

Poverty Traps, Aid, and Growth

Aart Kraay and Claudio Raddatz

The World Bank

First Draft: March 2005

Revised Draft: March 2006

Abstract

This paper examines the empirical relevance of the poverty trap view of underdevelopment. We calibrate simple aggregate growth models in which poverty traps can arise due to either low saving or low technology at low levels of development. We then assess the empirical relevance of these specific mechanisms and their consequences for policy. We find little evidence of the existence of poverty traps based on these two mechanisms. When put to the task of explaining the persistence of low income in African countries, the models require either unreasonable values for key parameters, or else generate counterfactual predictions regarding the relations between key variables. These results call into question arguments in favour of a large scaling-up of aid to the poorest countries that are premised on the existence of such poverty traps.

1818 H Street N.W., Washington, DC 20433, akraay@worldbank.org, craddatz@worldbank.org. This paper was prepared as a background paper for the 2005 Global Monitoring Report. We would like to thank Andy Berg, Carl-Johan Dalgaard, Oded Galor, Carlos Leite, Omer Moav, Zia Qureshi, seminar participants at the World Bank and International Monetary Fund, and two anonymous referees for helpful comments; and Ana Margarida Fernandes for generously sharing her data on Colombian plants. The views expressed in this paper are the authors' only and do not necessarily represent those of the World Bank, its executive directors, or the countries they represent.

1. Introduction

The idea of poverty traps has captured the imagination of development economists and policymakers alike for many years.¹ It is not hard to see why. There are many very plausible self-reinforcing mechanisms through which countries (or individuals) who start out poor are likely to remain poor. There is also an abundance of empirical evidence that most countries that were relatively and absolutely poor in the middle of the 20th century remain so today. Recent calls for across-the-board debt relief and a major scaling up of aid to help poor countries achieve the Millennium Development Goals have been significantly influenced by the idea that these countries are stuck in poverty traps and that major pushes are required to break free of these traps (see for example Sachs et. al. (2004, 2005)).

In this paper we examine the empirical evidence in support of the poverty trap view of underdevelopment. We do this by calibrating specific theoretical macroeconomic models of poverty traps in order to assess both their empirical relevance and their consequences for policy. We focus on two canonical mechanisms: low saving and low levels of productivity at low levels of development. It is not hard to see how either of these can generate a poverty trap. If either saving or productivity is low at low levels of development, investment will be low and countries will converge to an equilibrium with low capital and output per capita. If over some range of income levels saving rates and/or productivity increase sharply, then if countries can get to this point they might also be able to converge to an equilibrium with high capital and output per capita. We focus on these mechanisms because they capture some of the most popular explanations for the existence of poverty traps, and because they are often used as an explicit motivation for foreign aid to help countries escape from poverty traps.

We find little evidence of the existence of poverty traps based on these two broad mechanisms. Both mechanisms can theoretically generate poverty traps, or at least very persistent poverty. However, when put to the task of explaining the persistence of low income in African countries, the models require either unreasonable values for key parameters, or else generate counterfactual predictions regarding the relations between key variables. These results call into question poverty-trap arguments based on these mechanisms that a large scaling-up of aid to the poorest countries will result in sharp and sustained increases in growth.

¹ In an early contribution to the theory of poverty traps, Nelson (1956) writes that even 50 years ago "... the concept of low-level stagnation is scarcely new and different".

In the case of the savings mechanism, generating a poverty trap consistent with the experience of poor African countries with a simple Solow growth model with a variable saving rate requires this rate to be a steep s-shaped function of capital per capita. Savings should be low at low levels of capital per worker, increase significantly at some intermediate levels and level off at high levels. The actual cross-country relationship between saving rates and capital per capita observed in the data is far from meeting these conditions. If anything, saving rates seem to be increasing at low levels of capital per worker, flat at intermediate levels and increasing again at high levels.

The inability of the savings-based poverty trap to account for the experience of African countries is not the result of the exogenous behavior of the saving rate imposed by the Solow model. We also calibrate a model of subsistence consumption with optimizing agents, which provides a standard formalization of the inability of countries to save at low levels of income. We find that for this model to account for a country's persistent low level of income, the country must be very close to the level of subsistence consumption. Given the high dispersion of per capita consumption levels across countries in Sub-Saharan Africa, accounting for slow growth in these countries would require us to rather implausibly argue that subsistence consumption levels are country-specific and vary a lot from one country to the next. If we instead more reasonably assume a common level of subsistence consumption, the model predicts that most African countries should be rapidly converging to their stable equilibrium. The model also predicts that, by rapidly relaxing the constraints that subsistence consumption imposes on capital accumulation, foreign aid should quickly have an important impact on growth and savings. We perform a series of very simple regressions and review the existing literature to demonstrate that this prediction is not borne out by the data.

Technology-based explanations of poverty traps based on aggregate models do not fare better. We consider a broad class of models of optimizing agents facing a technology with increasing returns that are external to the firm. Within these models, we show that the existence of stable poverty trap equilibrium requires the external increasing returns to be sufficiently large relative to the diminishing returns in the production function. We find that such strong increasing returns are inconsistent with existing estimates from the vast empirical literature on the estimation of production functions, as well as some simple estimates of our own.

Based on the previous findings, we conclude that the simple -- but popular -- aggregate models of poverty traps we consider are not promising candidates to account for the experience of African countries. To the extent that these models capture some important features of reality, the

results suggest that a large expansion of aid is not a necessary condition to improve the prospects of these countries. The available data suggest that, at least at the aggregate level there is not enough non-convexity in underlying fundamentals to expect that large-scale increases in aid will be proportionately more effective than more moderate amounts of assistance. This does not mean that increasing aid to Africa (or other poor countries) is a bad idea. Rather, it suggests to us that we should not expect any disproportionate growth effects of large-scale aid.

We do not claim that this paper provides comprehensive coverage of the possible mechanisms of poverty traps, and the simple aggregate models we consider here are unable to capture some possible mechanisms mentioned in the literature. A significant fraction of this literature is based on highly stylized models, and on non-convexities at the micro-level, which are not easily amenable to calibration using aggregate data. Also, we face the standard trade-off of breadth versus depth. We have chosen to analyze a set of broad models that rely on aggregate mechanisms without delving deeply into the possible micro-mechanisms that lie behind our aggregate relations. A possible alternative approach would be to take seriously a specific micro-mechanism of poverty traps and make an effort to calibrate it using micro and macro data. Nevertheless, we consider that there are some important general lessons that can be derived from our analysis. The mechanisms that we analyze are usually mentioned in the discussion on poverty traps, and the models we use are standard in the growth literature. Also, under some conditions these models can be understood as reduced form representations of the richer underlying micro-mechanisms. And of course we recognize that theoretical literature on poverty traps has generated an abundance of potential poverty trap mechanisms other than the ones we study here. Our rather negative results regarding the mechanisms we study here suggest that it is important to subject other potential poverty trap mechanisms to similar empirical scrutiny before using them as a justification for increased aid to developing countries.

The rest of this paper is organized as follows. In the next section we review some of the existing evidence on poverty traps. In Section 3 we discuss saving-based poverty traps, and consider the role of foreign aid in helping countries to escape from such traps. In Section 4 we consider the evidence for technology-based poverty traps. Section 5 concludes.

2. Existing Empirical Evidence on Poverty Traps

In this section we briefly review some of the few existing empirical studies relating to poverty traps. One strand of the literature provides very reduced-form evidence based on cross-country data. If poverty traps are an important feature of growth dynamics, then over time one would expect to observe a bimodal distribution of per capita income, with a group of poor countries clustered around the low-level poverty-trap equilibrium, and a group of rich countries clustered around the high equilibrium (see Azariadis and Stachurski (2004) for details). In a series of papers Quah (1993a, 1993b, 1996, and 1997) has used non-parametric methods to estimate the evolution of the distribution of per capita income across countries. His results suggest the emergence of such a bimodal distribution of income. This finding is supported by a recent paper by Bloom, Canning, and Sevilla (2003) which provides evidence that, after controlling for countries' geographic characteristics, the evolution of income across countries is consistent with a bimodal ergodic distribution. There is some controversy, however, regarding the bimodality of the long-run income distribution: Kremer, Stock, and Onatski (2001) argue that the dynamics of the world income distribution is better characterized by a prolonged transition, during which inequality may increase, towards a single-peaked long-run distribution. A similar result is obtained in Azariadis and Stachurski (2004b), who uses a parametric model of income dynamics derived from an explicit growth model based on Azariadis and Drazen (1990). They also find that the long-run income distribution is unimodal, but that bimodality will appear during the transition. The last two papers are therefore only partly consistent with the idea of poverty traps, in the sense that income disparities may be very persistent and appear as poverty traps, but eventually, all countries would converge to the same equilibrium.²

² Of course, as pointed out by Quah (2001) the distinction whether poor countries are expected to remain poor for decades or centuries is nitpicking, and should not deter us from thinking about ways of accelerating the transition. On a different front, a recent paper by Feyrer (2003) raises some questions about the underlying causes behind the finding of bimodal distributions. The paper suggests that the bimodality of income distribution found by Quah results from a bimodal distribution of TFP, and therefore from productivity differences. Whether the bimodality of TFP results purely from increasing returns or from barriers to technology adoption may have different implications about the validity of the poverty traps story to account for the bimodality of the income distribution.

At a deeper level, the main drawback of this body of evidence is that, for the most part, their empirical analysis of the evolution of income distribution is non-parametric and unrelated to any underlying growth model, in particular, to any poverty trap story. The evidence is therefore at most consistent with a model of poverty traps. It does not provide proof of the validity of these models, nor does it address whether the magnitudes of the underlying mechanisms are empirically plausible. An exception is the paper by Azariadis and Stachurski (2004b) cited above, which assumes the presence of increasing returns external to the firm as the underlying cause of non-convexities. While this model can generate bimodality during the transition, the parameters estimated for the production function are somewhat unreasonable, suggesting incredibly large degrees of increasing returns at some levels of income. As we will discuss in section 4, these levels of increasing returns are not supported by the data.

If poverty traps are important, one would expect also to occasionally see growth spurts as countries manage to escape from such traps. In a recent paper, Hausmann, Pritchett, and Rodrik (2004) empirically examine such “growth accelerations”, which they define as increases in growth of at least two percentage points that are sustained for at least eight years. They find that there are surprisingly many such accelerations, and that while they are in general quite difficult to predict, they are more likely to occur in poorer countries. However, as in the previous case, this evidence is merely consistent with escapes from poverty traps and it is not conclusive. As the authors note, the standard dynamics of convergence in a growth model with shocks would also suggest that large jumps in growth rates are more likely in poor countries than in rich countries. Related reduced-form evidence is provided by Easterly (2005) who documents that even the poorest countries in the world consistently experience positive -- although slow -- growth on average, which is inconsistent with the existence of poverty traps with zero growth. Like Hausmann, Pritchett, and Rodrik (2004) Easterly (2005) also studies growth takeoffs but concludes that these are rare based on the filter that he uses.

