


This article was downloaded by: [Australian National University]
On: 15 May 2012, At: 00:49
Publisher: Routledge
Informa Ltd Registered in England and Wales Registered Number: 1072954
Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH,
UK



Journal of Development Studies

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/fjds20>

The economics of poverty traps and persistent poverty: An asset-based approach

Professor Michael R. Carter ^a & Christopher B. Barrett ^b

^a University of Wisconsin, Madison, USA

^b Cornell University, Ithaca, NY, USA

Available online: 19 Aug 2006

To cite this article: Professor Michael R. Carter & Christopher B. Barrett (2006): The economics of poverty traps and persistent poverty: An asset-based approach, *Journal of Development Studies*, 42:2, 178-199

To link to this article: <http://dx.doi.org/10.1080/00220380500405261>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.tandfonline.com/page/terms-and-conditions>

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae, and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand, or costs or damages

whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

The Economics of Poverty Traps and Persistent Poverty: An Asset-Based Approach

MICHAEL R. CARTER* & CHRISTOPHER B. BARRETT**

*University of Wisconsin, Madison, USA, **Cornell University, Ithaca, NY, USA

ABSTRACT *Longitudinal data on household living standards open the way to a deeper analysis of the nature and extent of poverty. While a number of studies have exploited this type of data to distinguish transitory from more chronic forms of income or expenditure poverty, this paper develops an asset-based approach to poverty analysis that makes it possible to distinguish deep-rooted, persistent structural poverty from poverty that passes naturally with time due to systemic growth processes. Drawing on the economic theory of poverty traps and bifurcated accumulation strategies, this paper briefly discusses some feasible estimation strategies for empirically identifying poverty traps and long-term, persistent structural poverty, as well as relevant extensions of the popular Foster-Greer-Thorbecke class of poverty measures. The paper closes with reflections on how asset-based poverty can be used to underwrite the design of persistent poverty reduction strategies.*

I. Persistent Poverty and the Challenge of Forward-looking Poverty Measurement

Much empirical poverty analysis of poverty is dedicated to defining, measuring or locating who is recently poor. Such analysis is almost unavoidably backward looking in the sense that it creates a portrait of who was poor at the time survey data were collected. However, the observation of persistent or chronic poverty motivates a more forward-looking question: Who will likely remain poor into the future?¹

The empirical papers that comprise this special issue contribute to answering this question by trying to understand the structural reasons that underlie poverty's persistence, asking when and why poverty reproduces itself over time, 'laying eggs' as one of the informants to the *Chronic Poverty Report* describes it (CPRC, 2004: 3). Some of the papers try to understand who among the poor is structurally positioned to take advantage of new economic opportunities when they appear (Adato et al., Barrett et al., Krishna et al., Peters and Whitehead).² Other papers ask who is positioned to ride out the negative shocks that destroy opportunities, while still maintaining a viable basis for future advance (Hoddinott and Little). In their search

Correspondence Address: Professor Michael Carter, University of Wisconsin, Madison, USA.
Tel: +1 608 263 2478; Fax: +1 608 262 4376; Email: mrcarter@wisc.edu

for viable and less viable structural positions, these studies implicitly define an asset-based approach to persistent poverty, where the term asset is understood to broadly include conventional, privately held productive and financial wealth, as well as social, geographic and market access positions that confer economic advantage. The goal of this paper is to frame this approach by developing the conceptual foundations for asset-based poverty measures that permit a forward-looking approach and help identify and ultimately understand the structure and persistence of poverty.

The remainder of this paper is organised as follows. Section II motivates an asset-based approach by noting the limited ability of conventional poverty measures to deal with time and poverty transitions. Section III then takes a first step towards a forward-looking, asset based approach to poverty by developing the concept of the (static) asset poverty line. A corresponding family of measures based on this line – modelled on the familiar Foster-Greer-Thorbecke (1984, hereafter FGT) logic – are then presented. These measures provide information on the depth of structural poverty given the current distribution of assets (and the returns to those assets).

