

Arndt, Channing; Jones, Sam; Tarp, Finn

Working Paper

What is the aggregate economic rate of return to foreign aid?

WIDER Working Paper, No. 2014/089

Provided in Cooperation with:

United Nations University (UNU), World Institute for Development Economics Research (WIDER)

Suggested Citation: Arndt, Channing; Jones, Sam; Tarp, Finn (2014) : What is the aggregate economic rate of return to foreign aid?, WIDER Working Paper, No. 2014/089, ISBN 978-92-9230-810-0, The United Nations University World Institute for Development Economics Research (UNU-WIDER), Helsinki,
<https://doi.org/10.35188/UNU-WIDER/2014/810-0>

This Version is available at:

<https://hdl.handle.net/10419/102987>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



UNITED NATIONS
UNIVERSITY
UNU-WIDER

WIDER Working Paper 2014/089

What is the aggregate economic rate of return to foreign aid?

Channing Arndt,¹ Sam Jones,² and Finn Tarp¹

May 2014

Abstract: Does foreign aid promote aggregate economic growth? In contrast to widespread perceptions, academic studies of this question have been rapidly converging towards a positive answer. We employ a simulation approach to (i) validate the coherence of recent empirics and (ii) calculate plausible ranges for the rate of return to aid. Our results highlight the long-run nature of aid investments and indicate the return to aid falls in ranges commonly accepted for public investments. We find no basis for the view that aid has a pernicious effect on productivity.

Keywords: foreign aid, rate of return, heterogeneity, simulation

JEL classification: C63, F35, O4

Acknowledgements: The authors are grateful for comments and advice on earlier drafts from Tony Addison, Richard Carey, John Hudson, Richard Manning, Lant Pritchett, Roger Riddell, L. Alan Winters, and Adrian Wood. The usual caveats apply.

Tables and figures appear at the end of the paper.

¹UNU-WIDER, ²University of Copenhagen; corresponding author: tarp@wider.unu.edu

This study has been prepared within the UNU-WIDER project ‘ReCom–Foreign Aid: Research and Communication’, directed by Tony Addison and Finn Tarp.

Copyright © UNU-WIDER 2014

ISSN 1798-7237 ISBN 978-92-9230-810-0

Typescript prepared by Lisa Winkler at UNU-WIDER.

UNU-WIDER gratefully acknowledges specific programme contributions from the governments of Denmark (Ministry of Foreign Affairs, Danida) and Sweden (Swedish International Development Cooperation Agency—Sida) for ReCom. UNU-WIDER also gratefully acknowledges core financial support to its work programme from the governments of Denmark, Finland, Sweden and the United Kingdom.

The World Institute for Development Economics Research (WIDER) was established by the United Nations University (UNU) as its first research and training centre and started work in Helsinki, Finland in 1985. The Institute undertakes applied research and policy analysis on structural changes affecting the developing and transitional economies, provides a forum for the advocacy of policies leading to robust, equitable and environmentally sustainable growth, and promotes capacity strengthening and training in the field of economic and social policy-making. Work is carried out by staff researchers and visiting scholars in Helsinki and through networks of collaborating scholars and institutions around the world.

UNU-WIDER, Katajanokanlaituri 6 B, 00160 Helsinki, Finland, wider.unu.edu

The views expressed in this publication are those of the author(s). Publication does not imply endorsement by the Institute or the United Nations University, nor by the programme/project sponsors, of any of the views expressed.

1 Introduction

After years of sharp debate, controversy over the effectiveness of foreign aid as a tool to promote social and economic progress in developing countries is slowly receding, at least in the empirical academic literature. As noted early on by Mosley (1986), evaluations of foreign aid at the microeconomic and meso levels have consistently pointed towards positive impacts on average. This view has been bolstered by recent studies that consider sub-components of aid, as well as outcomes in specific regions or sectors (e.g., see Michaelowa 2004; Mishra and Newhouse 2009; Arndt et al. 2014). Broadly speaking, these evaluations indicate a positive average impact of foreign aid, recognizing that heterogeneity of impact is prevalent.

Controversy has been most acute at the macroeconomic level. Even so, the large majority of the up-to-date empirical studies in the economics literature find positive impacts. The full range of independent studies published since 2008 based on cross-country growth regressions report comparable results for the marginal effect of aid on growth.

In rough terms, these studies suggest that receipt of foreign aid equal to 10 per cent of GDP over a sustained period is expected to boost growth by approximately one percentage point on average. While not all findings are statistically significant at conventional levels, in part reflecting noisy and sparse data, most are significant. Indeed, the broad magnitude and direction of these results are sufficiently similar to merit attention. This represents our point of departure.

In this paper, we add to the literature and ask two specific questions: (i) Are results from recent studies regarding the aggregate effect of aid on growth numerically coherent? (ii) If so, what do they imply about the economic rate of return to aid? To answer these questions, we run numerical simulations of a dynamic Solow-type growth model augmented with foreign aid. The first question is motivated by the notorious difficulty in pinning down causal effects in macroeconomic data. For assessment of foreign aid, these difficulties are compounded by the relatively low quality of data from developing countries regarding income growth and foreign aid volumes. Numerical simulation offers a simple and transparent means to validate the overall plausibility of empirical results. For instance, if estimates from these simulations were to indicate upper bounds for the effect of aid on growth that are systematically lower than estimates encountered in empirical studies, this would provide a basis to question whether these empirical results are upwards biased. Also, numerical simulations can help think through the empirical implications of specific modelling challenges, such as the suitable time frame over which aid impacts on growth.

The second question focuses on the comparative benefits of providing aid versus the financial costs of its provision. A positive long-run impact of aid on growth does not automatically imply that aid generates an acceptable return on investment when viewed over the lifecycle of a given project. At the same time, if returns to aid are found to be high, this might suggest there is scope to provide a larger share of development finance on non-concessionary terms. Even though rates of return represent a standard criterion for evaluating investments at the project level, these issues have not been addressed in the recent aid literature.¹ Our simulation approach helps to identify a plausible domain for economic rates of return to aid at the macroeconomic level.

¹ See Dalgaard and Hansen (2005) for estimates of the marginal productivity of investment associated with aid.

The remainder of this paper is structured as follows: Section 2 reviews recent published studies regarding the macroeconomic impact of aid on growth. Section 3 presents our baseline model and associated simulation protocol. Section 4 derives insights from simulations of the effect of aid on growth. Section 5 reports associated measures of economic returns. Section 6 considers an extension to the baseline model, allowing aid to induce permanent productivity effects. Section 7 summarizes and interprets our results.

2 Recent aid-growth studies

This section highlights the principal findings of recently published empirical studies that focus on the aid-growth relationship.² A meaningful starting point is Rajan and Subramanian (2008, hereafter RS08) who introduced a pair of influential innovations. First, they signaled a movement away from a reliance on cross-country dynamic panel data methods. Rather, they preferred a strategy involving a long-run cross-section regression in which both Aid/GDP and growth are taken as averages over relatively extended periods (up to 40 years). This responds to the insight that aid given at time t may only have a growth impact after $t + n$ years, and that this impact may yield benefits over an extended period. Second, to address the endogeneity of aid, RS08 deploy external instrumental variables rather than the internal instruments commonly deployed in dynamic panel estimators.

Table 1 summarizes core results from recent papers that address aid and growth.³ To the best of our knowledge, the table covers the full population of studies that meet the following criteria: they (i) refer to an average aggregate aid-growth relation for developing countries as a group; (ii) include data spanning at least 30 years; (iii) attempt to address the endogeneity of aid; and (iv) are published in a peer-reviewed economics journal since 2008.⁴ As the various papers included in the table use alternative specifications, an attempt has been made to select estimates from comparable models. In some instances non-linear specifications involving a squared aid term are included. For these we report the marginal effect of Aid/GDP on growth evaluated by fixing Aid/GDP at 2.5 percent. In the final column, we report the mass in the relevant tail of the standard normal distribution for the z-statistic calculated from the reported point estimate (beta) and its standard error.

In all but two cases, the beta coefficients in Table 1 are positive.⁵ Three quarters are approximately significant at the 10 per cent level and just two are insignificant at the 20 per cent level. The simple average of the point estimates for the average marginal effect of Aid/GDP on growth is 0.19; weighting by the logarithm of the inverse of their estimated variances yields an estimate of 0.12 with a standard error of 0.02. The finding of a statistically significant average effect is consistent with a formal meta-analysis of 68 earlier aid growth studies found in Mekasha

² More extensive overviews of the aid effectiveness literature can be found elsewhere. For example, see Clemens et al. (2012); Arndt et al. (2010, 2011, 2014); Roodman (2007); Dalgaard et al. (2004).

