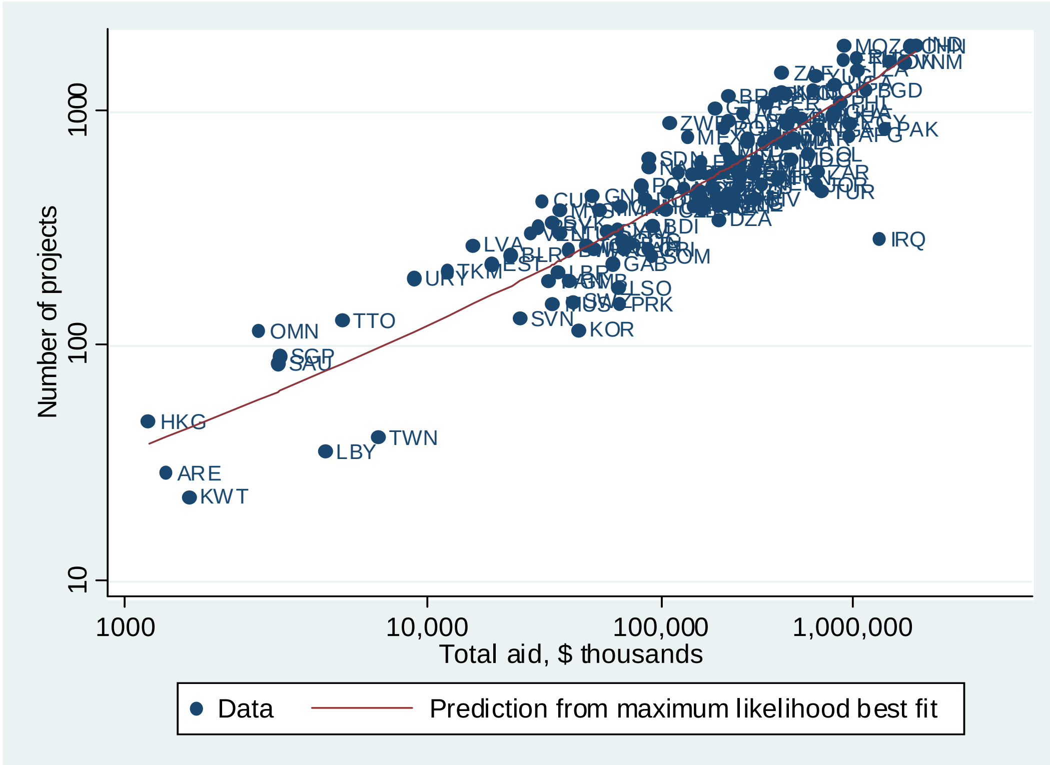

A simulation with sunk costs

This section develops the earlier examples in two important ways. It introduces a notion of sunk cost. And it adapts the model to the continuous setting. It then reports the results of simulations partially calibrated using the empirical patterns just described.

All of the earlier examples assume that the first penny the recipient invests in aid projects increases output of development and throughput. But it seems likely that aid projects have significant sunk costs that the recipient must cover for the project to proceed—meeting with donors in the capital, taking them on field visits, filing reports, and so on. We might expect sunk cost to rise with project size, but not as fast, according to a sunk cost function $s(A_i)$. This is captured in the following modified Cobb-Douglas production functions:

$$D_i = \begin{cases} 0 & \text{if } R_i \le s(A_i) \\ A_i^{\alpha_D}\left(R_i - s(A_i)\right)^{\rho_D} & \text{if } R_i > s(A_i) \end{cases}$$
$$B_i = \begin{cases} 0 & \text{if } R_i \le s(A_i) \\ A_i^{\alpha_B}\left(R_i - s(A_i)\right)^{\rho_B} & \text{if } R_i > s(A_i) \end{cases}$$
$$s \ge 0, s' \ge 0, s'' \le 0 \text{ everywhere}$$

Unlike in the earlier examples, there will be cases in which the recipient can maximize its utility by not funding certain projects. Indeed, it may not be able to *afford* to put resources into all projects. In general, the recipient will invest either 0 or more than $s(A_i)$ in each project, since investing less than the sunk cost would produce no output and waste resources. Among the set of recipient-funded projects, F, the recipient will allocate its resource much as in the case where there are no sunk costs, such that marginal utility is equalized across projects. For example, if the recipient is purely developmentalist, the optimal solution (compare to (1)) is:

$$\hat{R}_i = \begin{cases} \left(R - \sum_{j \in F} s(A_j)\right) \dfrac{A_i^{\gamma_D}}{\sum_{j \in F} A_j^{\gamma_D}} + s(A_i) & \text{if } i \in F \\ 0 & \text{if } i \notin F \end{cases} \tag{5}$$

The $R - \sum_{j \in F} s(A_j)$ term represents the piece of the recipient's budget that is not consumed by sunk costs, and takes the place of R in (1). As in (1), the proportion $\dfrac{A_i^{\gamma_D}}{\sum_{j \in F} A_j^{\gamma_D}}$ determines the allocation of that piece among funded projects.

But which projects will the recipient fund? In general, F can be characterized as set of size ranges, within each of which the recipient funds all projects. One range, however, could have an upper "bound" of infinity. In many cases, F consists of a single range.

Table 4 illustrates. It takes the case where there are five projects, of sizes 1, 2, 3, 4, and 5. It assumes that the recipient is purely developmentalist, that sunk cost rises with an elasticity of 0.8 with respect to aid $\left(s(A_i) = A_i^{0.8}\right)$, and that development is $D = A_i^{0.5}\left(R_i - s(A_i)\right)^{0.5}$. The table shows how the recipient's resource allocation evolves as its budget envelope expands. When the resource budget, R, is just 1, the recipient cannot exceed the sunk cost of even the smallest project, project 1. As R increases, the recipient can afford project 1, and funds it. With further increases, it shifts to larger projects. As R rises above 10, it becomes possible and optimal for the recipient to fund two projects, then three, and eventually all the projects. F is usually contiguous.