Section IV then allows for the possibility of asset accumulation/decumulation dynamics. After a brief review of the economics of poverty traps, Section IV develops forward-looking measures of dynamic asset poverty based the concept of a dynamic asset poverty threshold (which we label the ‘Micawber’³ threshold) that potentially separates those able to move to a high (non-poor) asset position from those caught in a low-level equilibrium trap. Section V considers some of the econometric challenges that confront the identification of the Micawber threshold and forward-looking poverty measures, while Section VI concludes the paper with reflections on the policy implications and uses of these measures.

II. Why an Asset Based Approach to Poverty?

As a starting point for thinking about persistent poverty, it is useful to consider what standard expenditure-based poverty measures can and cannot tell us. Figure 1 schematically represents alternative approaches to measuring poverty. The most common (first generation) approach to poverty measurement relies on household expenditure (or income) data from a single point in time. Once a money metric poverty line is defined, the population can be divided into poor and non-poor categories, and the standard suite of headcount and other FGT measures can be calculated to gauge the extent and depth of poverty within an economy. Application of these first generation poverty analysis methods to repeated cross-sectional surveys allows insight into the evolution of poverty within a society.

However, as numerous authors have remarked, cross-sectional poverty measurement is unable to distinguish between two very different patterns of poverty, each with a very different meaning. Consecutive cross-sectional findings of, say, a 33 per cent poverty headcount ratio could reflect a society in which the same one-third of individuals are persistently poor, period after period. In such a society, poverty would be experienced by only a minority, but intensely and indefinitely for those unlucky few. Alternatively, repeated observations of the same headcount ratio could reflect a reality in which poverty is a purely transitory phenomenon in which individuals routinely swap places on the basis of random outcomes, or perhaps based on age or other demographic process. Over time, all households would be

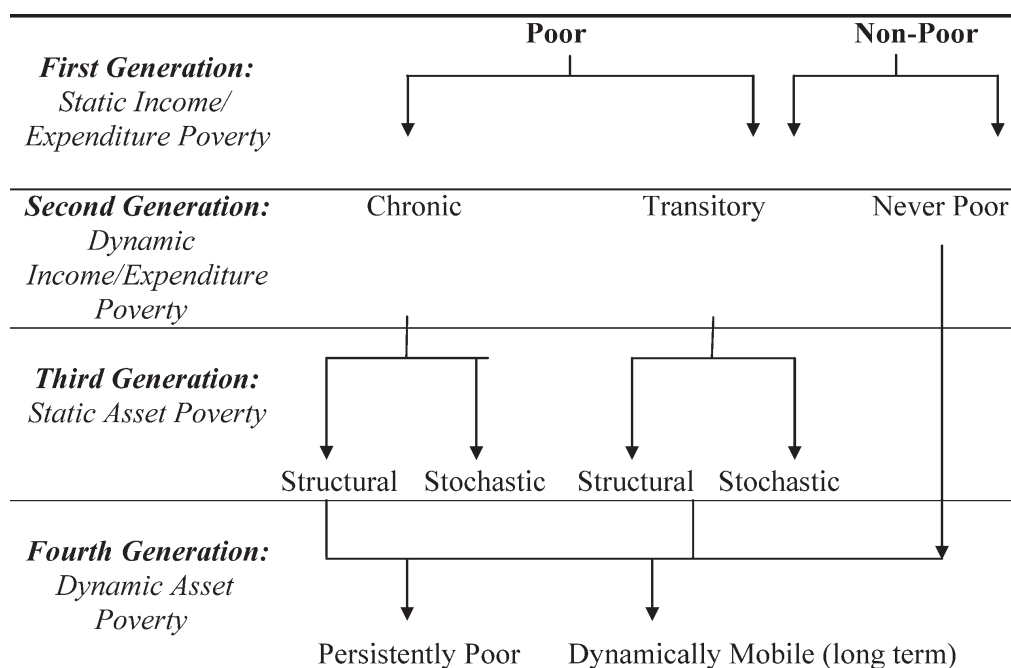


Figure 1. Alternative approaches to poverty measurement

poor one-third of the time, thus all would share the burden of poverty equally and only for a minority of the time.