Overall, this reduced-form cross-country evidence is at most suggestive of the existence of poverty traps, but it is hardly conclusive. Moreover, this type of evidence tells us little about the specific mechanisms generating the poverty trap or the policy interventions required to escape from that poverty trap. More useful in this respect is a line of microeconomic evidence that focuses on finding direct evidence of the mechanisms underlying specific models of poverty traps. For example, McKenzie and Woodruff (2004) use data from microenterprises in Mexico to search for evidence of poverty traps based on non-convexities in the production function. They take seriously the argument that poverty traps might exist if there are large fixed costs to starting

a business. If capital markets are imperfect and potential entrepreneurs are credit-constrained by their lack of collateral, then a poverty trap can exist since individuals who start out with low wealth are unable to finance potentially profitable investments in new businesses. They find that the fixed costs involved in starting up a small enterprise in Mexico are typically very low, in some sectors less than half the monthly wage of a low-wage Mexican worker. They also find that the returns to capital are very high even at very low levels of the capital stock, and cannot reject the hypothesis that returns are decreasing over the entire range of observed capital stocks. They conclude that their evidence is not consistent with this particular mechanism of poverty traps.

In contrast, a number of papers have found microeconomic evidence of spillovers or other externalities that could form the basis of poverty traps. For example, Jalan and Ravallion (2004) argue for the existence of spatial or geographic poverty traps. In a panel of households from China, they find that consumption growth at the household level increases with the local availability of “geographical capital” understood as the availability of roads, the local level of literacy, etc. Their evidence suggests that these aggregate factors (at the local level) increase the returns to capital faced by households. However, the main focus of the paper is to determine the impact of locally aggregate variables on household growth. They do not spell out a feedback mechanism from household growth to the evolution of the geographical capital that could generate a poverty trap. A direct interpretation of their evidence is that it shows that the factors captured in the geographical capital are a constraint to household growth. If we assume that household growth feeds back into these variables, it might be possible to generate a poverty trap, but the degree of feedback necessary to induce such a trap is not specified. A different type of microeconomic evidence can be found in Lokshin and Ravallion (2004), who empirically estimate non-linear dynamics in household income in two transition economies. While they find that adjustment to income shocks is nonlinear, they find no evidence of non-convexities that would cause temporary adverse shocks to permanently lower household income, as would be the case in a variety of models of poverty traps.

Finally, two recent papers take a calibration approach as we do here. Graham and Temple (2006) use a static two-sector variable-returns-to-scale model and show how, with minimal data requirements, it is possible to calibrate two steady-states for each country in the world. They then use these calibrations to document the contribution of the differences between these equilibria to cross-country income differences. Their model assumes that two different productive sectors co-exist within a country, a traditional sector (agriculture) producing under diminishing returns to all its factors, and a modern sector (non-agriculture) that exhibits increasing returns that are external

to the firm. They find that, depending on the degree of increasing returns assumed for the non-agricultural sector, their model can account from between 15% and 25% of the observed variation in per capita incomes across countries. The elegance and simplicity of the model, which allows it to determine the underlying alternative equilibrium for each country in the world with minor data requirements, comes at two costs. First, in the simplest version of the model the low equilibrium is unstable, so the authors have to resort to some ad-hoc assumptions about the presence of a fixed cost of switching sectors. Second, the model is static and focuses on the comparison of the different possible steady states of a country. So, the model is mute about how the equilibrium is selected for each country, and, most importantly, it does not allow us to analyze the dynamics of a country's income, the impact of policy interventions, or the size or type of policies required to move a country away from its low equilibrium.

The second paper that adopts a calibration approach is Caucutt and Kumar (2004). They calibrate models with multiple equilibria that capture four main mechanisms: coordination failures, inappropriate mix of occupations, insufficient human capital accumulation, and political-economy considerations. They argue that while it is possible to find plausible parameter values that deliver multiple equilibria in these models, the existence of this multiplicity is also quite sensitive to small variations in these parameters. Based on this they caution against policy prescriptions for one-time interventions designed to jump countries from a bad equilibrium to a good one. Although both the methodological approach and the conclusions of this paper are quite similar to ours, our paper differs in two respects. First, we focus on different mechanisms generating poverty traps than they do. Second, and perhaps more important, Caucutt and Kumar (2004) are primarily interested in understanding whether necessary conditions for the *existence* of multiple equilibria are satisfied for plausible parameter values. They do not ask, as we do, whether the magnitudes of the income differences attributable to these multiple equilibria are quantitatively reasonable.

3. Saving-Based Models of Poverty Traps

In this section of the paper we discuss one of the most popular models of poverty traps in which the source of the trap is inadequate saving at low levels of development. In these models, aid which augments domestic saving and finances accumulation can play a role in freeing countries from poverty traps. We first illustrate the basic mechanism and our calibration approach

using a very simple Solow growth model in which saving increases exogenously with the level of development. We next take seriously an important theoretical mechanism why saving rates might increase with income: the influence of subsistence consumption. We present a Ramsey growth model with subsistence consumption and show that, while it does not have a stable low-level equilibrium like the poverty trap in the Solow model, it can exhibit poverty-trap-like features such as persistent slow growth for very long periods of time. Finally, we go to the data and ask whether there is any evidence that (a) saving rates exhibit the sort of nonlinear relationship with income required for the existence of a poverty trap, and (b) the historical relationship between aid, saving, and growth is consistent with escapes from poverty traps.

The Basic Mechanism

We begin by using the Solow growth model to illustrate a saving-based poverty trap and to perform some simple calibration exercises. With a Cobb-Douglas technology the familiar dynamics of the capital stock per worker in the Solow model are given by:

$$(1) \quad \dot{k} = s(k) \cdot A \cdot k^\alpha - (n + \delta) \cdot k$$

where k is the capital stock per capita, A is the exogenous level of technology, α is the output elasticity of capital, and n and δ are population growth and depreciation rates. The key feature of the Solow model is its exogenous and constant saving rate.³ This is an important simplification which we will relax shortly. However, this assumption is very useful because it allows us to simply illustrate the existence of a saving-based poverty trap. In particular we allow the saving rate to be an exogenous function of the capital stock per capita, $s(k)$, and assume that the saving rate is constant at some low rate until a threshold value of the capital stock, \tilde{k} , is reached, and then it jumps to a constant higher rate, i.e.

$$(2) \quad s(k) = \begin{cases} s_L & , \quad k \leq \tilde{k} \\ s_H & , \quad k > \tilde{k} \end{cases}$$

³ At least, this is a key feature of the "textbook" Solow model. Solow (1956) contains a discussion of a version of the model in which the saving rate is assumed to be a declining function of capital per worker.

Figure 1 shows the familiar Solow diagram for this economy. There are two steady states, indicated by k_L and k_H . The lower one can be thought of as a poverty trap, since if a country starts out below k_L , it will grow until it reaches the low steady state, and then remains there forever. On the other hand, if a country starts out between \tilde{k} and k_H , and therefore has a high saving rate, it will grow steadily until it reaches the high steady state. The basic intuition for the poverty trap is simple: at low levels of income, investment is so low that it can only sustain a very small capital stock per capita.

We have calibrated Figure 1 so that it captures what we think are some key features of reality. We have drawn the low equilibrium at a capital stock per capita of \$1500 in 1996 PPP-adjusted terms, which is approximately equal to the population-weighted average capital stock per capita in Sub-Saharan Africa in 2000.⁴ The sum of the depreciation rate and population growth rate was set at 10%, which is standard in the literature. The observed average saving rate in Sub-Saharan Africa is 7.3 percent of GDP.⁵ The calibrated capital-output ratio in the low steady state is given by $s_L / (n + \delta)$, which at 0.73 is just a bit lower than what we observe for Sub-Saharan Africa as a whole. We then choose the level of technology, A , to ensure that the steady state level of the capital stock at the low equilibrium, i.e. $k_L = \left(\frac{s_L \cdot A}{n + \delta} \right)^{1/(1-\alpha)}$ is equal to \$1500. For a benchmark value of $\alpha=0.5$, which is the capital share in value added among low income countries from Purdue University's Global Trade Analysis Project (2005) database

⁴ We construct capital stocks by cumulating flows of investment at PPP, and assuming a depreciation rate of 6 percent per year.

⁵ We measure the saving rate as the investment rate at PPP from the Penn World Table Version 6.1, plus the current account in constant US dollars at market exchange rates divided by real GDP at PPP. This follows the PWT convention of assuming that international transactions occur at international prices. Since the relative price of investment goods tends to be higher in poor countries than in rich countries, saving and investment rates are substantially lower when evaluated at PPP in this way than in local currency terms. The use of PPP-adjusted saving and investment rates is however appropriate for our purposes since we are interested in comparisons of real quantities across countries. This is because a high saving rate evaluated at domestic relative prices will not translate into a high rate of real capital accumulation if the relative price of capital goods is high. We also note that our main empirical result on the observed nonlinearity of the relationship between saving rates and the level of development does not depend on whether we measure saving rates at international relative prices or at domestic relative prices.

(GTAP), this give us a calibrated value of $A=1.68$. The high steady state has been calibrated for a saving rate that is twice as high as we observe, at 15 percent of GDP, but with no other changes in any of the parameters of the model. Since the ratio of the capital stocks in the high and low steady states is $k_H / k_L = (s_H / s_L)^{1/(1-\alpha)}$, our parameterization implies that the calibrated capital stock in the high steady state is a bit more than four times that in the low steady state, at $k_H = \$6,333$.

Although Figure 1 shows that we can generate a growth model with a poverty trap that roughly matches Africa's experience, it is important to note that the low poverty trap equilibrium exists only by assumption. Suppose for example that the point at which saving jumps to the higher rate occurs to the left of what is now shown as the poverty trap equilibrium. In this case, there would be no poverty trap, and a country starting out at the low capital stock would steadily grow until it reaches the high steady state. Similarly, if population growth or depreciation were slightly lower, the low equilibrium can also disappear as the diagonal line rotates downwards, leaving only a single stable high equilibrium. Moreover, if we were to introduce exogenous technical progress into the model, then the poverty trap equilibrium would also disappear, provided that the threshold level of capital per capita \tilde{k} is fixed in absolute terms. If this is the case, thanks to exogenous technical progress all countries will eventually cross this threshold and the poverty trap vanishes as all countries will have the high saving rates.

Suppose despite all this that the poverty trap exists, and that some countries find themselves stuck in it. How can aid help such a country to break free of the poverty trap shown here? In the model, the reason countries are in a poverty trap is because their saving (and hence investment) rates are low. The role for aid in such a model is to augment saving sufficiently to allow the country to accumulate capital and grow to the point where the domestic saving rate jumps to the higher level.⁶ After the saving rate has increased, no further aid is required and the country grows steadily until it reaches the high equilibrium. To take a specific numerical example, suppose that we assume that the country in the low equilibrium receives a fixed fraction of its GDP in aid for a ten year period, and assume further that all of this aid is used to finance

⁶ Of course, the role of aid is very limited in this stylized model -- it simply augments domestic saving. If in contrast low levels of development were due to a very low productivity of investment, then more aid would not accelerate growth.

capital accumulation.⁷ We can then calculate how large the aid inflow must be to ensure that the country just reaches a per capita capital stock of \tilde{k} at the end of ten years. After this point, the country will embark on sustained growth and eventually reach the high steady state without requiring further aid.

The results of this exercise are reported in Table 1. Clearly, the amount of aid required depends on how far the threshold level of capital, \tilde{k} , is from the capital stock in the poverty trap equilibrium. The columns of the table correspond to different values of \tilde{k} , and the rows correspond to different assumptions on the degree of diminishing returns, α . For the benchmark value of $\alpha=0.5$, aid inflows ranging from 3 percent to 19 percent of GDP would be required for a ten-year period in order to break the country out of the poverty trap. Obviously, the higher the threshold level of capital at which saving rates increase, the greater is the aid required. In addition, the stronger are diminishing returns, i.e. the lower is α , the more aid is required.