Table 4. Optimal allocation of recipient resource of among 5 projects of size 1, 2, 3, 4, and 5 as resource budget *R* rises

	Optimal allocation				
R	R_1	R_2	R_3	R_4	R_5
1	0	0	0	0	0
2	2.00	0	0	0	0
3	0	3.00	0	0	0
4	0	0	4.00	0	0
5	0	0	0	5.00	0
6	0	0	0	0	6.00
7	0	0	0	0	7.00
8	0	0	0	0	8.00
9	0	0	0	0	9.00
10	0	3.07	0	0	6.93
11	0	0	4.27	0	6.73
12	0	0	0	5.41	6.59
13	0	0	0	5.85	7.15
14	0	0	0	6.30	7.70
15	0	0	0	6.74	8.26
16	0	0	0	7.18	8.82
17	0	0	4.39	5.68	6.93
18	0	0	4.64	6.01	7.35
19	0	0	4.89	6.34	7.76
20	0	0	5.14	6.68	8.18
21	0	0	5.39	7.01	8.60
22	0	3.34	4.81	6.23	7.62
23	0	3.48	5.02	6.52	7.98
24	0	3.63	5.24	6.80	8.34
25	0	3.77	5.45	7.09	8.69
26	1.95	3.63	5.25	6.82	8.36
27	2.01	3.77	5.45	7.08	8.69
28	2.08	3.90	5.65	7.35	9.02
29	2.15	4.03	5.85	7.62	9.36
30	2.21	4.17	6.05	7.88	9.69

A second elaboration needed for the simulations is a translation of the model into the continuous setting. As developed more fully in the Appendix, this is done by replacing the finite set of projects, $\{A_i\}$ with a differentiable function $h()$ that gives the density of projects of a given size, and representing the recipient's resource allocation rule with a distribution $r(x)$. Development and throughput are then given by:

$$D = \int_F x^{\alpha_D} \left(r(x) - s(x)\right)^{\rho_D} h(x) dx$$
$$B = \int_F x^{\alpha_B} \left(r(x) - s(x)\right)^{\rho_B} h(x) dx$$

Moreover, as shown in the Appendix, except in degenerate cases, F, the optimal set of funded projects, *is* always contiguous in this new setting (with Cobb-Douglas production). In fact, it always takes one of two forms: all projects below a certain size, or all projects above a certain size. So as a matter of algorithm, the maximizing recipient investigates two strategies—funding projects below or above a certain size. In each strategy, it seeks the size threshold, if any, that achieves a local maximum in utility, assuming that it allocates its resource among projects in F to equate marginal utility. It then determines which local optimum is the global one.

The recipient's choices of strategy and size cut-off of course depend on the parameters. So the idea of the simulations is to watch how the recipient's resource allocation shifts as the distribution of aid projects on offer evolves, and to see how this affects development. Projects are simulated as lognormally distributed and completely determined by a single parameter, total aid, according to (4). In particular, using the empirical parameters in column 1 of Table 3, it is assumed that

$$\mu = \mu_0 + c_\mu \ln A = 0.0546 + 0.2707 \ln A$$
$$\sigma = \sigma_0 + c_\sigma \ln A = 0.9580 + 0.1058 \ln A.$$

where A is in thousands of dollars. Recall that as A increases, according to the empirical pattern the number of projects also increases.

Letting x represent project size, the simulations further assume:

$$D(x) = x^{0.6}(r - s(x))^{0.6} \text{ or } x^{0.4}(r(x) - s(x))^{0.4}$$
$$B(x) = x^{0.8}(r(x) - s(x))^{0.2}$$
$$D = \int D(x)h(x)dx$$
$$B = \int B(x)h(x)dx$$
$$u(D, B) = B$$
$$s(x) = 0.5x^{0.8}$$
$$R = 200{,}000.$$

In words, development can have increasing returns to scale in one variant (0.6 + 0.6 > 1), once past sunk costs, and decreasing returns in an other (0.4 + 0.4 < 1). The recipient cares only about throughput, which is produced with constant returns to scale (0.8 + 0.2 = 1). The variables $r(x)$, and $s(x)$ are taken in thousands of purchasing-power-parity (PPP) dollars while A and x are taken in thousands of exchange-rate dollars. Sunk costs are thus \$500 PPP for a \$1,000 (exchange rate) project and rise with an elasticity of 0.8 with respect to size—so that, for example, those for a \$1,000,000 project (exchange rate) are about \$250,000 PPP. The resource budget is \$200 million, 1% of the PPP GDP of Tanzania, a country often pointed to as suffering from project proliferation. These parameter choices are meant to be reasonable and minimally arbitrary.

The simulations are not Monte Carlo. Representative project distributions are not generated. Rather, the problem is tackled analytically with the mathematics in the Appendix, the core challenge being to find the minimum or maximum project size that the recipient will fund in order to maximize throughput.[6]

Figure 7 shows how key variables evolve as aid rises in the variant with increasing returns to scale in development. Note first that the optimized value of B, called $\hat{B}$, rises monotonically with A—as it should since this is what the recipient is maximizing, and increasing

[6] The simulation is performed using an object-oriented Visual Basic for Applications program and accessed via user-defined functions in Microsoft Excel. The search algorithm for determining the minimum or maximum project sizes funded by the recipient is the "dbrent" routine in Press et al. (1988). The file is available from the author.

A only adds new production possibilities. Similarly for the maximum achievable value of D, called $\hat{D}$. But the actual value of D lags behind the ideal value, since the recipient is not maximizing it. For low aid levels, the ratio $D/\hat{D}$ is close to unity (graphed against the right axis), but around \$100 million, the ratio plunges dramatically, bottoming out after \$1 billion in total aid. Meanwhile, growth of development slows with respect to growth in aid, but never quite goes negative. For comparison, in this data set, Tanzania received \$1.3 billion in project aid in 2003.

What is behind the divergence between D and $\hat{D}$, potential and actual development? At low aid levels, the recipient's budget is ample enough that it is optimal to fund essentially all projects—regardless of whether the recipient is maximizing development or throughput.[7] The recipient does not allocate its resource among these projects quite the way it would if it were maximizing development (it favors smaller projects, relatively), though the difference does not have a large impact. Development and throughput are essentially complements. But as aid rises and projects proliferate, and because the recipient cares only about throughput, it is more reluctant to de-fund small projects with high transaction cost than it would be if it were intent on exploiting the scale economies of larger projects to maximize development. The gap becomes noticeable when total aid reaches some \$130 million. At this point, the optimal minimum project size for the recipient to fund is \$300,000 in aid—optimal, that is, if it is maximizing development. But from the recipient's throughput point of view, the optimal minimum project size is still only *half a penny*. (See Figure 8.) However, as total sunk costs continue rising (the bottom line in Figure 7), they eventually force even the throughput-minded recipient to de-fund small projects. The throughput-optimized size minimum then enters a catch-up period with respect to the development-optimized size minimum. But it never fully catches up, and so the

[7] "Essentially all" means that the minimum funded project size is a penny or less.

ratio between actual development and potential development is permanently lowered.

The second simulation differs from the first in that the exponents in the Cobb-Douglas function for development are 0.4 instead of 0.6. Now there are diminishing returns to scale. Figure 9 shows that because the development production function again differs from the throughput production that the recipient maximizes, aid is still not deployed optimally, and an absorption constraint appears.. However, the divergence this time is even sharper once it begins, and shows no signs of stopping even as aid passes \$10 billion, in that $D/\hat{D}$ continues declining. In fact, development declines in absolute terms. Why? This time the recipient departs from the development-optimal path not just in parameter choice, but in strategy. Because of the scale diseconomies, development is maximized when the recipient chooses to de-fund the largest projects. But the throughput-maximizing recipient instead de-funds the smallest projects, as shown in Figure 10.