³ Papers in Table 1 were collated from those listed on Google Scholar using the search expression: 'growth 'foreign aid' regression long-run' (accessed 11 February 2014).

⁴ It follows that 'naive' OLS estimates, where reported by the authors, are not included in Table 1. Equally, papers that focus exclusively on selected regions, disaggregated aid measures or alternative outcomes are not included in the table.

⁵ The data underlying the first of these negative results (NDHKM12) has been re-examined and found wanting by Lof et al. (2014). The second negative result (Herzer and Morrissey 2103, HM13) derives from an estimation model that controls for aggregate investment, implying the estimated effect of aid on output is restricted to non-investment channels. Note that the aid-investment channel is deemed to be important in a range of studies, including Juselius et al. (2014); Arndt et al. (2014).

and Tarp (2013); there the authors find a positive and significant weighted average partial correlation between aid and growth, which is robust to adjustments addressing publication bias and moderator effects.⁶

Interestingly, this range of results is consistent with the empirical estimates originally found in RS08, which sparked the recent wave of literature. Although the latter result is not statistically significant at standard levels, Arndt et al. (2010) correct the treatment of unreported values for aid flows from missing to null and show that plausible modifications to the authors' empirical strategy yield positive and significant estimates with the same data set.⁷ Furthermore, Arndt et al. (2014) find comparable and significant results using data updated to 2007 (Table 1, row 16). It also bears remarking that alternative econometric methods suggest similar conclusions. For example, drawing on country-specific (cointegrated) time series analysis for 36 African countries, Juselius et al. (2014) find that foreign aid had a significant and positive effect on investment or GDP (or both) in 27 cases. In seven cases the effect on investment or GDP is positive but insignificant and in only two cases does the impact of aid appear to be negative and significant.

In sum, recent empirical studies provide consistent support for the view that aid has had a positive average effect on growth when viewed over an extended time frame. The view that aid is ineffective finds much weaker support; and the notion that aggregate aid is actively harmful on average (e.g., Moyo 2009) finds no endorsement when recent studies are carefully scrutinized.

3 A numerical simulation

3.1 Baseline model

The previous section highlighted a notable degree of consistency across the full range of empirical studies published since 2008. An open question is whether these results are plausible from a theoretical and numerical-simulation perspective. That is, given what we know about aid volumes, is the distribution of these empirical estimates in accordance with standard assumptions about how aid affects growth? A preliminary answer is provided by RS08. Using a straightforward Solow-Swan framework, the authors derive an expression for the expected marginal effect of aid on growth, which is equal to the product of the share of aid invested, the share of capital in income and the output-capital ratio (K/Y). Based on rough estimates of the latter for developing economies, RS08 conclude that a plausible range for the aid-growth effect is from 0.08 to 0.16, but could be higher (lower) if aid enhances (undermines) productivity growth.

Three extensions to the framework suggested in RS08 are appropriate from the perspective of running more detailed numerical simulations.

First, explicit consideration of the dynamic structure of the effect of aid through the macro-economy is necessary. For instance, using RS08's assumption that aid only influences growth through expansion of the capital stock, it follows that the marginal return to capital may vary over time, and depreciation charges in future years must be taken into account.

⁶ Note that the estimates reported in Mekasha and Tarp (2013) are partial correlations (effect sizes) and thus are not directly comparable in magnitude to the estimates in Table 1.

⁷ In fact, the correction from missing to null is sufficient to achieve statistical significance.

Second, as already mentioned, there is a potential lag between receiving aid and its effect on growth. Different assumptions about the duration of this lag have important implications for the calculations of the effect of aid on growth and for returns to aid, especially over shorter time horizons.

Third, from a conceptual viewpoint, some distinction between longer-run and shorter-run effects due to aid may be warranted. For instance, one concern raised in the literature is that large aid flows may appreciate the real exchange rate, thereby weakening growth by disfavoring exporting firms that may have a high growth potential. Equally, channeling external savings to the public sector could either crowd-out or crowd-in domestic savings. Apart from the lagged effects of aid on the capital stock, these considerations point to the potential for more immediate developmental consequences associated with aid running through different channels.

Taking the static model of RS08 as a starting point, our simulation model addresses the above issues in a stylized economy. Our model is summarized by two principal equations:

$$Y_t = \theta_t K_t^\alpha L_t^{1-\alpha} \quad (1)$$

$$K_t = A_{t-\mu} + \gamma A_{t-1} + (1 - \delta)K_{t-1} + \delta K_0 \quad (2)$$

where t denotes time (in years; $t = 0$ is the initial or base year); Y_t denotes real national income; K_t is an aggregate measure of (physical) capital stock; L_t is the work force which, by assumption, is normalized to one and held constant ($L_t = L_0 = 1$); α is capital's share of income; A_t is the volume of foreign aid; and θ_t is total factor productivity (TFP); parameter δ refers to the rate by which capital depreciates ($0 < \delta < 1$); parameter γ captures an immediate crowding-out/in effect of aid on private investment; and parameter μ is an integer that captures aid timing effects ($\mu > 0$). Prices are presumed constant, normalized to one. Throughout we assume there is no pre-existing aid in the system (i.e., $A_0 = 0$).

Four main aspects of the above model merit comment. First, in keeping with RS08, our baseline assumption is that TFP is exogenous and constant over time ($\theta_t = \bar{\theta}$). This implies that foreign aid only influences aggregate welfare (income) through changes in the capital stock. This is the same as saying that all aid is invested. We could also introduce a parameter to capture the proportion of aid spent on imported consumption goods (a direct leakage from the domestic economy); yet, this possibility is implicit in choices for the parameter γ . The assumption that aid exclusively affects income via the investment channel implies, in theory and for the purposes of the main simulation exercise, there can be no persistent or permanent effects on income. All aid given in the past, as well as any indirect effects on private investment, will depreciate over the long run. Admittedly, this is a strong assumption; however, it provides analytical tractability to our baseline simulation. Furthermore, we consider extensions to this setting below.

Second, in addition to holding TFP fixed, the labour stock is, as noted, held constant throughout. As a result, we focus uniquely on the impact of changes in the capital stock on (per capita) income. This is equivalent to assuming that there are no changes to the quality of labour over time and that the population grows at the same rate as the work force. Thus, changes in income per capita, relative to a 'no aid' counterfactual path, can only derive from changes in the capital stock (per worker). This is a useful simplification. It circumvents the need to split the impact of aid across different factors of production. Moreover, we assume the economy begins in a steady state in which the real value of the capital stock is constant over time.

Third, the lag between aid delivery and its contribution to production (parameter μ) reflects delays that typically arise in investments financed by aid such as large-scale infrastructure, education, and health. Where the economy is already growing, such delays are expected to reduce the real impact of aid on growth as its contribution becomes smaller as a proportion of GDP. To see this, we derive an explicit expression for growth in terms of capital stock dynamics. Deploying our baseline assumptions ($L_t = 1, \theta_t = \bar{\theta} = 1$) implies that the definition for national income simplifies to $Y_t = K_t^\alpha$. Furthermore, assuming aid is provided in each period as a constant share of GDP: $A_t = \lambda Y_t, \lambda \geq 0$, the growth rate becomes:

$$\dot{Y}_t = \left[\frac{1}{K_{t-1}} (\lambda K_{t-\mu}^\alpha + \lambda \gamma K_{t-1}^\alpha + \delta K_0) + 1 - \delta \right]^\alpha - 1 \quad (3)$$

Equation (3) indicates the principal factors that determine the contribution of aid to growth at time t . They are the share of aid in GDP (λ), the short-run effect on private investment (γ), the lag through which aid fully adds to the capital stock (μ), the rate of capital depreciation, and the share of capital in income (α). As we assume no exogenous changes to the rate of return to capital, the model's counterfactual growth rate – i.e., the rate of growth that would have occurred in the absence of aid – is zero.⁸

To initialize the capital stock in each simulation we make two further assumptions: (i) the share of capital in aggregate income remains constant over time; and (ii) the marginal product of capital at $t = 0$ is equal to its current rental rate, denoted r_0 . The first assumption is a stylized empirical fact that is a direct corollary of Bowley's Law (Krämer 2011). Caselli and Feyrer (2007) describe the minimal conditions under which the second assumption is applicable. From these and previous assumptions, it follows that for chosen values of r_0 and α , the initial capital stock can be derived as follows:

$$K_0 = \left[\frac{r_0}{\alpha} \right]^{\frac{1}{\alpha-1}} \quad (4)$$

As noted above, our baseline model assumes that aid has no net effect on aggregate productivity. Views on this issue differ. On the one hand, some aid critics maintain that aid has long run pernicious effects on productivity. On the other hand, foreign aid often explicitly targets enhanced productivity growth through investments in education, health, technology, and institutions.⁹ Thus, if aid permanently augments TFP (e.g., via knowledge acquisition or introduction of new technologies), then steady state income levels would be raised permanently. There remains little consensus as regards the existence and/or empirical magnitude of any persistent effect of aid on economy-wide productivity. Consequently, we consider this possibility separately in an extension to the model (Section 6). Here, we note that this does not fundamentally change our conclusions.