This very simple example illustrates what we think are some of the main messages of the paper. First, while it is clearly possible for poverty traps to exist, the mere fact that saving increases with income is not sufficient for the existence of a poverty trap – it needs to do so in just the right way. As we will see shortly cross-country evidence is not very compelling that the required nonlinearities in saving rates exist. Second, even if a poverty trap does exist, the size of the policy intervention required to get countries out of poverty traps is very sensitive to the parameters of the specific model of poverty traps, in this case, the precise point at which saving rates jump, as well as the extent of diminishing returns, α .

Subsistence Consumption

An important limitation of the example above is that we simply assumed that saving rates increase at some level of development. In a fully-specified model, saving rates are endogenously determined by optimizing agents. The question therefore is what types of preferences can generate saving functions whose behavior is consistent with the existence of poverty traps. One important explanation offered in the literature is based on the role of subsistence consumption (see Azariadis, 1996; Ben-David, 1998; Sachs, 2004). The idea behind this explanation is that

⁷ A ten-year period can be motivated by the time between when this paper was written and 2015 when the Millenium Development Goals are supposed to be attained.

poor households do not save because they have to use their income to meet basic consumption needs, but once these needs are met households may save quite a lot of their income. We now present and calibrate a model that formalizes this intuition by explicitly introducing a minimum consumption level in the preferences of the representative consumer in an otherwise standard Ramsey model and deriving the optimal consumption and saving behavior under these conditions.⁸ The model is closely related to Ben-David (1998), which analyzes the implications of a model with subsistence consumption under exogenous growth.

The representative household chooses a path for consumption to maximize its inter-temporal utility:

$$(3) \quad U = \int_0^{\infty} u(c_t) e^{-\rho t} dt$$

where c is the consumption per member of the household and the utility function is:

$$(4) \quad u(c) = \frac{(c - \bar{c})^{1-\theta} - 1}{1-\theta}$$

The minimum subsistence consumption level is captured by \bar{c} . The representative household is endowed with a constant returns production technology $F(K, L)$ and with an endowment of labor that grows exogenously at rate n . Capital depreciates at a constant rate δ . Under these conditions, the evolution of the capital per worker is given by

$$(5) \quad \dot{k} = f(k) - c - (n + \delta)k,$$

⁸ This is of course not the only way to generate a saving rate that increases with the level of development. For example, Chakraborty (2004) develops a two-period overlapping generations growth model in which the probability of survival to the second period is endogenously determined as a function of public health care expenditures. The model exhibits an (unstable) low-level equilibrium in which the capital stock, income, and health spending are low. This in turn means that the probability of survival is low, and individuals rationally respond to this by saving less for their old age. Galor (1996) generates an increasing saving rate by assuming different propensities to save out of labour and capital income in a two-period OLG model, combined with a production function in which the factor shares vary with the level of development. Dalgaard and Hansen (2005) provide a capacity-utilization-based foundation for this argument in a Solow model.

where k denotes capital per worker, and $f(k) = F(K/L, 1)$ represents the production function in intensive form. The problem of the representative household is therefore

$$(6) \quad \begin{aligned} \max_c U &= \int_0^{\infty} u(c_t) e^{-\rho t} dt \\ \text{s.t.} & \\ \dot{k} &= f(k) - c - (n + \delta)k \\ c &\geq \bar{c} \end{aligned}$$

which is the standard Ramsey problem, plus the additional restriction that consumption is above subsistence.⁹

The first-order conditions of this problem are given by the following two differential equations:

$$(7) \quad \begin{aligned} \dot{c} &= \frac{(c - \bar{c})}{\theta} (f'(k) - n - \delta - \rho) \\ \dot{k} &= f(k) - c - (n + \delta)k \end{aligned}$$

The equations show that in contrast to the standard Ramsey model, $c = \bar{c}$ is also a stable locus for the consumption dynamics. The implications can be seen in the phase diagram depicted in Figure 2. Compared with the standard Ramsey model, this model exhibits an additional unstable equilibrium at the point (\bar{c}, k_l) , and it also has the property that for capital stocks below k_l the high equilibrium (k^*, c^*) cannot be attained. So, an economy that starts at (\bar{c}, k_l) can remain indefinitely at the subsistence level, and an economy that starts with capital below k_l will maintain its consumption at the subsistence level and keep depleting its capital stock until eventually capital is exhausted and the economy ends up at zero consumption. Of course, even for the cases in which the high equilibrium of the standard Ramsey model is attained, the transitional dynamics of this model will differ substantially from those of the case with no subsistence consumption.

⁹ It is important to notice that the presence of subsistence consumption in the preferences does not invalidate the use of a representative consumer model. Caselli and Ventura (2000) study a model with heterogeneous individuals sharing Stone-Geary preferences and show that despite the rich heterogeneity the model admits a representative agent. This model would unambiguously display convergence in the within country cross-section of consumption so that all households would converge to the same average level of consumption as the economy moves to the steady state

It is important to notice that this model is one example of a class of models where the multiplicity of equilibria results only from the form of the preferences. In general, the first order condition for the evolution of consumption in a Ramsey model corresponds to $\sigma(c)(f'(k) - n - \delta - \rho) = 0$, where $\sigma(c)$ is the elasticity of intertemporal substitution (EIS). It is clear from this equation that with a standard production function and constant values for the rate of population growth, depreciation, and discount factor there is only one possible equilibrium for the level of capital per worker (given by $f'(k) = n + \delta + \rho$) unless the EIS equals zero at some positive level of consumption. The Stone-Geary preferences we use have this property.¹⁰ An important distinction between this model and the Solow model of the previous section is that in this case the low equilibrium is unstable. As a result, countries cannot be “trapped” at this low equilibrium unless by some remarkable coincidence they begin just at the low equilibrium.¹¹ In other words, the model does not deliver a stable poverty trap. The reason is that, although the saving rate in the model is initially increasing with income, it does not have the s-shape (as a function of income) that is necessary to obtain multiple stable equilibria.¹²

Although our model does not deliver a standard poverty trap, its transitional dynamics may still generate something that looks like a poverty trap in the medium run. An economy that starts very close to the unstable equilibrium can exhibit consumption close to subsistence, low saving rates, and low growth for a significant length of time.¹³ In order to take this possibility seriously, we need to calibrate the key parameters of the model. We will assume that the

¹⁰ Bliss (2004) explores the role of a varying EIS for convergence and the possibility of poverty traps in OLG models. However, in the class of preferences he studies the EIS increases linearly with consumption, which means that it is only zero when consumption is also zero.

¹¹ Even this situation may not be feasible if the preferences are such that $\theta > 1$.

¹² The failure of this model to generate a “standard” poverty trap shows how difficult it is to obtain such a trap as a result of optimizing behavior without assuming functional forms that are dangerously close to simply assuming that poverty traps exist. For example, Azariadis (1996) presents a OLG model of subsistence consumption similar to ours that delivers a poverty trap by assuming that the future level of subsistence consumption is a step function of current wealth, hence making richer households effectively more patient. So, the way this model generates a poverty trap is by replacing the assumption of an s-shaped saving function with the assumption of an s-shaped discount factor.

¹³ A similar point was made by Steger (2000).

production function is Cobb-Douglas $f(k) = Ak^\alpha$, and specify values for the parameters $\alpha, \delta, n, \rho,$ and θ . For the benchmark specification we will set the share of capital $\alpha = 0.5$, which is larger than the value of 0.36 typically used for developed countries (e.g. Kehoe and Perri, 2002) and corresponds to the average share of capital in value added in low-income countries as noted earlier. For the robustness analysis we will vary the capital share between the value of 0.4 considered by Coe, Helpman, and Hoffmaister (1995), and 0.6.¹⁴ We set the sum of the depreciation rate δ and population growth n at 0.1, and the rate of intertemporal preference ρ at 0.05 which is consistent with the standard quarterly value of 0.99 for the discount factor used in the real business cycle literature (see for example Kehoe and Perri, 2002). Finally, we set the coefficient of relative risk-aversion (or the inverse of the elasticity of intertemporal substitution) θ at 2, which corresponds to the intermediate range of values obtained by Eichenbaum et al. (1998) for developing countries and is consistent with the EIS of 0.5 reported also for low income countries by Agenor (1999).¹⁵ The parameter A was chosen to calibrate the unstable low equilibrium to the level of consumption per capita and capital per-worker of the poorest African country in 1996 (Democratic Republic of Congo with a consumption per capita of about \$300 and capital per worker of \$330). This seems like a reasonable lower bound for a subsistence level of consumption, and it is slightly smaller than the \$300 in 1980 dollars computed by Ben-David (1998) as the least cost requirement for sustaining an individual's minimum dietary needs.

The bottom panel of Figure 2 shows the $\dot{c} = 0$ and the $\dot{k} = 0$ locus in the (c, k) space for the benchmark economy. It also shows the stable manifold of the economy, which was computed numerically using the time-elimination method of Mulligan and Sala-i-Martin (1991). In this parameterization, the levels of consumption and capital per person at the high stable equilibrium correspond to \$747 and \$3,434 respectively. These levels are very low. So, a first quantitative implication of the calibration exercise is that the high equilibrium consistent with the poorest African country being at its subsistence level implies that, even after reaching the high

¹⁴ We cannot increase the value of α further and simultaneously being able to calibrate the unstable equilibrium of the economy with the values observed for low income countries keeping the rest of the parameters fixed.

¹⁵ Standard values for θ range between 1 and 3. As it will be pointed out below, the results are not significantly affected by varying θ within this range.

equilibrium, per capita consumption will still be below that of the median African country, and well below that of middle income countries. Of course these magnitudes are sensitive to the specific parameterization, but, as will be shown below, taking them to more reasonable levels would require unlikely magnitudes for key parameters of the economy.

The effect of subsistence consumption on the pattern of growth for an economy that starts at an initial level of capital 1 percent above the low equilibrium is shown in Figure 3. For comparison, the figure also presents the evolution of capital per worker of an economy that starts with a much higher level of capital equal to the one observed for the median African economy, and of a standard Ramsey economy with the same parameters. The top-left panel shows that the convergence of the capital stock to the steady state is significantly slower for the benchmark economy. The remaining panels show that the persistence of low levels of capital per worker for the benchmark economy also translates in a low and persistent level of consumption, growth, and saving rates. During the first 20 years, average consumption and capital per worker are \$302 and \$339, which are barely above the subsistence levels, the saving rate averages 11 percent, and the average growth rate is 0.5 percent. These are magnitudes that roughly fit the experience of poor countries.¹⁶

The previous figures suggest that persistent low output and slow capital accumulation observed in African countries could be explained by assuming that these countries are close to their subsistence level of consumption. While these economies would not be “trapped” in the usual sense because they would eventually reach the high equilibrium, they would nevertheless display slow growth performance for an extended period of time. Over time, however, the growth rate accelerates as subsistence consumption becomes less important, and finally the growth rate declines as the economy approaches the stable high equilibrium.