As a sensitivity test, Figure 11 and Figure 12 show how the path for *D* changes when two crucial parameters are varied. Figure 11 shows the effect of halving or doubling c_μ, the coefficient on total aid in the equation for μ, relative to its empirically derived value of 0.2707. Higher values of c_μ lead to larger, fewer projects. Figure 11 does the same for c_σ, whose empirical value is 0.1058. Not surprisingly, increasing c_μ raises the aid level at which the development curve starts to bend. And when there are scale economies in development, it also substantially raises development for a given level of aid, especially once above that bend. However, when there are diseconomies of scale, varying c_μ has much less impact. It raises the aid level at which the marginal impact goes negative, but, strikingly, lowers the peak value for development. This is because the benefits from larger projects of lower sunk costs and fewer de-funded projects are offset by scale diseconomies. The overall picture is one of a hard ceiling of development impact when project economics favor small projects; this makes sense if, intuitively, the recipient government can only handle a limited number of projects. The picture is similar for c_σ, for in a lognormal distribution σ also influences average project size according to the formula $e^{\mu+\sigma^2/2}$.

Figure 7. Simulation with economies of scale in development

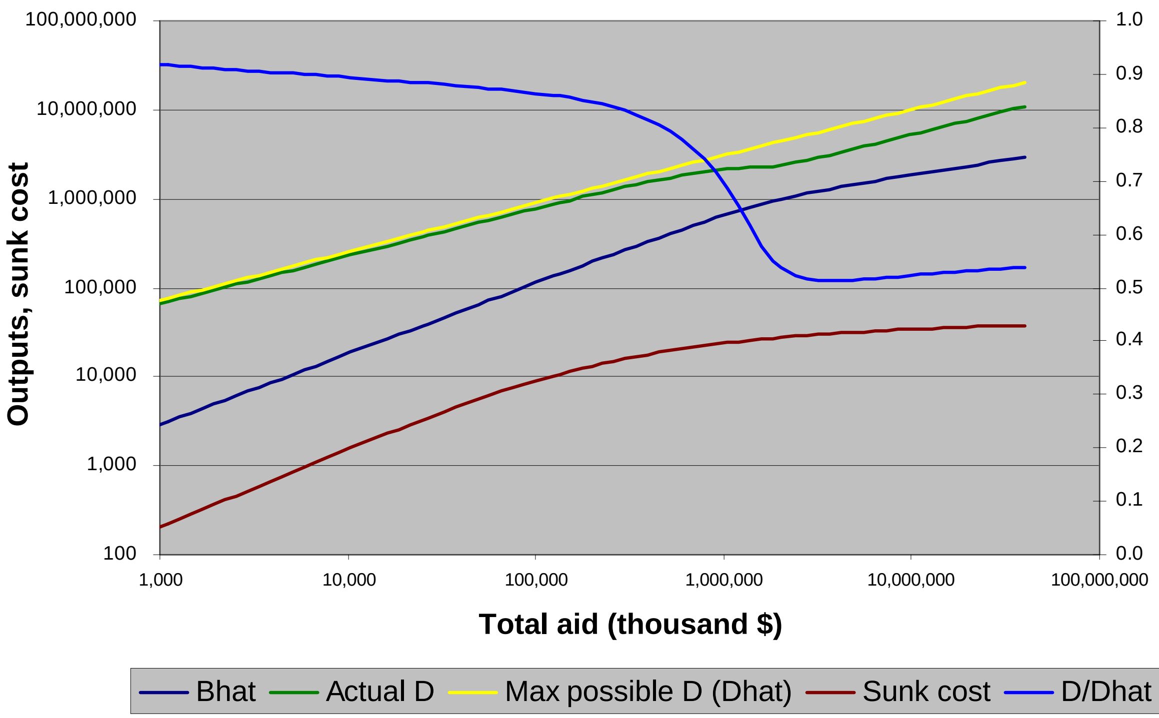

Figure 8. Simulation with economies of scale in development: minimum size of projects the recipient funds