Clearly a society typified by the first reality would be a much more polarised society, one vulnerable to hopelessness among a large subpopulation – and perhaps inter-class strife – and thus quite different from the one typified by the second poverty process. Unfortunately, first generation poverty measures are incapable of distinguishing between these starkly different poverty processes.

Interest in distinguishing between these two very distinctive situations has motivated a second generation of poverty analysis based on longitudinal or panel data that offer repeated observations over time on a single cohort of individuals or households. Grootaert and Kanbur (1995) offer an early and influential example of panel data-based expenditure or income poverty analysis. As illustrated in Figure 1, panel data permit a further decomposition of households into three categories: the always or chronically poor, the sometimes or transitorily poor, and the never poor.

In a summary of initial studies of panel data studies of poverty, Baulch and Hoddinott (2000) report on detailed studies of poverty dynamics based on panel data from ten countries. Updating that effort, Hoddinott (2003) found that the number of panel studies of African poverty had risen substantially. A common finding across all of these studies is that transitory poverty comprises a rather large share of overall poverty. The large share of transitory poverty based on income or expenditure underscores the inherent stochasticity of flow-based measures of welfare. People are better off one period than another without any significant or lasting change in their underlying circumstances, particularly the stock of productive assets under their control, due solely to random price and yield

fluctuations and irregular, stochastic earnings from remittances, gifts, lotteries, and so forth.⁴

The Achilles heel of these informative, second generation poverty measures is that they cannot distinguish between very distinctive sorts of poverty transitions. Individuals may appear to be transitorily poor in a standard panel study, moving from the poor to the non-poor state over time due to either of two markedly different experiences. Some may have been initially poor because of bad luck. Their transition to the non-poor state simply reflects a return to an expected non-poor standard of living (a stochastic poverty transition). For others, the transition may have been structural, due to the accumulation of new assets, or enhanced returns to the assets that they already possessed.

Similarly, those transitorily poor individuals who move from being non-poor to poor, can represent a mix of experiences. For some, it could represent a return to an expected standard of living, after a brief non-poor hiatus afforded by a spell of good luck. For others, it could be a temporary transition caused by bad luck in a later survey period. Finally, for yet others, it could be a structural move caused by the loss of assets (due to illness, natural disaster or theft), or by a deterioration in returns to their assets brought on changes in the broader economy (for example, unemployment or declining terms of trade).⁵

In short, the inability of second generation poverty analysis to distinguish structural from stochastic transitions limits its ability to describe how well an economy works for its least well off members. Does an observed amount of upward mobility reflect an economy that functions for at least some of the poor, facilitating asset accumulation and increased returns to the assets held by the poor? Or, does it reflect a large amount of a large amount of structural stasis and hopelessness that is masked by the churning of households that already possess assets and enjoy expected returns that predict a non-poor standard of living on average?⁶

To overcome these limitations of second generation poverty measurement, this paper reformulates poverty measurement in asset space. Section III uses the work of Carter and May (1999, 2001) to identify an asset poverty line as a natural extension of the familiar flow-based concept of an expenditure or income poverty line. This asset poverty line can be used to distinguish stochastic from structural transitions, making it possible to decompose poverty transitions, as shown in Figure 1. The asset poverty line can also be used as the basis for a suite of structural poverty indicators that provide a snapshot of structural poverty, cropped of the influence of stochastic transitions.