However, it is important to note that this poverty trap-like behavior is very sensitive to initial conditions. In particular, the length of time that low growth persists depends crucially on how close the initial level of capital is to the level corresponding to the low unstable equilibrium k_l . Table 2 illustrates this point by reporting the time required to close half of the distance to the

¹⁶ Under this parameterization, the steady-state saving rate is about 33%. Although this magnitude is high, it can be easily reduced by increasing the discount rate ρ without significantly affecting the properties of the solution. For example, in order to induce a steady-state saving rate of 25% we need to double the discount rate to 0.1, which correspond to an underlying annual discount factor of about 0.9.

steady state level of capital, assuming that the level of capital observed in low income countries is ten, one, one-tenth, and one-hundredth percent above the low unstable equilibrium. For the benchmark economy halving the distance to the steady state would take 54 years, as compared with just 12 years for the no-subsistence-consumption economy. However, a country that starts out 10 percent above the capital stock in the low equilibrium would take 41 years, and the median (un-weighted) African economy with a capital stock of \$934 per capita would take only 17 years. In fact, it is apparent from the figures that, under the benchmark parameterization, the median African country should behave in a similar way to the standard Ramsey economy. In other words, the existence of subsistence consumption should be practically irrelevant for the dynamics of the median African country, and even more so for countries richer than the median.

Table 2 also shows that the main quantitative implications of the model for growth near subsistence levels of consumption are robust to the capital share α . We consider a range of values for this parameter, and in each case re-calibrate the constant A to make the low unstable equilibrium correspond to the values of consumption and capital observed for the poorest African country.¹⁷ Although the level of the high equilibrium varies significantly, the pattern of convergence looks similar in all three cases. In particular the times required to close half the gap to the steady state are 59 years ($\alpha = 0.4$), 54 years ($\alpha = 0.5$), and 56 years ($\alpha = 0.6$). Comparing these differences with those obtained for economies that start at different distances from the low unstable equilibrium makes clear that the main determinant of the speed of convergence is the position of the initial conditions with respect to k_l .

In summary, the results from these simulations suggest that the model with subsistence consumption has the potential to generate something very similar to a poverty trap if we are willing to assume that poor African countries are very close to subsistence. However, the flipside of this assumption is that the median African country (and those above the median) should have a very high saving rate and be close to its high long run equilibrium. One possible way of reconciling the predictions of the model with the data would be to assume that the subsistence consumption levels are country specific, so that most low income African countries are close to their own subsistence levels, but this would require us to argue that the subsistence consumption level can vary significantly across countries. The consumption of the poorest African country is

¹⁷ Notice that because of the calibration of the constant A the effect of α on the steady state is in principle ambiguous. In practice, however, the standard increasing effect dominates.

one-third of the consumption of the median African country (PPP adjusted), and one-ninth of the consumption of the richest African country. Analogously, assuming a common but higher level of subsistence consumption, for example the consumption of the median African country, implies that half of the countries below that level would be drifting towards the zero consumption equilibrium, which is equally unappealing and even unfeasible under the benchmark value of θ . These quantitative implications of the model suggest that the subsistence consumption story as a source of poverty traps has serious problems matching the data. Nevertheless, in what follows we will maintain the assumption that the persistence of low levels of capital and consumption in poor countries results from the fact that they are close to the subsistence level, and ask how these countries can accelerate their escape from this situation.

What is the role of aid in this economy? As in the previous example, aid that augments saving and domestic capital accumulation can bring the country toward the high steady state. Unlike in the previous example, however, in the absence of aid countries would eventually get to the high equilibrium on their own: all aid can do is accelerate this process. To get a sense of the magnitude of the interventions involved, we compute the annual aid transfer as a fraction of GDP necessary to double and quadruple the initial level of capital per person in 10 years.¹⁸ Since saving rates are determined endogenously, we also need to make assumptions on how aid is delivered. We consider two cases: (i) all transfers are exogenously directed to capital accumulation; (ii) the allocation of transfers is left to the recipient country, but they are perceived as permanent. These two cases give us a lower and upper bound for the more realistic case in which the transfers have some bias towards capital accumulation and are at least partially perceived as transitory.

The values obtained for the transfers are reported in Table 3 for the benchmark economy, the median African country, the Ramsey case, and for alternative values for the capital share ($a = 0.4, 0.6$). For the benchmark economy, the annual transfer, as a fraction of GDP, required to double the capital per worker in a 10 year period is about 4 percent in the case in which all the transfer goes to capital accumulation and 9 percent if the transfer goes to income but is perceived

¹⁸ One can motivate this particular experience again with reference to the Millenium Development goals, the first of which is to halve the headcount measure of poverty. For a very poor country with per capita income of \$300 and a Gini coefficient of 40, halving the headcount measure of poverty with a 1\$ per day poverty line requires a doubling of income (assuming an unchanged lognormal distribution of income). This in turn requires a quadrupling of the capital stock for our benchmark case of $\alpha=0.5$.

as permanent (see columns (4) and (5)). As mentioned above, these two values represent a lower and upper bound for the actual transfers required assuming different degrees of anticipation of the end of the policy and biases towards capital accumulation. For this economy, doubling the initial capital stock would have taken almost 30 years if left to itself (see column (6)). Doubling the capital stock in 10 years therefore takes approximately 20 years from the development process of the baseline economy. Quadrupling the capital stock however requires a much larger effort, with transfers somewhere between 13 and 33 percent of GDP for 10 years.

Figure 4 shows the effects of aid on saving and growth in the first scenario of doubling the per capita capital stock. Remember that, in each graph, the two exercises represent a lower and upper bound for the true effects that should be observed. Given the apparent similarity of the evolution of saving rates and growth rates in both scenarios the figure gives us a pretty good idea of the likely evolution of these variables. In both cases we observe a large initial increase in saving and growth rates and a considerably faster convergence to steady state than in the baseline case.

There is a striking difference in the magnitude of required aid inflows in this example compared with the previous example based on the Solow model. To see this, note that in Table 1 the case of attaining a threshold capital stock of \$3000 in ten years corresponds to a doubling of the initial capital stock. In the Solow economy doubling the capital stock required 7.7 percent of GDP in aid for 10 years for the benchmark value of $\alpha=0.5$. In contrast, in Table 3 we see that doubling the capital stock in the Ramsey economy with subsistence consumption requires only 3.9 percent of GDP in aid for the same capital share when aid goes directly to investment. The reason for this difference is that in the Solow economy, the domestic saving rate by assumption remains constant and does not respond to the aid inflow. In contrast, in the subsistence economy model there is a very strong increase in saving. In the absence of aid saving rates would have averaged 10% of GDP in the first 10 years, while with aid saving rates increase sharply and average 16% of GDP during this period. This additional aid-induced domestic saving and investment has a strong effect on growth over and above the direct effects of aid.

In summary, the model with subsistence consumption could be used to rationalize the persistence of low levels of consumption and capital observed in low income African countries by assuming that they are close to the subsistence level, but this requires either to assume that there is significant variation in the levels of subsistence consumption across African countries or that half of these countries should exhibit high saving and growth rates. Also, taking the model seriously would imply that relatively modest transfers, well within the magnitudes we observe

historically should have an important effect on the performance of these countries. We next ask whether this has been the case empirically.

Empirical Assessment

In this sub-section we assess the empirical relevance of the saving based poverty trap models discussed above. We first ask whether the cross-sectional and time-series behavior of savings are consistent with models of poverty traps. We then ask whether there is any evidence that in the past aid has helped countries to escape from saving-based poverty traps.

Consider first the nonlinear relationship between saving and the level of development that is required for a poverty trap to exist in the Solow model. In the left-hand panels of Figure 5 we plot the empirical counterpart of this relationship. In each panel we have saving rates averaged over all available years between 1970 and 2000 on the vertical axis, and on the horizontal axis we have the capital stock per capita in 2000. The top panel plots data for all countries, the second only for low-income countries, and the third only for countries in Sub-Saharan Africa. In all these graphs we have fitted a third-order polynomial relationship to pick up any non-linearities that might be present in the data. These non-linearities do appear to be present in the cross-sectional relationship between investment rates and the level of development. In the top panel we find sharp increases in saving at very low capital stocks, then a fairly flat section, and then finally again an increasing section.¹⁹ These nonlinearities are statistically significant in the sense that the quadratic and cubic terms enter significantly. In the middle and bottom panel we find that saving rates increase at very low levels of the capital stock per capita and then flatten out and even decline slightly around \$5000. In these cases we do not find significant evidence of nonlinearities, although this may be due to the smaller sample size and smaller range of variation in capital stocks per person in these sub-samples.

Despite the presence and at times significance of these nonlinearities, the patterns shown in the left side of Figure 5 are not very promising for the poverty trap story. Recall that for the poverty trap story to work, we need to find low saving rates at low levels of development, and then a sharp increase over some intermediate range before leveling out at a high rate. This is just the opposite of what we see in the data, where saving rates increase sharply at very low income

¹⁹ We have reproduced the results in this figure using saving rates evaluated at domestic, rather than international prices. The results are qualitatively similar, in the sense that saving rates tend to increase sharply at low levels of income, then flatten out over an intermediate range, and finally increase again.

levels and then flatten out. The consequences of these saving patterns for the Solow model are shown in the right panel of Figure 5. In the top and middle panels the fact that saving rates are increasing and concave over the relevant range implies a unique stable equilibrium around a capital stock of \$1500 per capita. For the Sub-Saharan African sample there is an unstable equilibrium around \$2000 and a stable equilibrium around \$5000. In none of these cases do we find a sufficiently S-shaped pattern in saving rates to generate a stable low-level equilibrium that might be thought of as a poverty trap, together with a stable high equilibrium that countries might be able to attain with appropriately large amounts of aid. In what follows we focus on stable equilibria because we find it hard to interpret unstable equilibria as the basis for persistent differences in incomes.²⁰

The bottom line here is that the cross-sectional patterns in saving rates do not appear to be particularly supportive of the subsistence consumption-based poverty trap mechanism we have examined. This is not because saving rates do not increase with the level of development. Rather the data suggest that they do not increase in the right nonlinear fashion required for the existence of a poverty trap. Moreover, this empirical regularity suggests that other poverty trap mechanisms that generate multiple equilibria through a nonlinear response of saving may not fare well when confronted with the data.

We now turn to the empirical predictions of the Ramsey economy with subsistence consumption. A key prediction of this model, which can readily be seen from Figure 3, is that countries starting out near subsistence levels will eventually experience sharp increases in saving rates and sharp accelerations in growth as they reach a level of development where subsistence considerations are no longer important. For the baseline calibrations a country starting off near subsistence will see its saving rates roughly double from 10 percent to 20 percent during the first 30 years, and growth will accelerate from near zero to 5 percent per year.²¹ A first question to ask of the data is whether jumps in saving rates of this magnitude are ever observed. In the top panel

²⁰ At the same time, we recognize that it is often possible to transform unstable equilibria into locally stable ones with suitable additional assumptions about the dynamics of adjustment. However, we think it is useful to fully specify such extensions and evaluate them empirically on a case by case basis -- something which is beyond the scope of the exercise we provide here.

²¹ The jump in saving would be large even if we had a higher discount rate and therefore a lower steady-state saving rate. For example, in the case in which we double the discount rate, which induces a steady-state saving rate of 25%, savings would still increase from 10% to 17% in 10 years.

of Figure 6 we graph the change in decadal average saving rates between the 1970s and the 1990s on the vertical axis, and the initial log-level of per capita GDP in 1970 on the horizontal axis.

The first thing to notice from this graph is that such large increases in saving rates, even over a 30-year period, are very unusual. Only six countries (Lesotho, Gambia, Mauritania, Chile, Singapore, and Switzerland) have seen increasing in saving rates of more than 10 percent of GDP over this 30-year period. This suggests that accelerations in saving of the magnitude suggested by subsistence consumption models are rare. Second, there is no obvious pattern in the data suggesting that accelerations in saving are more likely at low income levels. There is only a small and insignificant negative correlation between changes in saving and initial income. Moreover the large increases in saving above 10 percent of GDP are fairly evenly distributed across the entire income distribution in 1970, from Lesotho at \$900 to Switzerland at \$20,000. In contrast, the benchmark calibration of the model implies that large increases in saving rates occur in a range much closer to subsistence levels. In fact, according to the calibration, the saving rate reaches 15% at about \$424 of capital per capita (which according to the other parameters implies a level of income per capita of about \$350).