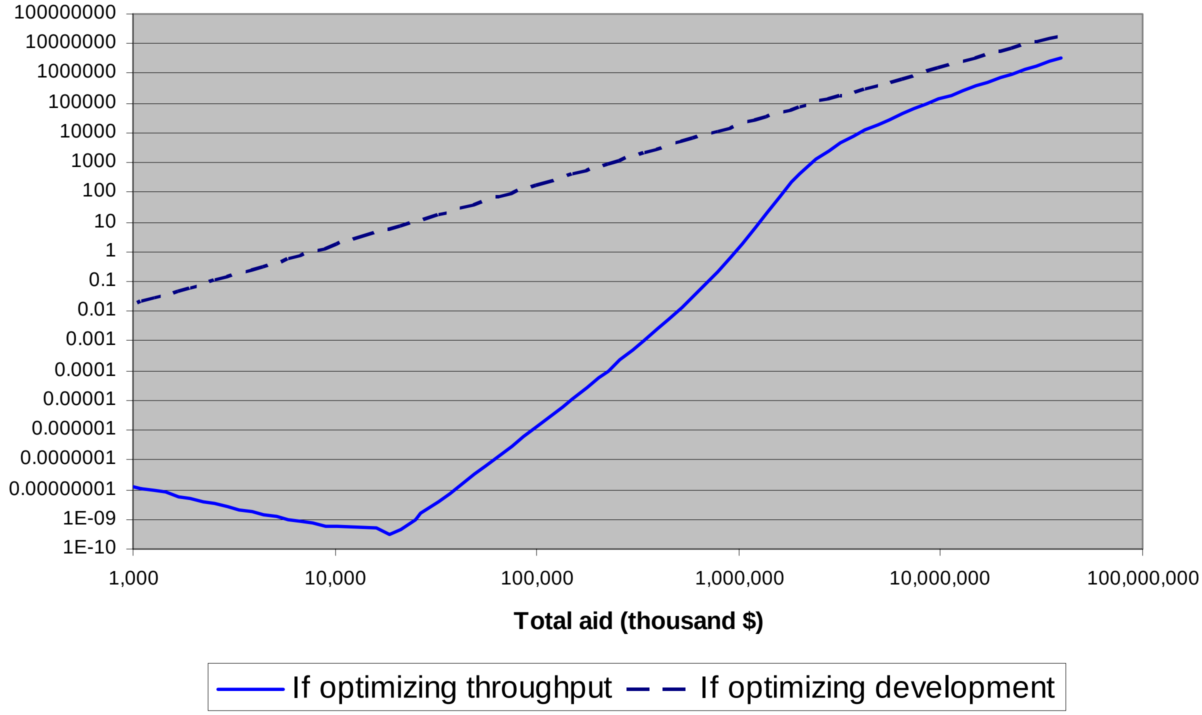

Figure 9. Simulation with diseconomies of scale in development

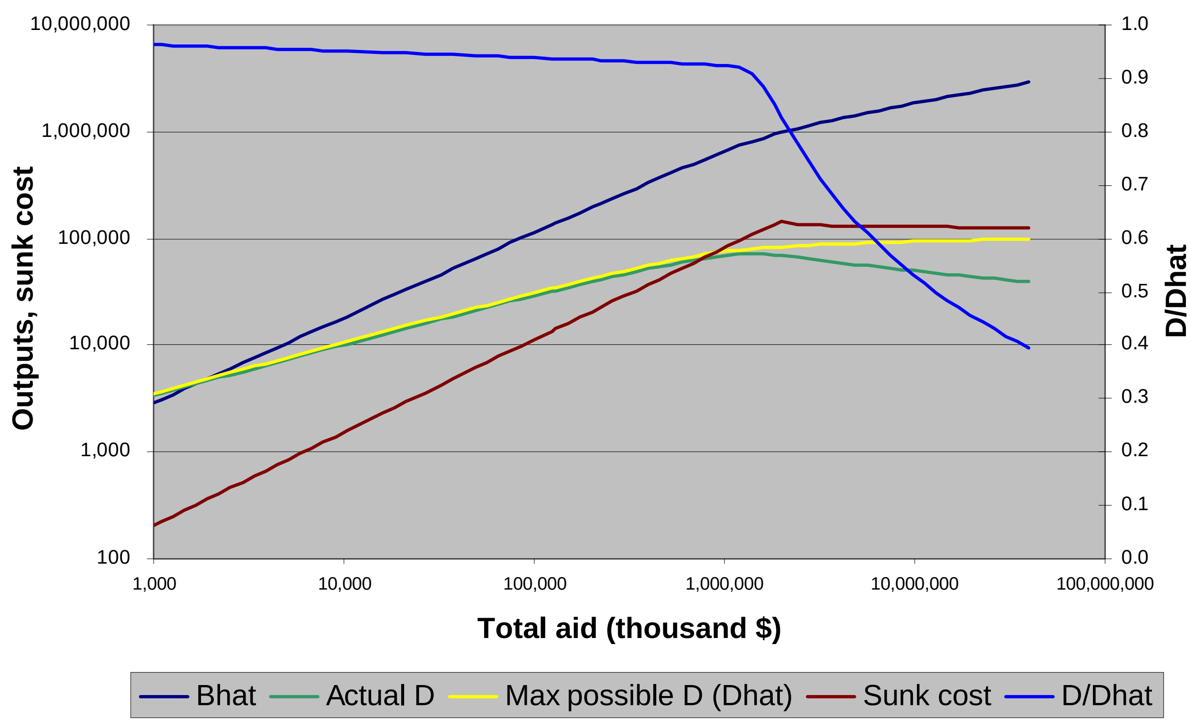

Figure 10. Simulation with diseconomies of scale in development: threshold size of projects the recipient funds

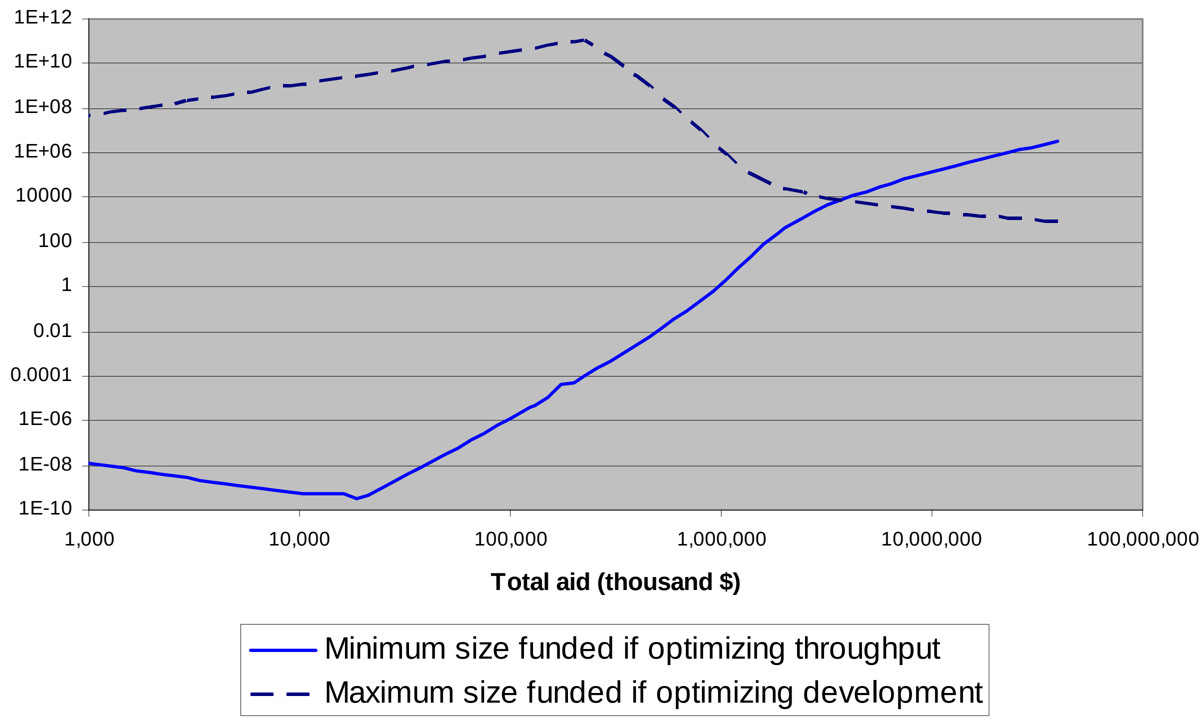

Figure 11. Sensitivity analysis: varying c_μ, the coefficient on total aid for average log project size,