Finally, whilst we assume that all aid is invested directly in the physical capital stock, we do not permit income gains from this additional capital to be reinvested domestically. Although in some

⁸ This is obtained trivially from equation (3) by setting $\lambda=0$ and allowing the current capital stock to be at its steady state level ($K_{t-1} = K_t = K_0$).

⁹ Alternatively, as low income countries are almost surely situated below their steady state income levels, successful investments of this kind would at least be expected to boost (speed-up) a process of convergence.

cases such reinvestment may represent an important indirect effect due to aid, the extent to which this occurs is likely to vary (and is not observed). To avoid arbitrary specifications of the model in this regard, we retain a conservative approach and assume that all income gains are either consumed or used to repay aid-related financing costs (see below).

3.2 Simulation protocol

The primary objective of the simulation exercise is to estimate the marginal effect of aid on growth under alternative combinations of the key parameters. Following the previous discussion, for any given set of these parameters, the effect of aid on growth is directly obtained when we hold all aspects of the model economy fixed through time, with the exception of the aid injection. Since we are interested in the long-run effect of aid on growth, in each simulation we assume a fixed share of income is received in foreign aid over a 30-year period – i.e., from $t = 1$ to $t = 30$.¹⁰ On this basis, and under previous assumptions, it follows that after 30 periods (plus the maximum lag on aid), the only changes to the capital stock will be depreciation charges.

The rationale for undertaking this simulation is to explore the range of values taken by the marginal effect of aid on growth, as well as the aggregate economic returns to aid, over a plausible domain of the model’s parameters. For each simulation, estimates of $J^{-1} \sum_{j=1}^J \dot{Y}_{t+j} / \lambda$ give the average marginal effect of Aid/GDP on growth over a window of J periods. We define a realization of each simulation, indexed by i , as a run of the model for a vector of exogenous parameters $s_i = \{\alpha, \lambda, \mu, \delta, r_0, \gamma\}_i$. If we could specify a distribution, Ω , associated with the (stochastic) vector s_i , then by taking repeated random draws from Ω through the simple economy outlined in equations (1) and (2), we would obtain a very good approximation to the distribution of returns to aid. In turn, this Monte Carlo method allows us to report moments of the estimated distribution of the marginal growth effects of aid (and associated estimates of returns to aid).

Insufficient prior information exists to completely specify Ω . In particular, covariances and moments greater than two are frequently unknown. We therefore use existing literature to identify a plausible range of values for each parameter and presume an agnostic independent uniform distribution for most parameters (see below). This corresponds to Laplace’s principle of insufficient reason, often adopted in Bayesian statistics to construct prior probability distributions. We proceed by drawing randomly from these distributions and calculating the associated marginal effects and returns to aid for each draw, as per a Monte Carlo analysis. The analysis of results departs from the formal Monte Carlo approach, however, by emphasizing the domain of results for different parameter permutations. That is, the intention is to highlight a plausible range of values for each parameter as well as any observed tendency for probability mass to concentrate within particular subspaces of their joint distribution.

Each of the exogenous parameters in s_i is chosen according to the following rules:

$$\alpha_i = 0.3 + \psi_{\alpha i}(0.7 - 0.3)$$

$$\lambda_i = \min\{0.005, 0.25 \cdot B_i(1.5, 5)\}$$

¹⁰ This amounts to fixing $\epsilon = 30$ in equation (3). This period could also be varied; however, it is not important from the current analytical perspective.

$$\mu_i = \text{int}(0.5 + 10 \cdot \psi_{\mu i})$$

$$\delta_i = 0.02 + 0.25 \cdot (\psi_{\delta i} - 0.02)$$

$$r_{0i} = N_i(0.27, .09) + [\delta_i - 0.25 \cdot (\psi_{r_i} - 0.02)]$$

$$\gamma_i = (\psi_{\gamma i} - 0.5)/2$$

where $B_i(a, b)$ indicates an independent random draw from the Beta distribution with shape parameters a and b ; $N_i(c, d)$ denotes a random draw from a normal distribution with mean c and standard deviation d ; and ψ_j is a random draw from the standard uniform distribution.¹¹ These choices are meant to capture a wide but plausible prior over the joint distribution of parameters. For instance, capital's share of income ranges from 30 to 70 percent; the delay from aid delivery to it making a full contribution to the capital stock is an integer uniformly distributed between 1 and 10 periods (years). Similarly, the share of aid in income is bound between 0.5 and 25 percent, but is distributed with a long right-hand tail with a median of approximately 5 percent. Broadly, this corresponds to what we find in the existing aid data.

Choices regarding the rate of depreciation and marginal returns to capital are also important and controversial. Few rigorous estimates of depreciation rates are found in the literature, much less for developing countries. One study is Bu (2006), who employs data from enterprise surveys. With the exception of Indonesia, where large tax incentives appear to drive very rapid effective depreciation rates, the author estimates aggregate capital stock depreciation rates of between 9 and 23 per cent for a range of developing countries (Bu 2006: Table 10). These relatively high values compare to rates of between 2 and 10 per cent employed in company accounts in the same countries (Bu 2006: Table 8), which are equivalent to rates used in developed countries and many cross-country capital stock estimates. Faced with these ranges, we take an agnostic stance and assume depreciation rates are uniformly distributed between 2 and 25 percent.

For rates of return, we rely primarily on the calculations in Caselli and Feyrer (2007). Their baseline estimates point to an average marginal product of capital, which is equivalent to the rental rate under the assumption of constant returns to scale and competitive capital markets of 0.27 with a standard deviation of 0.09 among low income countries. Notably, this estimate does not take into account 'correction factors' such as differences in the prices of capital goods or the absence of complementary factors. To some extent, these factors can be considered to contribute to the effective depreciation rate; as such, we impose a positive correlation between the rental rate and the rate of depreciation. For the simulation sample this correlation equals 0.53 and the net rate of return ($r_0 - \delta$) takes a mean of 15.4 per cent and standard deviation of 11.2 percent, which is closer to the corrected rates of return estimated by Caselli and Feyrer (2007).

Finally, in the absence of clear empirical guidance, the short-run effect on private investment, γ , is assumed to be uniformly distributed over the range $[-0.25, 0.25]$. For simulations where this parameter is negative, the notion is that harmful effects due to aid are material. Where it takes positive values, the notion is that aid has a crowding-in effect on private investment; this might be the case where expectations of future new infrastructure investments stimulate complementary private activities.

¹¹ In practice, for these uniform draws we employ a 10,000 x 5 matrix drawn via Latin Hypercube Sampling. This helps ensure systematic (as opposed to patchy or clustered) coverage of the entire search space.

4 Aid-growth results

To investigate the macroeconomic impact of aid for plausible combinations of parameters, we explore the domain of the parameter space via 10,000 simulations, whereby each simulation can be thought of as reflecting the behaviour of a hypothetical aid recipient. Before looking at these results, a simple visual illustration helps demonstrate the behaviour of the model. Figure 1 plots the time paths of key variables for a single simulation, wherein each parameter is fixed at the median of the corresponding generated distributions used in the simulations discussed below. As such, the figure captures the behaviour of the median aid recipient in our generated dataset. Panel 1(a) shows the path of aid inflows as a share of GDP, fixed at around 5 per cent from $t=1$ until $t=30$ and zero thereafter. Panel 1(b) indicates the volume of additional capital stock attributable to these aid inflows (net of depreciation), shown as a per cent of its base (steady state) level. Due to the lag between receiving aid and additional capital entering productive use, positive values are registered first at $t=7$. Similarly, although there are no aid inflows from $t=31$ onwards, new aid-financed productive capital continues to be added to the economy until $t=36$. Following equation (2), it is relevant to note that depreciation charges are made to the capital stock in each period. Since there are no further aid-financed additions to the capital stock from $t=37$ onwards, these depreciation charges then act to ensure that the economy gradually returns to its steady state position.