While defining and measuring poverty based on the asset poverty line provides important information on the structural foundations of poverty, it does not speak to the long-term persistence of structural poverty. As illustrated in Figure 1, analysis based on the asset poverty line cannot by itself identify whether the currently structurally poor are likely to remain poor over the longer term, caught in a poverty trap, or indeed whether a subset of the structurally non-poor can sustain their positions over the longer term. To further decompose these groups according to their long-term, persistent poverty status requires a fourth generation approach to poverty based on an understanding of underlying patterns of asset dynamics. As discussed in Section IV, identification of the dynamic asset poverty threshold is the key to decomposing current structural poverty into its persistent and more transitory components.

III. The Asset Poverty Line and Measures of Structural Poverty

This section takes first steps toward developing forward-looking poverty measures that are informative about the nature of long-term, structural poverty. After defining the concept of an asset poverty line, this section shows how the asset poverty line can be used to identify those households who lack the assets that, on average, generate a non-poor level of expenditure or income. Such measures are informative about the likely prospects of a household possessing a given asset portfolio, given past asset productivity, much like increasingly-popular vulnerability measures (Christiaensen and Boisvert, 2000; Christiaensen and Subbarao forthcoming).

(a) Using the Asset Poverty Line to Decompose Poverty Transitions

Distinguishing between stochastic and structural transitions requires information on assets and expected levels of well-being. Conceptually, this is a relatively straightforward exercise, as indicated by Figure 2, adapted from Carter and May (2001). The vertical axis measures a standard flow indicator of achieved material well-being (or utility), typically measured as income or expenditure. The conventional money metric poverty line measured in this dimension is denoted \underline{u} . The horizontal axis measures the assets that generate a household's livelihood.




Figure 2. Single period income and asset poverty lines

While these assets are multi-dimensional, tangible and intangible, we assume here for illustrative purposes that assets are one-dimensional, or that we have non-problematically aggregated them into a one-dimensional index measure. We briefly discuss methods for building such an asset index in Section V.

One can map the relationship between assets and income, expenditures or some other flow measure of well-being, as illustrated by the (expected) livelihood function graphed in Figure 2.⁷ The asset poverty line is then simply the level of assets (denoted \underline{A} in Figure 2) that predicts a level of well-being equal to the poverty line, \underline{u} . Purely for expositional purposes, assume for the moment that the livelihood function does not change over time.⁸ Then in any time period, a household is stochastically poor if it holds assets worth at least \underline{A} yet its realised income or expenditure falls stochastically below \underline{u} . Conversely, the household is structurally poor if its stock of assets is less than \underline{A} and its realised income or expenditure level falls, as expected, below \underline{u} .

Panel data permit estimation of the asset poverty line and enable the third generation decomposition of poverty transitions shown in the third row of Figure 1. A household that moved over time from above to below the standard expenditure-based poverty line could be said to have made a stochastic transition back to its expected status if the household's assets still mapped into an expected standard of living below the poverty line. In Figure 2, this transition is illustrated as the movement from point C back to the point $\hat{u}(A')$. Alternatively, a household that moves from $\hat{u}(A'')$ to $\hat{u}(A')$ would have made a structural transition below the poverty line due to a loss of assets from A'' to A' .

Similarly, a household that made the opposite observed expenditure transition (from below to above the poverty line) could be said to have made a structural transition if household assets predicted expenditure initially below the poverty line, at $\hat{u}(A')$, but in the subsequent period assets yield expected expenditures above the poverty line. Such a shift could occur either because of asset accumulation that moved the household to point $\hat{u}(A'')$, or because of improved returns on the household's stock of assets, which shifted the livelihood function from $\hat{u}(A)$ to $\tilde{u}(A)$, bumping expected and observed expenditures from $\hat{u}(A')$ to point C in Figure 2. Finally, in Figure 2, the stochastic transition out of poverty would be manifest as a movement from point B to $\hat{u}(A'')$, which merely reflects a return to a household's expected welfare level given its asset holdings and the livelihood function mapping assets into expenditures.