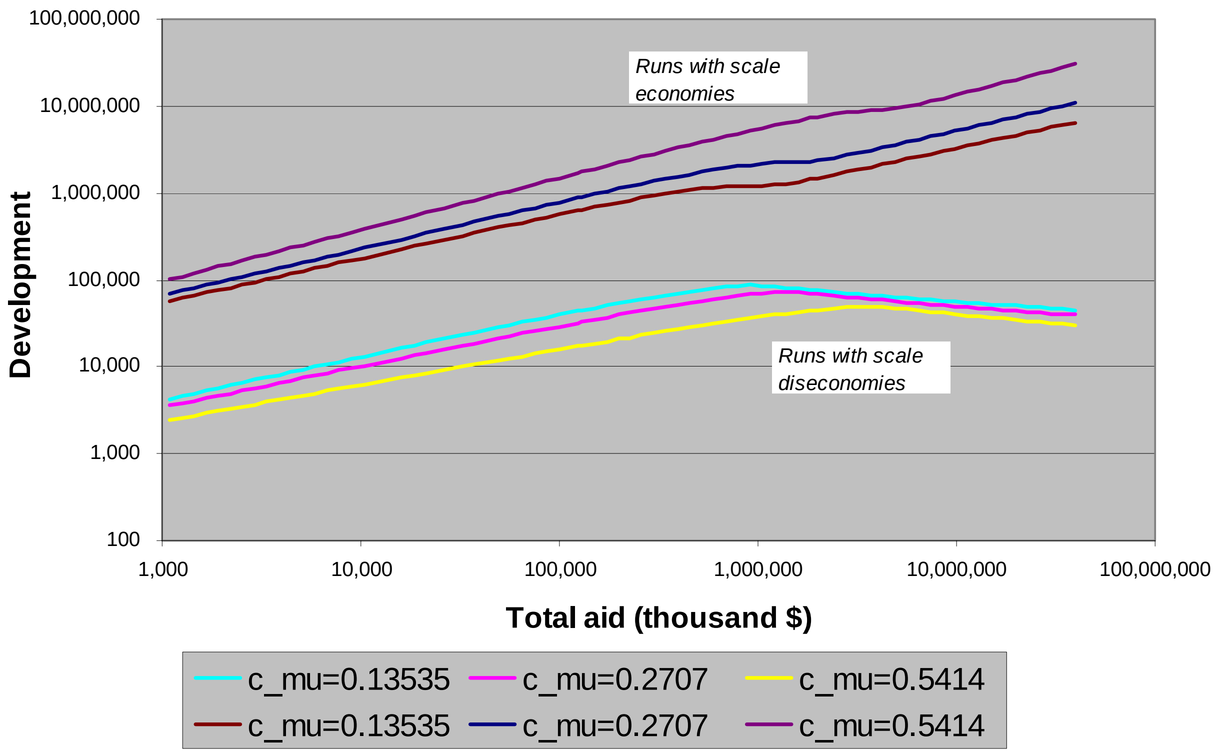

Figure 12. Sensitivity analysis: varying c_σ, the coefficient on total aid for standard deviation of log project size

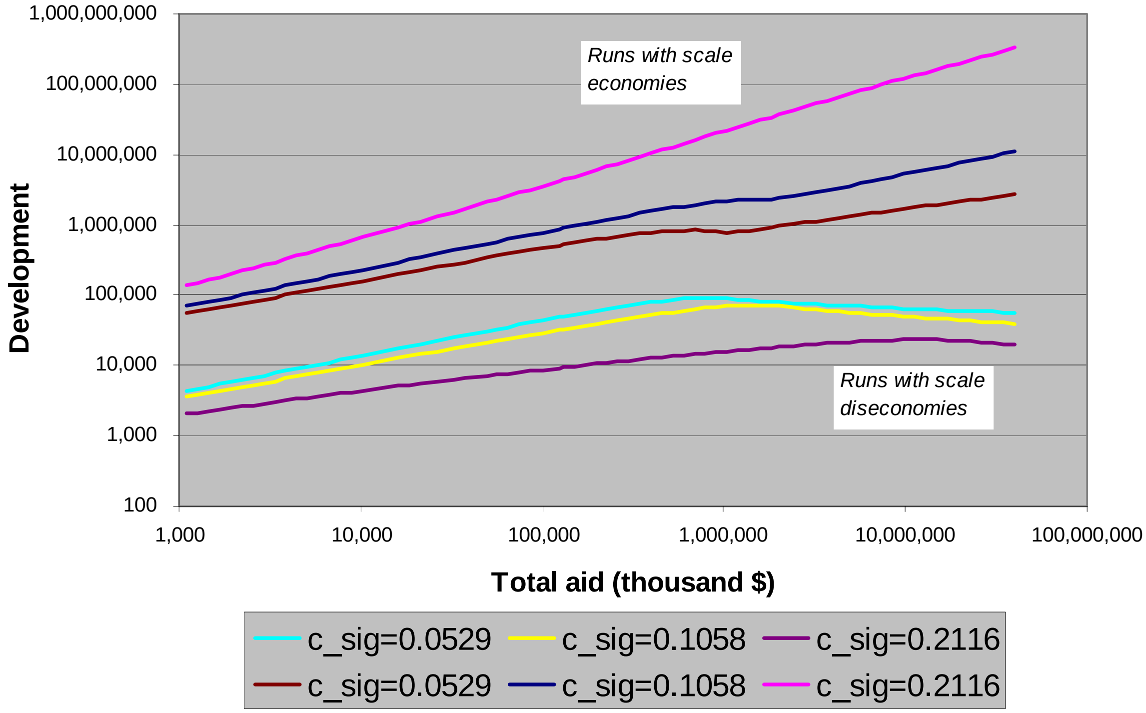

Conclusion

The above simulations assume for simplicity that the recipient cares only about throughput. If we assume, more realistically, that the recipient cares about both throughput and development, then the same pattern would appear, but less strongly. What is important is that three features of the model—the recipient not being purely developmentalist, differences in the throughput and development production functions, and sunk costs—interact to generate a notion of absorptive capacity for aid. Above a certain threshold level, the marginal productivity of aid declines precipitously—assuming, that is, that aid increase goes hand in hand with a pattern of project proliferation that is typical in the current cross-section of countries. The deep source of this threshold character is the discontinuous nature of conditionality. If recipients just cover sunk costs, projects go forward. If they fall just short, donors shut projects down.

If donors break out of current proliferation patterns, they may be able to raise development impact in countries suffering from proliferation *if* there are scale economies to be exploited. Descending somewhat from the abstract world of the model, the key may be for donors to emphasize sectors where scale economies are more likely, such as infrastructure.

Scholarly examination of aid project proliferation, like the economics of aid administration generally, is in its infancy. Concern about the extent of project proliferation, especially in low-income "donor darlings," is itself proliferating faster than the body of sharp theoretical and empirical analysis of the topic. Since this paper is primarily theoretical, its main value may lie in offering a mathematical paradigm. The aid process is conceived of as a set of production activities, lognormally distributed by size, each taking inputs from both donor and recipient. The complexity arises from differences in the production functions for two outputs and the differences in utility functions of the two agents, all of which ought to be tested empirically.

Indeed, this paper could serve as a point of departure for additional empirical research as well as theory. How great are sunk costs, what are they, and how do they vary with project size? How are aid projects best represented as production functions? Perhaps a more flexible form such as constant-elasticity-of-substitution functions is needed. Are recipient or donor resources the dominant factors in aid project production? How is the recipient resource *R* best thought of? Can it be measured? How does it vary by country? When do recipients refuse to invest resources in some aid projects even at the expense of their termination? Potential future directions in the theory include introducing alternative production functions, endogenizing donor behavior, allowing multiple donors, modeling interactions between projects, and repeated game aspects.

The model suggests that the question of when recipients refuse to invest resources in some aid projects is particularly important because the sharp drop in the marginal effectiveness of aid occurs at the point where it becomes optimal for development for the recipient to de-fund projects—and yet does not. The model does say that as aid continues to rise the recipient *will* eventually begin de-funding projects, but by the time there are such signs of trouble, the marginal and average productivity of aid will have already fallen substantially.

This suggests that especially when donors are contemplating a double or tripling of aid to Africa, they need to be sensitive to indications, such as in Tanzania, that relatively developmentalist recipients are resisting the proliferation of projects and the associated administrative burden. While donors' desire for monitoring and accountability is understandable, pushing these desires too hard in recipient countries with limited budgets can push the aid delivery process against structural limits and undermine overall effectiveness. In such countries donors should therefore contemplate varying the parameters within their control, funding fewer, larger aid activities. That in turn may call for careful analysis of which sectors offer the great scale economies for aid projects.