In the bottom panel of Figure 6 we ask whether large increases in saving tend to be accompanied by growth accelerations, as the model with poverty traps would suggest. Again consider the change in decadal average saving rates in the 1990s relative to the 1970s, and we look at the change in average growth rates over the same period. There is a small and statistically insignificant relationship between changes in investment rates and changes in growth rates, contrary to the implication of the theory that saving and growth accelerations should go together.

Finally, we have seen that even fairly modest aid inflows can have very strong effects on growth in the subsistence economy model. This reflected both the direct effects of aid on capital accumulation, as well as the indirect effect of sharp increases in saving as aid makes the subsistence constraint less binding. In particular, the model predicts that aid should increase investment substantially, and that this in turn leads to a sharp increase in growth. Depending on the specific assumptions on the use of aid flows (whether they go completely to investment or to general expenditure) the increase in investment can be even more than one for one. At the very least, increases in investment should be roughly of the same magnitude than the aid flows (as a percentage of GDP). In Figure 7 we explore the links between aid and investment. In the top panel we plot the relationship between aid as a fraction of GDP and the investment rate, both averaged over the 1990s, for all low-income countries. The first thing to note is that there have been cases of quite poor countries that have received very substantial aid inflows over 10-year

periods, for example Sao Tome and Principe at 30% of GDP, Guinea-Bissau at 17 percent of GDP, or Equatorial Guinea at 10 percent of GDP during the 1990s. Note also that these aid ratios have GDP adjusted for purchasing power parity in the denominators. Had we used market exchange rates in the denominator the aid ratios would be much higher.²² In short, there have been quite a few cases of poor countries receiving aid flows on the order of magnitude that our calibrations suggest would be required to get out of poverty traps.

Unfortunately, however, the poverty trap models also imply very large increases in investment should occur with aid. Yet the simple correlation between aid ratios and investment rates is only very weakly positive, and the marginal significance of this relationship is entirely due to the inclusion of two extreme observations, Sao Tome and Principe and Guinea-Bissau. Excluding these countries the slope of the relationship between investment and aid is zero. In the bottom panel we show the same relationship, but now in changes for the decade of the 1990s relative to the 1970s. Again the correlation is not significantly different from zero. A particularly extreme example is the case of Sao Tome and Principe. Our data suggest that aid increased by 20% of GDP between the 1980s and the 1990s, yet investment rates increased by only 8%.

Another interesting piece of evidence comes from taking seriously the calibrations of the effect of aid in the Ramsey model shown in Table 3. There we saw that for countries near the subsistence level of consumption, aid flows between 4 and 9 percent of GDP for a 10-year period would be enough to double per capita capital stocks. In our data there are 25 countries with aid as a share of GDP greater than 10 percent during the 1990s. For this group of countries the median increase in the per capita capital stock during the 1990s was just 20 percent, and the highest increase was 50 percent. The fact that these increases in per capita capital stocks were so low in countries with very large aid inflows suggests to us that a subsistence-related poverty trap is probably not very empirically relevant.

²² Which of the two is more appropriate? This depends on whether aid dollars are used to buy goods and services in rich countries (for example, advanced equipment, or consultant services) or in the poor aid recipient (for example on locally-produced building materials). We are not aware of estimates of the division of aid along these lines. Here we report the PPP aid ratios simply to be conservative about the magnitude of aid ratios that we actually observe. Aid ratios expressed as a share of GDP at market exchange rates are much higher, which would suggest that even more countries have already received enough aid to break them out of the kind of poverty traps we study. We have also verified that our results on the correlations between aid and investment hold for both measures of aid.

None of these simple stylized facts should be especially surprising. Rodrik (2000) provides a careful analysis of episodes of “saving transitions” which he defines as increases of five percent or more in saving rates. He finds only 20 such episodes during the period 1960-1995. This suggests that escapes from poverty traps must be quite rare since their hallmark is a sharp increase in saving rates. It is also interesting to note that the saving transitions he identifies occur at a very wide range of per capita income levels, ranging from around \$800 in China to over \$6000 in Mauritius. It seems difficult to interpret the saving transitions occurring at high income levels as resulting from the escape from a poverty trap based on subsistence consumption simply because these countries are so rich that subsistence considerations matter little. Finally, Rodrik also documents that saving transitions do not tend to be accompanied by higher growth. Although these sharp increases in saving tend to persist for long periods of time, growth accelerates only for a relatively short period and quickly returns to its pre-transition level. This again seems inconsistent with the poverty trap story.

The other observation on the weak relationships between aid, investment and growth has been amply documented in a series of recent papers. Easterly (1999) finds that in only six of 88 countries for which data are available does aid raise investment at least one by one, and in only 17 countries is the effect positive and statistically significant. He also finds virtually no evidence that investment and growth are correlated within countries. For Africa in particular, Easterly and Dollar (1999) find that in none of 34 African countries aid increases savings at least one by one. Similarly, Devarajan, Easterly and Pack (2001, 2003) document that the effect of public investment on growth is not significantly different from zero, and that the effect of private investment on growth is significant only if Botswana is included in the sample. Based on this as well as a variety of micro evidence they conclude that the productivity of investment in Africa is extremely low.

In summary, while models of saving based poverty traps are quite intuitive and are frequently invoked as a motivation for aid, we fail to find any clear evidence in support of them. We have seen that there is little evidence in support of the nonlinearities in saving rates across countries required to generate poverty traps. We have also seen that there are very few instances of saving and growth accelerations that would be consistent with calibrated models of poverty traps. And finally, as many others have noted, there appears to be very little evidence that aid has an at least proportional effect on investment as these models would suggest. This is the case despite the fact that there have been quite a few episodes of countries receiving very large aid inflows that our calibrations suggest would be sufficient to break these countries out of poverty

traps. Overall this evidence leads us to doubt the empirical relevance of saving-based models of poverty traps.

4. Technology-Based Models of Poverty Traps

We now turn to another popular argument for the existence of poverty traps: the idea that there are sharp increases in productivity once a certain threshold level of development is attained.²³ This idea can be generated in a variety of models in which indivisibilities in production, externalities across firms, public infrastructure, and many other factors play a role. In this section we will not restrict ourselves to a particular mechanism generating the acceleration in productivity. Rather we will review some fairly general requirements for the existence of a technology-based poverty trap, which we interpret as the existence of a stable low-level equilibrium, and then we will show that these requirements do not appear to hold in cross-country and firm-level data. We interpret here “technology” in a broad sense, and our argument also applies to other mechanisms linking total factor productivity with the level of development.

Theoretical mechanisms

To be consistent with perfect competition, theoretical models usually assume that increasing returns are external to the firm. As surveyed by Caballero and Lyons (1990), standard explanations for these economies of scale that are external to the firm but internal to the industry or country include advantages of within-industry specialization, agglomeration, indivisibilities and public intermediate inputs such as roads or other infrastructure. Some papers along these lines include Bryant (1983), Weil (1989) and Durlauf (1991) who introduce some form of externality in the production process, and Diamond (1982), Howitt (1985), and Howitt and McAfee (1988) who analyze models of search and matching with thick market externalities. Most of these models, however, are highly stylized and not amenable to calibration. One exception is the model of production externalities calibrated by Graham and Temple (2004) discussed in section 2. Nevertheless, as mentioned above all these models share the common feature that the increasing returns are assumed to be external to the firm, so that, in a model of homogeneous firms, the technology of a representative firm j has the form:

²³ Formalizations of this idea go back at least as far as Solow (1956) who illustrates the case of technologically-induced stable low and high equilibria (Figure II in that paper).

$$(8) \quad Y_j = A(E)F(K_j, L_j)$$

where Y_j , K_j , and L_j are the total output, capital, and labor of firm j , and $A(E)$, the scale factor capturing total factor productivity, depends on a measure E of the average or aggregate level of activity. This dependence is taken as given by the atomistic firm. As usual, the function $F()$ is assumed to have constant returns to scale.

The crucial aspect of the specification above is the dependence of the scale factor on a measure of aggregate activity. Different models take different stances on what is the source of the external economies, which will determine what is the relevant variable that affects A . We will consider three possibilities. In the first specification, the scale factor depends on the level of capital per worker. This approach was introduced by Azariadis and Drazen (1990) in their model of threshold externalities, and, as noted by Barro and Sala-i-Martin (1991), it has the advantage of eliminating the scale effects that appear in models like the ones presented by Romer (1986) and Barro (1990). In the second specification, we assume that the scale factor depends on the aggregate level of output. This approach is similar to the one taken by Caballero and Lyons (1990, 1992, 1994), and Cooper and Haltiwanger (1996). Finally, we also consider the case in which the scale parameter depends on the aggregate level of capital, in the spirit of the literature on capital-embodied knowledge started by Romer (1986). For each of these specifications, we explore the restrictions that the existence of a stable low equilibrium imposes on the technological parameters and then we check whether these restrictions are actually borne out by the data.

The simplest specification for the case in which the scale factor depends on the level of capital per worker is the one considered by Azariadis and Drazen (1990), where there is a threshold level of capital at which there is a “technological jump”. That is

$$(9) \quad A(k) = \begin{cases} \underline{A} & \text{if } k \leq \tilde{k} \\ \bar{A} > \underline{A} & \text{otherwise} \end{cases}$$

The ability of a technology like this to account for the differences in income observed in low income countries will of course depend on the threshold capital level \tilde{k} . In a Solow framework, the analysis will be very similar to the one presented in Section 2. This is because, in the Solow model, the role of jumps in saving and jumps in technology are more or less interchangeable. In particular, all of the discussion in the first part of Section 2 can be re-interpreted as a story in which there is an exogenous improvement in productivity once capital per worker reaches a certain level. In this framework a poverty trap featuring low productivity can exist since for a

given saving rate, low productivity means that investment is insufficient to generate sustained growth in the capital stock.

As we shall see shortly in the data, it seems difficult to find empirical evidence for such a discrete jump in the level of technology. However, it is possible to have a poverty trap with a less dramatic nonlinearity in technology. In fact, we can use the Ramsey economy of the previous section to ask what degree of increasing returns would be necessary over some range of the production function to obtain a stable low equilibrium corresponding to a poverty trap. Assume that TFP (A) is a function of aggregate capital stock per capita ($A(k)$) but that firms do not internalize the presence of the external effect in their production decisions. Under these assumptions, the two equations that govern the dynamics of the Ramsey economy are:

$$(10) \quad \begin{aligned} \frac{c}{c} &= \frac{1}{\theta} (A(k)f'(k) - n - \delta - \rho) \\ k &= A(k)f(k) - c - (n + \delta)k \end{aligned}$$

As in the standard Ramsey case, the equilibrium levels of capital are pinned down by the dynamics of consumption. In a steady state with no consumption growth the capital level has to satisfy the condition $A(k)f'(k) = n + \delta + \rho$. As $A(k)$ and $f'(k)$ are an increasing and decreasing function of k respectively it is easy to see that it is possible for this equation to have multiple solutions, depending on the specific form of the $A(k)$ function. Assuming, as is reasonable, that $\lim_{k \rightarrow 0} A(k) < \infty$ and $\lim_{k \rightarrow \infty} A(k) < \infty$, the only way of inducing multiple equilibria is for $A(k)$ to be significantly steeper than $f'(k)$ in some range. The situation is depicted in Figure 8. The exact condition is:

$$(10) \quad \frac{A'(k)k}{A(k)} > -\frac{f''(k)k}{f'(k)},$$

that is, there must be a range where the elasticity of $A(k)$ is higher than the elasticity of $f'(k)$. Assuming that $f(k) = k^\alpha$, the condition translates into

$$(11) \quad \gamma(k) = \frac{A'(k)k}{A(k)} > (1 - \alpha).$$

The intuition behind this condition is that the stronger the diminishing returns of the production function $f(\cdot)$, the larger has to be the scale effect on TFP to offset the tendency of the aggregate returns to decrease. Assuming that α takes values somewhere between 0.4 and 0.6, the

condition requires values for the elasticity of productivity with respect to the capital stock per worker to be between 0.4 and 0.6.