The income growth effects attributable to aid inflows are illustrated in Panel 1(c). The sharp discontinuities at $t=7$ and $t=37$ reflect non-linear shifts in the productive capital stock, reflecting discontinuities in the supply of aid. Specifically, the jump in growth at $t=7$ is driven by the first aid-financed addition to the productive capital stock (see above). From $t=7$ to $t=36$, the growth effect due to aid is positive but declining, which largely reflects the fact that depreciation costs are increasing relative to annual additions to the capital stock. At $t=36$ the capital stock is at its maximum and, following previous discussion (also, Panel 1b), suffers a large write-down due to depreciation charges in $t=37$, in turn generating negative income growth.¹² Finally, panel 1(d) illustrates the cash flow associated with aid, which is treated separately in Section 5.

To investigate the full set of simulations, for each run we calculate the marginal effect of aid on growth over 5, 10, and 30 year periods (starting at $t=t$), which correspond to alternative windows over which the aid-growth effect has often been assessed empirically. Table 2 summarizes these results. In order to provide a sense of the impact of the various parameter choices on these outcomes, for each parameter draw we create a dummy variable that takes a value of one if the parameter is larger than its median (over all simulations) and zero if not. A vector of ones (row 16 of the table) therefore would refer to simulations in which all parameters are above their respective medians. In essence, whilst results for each simulation can be thought of as referring to a hypothetical aid-recipient (country), the table reports the central tendency in distinct, non-overlapping sub-regions of the parameter space. For ease of exposition, all coefficients are multiplied by ten and thus indicate the expected growth response from a 10 per cent share of aid in GDP.

Three main conclusions can be drawn from the table. First, the marginal aid-growth results are sensitive to the length of the assessment window. The global mean of the distribution of aid-growth results is lowest when calculated over the shortest window at 0.49; however, for the 10- and 30-year periods the corresponding means are 0.87 and 0.70. Critically, the standard

¹² Alternative aid supply rules could be employed that induce a less volatile income growth response. However, volatility is not in focus here and therefore a simple time delimited fixed share aid supply rule is used.

deviations of these effects decline as the time horizon increases. For the 5- and 10-year windows, the standard deviations are 0.79 and 0.64 respectively, yielding coefficients of variation of 1.63 and 0.74. For the 30-year window, the distribution of effects is tighter, yielding a coefficient of variation equal to 0.51. These results imply that the estimated effect of aid is most sensitive to alternative parameter choices over shorter time frames. This is consistent with the comparatively mixed findings from empirical aid-growth studies that focus on relatively short time frames (e.g., 5-year panels).

The importance of assessing aid over an appropriate time frame is illustrated from boxplots of the simulated distributions of the marginal effect of aid on growth (for units equivalent to a 10 per cent Aid/GDP ratio) assessed after 5, 10, and 30 years, shown in Appendix Figure A1. A key point to note is that whilst 30 per cent of all aid-growth estimates are in the negative domain when viewed over a 5-year horizon, *none* of the 30 year estimates fall below zero. Thus, according to the model, the cumulative effect of aid in contributing to the capital stock outweighs possible short-run distortionary effects when viewed over longer horizons.

A second observation from Table 2 is that for windows of 5 and 10 years, a principal driver of differential effects is the choice of the lag length (μ_i). In the lower half of the table, where the lag ranges from 6 to 10 years (rows 9-16), the average marginal aid-growth coefficients are 0.00 and 0.48 for the 5 and 10 year windows, compared to 0.67 for the 30-year window. In these cases the full capital stock contribution only arrives from $t = 6$ or later, meaning that in the 5-year window the only possible effect of aid comes through the (mean zero) short-run investment effect. This explains why, for the 5-year window, the estimated marginal aid-growth effect is negative when the lag is long and the short-run investment effect is negative ($\gamma_i < 0$).

Third, the global average estimate of the marginal aid-growth effect over 30 years (equal to 0.07 when evaluated for units of 1 per cent aid in GDP) is only slightly lower than the estimated weighted effect found in recent empirical studies (Table 1), as well as estimates from meta-analyses that consider a wider range of studies (Mekasha and Tarp 2013). However, it bears recalling that our baseline model excludes any persistent effects of aid on productivity. Also, the 95 per cent confidence interval for the 30-year marginal aid-growth effect ranges from 0.015 and 0.158, which substantially overlaps with the distribution of empirical estimates in Table 1. Moreover, our results suggest that benefits due to aid in terms of additional national income are material. On average across the simulations, per capita income is 12.8 per cent higher at $t=30$ than it would have been without assistance, with a 95 per cent confidence interval of 1.3 to 35.7 percent.

5 Economic returns to aid

The previous section demonstrated that our numerical simulation model delivers estimates of aid-growth effects that are comparable to those in the recent empirical literature. Here we employ the same simulation results to estimate the returns to aid at the economy-wide level. The simulations are particularly well placed for this exercise because they can be used to deliver a time path for the cash flows associated with aid. This can be evaluated as follows:

$$CF_t = Y_t - (Y_0 + A_t) \tag{5}$$

which says that the cash flow associated with aid in a given year (i.e., its net value added) is the difference between per capita income at time t minus the counterfactual income (without aid)

plus the effective cost of aid in the same period. Under the assumption (applied throughout the simulations) that the only source of capital accumulation comes through aid, Y_0 is used as the counterfactual. The time path for the cash flow from the example simulation discussed previously is shown in Figure 1, panel 1(d). This profile is dissimilar to that of the marginal growth effect in panel 1(c) – the cash flow is negative for the first 10 periods; also it spikes upwards at $t=37$ when no further aid inflows occur but income remains higher than in the ‘no aid’ counterfactual due to accumulated additions to the capital stock.

From a project finance perspective, the returns to a given investment depend on the magnitude and profile of its associated cash flows. Typically, although cash flows are expected to be negative in early years, viable investments will yield a positive net present value (NPV) of cumulative flows over the life-time of the project. To calculate the NPV a discount factor must be specified ex ante, often referred to as the hurdle rate. Where this is unknown, the internal rate of return (IRR) can be used to identify the discount factor which yields an NPV of zero. Thus, the IRR indicates the maximum cost of capital required for the project to break-even. In this light, IRR calculations for each aid simulation indicate the maximum effective financing costs of foreign aid inflows (incurred by recipients) that generate no net losses. This interpretation fits our simulation model as we assume all income gains associated with aid are available for consumption or payment of financing costs. For instance, if the IRR is found to be 10 percent, then the effective costs of financing aid over the lifetime of the project must be below 10 per cent for such aid to have a positive overall net income effect.

The IRR has been subject to criticism in both the development and corporate finance literatures (e.g., Kierulff 2008). This stems from the fact that, like the NPV, it is primarily designed to evaluate projects without interim cash flows – i.e., projects with a single initial outlay and a final return. Where interim cash flows occur (see Figure 1d) there is debate as to how these should be discounted. The implicit assumption of the IRR is that interim cash flows (positive and negative) attract a rate of return equal to the IRR. For our purposes this assumption is reasonable since the notion of a break-even average cost of capital, applied to positive and negative cash flows, provides a useful summary of the corresponding macroeconomic returns to aid.¹³ Indeed, if the IRR on aid were found to be very high (say, over 20 percent), this might raise questions as to why concessional rates of financing were necessary (see the discussion of Lucas 1990).

To complement the IRR measure, two additional returns indicators are calculated for each simulation. The first (CFY) is the average ratio of net cash flow to initial income (cash flow income (income= Y)). This is equivalent to the NPV calculated with a discount factor of zero, divided by the product of the initial income level and the number of periods in the simulation. Values greater than zero imply that the (undiscounted) cumulative cash flow from the project is positive – i.e., the real value of income generated by aid is greater than the value of aid received, regardless of the time profile of positive and negative cash flows (to which donors may be indifferent). The second indicator (PCCF) reports the first time period when cumulative cash flows are positive (positive cumulative cash flows), which gives the time horizon necessary for aid to generate cumulative income gains equal in value to the amount of aid given, again ignoring

¹³ An alternative to the IRR is the modified IRR (MIRR), which explicitly allows interim cash flows to attract different rates, depending on whether they are negative or positive. However, in the present context where the ‘project’ is active over a long period (at least 30 years) and there is no large terminal cash flow, the MIRR is sensitive to assumptions regarding interim rates. Moreover, correctly specifying such rates is likely to be controversial. For these reasons, the MIRR metric is not employed.

time discounting or financing costs.¹⁴ This captures the extent of patience required by donors to observe positive overall returns to aid outflows.