Appendix 1. A general model of aid projects

The core model

Write the development and throughput technologies as

$$D_i = D_i(A_i, R_i)$$

$$B_i = B_i(A_i, R_i)$$

where i indexes projects. D_i is development, B_i is "throughput" as a side-benefit, A_i is aid quantity going into the project, and R_i is the recipient-side resource. A_i will also be called the "size" of project i. To keep the notation compact, draw the equations together into vector-valued functions **D**(**A**,**R**) and **B**(**A**,**R**). Also, let $D = \sum D_i$ (total development produced) and define analogous symbols for the B_i, R_i, and A_i (except that we will treat $R = \sum R_i$ as a binding budget constraint rather than a definition). We assume **A, R** ≥ **0**—inputs are never negative.

The recipient's utility is $u_R(D,B)$ and the donor's, $u_A(D,B)$. Thus, both utilities depend only on the simple sums of the two kinds of outputs from individual projects. The recipient's problem, given the donor's aid allocation **A** and the recipient resource budget R, is

$$\max_{\mathbf{R}} u_R(D,B)$$

$$\text{such that } \boldsymbol{\iota}'\mathbf{R} \le R \text{ and } \mathbf{R} \ge 0$$

where $\boldsymbol{\iota}$ is a column of 1's. If u_R satisfies the appropriate regularity conditions when considered as a function of **A** and **R**, we can analyze this problem using a Lagrangian:

$$\mathcal{L}(\lambda, \mathbf{R}) = u_R\left(\sum D_i(A_i,R_i), \sum B_i(A_i,R_i)\right) - \lambda(\boldsymbol{\iota}'\mathbf{R} - R).$$

Imposing $\nabla \mathcal{L} = 0$ yields the first-order conditions for a maximum:

$$\boldsymbol{\iota}'\mathbf{R} = R$$
$$(\nabla_{\mathbf{R}} u_R)' = \lambda \boldsymbol{\iota}$$

where $\nabla_{\mathbf{R}}$ is the gradient operator with respect to **R**. In words, the recipient consumes its budget and allocates so that the marginal utility of its resource is equalized across projects to λ. The second-order condition is that

$$\nabla^2 \mathcal{L} = \begin{bmatrix} 0 & -\boldsymbol{\iota}' \\ -\boldsymbol{\iota} & \nabla^2_{\mathbf{R}} u_R \end{bmatrix}.$$

is negative semi-definite.

To describe the dynamics as **A** varies, let $\hat{\mathbf{R}}(\mathbf{A})$ be the vector-valued "indirect demand function" that maps the donor's allocation of aid among projects to the (optimizing) recipient's allocations of its resource. Define the "indirect supply functions" and indirect utility functions

$$\begin{aligned}
\hat{\mathbf{D}}(\mathbf{A}) &= \mathbf{D}\left(\mathbf{A}, \hat{\mathbf{R}}(\mathbf{A})\right) \\
\hat{\mathbf{B}}(\mathbf{A}) &= \mathbf{B}\left(\mathbf{A}, \hat{\mathbf{R}}(\mathbf{A})\right) \\
\hat{u}_R(\mathbf{A}) &= u_R\left(\hat{\mathbf{D}}(\mathbf{A}), \hat{\mathbf{B}}(\mathbf{A})\right) \\
.\hat{u}_A(\mathbf{A}) &= u_A\left(\hat{\mathbf{D}}(\mathbf{A}), \hat{\mathbf{B}}(\mathbf{A})\right)
\end{aligned}$$

These describe how output of development and throughput and the utility derived from them change as the donor varies its allocation of aid among projects and the recipient adapts its own allocation. Let $\hat{D}_i(\mathbf{A})$ and $\hat{B}_i(\mathbf{A})$ be the i^{th} components of $\hat{\mathbf{D}}(\mathbf{A})$ and $\hat{\mathbf{B}}(\mathbf{A})$, and $\hat{D}(\mathbf{A})$ and $\hat{B}(\mathbf{A})$ be the sums of these components, or "total indirect supply functions." These give total development and throughput for a given donor portfolio of projects of various sizes, assuming optimizing behavior on the part of the recipient. All these variables actually depend on the recipient resource budget, *R*, too, but we suppress this argument for clarity. A vector of central interest is the gradient $\nabla\hat{D}$, that is, the derivative of total development with respect to the donor's allocation of aid across projects. The chain rule gives

$$\nabla\hat{D} = \nabla_{\mathbf{A}} D + \nabla_{\mathbf{R}} D \cdot \nabla\hat{\mathbf{R}}. \tag{6}$$

One special case is worth noting. If the recipient is "purely developmentalist," i.e., $u_R = D,$ then $\nabla_{\mathbf{R}} D = \nabla_{\mathbf{R}} u_{\mathbf{R}}$, which by the first order condition is just $\lambda \boldsymbol{\iota}'$ because the marginal development impact of the recipient's resource is equalized across projects. If R is fixed, then the elements of $\nabla \hat{\mathbf{R}}$ sum to zero, meaning that $\nabla \hat{\mathbf{R}}$ is orthogonal to $\lambda \boldsymbol{\iota}'$. As a result, the second term of (6) drops out:

$$\nabla \hat{D} = \nabla_{\mathbf{A}} D. \tag{7}$$

If the recipient is a development-optimizer, then the marginal impact of a change in aid, factoring in how the recipient will reallocate its own resources at the margin, equals the marginal impact that would occur if the recipient made no reallocation. Although the recipient may in fact reallocate its resource at this margin, it has no effect on development, since at the margin the development impact of its resource is equalized across projects. This is an example of the Envelope Theorem (Varian 1992).

As discussed in text, we can introduce a notion of sunk cost, which the recipient must cover out of its budget for the project to go forward. It rises with project size, but not as fast, according to a sunk cost function $s(A_i)$:

$$\begin{gathered} D_i = \begin{cases} 0 & \text{if } R_i \le s(A_i) \\ D_i\left(A_i, R_i - s(A_i)\right) & \text{if } R_i > s(A_i) \end{cases} \\ B_i = \begin{cases} 0 & \text{if } R_i \le s(A_i) \\ B_i\left(A_i, R_i - s(A_i)\right) & \text{if } R_i > s(A_i) \end{cases} \\ s \ge 0, s' \ge 0, s'' \le 0 \text{ everywhere} \end{gathered} \tag{8}$$

And we can translate the model into the continuous setting by replacing the finite set of projects, $\{A_i\}$ with a differentiable function $h()$ that gives the density of projects of a given size. Assuming that the recipient puts the same amount of resource into each project of any given size, which it should by symmetry, its resource allocation rule can be expressed as function of project size: $r(x)$. Equations (8) translates to:

$$D = \int_F D(x, r(x) - s(x))h(x)dx$$
$$B = \int_F B(x, r(x) - s(x))h(x)dx$$

where F is the set of projects that the recipient chooses to fund, $F = \{x \mid r(x) \geq s(x)\}$.