The above discussion is not significantly affected by assuming that the scale factor depends on aggregate output ($A=A(Y)$) or on aggregate capital ($A=A(K)$) instead of depending on capital per worker. A little bit of algebra shows that the existence of a stable low equilibrium—of the kind consistent with a poverty trap story – imposes similar constraints on the elasticity of the scale factor with respect to its corresponding argument. A necessary condition to have a low equilibrium is that the implied elasticity of TFP with respect to the relevant variable (aggregate output or aggregate capital) has to be strictly larger than $1 - \alpha$ over some range.

Empirical Assessment

A natural place to begin looking for evidence of these scale effects on productivity is in the large empirical literature on estimating returns to scale. In this literature, the degree of returns to scale observed in an industry (firm) is obtained by estimating a relationship between the growth rate of value added (or output) and the growth rates of the different inputs:

$$(12) \quad \hat{Y} = \beta_0 + \beta_1 \hat{K} + \beta_2 \hat{L} + \varepsilon$$

where the hat denotes the percentage change of a variable. The sum of the coefficients β_1 and β_2 denotes the degree of returns to scale of the production function, and this sum is usually the value reported for the papers that empirically estimate the degree of returns to scale.²⁴

Under our assumed functional form for the production function, the growth rate of aggregate output is given by

$$(13) \quad \hat{Y} = \alpha \hat{K} + (1 - \alpha) \hat{L} + \hat{A}$$

where α is the share of capital on total factor costs. For the cases in which A depends on aggregate output and aggregate capital, the growth rate of the scale factor can be expressed as a function of its elasticity γ and the growth rate of the underlying variable, so that the expression for the growth rate of output corresponds to

²⁴ This estimation is not as simple as it seems. There are many technical complications in the estimation of the parameters because of the endogeneity of the input choice, and also because of the different implications of the aggregation process in the presence of increasing returns and monopoly power. The interested reader is referred to the discussion in Basu and Fernald (1995).

$$(14) \quad \hat{Y} = \frac{\alpha}{1-\gamma} \hat{K} + \frac{1-\alpha}{1-\gamma} \hat{L}$$

in the case in which the scale factor is a function of aggregate output, and to

$$(15) \quad \hat{Y} = \alpha \hat{K} + (1-\alpha) \hat{L} + \gamma \hat{K}$$

in the case in which the scale factor is a function of aggregate capital.

These two expressions allow us to compare the magnitude of the returns to scale estimated in the literature with those required for the existence of technology-based poverty traps. Recall that for a poverty trap to exist we require $\gamma > 1-\alpha$ over some range. For the case in which the scale factor depends on total output, equation (14) tells us that the sum of the estimated coefficients on capital and labor equals $(1-\gamma)^{-1}$, which is increasing in γ . Assuming that the capital share ranges from $\alpha=0.4$ to $\alpha=0.6$, this implies that estimates of returns to scale must range from 1.7 to 2.5 in order for a poverty trap to exist. When the scale factor depends on total capital, Equation (15) tells us that the sum of the estimated coefficients on capital and labor is equal to $1+\gamma$. This in turn means that estimated returns to scale must range from 1.4 to 1.6 over some range of the production function.

These required levels of returns to scale ranging from 1.4 to 2.5 are much higher than what is typically found in the literature, which are summarized in Table 4. The table shows that the only study that approaches the lower end of this range is Levinsohn and Petrin (1999), who look at data from Chilean plants. However, most other available studies only find constant to moderate increasing returns. In particular, the recent paper by Van Biesebroeck (2005) that focuses on Sub-Saharan African countries situates the degree of returns to scale between 1 and 1.1.

The absence of empirical evidence of strong increasing returns casts doubt on these technology-based explanations for poverty traps. However, it is important to note that technology-based poverty traps could exist even in the absence of increasing returns in the aggregate production function. In the first example we discussed, we assumed that total factor productivity was a function of capital per capita. In this case, the aggregate production function has constant returns to scale, but a poverty trap could nevertheless exist. Moreover, in this case existing evidence on the degree of returns to scale is not helpful in assessing the empirical relevance of this particular mechanism. In the remainder of this section, we directly examine the

relationship between estimates of total factor productivity and capital per worker using firm-level and cross-country data.

The three panels of Figure 9 show the relation between (log) TFP and (log) capital per worker, total capital, and output for a cross-section of countries in 2000. Total factor productivity is obtained as residuals from a Cobb-Douglas production function in capital and labor, with a capital share of 0.5.²⁵ Each panel also shows the estimated relation between the variables assuming both a simple linear specification and a cubic polynomial to allow for the nonlinearities required to generate multiple equilibria. In the first panel we observe a linear relation between (log) TFP and (log) capital per worker. It is apparent that the polynomial specification does not offer a significant improvement on the fit with respect to the linear specification. Overall, the estimated relationship suggests a constant elasticity γ of about 0.2. Both the finding of a constant elasticity and its magnitude are inconsistent with the requirements of multiple equilibria. A constant elasticity γ implies that the function $Af'(k)$ is either always decreasing or always increasing, which implies a unique equilibrium (or none). Also, as determined above, the lower bound for γ (obtained by assuming a very large capital share) is 0.4, which is almost twice the estimated value. The relation of TFP with total capital, shown in the second panel, does not fare better. In fact, the figure shows practically no relation between the two variables. Even if we take the econometric results at face value, they also suggest a constant and small elasticity (about 0.06). As explained above, a constant elasticity is inconsistent with multiple equilibria, and the implied value of the elasticity is way below the magnitude required for multiplicity. Finally, the bottom panel shows the relation of TFP with total output. It is apparent that the situation is almost identical to that obtained for total capital.

The panels of Figure 10 repeat the exercise performed in Figure 9, but this time using panel data from Colombian plants.²⁶ The graphs plot the relations between the log of TFP at the plant level and the log of the different measures of activity at the industry level after partialling-

²⁵ Results are similar using other values for the capital share. Note also that we make no effort to adjust for the quality of labour input, as is often done in development accounting exercises. It could be that the nonlinearity in productivity comes precisely from sharp improvements in worker skills at some level of development, and so we would like our estimates of productivity to capture this.

²⁶ The data on Colombian plants comes from Fernandes (2003). Industry level data was obtained from UNIDO (2002).

out plant, industry, and year fixed effects. The conclusion that emerges from the three figures is similar, and suggests that the scale factor is independent of these aggregate variables. The linear specifications show a clearly flat relation between the log of firm level TFP and the aggregate variables, and the non-linear specifications are undistinguishable from the linear ones. Again, although these results should be taken with caution because of the caveats mentioned above, the message we want to emphasize is that the scant evidence available is far from supporting any increasing or non-linear increasing returns external to the firm.

5. Concluding Remarks

There is by now an abundance of theoretical models of poverty traps capturing a wide variety of self-reinforcing mechanisms through which poor countries might remain poor. What is lacking is a similar body of empirical evidence in support of such models. This is particularly unfortunate given the emphasis on poverty traps in recent discussions of aid policy. In this paper we have tried to take seriously the quantitative implications of two leading mechanisms generating poverty traps in aggregate growth models: low saving and low productivity at low levels of development. We do not find much evidence in support of the idea that these two mechanisms are empirically relevant. While saving rates and productivity do clearly increase with income levels, we find that they do not increase in the right nonlinear way required to generate stable low-level equilibria that can plausibly account for very persistent income differences across countries. We also find that a poverty-trap view of aid relying on these mechanisms leads to counterfactual predictions for the relationship between aid, investment, and growth. This leaves us skeptical of the popular idea that sufficiently large increases in aid will have disproportionate effects on economic growth in low-income countries. This does not mean that aid does not matter for growth, or for other development outcomes. Rather we do not find evidence of threshold effects based on these two leading explanations for poverty traps whereby sufficiently high levels of aid are necessary to “jump-start” a sustainable growth process.

An important limitation of this paper is that we have only considered a few of the many possible theoretical mechanisms generating poverty traps, and we have restricted ourselves to highly aggregative models of growth. More empirical work on calibrating specific mechanism using both micro and macro data may show that other self-reinforcing processes are empirically important as explanations for underdevelopment. It may also be the case that a variety of mechanisms operating simultaneously will deliver empirically relevant poverty traps. Results

along these lines may well provide a basis for the argument that large-scale increases in aid will have disproportionate effects on growth. This work however remains to be done.

References

Antweiler, Werner and Daniel Trefler, (2000). "Increasing Returns and All That", NBER Working Paper 7941.

Azariadis, Costas and Allan Drazen, (1990), "Threshold Externalities in Economic Development," *The Quarterly Journal of Economics*, 105(2), pages 501-26.

Azariadis, Costas (1996). "The Economics of Poverty Traps, Part One: Complete Markets". *Journal of Economic Growth*. 1(4)449-496.

Azariadis, Costas, and John Stachurski, (2004), "Poverty Traps", Forthcoming in the Handbook of Economic Growth, Aghion and Durlauf eds.

Azariadis, Costas, and John Stachurski, (2004b), "A Forward Projection of the Cross-Country Income Distribution", Manuscript, UCLA.

Barro, Robert J, (1990), "Government Spending in a Simple Model of Endogenous Growth," *Journal of Political Economy*, 98(5), pages S103-26.

Bartelsman, Eric J, Ricardo J Caballero, and Richard K Lyons, (1994), "Customer- and Supplier-Driven Externalities," *American Economic Review*, 84(4), pages 1075-84.

Basu, Susanto and John G. Fernald, (1995), "Are apparent productive spillovers a figment of specification error?," *Journal of Monetary Economics*, 36(1), pages 165-188.

Ben-David, Dan, (1998), "Convergence Clubs and Subsistence Economies", *Journal of Development Economics*, 55, 153-159.

Bliss, Christopher, (2004), "Some implications of a variable EIS", Working Paper 2004-W26, Nuffield College, Oxford University.

Bloom, David E, Canning, David and Sevilla, Jaypee, (2003). "Geography and Poverty Traps," *Journal of Economic Growth*, vol. 8(4), pages 355-78, December.

Bryant, John, (1983), "A Simple Rational Expectations Keynes-Type Model," *The Quarterly Journal of Economics*, 98 (3), pages 525-28.

Caballero, Ricardo J. and Richard K. Lyons, (1990), "Internal versus external economies in European industry," *European Economic Review*, 34(4), pages 805-826.

Caucutt, Elizabeth and Krishna Kumar (2004). "Evaluating Explanations for Stagnation". Manuscript, University of Western Ontario and USC Marshall School of Business.