Table 3 reports calculations for the three aid performance metrics (IRR, CFY, and PCCF). It follows the same format as Table 2. Considering all 10,000 simulations, the simple average IRR is equal to 7.3 percent; however, the average is less informative than the ranges.¹⁵ As Appendix Figure A2 illustrates, the simulated empirical distribution of the IRR metric spans a wide range, thereby confirming that returns to aid (as expected) are sensitive to parameter choices. Nevertheless, 90 per cent of simulations present a positive IRR and 28 per cent attract an IRR above a 10 per cent threshold. The 95 per cent confidence interval of the simulated distribution of returns ranges from -2.3 per cent to 19.9 percent. Put differently, Table 3 indicates that the average IRR in the most favorable sub-domain of parameters (row 4) is 18.7 percent; however, parameter draws located in the least favourable sub-domain present an average IRR of just 1.1 per cent (row 13).

What accounts for these relatively large differences in IRR outcomes? As noted, the model does not allow permanent productivity or indirect reinvestment effects from aid inflows (see Section 3). As a consequence, positive cash flows associated with aid decline relatively rapidly once aid flows cease as depreciation charges diminish the additional capital stock (see Figure 1). More importantly, while time discounting does not feature in the aid-growth estimates, it plays a critical role in the evaluation of economic returns. This dimension is fundamental due to the relatively long lags permitted between aid inflows and the start of significant income benefits, at more than five years on average.¹⁶ That is, the switch from considering long-period average growth rates (Section 4) to evaluating the IRR of associated aid-related cash flows means that the timing of positive cash flows becomes crucial. Relatedly, income gains from aid do not arrive as a lump sum. Rather, they arise as returns to a stock of additional capital and, consequently, are spread out over time. Thus, the NPV of each simulated cash flow is typically substantially lower than its undiscounted cumulative sum. From this perspective, the relatively moderate IRR values, which represent the bulk of results reported in Table 3, substantiate a view of foreign aid (on aggregate) as a long-term investment whose net benefits emerge in a cumulative fashion over a long period.

A complementary interpretation of these results is that they substantiate a role for foreign aid as a source of concessionary financing. A primary rationale for foreign aid is that recipients cannot raise financing on domestic or international markets (on reasonable terms). However, if macroeconomic returns to foreign aid were generally high, then the need for concessionary financing would be more ambiguous. The simulations suggest that whilst foreign aid can enhance growth (Section 4), the corresponding macroeconomic returns are often moderate due to the time profile of cash flows. In addition, the macroeconomic benefits of aid are likely to be diffuse, which implies that repayments would need to be financed via the tax system with an associated marginal cost of funds. In this light, concessionary terms of financing are likely to remain necessary.

The additional metrics reported in Table 3 are broadly consistent with the IRR estimates. Implicitly employing a discount factor of zero, the CFY results indicate that the cumulative net

¹⁴ If the project generates no positive cash flows then the last period is used.

¹⁵ A small number of outlier IRR values are excluded, defined as being over 50 per cent or below -50 percent.

¹⁶ For instance, a flow of 100 units received after six periods is only worth 56.4 in present terms based on a discount factor of 10 percent.

cash flow associated with aid is consistent with a 2.6 per cent income gain per period (relative to initial income levels) over the entire project lifecycle. Thus, for a median project span of 71 periods, the total net income gain due to aid is equal to 181 per cent of initial income on average. Mathematically, the same share of CFY values as IRR values (90 percent) are positive, reiterating that for the large majority of parameter choices positive income benefits due to aid are greater than or equal to the real value of aid received (ignoring time preferences). In turn, the PCCF results reinforce the conclusion that foreign aid must be evaluated over a long-run perspective.

6 Productivity effects

The analysis of the previous sections ignored the potential impact of aid on total factor productivity. As a result, the mechanics of the baseline model (and simulations) imply that over the very long run the effect on the economy of aid given in any specific period will be zero. This occurs because the aid-financed capital stock eventually depreciates and the economy converges back towards its assumed original steady state (assuming aid flows cease). However, many debates around foreign aid suggest it may have more permanent aggregate effects (e.g., Moss et al. 2006; Rajan and Subramanian 2007; Djankov et al. 2008); thus, exclusion of such productivity effects could be significant.

To address this gap, we extend the baseline model and allow for aid-induced productivity effects. At the outset, it is important to highlight that in doing so we enter controversial terrain. Despite wide-ranging claims, concrete empirical evidence regarding the nature and scale of aid-induced productivity effects is scarce. Moreover, precisely how these might enter a simulation model can only be considered tentative. Accordingly, the present section seeks to quantify the range of consequences of incorporating productivity effects for the outcomes of interest *relative* to our baseline simulations. This can be seen as a form of robustness check on the baseline results.

In order to incorporate permanent productivity effects, we assume that aid inflows over a sustained period gradually push the economy's aggregate TFP to a new level, given by:

$$\theta_i^* = \bar{\theta} + \tau_i \lambda_i, \tau_i = 2\psi_{\tau i} \quad (6)$$

This equation says that the aid-adjusted permanent level of TFP, which in turn defines a new steady state income level, is a deterministic function of the amount of aid received as a share of income (λ) multiplied by a scaled independent draw from the uniform distribution. Since $0 < \tau_i < 2$, we impose a positive correlation between Aid/GDP and the final level of productivity. A main reason for doing so is that IRR calculations become highly unstable in the absence of (any) positive cash flows. We recognize, however, that this assumption is open to question. Nonetheless, as shown below, because the first order effects of τ are symmetric around the origin (on average), we can use the simulation results to consider the (average) magnitude of τ necessary for aid-induced productivity losses to fully dominate capital stock accumulation effects.

Retaining our convention that aid begins to enter the economy at $t=1$ and continues for 30 periods, the time path for the TFP parameter (dropping simulation subscript λ), θ_t , is given by:

$$\theta_t = \begin{cases} \bar{\theta}=1 & \text{if } t \leq \mu \\ \theta^* & \text{if } t > \kappa + \mu \\ (\theta^*)^{(t-\mu)/\kappa} & \text{otherwise} \end{cases} \quad (7)$$

This says that the productivity parameter gradually increases to its specified final value starting at period $t = \mu + 1$ and ending at period $t = \kappa + \mu$. Parameter κ thus indicates the duration of the productivity phase-in period and is drawn from a discrete uniform distribution taking a minimum of ten and maximum of 30. This set-up implies that productivity grows at a compound rate of $(\theta^*)^{1/\kappa}$ per period during this phase-in period.

To simulate the extended model, the baseline simulation protocol is employed without alteration. Under the assumptions of the model, the volume of the capital stock will always converge to its original (steady state) value regardless of the chosen productivity parameter; thus, each simulation can be terminated in the same way as before. To make comparison between the extended and baseline models as clear as possible, we employ the same parameter values as the corresponding baseline simulations for each extended simulation. Thus, the only difference between each of the baseline and extended simulations reflects the chosen path for the productivity (TFP) parameter. In this way, differences in outcomes between the extended and baseline models reflect the unique contribution of aid-induced productivity effects.

Table 4 summarizes results from the 10,000 extended simulations. As before, we divide the parameter space into interesting sub-domains and show results separately for each sub-domain, as well as globally. In the table, we focus on the new parameters, τ and κ , both of which are split by their medians (1 and 18 respectively). The outcomes of interest (denoted by z) are the marginal effect of aid on growth (calculated over a 30-period window) and the IRR measure. For each of these, the table reports means of: (i) the raw outcome measures from the extended simulations, z_e ; (ii) the absolute difference between the extended and baseline simulation outcomes, $\Delta_e = z_e - z_b$, which is always greater than or equal to zero for the assumed parameter bounds; and (iii) a measure of the relative magnitude of these differences, calculated as:

$$\Lambda = 1 - 2(z_b/z_e) \quad (8)$$

The rationale for this final measure derives from the supposition of effect symmetry. If $\Delta_e > z_b$ then the difference between the extended and baseline simulation outcomes, which reflects the induced productivity effect due to aid, is absolutely larger than the pure capital accumulation effect. By symmetry, this means that if the corresponding value for τ were to take the opposite sign (negative) we would expect the outcome variable to be consigned to a negative domain (*ceteris paribus*).¹⁷ In other words: $\Delta_e > z_b \Leftrightarrow \Lambda > 0$, which provides a simple benchmark to evaluate the relative magnitude of the productivity effects.

As might be expected, Table 4 shows that when positive aid-induced productivity effects are incorporated, the estimated aggregate economic contribution of aid increases. Over all 10,000 simulations, the average simulated productivity increment is 5.7 per cent (i.e., the TFP parameter

¹⁷ A further rationale for this measure is that for $z_b > 0$, which covers virtually all simulations, we have: $\Lambda \in (-1,1)$. These bounds are helpful for subsequent analysis and visual illustration.

increases from 1 to 1.057 on average), with a minimum of 0 and a maximum of 38.0 percent. These choices are associated with an average marginal aid-growth effect of 1.0 percentage points (expressing Aid/GDP in units of 10 percent) and an average IRR of 11.9 percent, which are respectively 58 and 63 per cent larger than the overall baseline simulation averages. In addition, the variance of outcomes increases moderately when productivity effects are introduced, underlining the likely relevance of substantial heterogeneity in the macroeconomic effects of aid.