The Cobb-Douglas case

The simulations reported in the main text assume that D and B are both Cobb-Douglas. Here we derive the recipient's solution in this particularly tractable case. Let α_D and ρ_D, α_B and ρ_B be the exponents that characterize the two production functions. Then:

$$\begin{aligned} D &= \int_F x^{\alpha_D} (r(x) - s(x))^{\rho_D} h(x)dx \\ B &= \int_F x^{\alpha_B} (r(x) - s(x))^{\rho_B} h(x)dx \end{aligned} \qquad (9)$$

The mathematics of the recipient's problem are essentially identical whether it is a pure development or a pure throughput optimizer (the two possibilities that are easiest to analyze). So consider the case when the recipient is purely developmentalist. Its problem is

$$\max_{r(x)} D$$

$$\text{such that} \int_{\mathrm{F}} r(x)h(x)dx \leq R.$$

By analogy with the discrete case, the natural candidate solution is (compare to (5)):

$$\hat{r}(x) = \begin{cases} \left(R - \int_F s(y)h(y)dy\right) \dfrac{x^{\gamma_D}}{\int_F y^{\gamma_D} h(y)dy} + s(x) & \text{if } i \in F \\ 0 & \text{if } i \notin F \end{cases}. \qquad (10)$$

We must apply the calculus of variations to confirm that (10) is the optimum for a given

F. This is more easily done if the problem is recast so that the budget constrain becomes a boundary condition. So define the function $R(x)$ to be the total resource that the recipient invests in projects of size x or smaller. Then we have:

$$R(x) \equiv \int_0^x r(y)h(y)dy \tag{11}$$

$$\delta(x) \equiv x^{\alpha_D}\left(r(x)-s(x)\right)^{\rho_D} h(x) = x^{\alpha_D}\left(\frac{R'(x)}{h(x)}-s(x)\right)^{\rho_D} h(x)$$

Thus according to (9),

$$D = \int_F \delta(x)dx$$

The recipient's problem is then

$$\max_{R(x)} \int_F \delta(x)dx \tag{12}$$

such that $R(0) = 0$ and $R(\infty) = R$.

This is a calculus of variations problem in which the integrand can be expressed purely as a function of the derivative of the argument: $\delta(x)$ depends directly on $R'(x)$ but not $R(x)$. The version of the Euler equation that $R(x)$ must satisfy in this case is

$$\frac{d}{dx}\left(\frac{d\delta}{dR'}\right) = 0$$

(Chiang 1992, p. 37). Differentiating (11),

$$\frac{d\delta}{dR'} = \rho_D x^{\alpha_D}\left(\frac{R'(x)}{h(x)}-s(x)\right)^{\rho_D-1} = \rho_D x^{\alpha_D}\left(r(x)-s(x)\right)^{\rho_D-1}$$

Plugging in the trial solution for $r(x)$ in (10),

$$\frac{d\delta}{dR'} = \rho_D x^{\alpha_D}\left(\left(R - \int_F s(y)h(y)dy\right)\frac{x^{\gamma_D}}{\int_F y^{\gamma_D} h(y)dy} + s(x) - s(x)\right)^{\rho_D - 1}$$
$$= \rho_D \left(\frac{R - \int_F s(y)h(y)dy}{\int_F y^{\gamma_D} h(y)dy}\right)^{\rho_D - 1} x^{\alpha_D + \gamma_D(\rho_D - 1)}.$$

Since $\gamma_D = \alpha_D / (1 - \rho_D)$, $x^{\alpha_D + \gamma_D(\rho_D - 1)} = 1$. The above expression therefore does not vary with x, so it satisfies the Euler equation, as needed.

Equation (10) describes the recipient's solution given the set of funded projects, F, but does not include a recipe for F. To study F, we first substitute the formula for the recipient's optimal allocation within F (equation (10)) back into the equation for D in (9):

$$D = \left(\frac{R - \int_F s(x)h(x)dx}{\int_F x^{\gamma_D} h(x)dx}\right)^{\rho_D} \int_F x^{\alpha_D}\left(x^{\gamma_D}\right)^{\rho_D} h(x)dx = \left(\frac{R - \int_F s(x)h(x)dx}{\int_F x^{\gamma_D} h(x)dx}\right)^{\rho_B} \int_F x^{\gamma_D} h(x)dx$$
$$= \left(\int_F x^{\gamma_D} h(x)dx\right)^{1-\rho_D}\left(R - \int_F s(x)h(x)dx\right)^{\rho_D} \qquad (13)$$

(taking advantage in the second step of the fact that $\alpha_D + \gamma_D\beta_D = \gamma_D$).

It is impossible to write down a general, explicit solution for the F that maximizes this geometric average of integrals. However, in this Cobb-Douglas case, F can be tightly characterized, assuming sunk cost has constant elasticity with respect to project size.

Proposition. *If we can write* $s(x) = s_0 x^{c_s}$, *then the set F that maximizes (13) is almost always contiguous and takes the form* (0, MAX] *or* [MIN, ∞) *where* MIN, MAX > 0.

In other words, except in certain degenerate examples, the recipient's optimal strategy is to fund all projects below a certain size or all projects above a certain size—or, as a special case of the latter, all projects of any size (MIN = 0).

Proof. A priori, F could consist of multiple disjoint segments within the set of non-negative numbers. Let $S = [x_1, x_2]$ be such a segment. Then (13) can be expanded to

$$
\begin{aligned}
D &= \left(\int_F x^{\gamma_D} h(x)dx\right)^{1-\rho_D}\left(R - \int_F x^{c_s} h(x)dx\right)^{\rho_D} \\
&= \left(\int_{x_1}^{x_2} x^{\gamma_D} h(x)dx + \int_{F-S} x^{\gamma_D} h(x)dx\right)^{1-\rho_D}\left(R - s_0\left(\int_{x_1}^{x_2} x^{c_s} h(x)dx + \int_{F-S} x^{c_s} h(x)dx\right)\right)^{\rho_D}
\end{aligned}
\tag{14}
$$

Analysis of how to maximize (14) is easier if we work with its logarithm. Since x_1 and x_2 are parameters within the recipient's control, at a maximum, x_1 and x_2 are each either roots of a derivative of ln u or boundary solutions:

$$
\begin{aligned}
\partial(\ln D)/\partial x_1 &= 0 \ \text{ or } x_1 = 0 \\
\partial(\ln D)/\partial x_2 &= 0 \ \text{ or } x_2 = \infty.
\end{aligned}
$$

If x_1 is an interior solution, then differentiating the log of (14) leads to

$$
0 = \frac{\partial(\ln D)}{\partial x_1} = (1-\rho_D)\frac{-x_1^{\gamma_D} h(x_1)}{\int_{x_1}^{x_2} x^{\gamma_D} h(x)dx + \int_{F-S} x^{\gamma}{}_D h(x)dx} - \rho_D \frac{-s_0 x_1^{c_s} h(x_1)}{R - s_0\left(\int_{x_1}^{x_2} x^{c_s} h(x)dx + \int_{F-S} x^{c_s} h(x)dx\right)}.
$$