Caballero, Ricardo J. and Richard K. Lyons, (1992), "External effects in U.S. procyclical productivity," *Journal of Monetary Economics*, 29(2), pages 209-225.

Caselli, Francesco and Jaume Ventura, (2000), "A Representative Consumer Theory of Distribution", *American Economic Review*, 90 (4), 909-926.

Chakraborty, Shankha (2003). "Endogenous Lifetime and Economic Growth". *Journal of Economic Theory*, forthcoming.

Cooper, Russell and John Haltiwanger, (1996), "Evidence on Macroeconomic Complementarities," *The Review of Economics and Statistics*, 78(1), pages 78-93.

Dalgaard, Carl-Johan and Jes Winther Hansen (2005). "Capital Utilization and the Foundations of Club Convergence". *Economics Letters*. 87:145-152.

Devarajan, Shantayanan, William Easterly, and Howard Pack (2001). "Is Investment in Africa Too High or Too Low? Macro- and Micro-Evidence". *Journal of Africa Economies*. 10:81-108.

Devarajan, Shantayanan, William Easterly, and Howard Pack (2003). "Low Investment is Not the Constraint on African Development". *Economic Development and Cultural Change*.

Diamond, Peter A, (1982), "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy*, 90(5), pages 881-94.

Dollar, David, and William Easterly (1999), "The Search for the Key: Aid, Investment, and Policies in Africa", *Journal of African Economies*, 8 (4), 546-577.

Durlauf, Steven N, (1991), "Multiple Equilibria and Persistence in Aggregate Fluctuations," *American Economic Review*, 81(2), pages 70-74.

Easterly, William, (1999), "The ghost of financing gap: testing the growth model used in the international financial institutions", *Journal of Development Economics*, 60, 423-438.

Easterly, William (2005). "Reliving the 1950s: The Big Push, Poverty Traps and Takeoffs in Economic Development". Manuscript, NYU.

Fernandes, Ana Margarida, (2003), "Trade Policy, Trade Volumes, and Plant-Level Productivity in Colombian Manufacturing Industries", World Bank Policy Research Working Paper 3064.

Feyrer, James (2003), "Convergence by Parts", Mimeo, Dartmouth College.

Fuss, Melvyn A. and V. K. Gupta (1981) "A Cost Function Approach to the Estimation of Minimum Efficient Scale, Returns to Scale, and Suboptimal Capacity," *European Economic Review*, 15 (2), 123–135.

Galor, Oded (1996). "Convergence? Inferences from Theoretical Models". *Economic Journal*. 106(437): 1056-1069.

Global Trade Analysis Project database (GTAP) (2005), Purdue University, <https://www.gtap.agecon.purdue.edu/>

Graham, Bryan S. and Jonathan Temple, 2006. "Rich Nations, Poor Nations: How much can multiple equilibria explain?," *Journal of Economic Growth* (in press).

Griliches, Zvi and Vidar Ringstad (1971) *Economies of Scale and the Form of the Production Function: An Econometric Study of Norwegian Manufacturing Establishment Data*, Amsterdam: North Holland.

Hallward-Driemeier, Mary, Giuseppe Iarossi, and Kenneth Sokoloff (2002), “Exports and Manufacturing Productivity in East Asia: A Comparative Analysis with Firm Level Data”, NBER Working Paper 8894.

Harrigan, James (1999) “Estimation of Cross-Country Differences in Industry Production Functions,” *Journal of International Economics*, 47 (2), 267–93.

Hausmann, Ricardo, Lant Pritchett, and Dani Rodrik (2004). “Growth Accelerations”. Manuscript, Kennedy School of Government.

Howitt, Peter, (1985), “Transaction Costs in the Theory of Unemployment,” *American Economic Review*, 75(1), pages 88-100.

Howitt, Peter and R Preston McAfee, (1988), “Stability of Equilibria with Externalities,” *The Quarterly Journal of Economics*, 103(2), pages 261-77.

Kehoe, Patrick J. and Fabrizio Perri, 2002, “International Business Cycles with Endogenous Incomplete Markets.” *Econometrica*, 70(3), pages 907-928.

Levinsohn, James and Amil Petrin (1999) “When Industries Become More Productive, Do Firms? Investigating Productivity Dynamics.” Mimeo, University of Michigan.

Lokshin, Michael and Martin Ravallion (2004). “Household Income Dynamics in Two Transition Economies”. *Studies in Nonlinear Dynamics and Econometrics*. 8(3), Article 4.

Mulligan, Casey, and Xavier Sala-i-Martin, (1991), “A Note on the Time Elimination Method for Solving Recursive Dynamic Economic Models”, NBER Technical Working Paper 116.

McKenzie, David and Christopher Woodruff (2004). “Is There an Empirical Basis for Poverty Traps in Developing Countries?”. Manuscript, Stanford and UCSD.

Nelson, Richard (1956). "A Theory of the Low-Level Equilibrium Trap in Underdeveloped Economies". *American Economic Review*. 46(5):894-908.

Paul, Catherine J. Morrison and Donald S. Siegel (1999) "Scale Economies and Industry Agglomeration Externalities: A Dynamic Cost Function Approach," *American Economic Review*, 89 (2), 272–290.

Quah, Danny T, (1996). "Twin Peaks: Growth and Convergence in Models of Distribution Dynamics," *Economic Journal*, 106 (127), pages 1045-55.

Quah, Danny T, (1993a), "Empirical Cross-Section Dynamics in Economic Growth". *European Economic Review* 37, pp. 426–434.

Quah, Danny T, (1993b), "Galton's Fallacy and Tests of the Convergence Hypothesis". *The Scandinavian Journal of Economics* 95, pp. 427–443.

Quah, Danny T, (1997), "Empirics for Growth and Distribution: Stratification, Polarization, and Convergence Clubs". *Journal of Economic Growth* 2 (1997), pp. 27–59.

Rodrik, Dani (2000), "Saving Transitions". *World Bank Economic Review*. 14(3):481-507.

Romer, Paul M, (1986), "Increasing Returns and Long-run Growth," *Journal of Political Economy*, 94(5), pages 1002-37.

Sachs, Jeffrey, John W. McArthur, Guido Schmidt-Traub, Margaret Kruk, Chandrika Bahadur, Michael Faye, and Gordon McCord, (2004), "Ending Africa's Poverty Trap", *Brooking Papers on Economic Acitivity*, 2004 (1).

Sachs, Jeffrey (2005). "Investing in Development: A Practical Plan to Achieve the Millenium Development Goals. New York. UN Millenium Project.

Solow, Robert (1956). "A Contribution to the Theory of Economic Growth". *Quarterly Journal of Economics*. 70(1):65-94.

Steger, Thomas, (2000), "Economic Growth with Subsistence Consumption", *Jornal of Development Economics*, 62, 343-361.

Tybout, James R. (2000), "Manufacturing Firms in Developing Countries: How well do they do, and why?" *Journal of Economic Literature*, 38 (1), 11-44.

Tybout, James R. and M. Daniel Westbrook (1995) "Trade Liberalization and the Dimensions of Efficiency Change in Mexican Manufacturing Industries," *Journal of International Economics*, 39 (1-2), 53-78.

United Nations Industrial Development Organization (UNIDO) (2002). Industrial statistics database. CD-ROM.

Van Biesebroeck, Johannes (2005). "Exporting Raises Productivity in Sub-Saharan African Manufacturing Firms", *Journal of International Economics*, 67 (2), 373-391.

Weil, Philippe, (1989), "Increasing Returns and Animal Spirits", *American Economic Review*, 79, 889-894.

Table 1: Aid as Share of GDP Required to Escape From Poverty Trap

Capital Share (α)	Threshold Capital Stock (\$000's)				
	<u>2</u>	<u>3</u>	<u>4</u>	<u>5</u>	<u>6</u>
0.25	3.3%	9.4%	15.0%	20.3%	25.3%
0.5	2.9%	7.7%	11.7%	15.3%	18.6%
0.75	2.5%	6.2%	9.2%	11.6%	13.7%

Table 2: Persistence of Slow Growth

Initial capital (% above low equilibrium)	Time to close half the gap to steady state (years)
10%	40
1% (benchmark)	54
0.10%	68
0.01%	85
Median African (300%)	24
Ramsey economy	15
1% (capital share = 0.4)	59
1% (capital share = 0.6)	56

Table 3: Aid and Growth with Subsistence Consumption

Parameterization	Annual transfer to double initial capital in 10 periods (% GDP)	Annual transfer to double considering permanent transfers	Time to double capital by itself	Annual transfer to quadruple initial capital in 10 periods (% GDP)	Annual transfer to quadruple considering permanent transfers	Time to quadruple capital by itself
	(1)	(2)	(3)	(4)	(5)	(6)
Baseline	3.9	9.3	29	12.8	32.5	42
Median african	4.8	12.8	15	--	--	--
Ramsey ($\underline{c}=0$)	0.0	0.0	3	0.0	0.0	8
Low capital share ($\alpha = 0.4$)	5.4	15.5	38	17.4	51.1	69
High capital share ($\alpha = 0.6$)	2.3	6.2	22	8.7	18.6	30

Table 4: Estimates of increasing returns in literature

Source	Sample of Countries	Estimated degree of returns to scale
Harrigan (1999)	OECD	≤ 1
Paul and Siegel (1999)	US	1.3
Tybout (2000)	Asia, LAC, MENA	1.05-1.1
Tybout and Westbrook (1995)	Mexico	1
Van Biesebroeck (2005)	SSA	1-1.14
Levinsohn and Petrin (1999)	Chile	1.2-1.44
Hallward-Driemeier et al. (2002)	East Asia	1.01-1.18
Antweiler and Trefler (2000)	71 countries	1-1.2

Notes: LAC stands for Latin-America and the Caribbean, MENA for Middle East and North Africa, SSA for Sub-Saharan Africa, US for the United States of America, and OECD for the countries that belong to the Organization for Economic Cooperation and Development