What should we make of these results? First, the average value for τ is equal to one, meaning that on average we assume that a sustained inflow of aid equal to λ per cent of GDP over a 30-year period also increases the aggregate level of TFP by λ percentage points over the long run. At the means of the new parameter distributions, the implied growth in TFP over the phase-in period (mean of 20 years) equals 0.25 percentage points per period. In light of the assumed volume of aid relative to GDP this is a moderate but non-negligible effect size – for example, as Nordhaus (2001) points out, long-run estimates of annual labour productivity growth in the USA cluster around one percentage point on average.

Second, under our parameter assumptions, it is notable that the extended simulation results for the long-run marginal effect of aid on growth more closely approximate the distribution of results from recent empirical results. Recall that the weighted average of the regression estimates reported in Table 1 suggests that an inflow of aid equal to 10 per cent of GDP is expected to raise income growth by 1.25 percentage points on average, which is slightly higher than the global average of the extended simulations. Therefore we conclude that recent empirical estimates regarding the impact of aid (Table 1) point to moderate positive productivity effects due to aid.

Third, by the same token, the results suggest that for small to moderate values of τ , the induced productivity effects are insufficient to dominate capital accumulation effects. Instead, we only consistently find that $\Delta_e > z_b$ for values of τ that are substantially larger than one. This insight is illustrated in Figure 2, which is a scatterplot of τ against the relative magnitude indicator (Λ) for both the marginal aid-growth effect and IRR (each point represents one simulation). In both panels, the large majority of observations lie below the benchmark value $\Lambda = 0$ (i.e., 84 and 75 per cent respectively). The bivariate regression line, which is just the estimated linear prediction of Λ given τ , crosses zero above $\tau = 1.5$ in both panels. The implication is that substantial negative aid-induced productivity effects are required to fully undermine the positive investment effects. Moreover, for the macroeconomic effects of aid to be both negative *and* statistically significant, even larger negative aid-induced productivity effects would be necessary. From this perspective, recent empirical results are incompatible with the postulate that aid has generated large pernicious effects on aggregate productivity (on average) over the long run.

7 Summary and implication

7.1 Summary

This paper took as its point of departure the substantial degree of consistency across a range of recently published empirical studies that investigate the long-run marginal effect of aid on growth. Employing different methods and data sets, the very large majority of these studies place this effect in the positive domain. The weighted average result from these studies, where the weights correspond to the inverse of the standard error of the estimates, indicates that a

sustained inflow of foreign aid equivalent to 10 per cent of GDP is roughly expected to raise growth rates per capita by one percentage point on average.

The first objective of this study was to investigate the coherence of these results using numerical simulations. To do so, we set up a simple simulation model that captures key factors that are likely to influence the size and direction of the macroeconomic effects of aid. These included a lag between receipt of aid and its contribution to income, as well as shorter-term effects on private investment (positive or negative). By assumption, our baseline model ruled-out permanent income effects from productivity spillovers or indirect gains from reinvestment.

Based on 10,000 random draws from the parameters underlying the model (constrained over reasonable empirical ranges), we simulated the growth impact of aid and calculated project performance metrics such as the IRR. The results revealed that, when assessed over a short time horizon (five years), the marginal effect of aid on growth is negative over a significant share of observations and has large variance. This reflects both sensitivity to the lag structure through which aid effectively contributes to the capital stock, as well as the impact of short-run factors such as crowding-out effects. However, when assessed over the long-run (30 years), the marginal effect of aid is consistently positive and distributed more tightly. After 30 years of assistance, incomes in the simulated economies were 12 per cent larger on average. Moreover, the bulk of simulated marginal effects of aid lie in ranges comparable to findings from recent empirical studies. Thus, numerical simulations substantiate the need to evaluate aid over long periods and indicate that recent empirical results are coherent with plausible parameterizations of a simple yet general growth model. The simulations also highlight substantial heterogeneity in the aggregate effects of aid.

The IRR calculations, based on the same simulations, demand a more nuanced interpretation. The average IRR over all 10,000 model runs is equal to 7.3 percent, with 95 per cent of outcomes falling in the interval ranging from -2.4 to 18.7 percent. The simulated IRR is highly sensitive to the time profile of (net) cash flows due to foreign aid. On average, cumulative undiscounted net cash flows only turn positive after more than 32 periods (years). This reflects the lag structure of the model and that aid contributes positively to income via additions to the stock of capital, meaning that returns are spread out over time. Long lags between the injection of funds and realization of net benefits militate against high IRRs. This substantiates a macroeconomic view of aid as a long-term investment whose benefits cumulate slowly over relatively long periods. Additionally, rates of return of this magnitude confirm a valid role for aid as a concessionary source of financing.

As an extension to our baseline model, we allowed aid to have a permanent effect on aggregate productivity. The results suggested this is likely to be a relevant mechanism through which aid has macroeconomic impacts. For moderate choices of the partial correlation coefficient between aid and the final level of TFP (τ), we found that the simulated distribution of the long-run marginal aid-growth effect more closely approximated the weighted average from recent empirical studies (Table 1). The corresponding average IRR also was significantly higher, at 11.9 percent, but remains within a plausible domain.

7.2 Implications

The upshot of our analysis is that the overall picture given by the aid-growth empirics of Table 1 is reasonable. They are compatible with a range of plausible parameterizations of a simple yet general model of the growth process. The central tendency of our baseline simulations approaches the mean of recent empirical aid-growth studies and falls spot on once moderate

positive effects from aid on aggregate productivity are permitted. In sum, the latest simulation and empirical evidence suggest positive average growth returns to foreign aid when viewed over long time frames. The notion that aid has broadly harmed development performance receives essentially no empirical support.

Aid does not deliver the kind of high returns often implied by two gap model exercises, which underpinned optimistic projections of rapid economic development in the early days of foreign assistance (a point vividly made by Easterly 1999). Similarly, aid proponents have suggested that aid may deliver very large returns when applied to escape from poverty traps. As Kraay and Raddatz (2007) note, however, these hopes appear not to have materialized. The observed reality, which is consistent with our simulation results, is that the journey from low to middle income status typically remains long and arduous, even when supported by aid inflows.

The same results also suggest that neither the predictions of simplistic two gap modelling nor the existence of poverty traps is required to rationalize foreign aid. A sufficient criterion is its macroeconomic rate of return. The distribution of returns summarized in Table 2 (also Figure A2), which should be viewed as *ex post*, cluster around commonly used rates of return in *ex ante* project analysis such as the 10 per cent return cutoff applied by the World Bank and the seven per cent cutoff employed by the United States Office of Management and Budget (Pohl and Mihaljek 1992; Powers 2003). Thus, recent growth empirics combined with a simple but reasonable model of growth point to rates of return to aid in a range that is commonly associated with successful long-run public sector investments.

Why did it take so long to figure this out? Both aid volumes and their associated impacts are not so large as to be easily identifiable in macroeconomic data. The simulation modeling presented here underscores that long time frames are required to detect a growth impact. This reflects lags in the realization of benefits and the relatively moderate contribution of aid to the overall growth rate. In reality, detecting the contribution of aid is further complicated by large fluctuations in growth that have been an inherent part of the experience of nearly all developing countries. On top of this, and as noted earlier, observations of both the flow of aid funds to developing countries and the growth rates achieved are known to be imperfect. For these reasons, it is not surprising that the economics profession is only recently converging on the broadly similar empirics presented in Table 1.

While growth is clearly important, the return to aid in the form of enhanced economic growth is not the only metric that can or should be employed to evaluate aid. And, other metrics largely support the case for aid. For example, aid has been shown to contribute to accumulation of important elements of human capital, particularly improved educational attainment and improved health (Arndt et al. 2014 review this literature). These are expected to yield indirect growth benefits but can also be considered merit goods with intrinsic worth. Humanitarian assistance is given with the expectation of saving lives rather than making a growth contribution.¹⁸ Finally, under the assumption of diminishing marginal utility of income, the chasm that separates the living standards of citizens of recipient countries and the living standards of citizens of donor countries should enter the calculus. In income terms, our analysis


¹⁸ A contribution to growth may nevertheless occur. For example, Hoddinott et al. (2008) find that improvements in early childhood nutrition contribute to adult labour productivity and hence to long-term economic growth. To the extent that humanitarian assistance prevents major declines in early childhood nutrition, it may also improve the adult labour productivity of cohorts born during or shortly before times of crisis relative to a no assistance counterfactual.

indicates that the monetary gains to recipients are somewhat larger than the monetary costs imposed on donors (depending on the assessed opportunity cost of the aid to the donor). However, a utility lens would greatly magnify the assessment of the returns to aid.