Rearranging,

$$
x_1^{\gamma - c_s} = \frac{\rho_D}{1-\rho_D}\frac{\int_{x_1}^{x_2} x^{\gamma_D} h(x)dx + \int_{F-S} x^{\gamma_D} h(x)dx}{R/s_0 - \left(\int_{x_1}^{x_2} x^{c_s} h(x)dx + \int_{F-S} x^{c_s} h(x)dx\right)}.
\tag{15}
$$

Similarly, it works out that if x_2 is an interior solution,

$$
x_2^{\gamma - c_s} = \frac{\rho_D}{1-\rho_D}\frac{\int_{x_1}^{x_2} x^{\gamma_D} h(x)dx + \int_{F-S} x^{\gamma_D} h(x)dx}{R/s_0 - \left(\int_{x_1}^{x_2} x^{c_s} h(x)dx + \int_{F-S} x^{c_s} h(x)dx\right)},
$$

which is almost identical to (15). So if x_1 and x_2 are both interior, $x_1^{\gamma-c_s} = x_2^{\gamma-c_s}$. This equation only admits solutions in special cases, which are unrepresentative and degenerate, in the sense that an infinitesimal perturbation of the parameters will cause them to disappear. If γ exactly equals c_s, then the equation has infinitely many solutions. But if they differ even infinitesimally it has only one: $x_1 = x_2$. But if $x_1 = x_2$ is itself a degenerate solution in the continuous setting since the integrals bounded by x_1 and x_2 in (14) are 0. Thus in general x_1 and x_2 are not both interior. x_1 takes a boundary value (0) or x_2 effectively does (∞), or both.

The above argument shows that *S* takes the form asserted in the proposition. But *S* is only one piece of *F*. In general, since *F* is the disjoint union of segments in such forms, it could itself take one of those forms. Or it might take the form $(0, x_1] \cup [x_2, \infty)$, where $0 < x_1 < x_2$. But an argument almost identical to the one above shows that it cannot. In particular, such an *F* could only be a maximum if $x_1 = 0$ or $x_2 = \infty$ or $x_1 = x_2$, none of which allows *F* to be disconnected. □

So as a matter of algorithm, the maximizing recipient investigates two strategies: funding projects below or above a certain size. In each strategy, it seeks the size threshold, if any, that achieves a local maximum in utility. It then determines which local optimum is the global one. All of the forgoing applies mutatis mutandis if the recipient is a pure throughput rather than development optimizer.

The proof does not carry over to discrete distributions of projects such as in Table 4 of the main text. These can be thought of as distributions, $h(x)$, that are infinitely dense at some points, and non-differentiable there. In these distributions, the possibility rejected in the above proof that $x_1 = x_2$ is no longer degenerate, and all we can conclude is that *F* is a set of singletons, which is not revealing. However, the larger the set of projects and the better it can be approximated as a continuous distribution, the greater will be the tendency for *F* to be contiguous, a tendency seen in Table 4.

One final result needed to run the simulations in the main text is the formula for development when the recipient maximizes throughput. Substituting (10) into (9),

$$D = \left(\frac{R - \int_F s(x)h(x)dx}{\int_F x^{\gamma_B} h(x)dx} \right)^{\rho_D} \int_F x^{\alpha_D + \gamma_B \rho_D} h(x)dx. \tag{16}$$

References

Acharya, A., A. de Lima, and M. Moore (2006). 'Proliferation and fragmentation: Transactions costs and the value of aid'. *Journal of Development Studies*, 42 (1): 1–21.

ActionAid International (2005). *Real Aid: An Agenda for Making Aid Work*. London.

Aitchison, J., and J.A.C. Brown (1963). *The Lognormal Distribution*. Cambridge, U.K.: Cambridge University Press.

Chiang, A.C. (1992). *Elements of Dynamic Optimization*. New York: McGraw-Hill.

Clemens, M.A., and T.J. Moss (2005). 'Ghost of 0.7%: Origins and Relevance of the International Aid Target'. Working Paper No. 68. Washington: Center for Global Development.

Clemens, M., and S. Radelet (2003). 'The Millennium Challenge Account: How much is too much, how long is long enough?' Working Paper No. 23. Washington: Center for Global Development.

Clemens, M., S. Radelet, and R. Bhavnani (2004). 'Counting Chickens When They Hatch: The Short Term Effect of Aid on Growth'. Working Paper No. 44. Washington: Center for Global Development.

Collier, Paul and David Dollar (2002). 'Aid Allocation and Poverty Reduction'. *European Economic Review*, 45 (1): 1–26.

Development Assistance Committee (DAC) (2002). *Reporting Directives for the Creditor Reporting System*. Paris.

Dalgaard, C., H. Hansen, and F. Tarp (2002). 'On the Empirics of Foreign Aid and Growth'. Working Paper No. 02/08, Nottingham, U.K.: Centre for Research in Economic Development and International Trade, University of Nottingham.

Durbarry, R., N. Gemmell, and D. Greenaway (1998). 'New Evidence on the Impact of Foreign Aid on Economic Growth'. Working Paper No. 98/08, Nottingham, U.K.: Centre for Research in Economic Development and International Trade, University of Nottingham.

Hadjimichael, M.T., D. Ghura, M. Muhleisen, R. Nord, and E.M. Ucer (1995). 'Sub-Saharan Africa: Growth, Savings, and Investment, 1986-93'. Occasional Paper No. 118. Washington: International Monetary Fund.

Hansen, H. (2001). 'The Impact of Aid and External Debt on Growth and Investment: Insights from Cross-Country Regression Analysis'. Presented at WIDER Conference on Debt Relief, Helsinki, 17–18 August.

Hansen, H., and F. Tarp, 2000. 'Aid Effectiveness Disputed'. *Journal of International Development*, 12 (3): 375–98.

__________ (2001). 'Aid and growth regressions'. *Journal of Development Economics*, 64 (2): 547–70.

Morss, E.R. (1984). 'Institutional Destruction Resulting from Donor and Project Proliferation in Sub-Saharan African Countries'. *World Development*, 12 (4): 465–70.

Lensink, R., and H. White (2001). 'Are There Negative Returns to Aid?' *Journal of Development Studies*, 37 (6): 42–65.

Press, W.H., et al. (1998). *Numerical Recipes in C.* Cambridge, U.K.: Cambridge University Press.

Roodman, D. (2005). 'An Index of Donor Performance'. Working Paper No. 67. Washington: Center for Global Development.

U.N. Development Programme (2005a). *Human Development Report 2005: International Cooperation at a Crossroads: Aid, Trade and Security in an Unequal World*. New York.

U.N. Development Programme (2005b). *Investing in Development: A Practical Plan to Achieve the Millennium Development Goals*. London: Earthscan.
Varian, H.R. (1992). *Microeconomic Analysis, 3rd ed.* New York: W.W. Norton.
World Bank (2005). *Global Monitoring Report 2005: Millennium Development Goals: From Consensus to Momentum*. Washington.