Figure 1: Solow Model With Poverty Trap





Figure 2: Ramsey Economy with Subsistence Consumption

Phase Diagram



Calibrated Phase Diagram

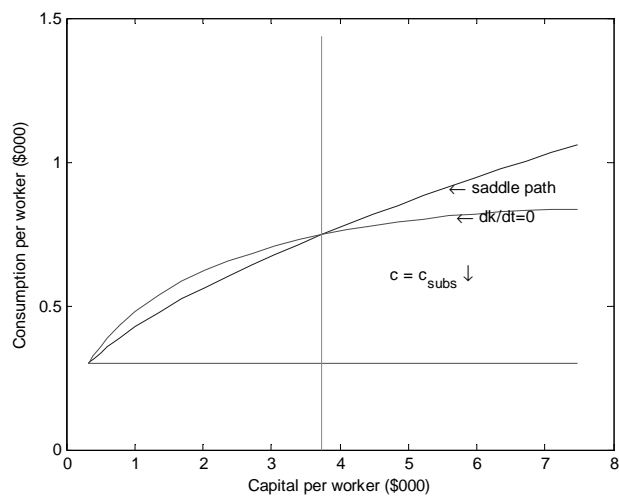




Figure 3: Growth With Subsistence Consumption


Capital Stock Per Capita



Consumption Per Capita



Growth Rate



Saving Rate





Figure 4: Effects of Aid

Saving



Growth




Figure 5: Cross-Country Relation Between Saving and Capital Stocks Per Capita

Empirical Evidence

Consequences for Solow Model

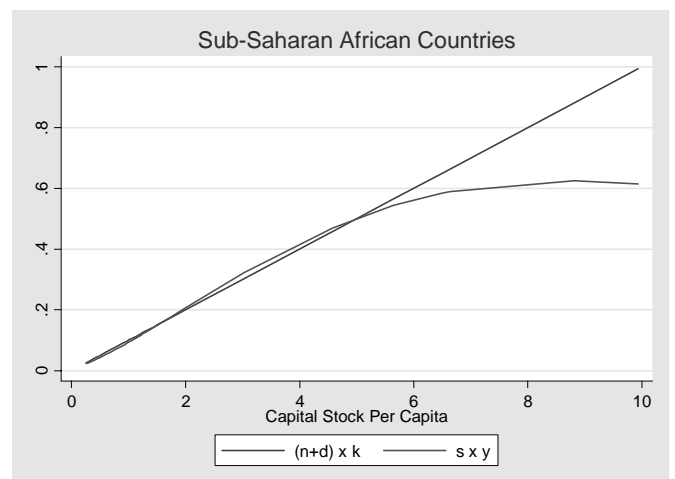
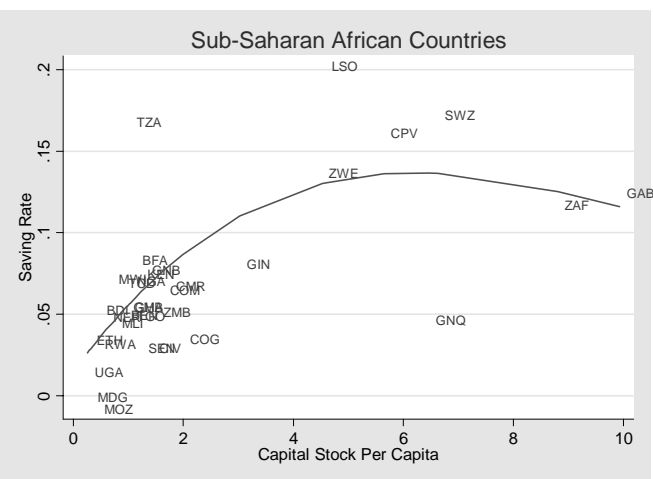
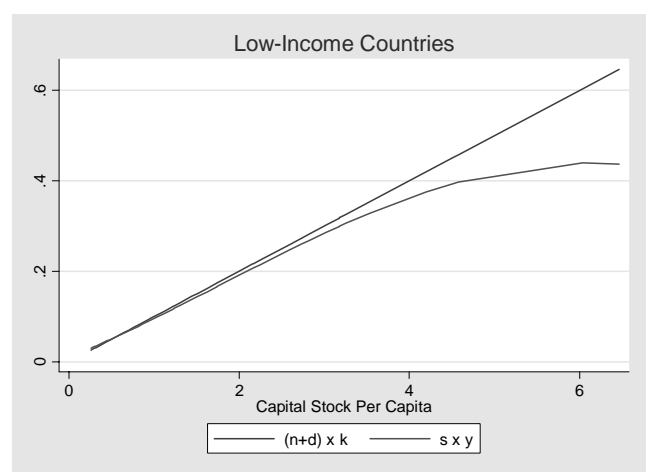
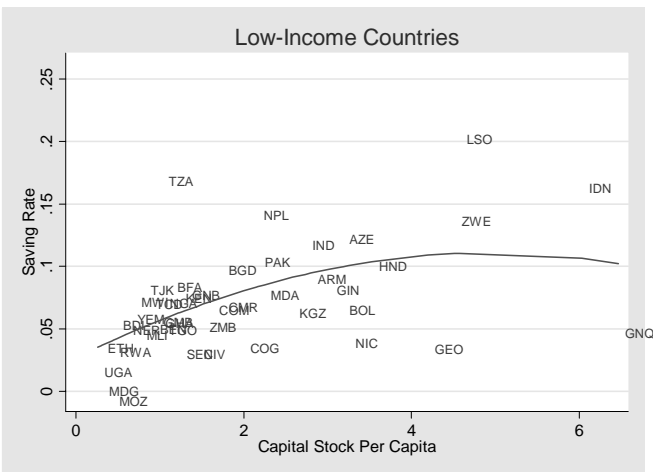
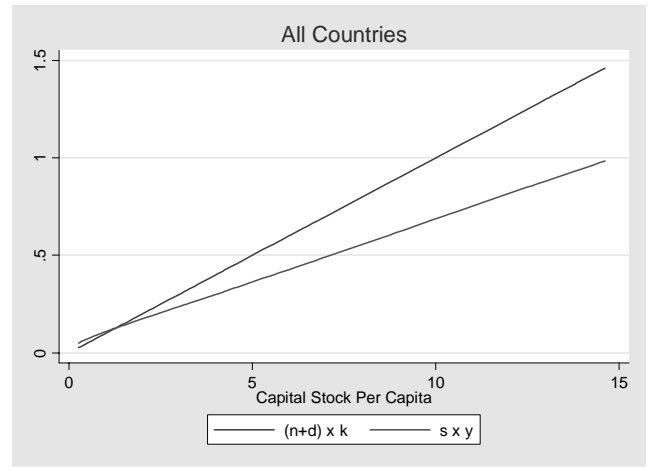
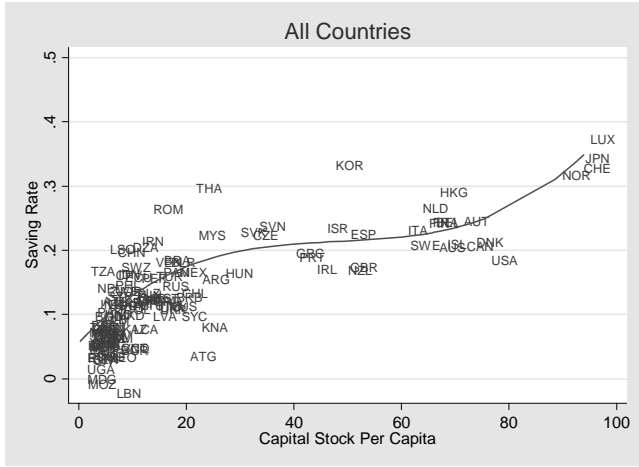



Figure 6: Changes in Saving and Growth, 1990s vs. 1960s

Saving Accelerations



Changes in Saving and Growth


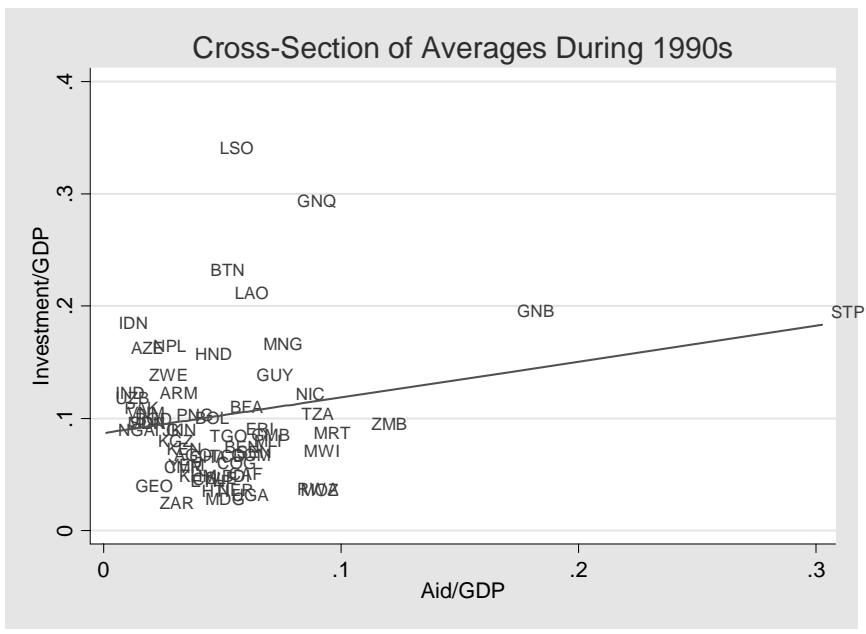


Figure 7: Aid and Investment

Cross-Section in 1990s



Changes in 1980s vs. 1990s

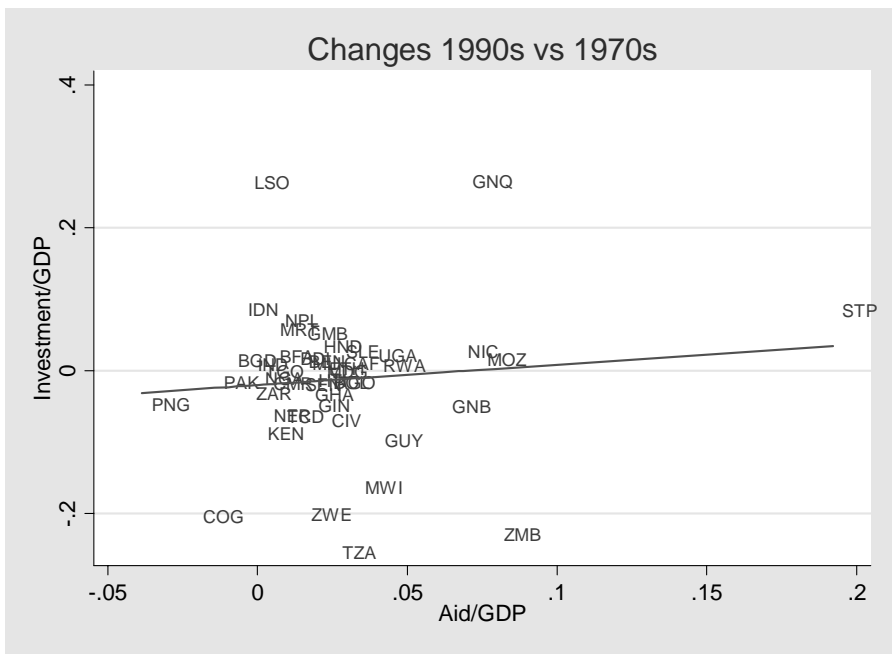


Figure 8: Technologically induced multiple equilibria



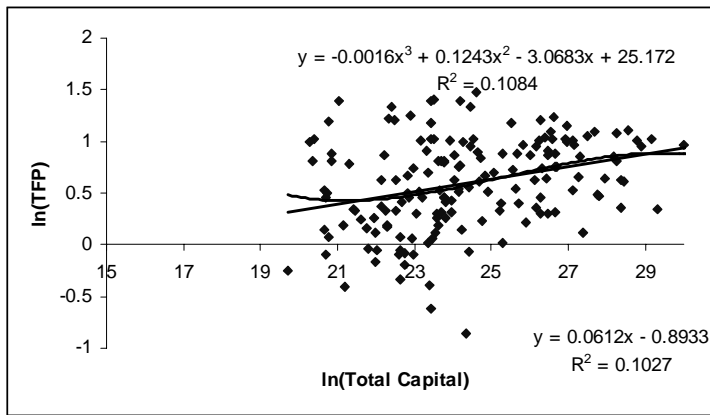


Figure 9: Cross-Country Relation Between TFP and Capital

Per Capita Capital



Total Capital



Total Output




Figure 10: Relation between TFP and measures of activity in Colombian firms

Figures show the partial relations between a firm's log TFP and the level of capital per worker, total capital, and total output of the industry where the firm's operates. Each figure also shows the results obtained in two regressions of the (log) firm's TFP and each corresponding measure of aggregate activity (also in logs). The first regression includes only the linear term while the second fits a cubic polynomial. Each regression also included a firm, industry, and year fixed effect