To close, we note that simulation is a powerful and flexible tool employed in a wide array of disciplines to shed light on empirical phenomena. In a widely cited article, Alroy (2001) employed simulation to better understand the end-Pleistocene mass extinction of large species (mega-fauna). He finds that across a wide array of plausible parameter values, human population growth and hunting almost invariably leads to a major mass extinction. Here, we employed simulation to better understand the returns to aid investments over the past 40 years. Across a wide array of plausible parameter values we find that the best evidence points to a material contribution of aid to achieving development objectives. We conclude that calls for the extinction of aid and its associated institutions on the basis of poor or negative returns are unjustified.

Appendix


Figure A1: Boxplot of simulated empirical distributions of the marginal impact of aid on growth evaluated over alternative horizons (5, 10, and 30 periods)



Notes: Aid/GDP is given in units of 10 percent; periods refer to the window over which the average marginal aid-growth effect is calculated (starting at $t=1$).

Source: authors' calculations based on 10,000 simulations.

Figure A2: Histogram and density plots of the simulated distribution of internal rates of return to foreign aid



Notes: IRRs are calculated over the full project lifecycle.

Source: authors' calculations based on 10,000 simulations.

References

- Alroy, John. 2001. 'A Multispecies Overkill Simulation of the End-Pleistocene Megafaunal Mass Extinction.' *Science* 292: 1893-96.
- Arndt, Channing, Sam Jones, and Finn Tarp. 2014. 'Assessing Foreign Aid's Long-Run Contribution to Growth and Development.' *World Development* (forthcoming), dx.doi.org/10.1016/j.worlddev.2013.12.016.
- Arndt, Channing, Sam Jones, and Finn Tarp. 2011. 'Aid Effectiveness: Opening the Black Box.' WIDER Working Paper 2011/44. Helsinki: UNU-WIDER.
- Arndt, Channing, Sam Jones, and Finn Tarp. 2010. 'Aid, Growth, and Development: Have We Come Full Circle?.' *Journal of Globalization and Development* 1(2), 1-29.
- Brückner, Markus. 2013. 'On the Simultaneity Problem in the Aid and Growth Debate.' *Journal of Applied Econometrics* 28(1): 126-50.
- Bu, Yisheng. 2006. 'Fixed Capital Stock Depreciation in Developing Countries: Some Evidence From Firm Level Data.' *Journal of Development Studies* 42(5): 881-901.
- Caselli, Francesco, and James Feyrer. 2007. 'The Marginal Product of Capital.' *The Quarterly Journal of Economics* 122(2): 535-68.
- Clemens, Michael A., Steven Radelet, Rikhil R. Bhavnani, and Samuel Bazzi. 2012. 'Counting Chickens when they Hatch: Timing and the Effects of Aid on Growth.' *The Economic Journal* 122(561): 590-617.
- Dalgaard, Carl-Johan, Henrik Hansen, and Finn Tarp. 2004. 'On The Empirics of Foreign Aid and Growth.' *The Economic Journal* 114(496): F191-F216.
- Dalgaard, Carl-Johan, and Henrik Hansen. 2005. 'The Return to Foreign Aid.' Working Paper 05-04. Copenhagen: University of Copenhagen, Department of Economics.
- Djankov, Simeon, Jose G. Montalvo, and Marta Reynal-Querol. 2008. 'The Curse of Aid.' *Journal of Economic Growth* 13(3): 169-94.
- Easterly, William. 1999. 'The Ghost of Financing Gap: Testing The Growth Model Used in the International Financial Institutions.' *Journal of Development Economics* 60(2): 423-38.
- Herzer, Dierk, and Morrissey, Oliver. 2013. 'Foreign Aid and Domestic Output in the Long Run.' *Review of World Economics* 149(4): 723-48.
- Hoddinott, John, John A. Maluccio, Jere R. Behrman, Rafael Flores, and Reynaldo Martorell. 2008. 'Effect of a Nutrition Intervention during Early Childhood on Economic Productivity in Guatemalan Adults.' *Lancet* 371: 411-16.
- Juselius, Katarina, Møller, Niels F., and Finn Tarp. 2014. 'The Long-Run Impact of Foreign Aid in 36 African Countries: Insights from Multivariate Time Series Analysis.' *Oxford Bulletin of Economics and Statistics* 76(2): 153-84.
- Kalyvitis, Sarantis, Thanasis Stengos, and Irene Vlachaki. 2012. 'Are Aid Flows Excessive or Insufficient? Estimating the Growth Impact of Aid in Threshold Regressions.' *Scottish Journal of Political Economy* 59(3): 298-315.
- Kierulff, Herbert. 2008. 'MIRR: A Better Measure.' *Business Horizons* 51(4): 321-329.
- Kraay, Aart, and Claudio Raddatz. 2007. 'Poverty Traps, Aid, and Growth.' *Journal of Development Economics* 82: 315-47.

- Krämer, Hagen M. 2011. 'Bowley's Law: The Diffusion of an Empirical Supposition into Economic Theory.' *Cahiers d'Économie Politique/Papers in Political Economy* 2(61): 19-49.
- Lessmann, Christian, and Gunther Markwardt. 2012. 'Aid, Growth and Devolution.' *World Development* 40(9): 1723-49.
- Lof, Matthijs, Tseday J. Mekasha, and Finn Tarp. 2014. 'Aid and Income: Another Time-series Perspective.' *World Development* (forthcoming), dx.doi.org/10.1016/j.worlddev.2013.12.015
- Lucas, Robert E. 1990. 'Why Doesn't Capital Flow From Rich to Poor Countries?.' *The American Economic Review* 80(2): 92-6.
- Mekasha, Tseday J., and Finn Tarp. 2013. 'Aid and Growth: What Meta-Analysis Reveals.' *Journal of Development Studies* 49(4): 564-83.
- Michaelowa, Katharina. 2004. 'Aid Effectiveness Reconsidered – Panel Data Evidence for the Education Sector.' HWWA Discussion Paper 264.
- Minoiu, Camelia, and Sanjay G. Reddy. 2010. 'Development Aid and Economic Growth: A Positive Long-run Relation.' *The Quarterly Review of Economics and Finance* 50(1): 27-39.
- Mishra, Prachi, and David Newhouse. 2009. 'Does Health Aid Matter?.' *Journal of Health Economics* 28(4): 855-72.
- Mosley, Paul. 1986. 'Aid-effectiveness: The Micro-Macro Paradox.' *IDS Bulletin* 17(2): 22-7.
- Moss, Todd, Gunilla Pettersson, and Nicholas van de Walle. 2006. 'An Aid-Institutions Paradox? A Review Essay on Aid Dependency and State Building in sub-Saharan Africa.' Technical Report 74. Washington, DC: Center for Global Development.
- Moyo, Dambisa. 2009. *Dead Aid: Why Aid Is Not Working aHow There is A Better Way for Africa*. London: Penguin Books.
- Nordhaus, William D. 2001. 'Productivity growth and the new economy.' NBER Working Paper 8096. Washington, DC: National Bureau of Economic Research. Available at: www.nber.org/papers/w8096.
- Nowak-Lehmann, Felicitas, Axel Dreher, Dierk Herzer, Stephan Klasen, and Inmaculada Martinez-Zarzoso. 2012. 'Does Foreign Aid Really Raise Per Capita Income? A Time Series Perspective.' *Canadian Journal of Economics/Revue Canadienne d'Économie* 45(1): 288-313.
- Pohl, Gerhard, and Dubravko Mihaljek. 1992. 'Project Evaluation and Uncertainty in Practice: A Statistical Analysis of Rate-Of-Return Divergences of 1,015 World Bank Projects.' *The World Bank Economic Review* 6(2): 255-77.
- Powers, Kyna. 2003. 'Benefit-Cost Analysis and the Discount Rate for the Corps of Engineers Water Resource Projects: Theory and Practice.' Congressional Research Service Report RL31976.
- Rajan, Raghuram G., and Arvind Subramanian. 2008. 'Aid and Growth: What Does the Cross-Country Evidence Really Show?.' *The Review of Economics and Statistics* 90(4): 643-65.
- Rajan, Raghuram G., and Arvind Subramanian. 2007. 'Does Aid Affect Governance?.' *The American Economic Review* 97(2): 322-27.
- Roodman, David. 2007. 'The Anarchy of Numbers: Aid, Development, and Cross-Country Empirics.' *The World Bank Economic Review* 21(2): 255-77.