This asset-based approach thus moves us considerably closer to being able to address the key questions surrounding households' longer-term prospects of being non-poor. The challenge in implementing these ideas results from the need to estimate a livelihood mapping between assets and expenditures (or income) statistically. Carter and May (2001) illustrate an application of this method to South African households, cautiously denoting a household as stochastically poor only if one can reject the statistical hypothesis that their assets are expected to yield flow-based welfare measures below the standard expenditure or income poverty line.

While one could quibble with components of their methodology, the Carter and May analysis nicely illustrates both the strengths and limitations of the asset poverty line. They estimate that less than half of the observed transitions out of

poverty in South Africa over the 1993–98 period are structural, as 60 per cent of the households who made the transitions had initial period assets that strongly predicted well-being in excess of the standard poverty line. In terms of downward mobility, Carter and May find that only a small fraction (15 per cent) clearly fell into poverty for stochastic reasons, while fully 51 per cent of those who fell behind suffered asset losses that left them structurally poor in the latter survey period.

In addition to helping unpack the nature of observed past transitions between poor and non-poor status, the asset poverty line can also distinguish households that have a current asset base that predicts a non-poor future standard living from those whose current circumstances predict a standard of living below the poverty line. This latter group is arguably of greater concern – and an appropriate target for intervention – as they would be expected to remain or become poor in the future, absent asset accumulation or further structural change in the economy.

(b) Structural Poverty Measures

The forward-looking, structural insights into poverty afforded by the asset poverty line can be used to create a class of structurally-based poverty measures based on the familiar FGT poverty measures. Recall that the FGT class of decomposable, single-period P_α poverty measures is defined as:

$$P_\alpha = \frac{1}{N} \sum_{i=1}^N I_i \left(\frac{u - u_i}{u} \right)^\alpha \quad (1)$$

where N is the sample size, u is the scalar-valued poverty line, u_i is the flow-based measure of welfare (income or expenditures), I is an indicator variable taking value one if $u_i < u$ and zero otherwise, and α is a parameter reflecting the weight placed on the severity of poverty. Setting $\alpha = 0$ yields the headcount poverty ratio P_0 (the share of a population falling below the poverty line). The higher order measures, P_1 and P_2 , yield the poverty gap measure (the money metric measure of the average financial transfer needed to bring all poor households up to the poverty line) and the squared poverty gap (an indicator of severity poverty that is sensitive to the distribution of well-being amongst the poor).

Consider now the FGT class of measures defined around the asset poverty line, \underline{A} :

$$P_\alpha^A = \frac{1}{N} \sum_{i=1}^N I_i^A \left(\frac{\underline{A} - A_i}{\underline{A}} \right)^\alpha \quad (2)$$

where A_i is asset stock of household i and the binary indicator variable $I_i^A = 1$ if $A_i < \underline{A}$ reflects whether the household i 's asset stock falls below the static asset poverty line. The order zero measure, P_0^A provides a head-count measure of the structurally poor, while P_1^A measures magnitude of the average asset transfer (or accumulation) needed to bring the structurally poor just up to the asset poverty line. Analogous to the flow-based FGT measures, higher order asset poverty measures ($\alpha > 2$) will be sensitive to the distribution of assets amongst the poor.

(c) The Temporal Dimension: Limitations of the Asset Poverty Line

As in the South African example discussed above, the asset poverty line gives a sharper picture of the nature of poverty dynamics. Analysis of P_{α}^A measures along with the standard flow-based FGT measures also promises a much more complete and forward-looking poverty portrait.

However, analysis based on the (static) asset poverty line suffers from two conceptual weaknesses. First, like the standard flow-based measures, its definition depends on an arbitrary living standard (\underline{u} in the notation above). Second, analysis based on the asset poverty line does not account for any predictable future changes in the assets of the poor (nor predictable changes in the future returns to those assets). Analysis based on the static asset poverty line therefore cannot reliably indicate whether structurally poor households are likely to remain so into the foreseeable future, or whether they are headed in the right direction, nor whether structurally non-poor households can be expected to remain non-poor indefinitely, that is, are they free and clear of the poverty line for good? Put differently, how many of the structurally poor are likely to be structurally mobile over the long term? Alternatively, how many are caught in a long-term trap of persistent poverty? Similarly, how many of the structurally non-poor are actually in a sustainable situation?