Table 1: Point estimates of the marginal effect of aid on growth from recent studies

Study ^a	Reference	Period	Specific ^b	Beta	Std. Error	≈ Prob.
RS08	Table 4, col. 1	1960-2000	Linear	0.06	0.06	0.30
RS08	Table 4, col. 2	1970-2000	Linear	0.10	0.07	0.17
MR10	Table 4, col. 1	1960-2000	Linear	0.08	0.03	0.01
AJT10	Table 6, col. 2	1960-2000	Linear	0.09	0.04	0.02
AJT10	Table 4, col. 4	1970-2000	Linear	0.13	0.05	0.01
CRBB12	Table 7, col. 6	1970-2005	Non-linear	0.15	0.06	0.01
CRBB12	Table 7, col. 10	1970-2005	Non-linear	0.31	0.17	0.07
CRBB12	Table 9, col. 9	1971-2005	Non-linear	0.27	0.13	0.04
CRBB12	Table 9, col. 9	1971-2005	Non-linear	0.42	0.20	0.04
KSV12	Table 2, col. 5	1970-2000	Linear	0.05	0.05	0.32
LM12	Table 3, col. 4	1960-2001	Linear ^c	0.85	0.43	0.05
NDHKM12	Table 1, col. 4	1960-2006	Linear	-0.02	0.01	0.14
B13	Table 3, col. 1	1960-2000	Linear	0.12	0.04	0.00
B13	Table 3, col. 1	1970-2000	Linear	0.18	0.07	0.01
HM13	Table 2, row 2	1971-2003	Linear ^d	-0.01	0.00	0.00
AJT14	Table 1, col. 2	1970-2007	Linear ^e	0.30	0.18	0.09
Mean	Unweighted			0.19	0.05	0.00
	Weighted			0.12	0.02	0.00

Notes: ^a. RS08 is Rajan and Subramanian (2008); MR10 is Minoiu and Reddy (2010); AJT10 is Arndt et al. (2010); CRBB12 is Clemens et al. (2012); KSV12 is Kalyvitis et al. (2012); NDHKM12 is Nowak-Lehmann et al. (2012); LM12 is Lessmann and Markwardt (2012); B13 is Brückner (2013), HM13 is Herzer and Morrissey (2013); and AJT14 is Arndt et al. (2014). ^b. For non-linear specifications (which involve a squared aid term) the marginal effect of Aid/GDP on growth is evaluated by fixing Aid/GDP at 2.5 percent; standard errors are also approximate for these cases. ^c. For comparability, this result is the marginal effect due to aid assuming no decentralization of government spending; however, the authors find that if the degree of decentralization exceeds 7 percent, the marginal effect of aid on growth is no longer statistically significant. ^d. This estimate controls for investment and is derived as an average from country-specific regressions. ^e. Beta coefficient and standard errors are adjusted to raw values from standardized values as reported in the study. Standard errors for unweighted and weighted means are derived from the reported distribution of beta coefficients in the table. Probability is based on the normal distribution.

Source: authors' collation based on Google Scholar references.

Table 2: Summary of the marginal effect of aid on growth from numerical simulations

Id	Parameter domains				Obs.	Window (periods)		
	μ	λ	$r_0 - \delta$	γ		5	10	30
1	0	0	0	0	603	0.55	0.76	0.43
2	0	0	0	1	621	0.83	1.01	0.55
3	0	0	1	0	611	1.00	1.47	0.91
4	0	0	1	1	631	1.54	1.94	1.14
5	0	1	0	0	665	0.52	0.72	0.42
6	0	1	0	1	626	0.85	0.99	0.53
7	0	1	1	0	648	1.01	1.40	0.84
8	0	1	1	1	595	1.53	1.81	1.01
9	1	0	0	0	593	-0.15	0.21	0.41
10	1	0	0	1	628	0.16	0.47	0.52
11	1	0	1	0	644	-0.30	0.36	0.82
12	1	0	1	1	669	0.30	0.88	1.02
13	1	1	0	0	622	-0.15	0.20	0.39
14	1	1	0	1	642	0.15	0.46	0.49
15	1	1	1	0	614	-0.31	0.39	0.77
16	1	1	1	1	588	0.30	0.87	0.95
Mean					625.0	0.49	0.87	0.70
St. dev.					24.30	0.79	0.64	0.36

Notes: marginal effects are calculated for Aid/GDP in units of 10 percent; parameter domains are dummy variables taking a value of one if the indicated parameter (or combination thereof) is greater than its median over all simulations; in column order the parameter domains refer to the lag length, the share of aid in GDP, the net return at time zero, and the short-run crowd-in effect.

Source: authors' calculations based on 10,000 simulations.

Table 3: Project performance metrics for foreign aid from numerical simulations

Id	Parameter domains				Obs.	Window (periods)		
	μ	λ	$r_0 - \delta$	γ		IRR	CFY	PCCF
1	0	0	0	0	603	2.42	0.34	44.31
2	0	0	0	1	621	6.29	0.75	36.57
3	0	0	1	0	611	12.80	1.94	16.11
4	0	0	1	1	631	18.70	2.72	11.65
5	0	1	0	0	665	1.91	0.83	53.53
6	0	1	0	1	626	5.60	2.09	38.04
7	0	1	1	0	648	11.76	5.17	17.38
8	0	1	1	1	595	17.26	7.67	12.63
9	1	0	0	0	593	1.36	0.32	52.31
10	1	0	0	1	628	3.53	0.69	39.88
11	1	0	1	0	644	6.27	1.69	27.51
12	1	0	1	1	669	10.01	2.52	19.75
13	1	1	0	0	622	1.06	0.79	56.98
14	1	1	0	1	642	3.14	1.86	47.04
15	1	1	1	0	614	5.96	4.92	28.88
16	1	1	1	1	588	9.39	6.98	21.17
Mean					625.0	7.34	2.56	32.76
St. dev.					24.30	6.93	2.99	40.65

Notes: parameter domains are dummy variables taking a value of one if the indicated parameter (or combination thereof) is greater than its median over all simulations; in column order the parameter domains refer to the lag length, the share of aid in GDP, the net return at time zero, and the short-run crowd-in effect. IRR is the internal rate of return (in percent). CFY is the ratio of the average net cash flow over all periods to initial income (in percent), with no time discounting. PCCF is the time horizon necessary for aid to generate cumulative income gains equal in value to the amount of aid given, ignoring time discounting.

Source: authors' calculations based on 10,000 simulations.


Table 4: Summary results from extended simulations with aid-induced productivity effects

Parameter domains		Aid-growth effect			IRR		
τ	κ	z_e	Δ_e	Λ	z_e	Δ_e	Λ
(0,1]	(10,20]	0.87	0.18	-0.56	10.37	3.11	-0.43
(1,2]	(10,20]	1.22	0.52	-0.09	14.64	7.36	-0.05
(0,1]	(20,30]	0.85	0.16	-0.59	9.68	2.37	-0.50
(1,2]	(20,30]	1.17	0.47	-0.14	12.87	5.52	-0.16
Mean	.	1.03	0.33	-0.34	11.94	4.66	-0.28
St. dev	.	0.42	0.19	0.36	7.17	3.67	0.46

Notes: each simulation of the extended model is a replicate of a baseline simulation but with different productivity parameters (only); parameter domains refer to categories of the new parameters of interest; τ is the coefficient on Aid/GDP used to derive the final level of TFP; κ is the duration of the productivity phase-in period; sub-domains for parameters are identified by their medians; the aid-growth effect is calculated over 30 periods, scaled to units of 10 percentage points; z_e refers to raw outcomes from the extended simulations; Δ_e is the absolute difference between the extended and baseline simulation results; Λ measures the relative magnitude of the aid-induced productivity effect vs. the baseline effect, as per equation (8).

Source: authors' calculations based on 10,000 simulations.


Figure 1: Path of aid's contribution to growth and associated net cash flow over a single simulation



Notes: parameters for this example are median parameter values based on all 10,000 draws employed to examine the properties of the model; Aid/GDP is given in units of 10 percent; cash flow is net of the cost of aid.

Source: authors' calculations.

Figure 2: Scatter plots of parameter τ (tau) against the relative magnitude of productivity-induced effects for marginal aid-growth estimates (30-year window) and internal rate of return (IRR)



Notes: relative magnitude (y-axis) is calculated according to equation (8); values greater than zero indicate that, for chosen model parameters, the productivity effect dominates the capital accumulation effect, implying (by symmetry) that the corresponding negative value for τ would consign the outcome variable to the negative domain.

Source: authors' calculations.