Answering these questions requires an approach to poverty based on asset dynamics. Similar to the way in which the single period asset poverty line can distinguish between stochastic and structural poverty transitions in the short term, the remainder of this paper argues that a dynamic asset poverty line can help distinguish households caught in a long-term structural poverty trap from those expected to follow an upward trajectory, that is, those who enjoy structural economic mobility. The next section develops the theoretical foundations for the dynamic asset poverty line, the threshold at which accumulation dynamics bifurcate, leading to multiple dynamic welfare equilibria, including the possibility of a poverty trap.

IV. Poverty Traps and the Dynamic Asset Poverty Threshold

Households that can steadily accumulate assets or who enjoy steady technical change or favourable shifts in their terms of trade will grow their way out of poverty. Among very poor populations, this growth could take some time, but movement nonetheless proceeds steadily in the right direction. For these households, time would be a dependable ally in the fight against poverty and would oversee a domestic process of convergence as poor households climb out of poverty and catch-up to their better-off neighbours. But does time work in favour of poor households, or is the case that the many of the poor ‘can’t get ahead for falling behind’ (Barrett and Carter, 2001–2)?

Analogous questions of convergence have figured prominently in the macro-economic debate over the growth of nations.⁹ While there are some critical differences between economic growth at national and household levels, macro-economic growth theory and its attendant convergence controversy provide some useful insights and language for thinking about poverty and growth within nations. After a brief consideration of key ideas that have emerged from the macro growth

and convergence literature, this section explores the microeconomics of household accumulation and poverty traps.

(a) Thresholds and Clubs: Insights from the Convergence Controversy over the Growth of Nations

The workhorse model of neoclassical economic growth relies on an assumption of diminishing returns to assets (that generate a stream of income) to hypothesise that poorer nations will tend to catch up over time, or converge, with the incomes of richer nations. However, there is overwhelming empirical evidence that income convergence does not accurately describe economic growth at the macro level of nation states. In the words of Lant Pritchett (1997), the observation of ‘divergence, big time’ has invited 20 years’ debate and new theorising about alternative frameworks that might fit the data better.¹⁰

Within the macro-growth literature, two alternatives to the neoclassical growth model have emerged to account for the observed pattern of divergence. The first is the idea of conditional or club convergence, meaning that groups of countries that share similar intrinsic characteristics tend to converge to a living standard that is unique to their group or club. While there is convergence within clubs, there can be divergence between clubs.

The idea of conditional convergence dates back at least to Baumol (1986) and DeLong (1988) writing on ‘club convergence’, wherein distinct subpopulations (of nations, in their case) appear to converge on different steady-state growth rates. Quah (1993, 1996, 1997) extended this notion to more general distribution dynamics to explore the mobility of countries across income levels. Theories of conditional convergence turn fundamentally on the existence of an exclusionary mechanism, an immutable intrinsic characteristic that keeps members of one group or club facing a lower level equilibrium from moving to another group or club with a higher level equilibrium. The extant macro literature offers only rather vague suggestions as to why such exclusionary mechanisms might exist, hypothesising about distance from sea ports, agro-ecological conditions and their impacts on health and agricultural productivity, natural resource endowments and their effects on incentives to industrialise, or the institutional legacies of colonial history, including intra-national ethnic diversity (Acemoglu et al., 2001; Bloom and Sachs, 1998; Bloom et al., 2003; Easterly and Levine, 1997; Masters and Macmillan, 2001; Sachs and Warner, 1997).

The second alternative posits thresholds and multiple equilibria. From this perspective, there is no unique equilibrium for a country. Instead, controlling for a country’s intrinsic characteristics, both high- and low-level equilibria are available. Whether the country reaches a high level equilibrium, or remains trapped at a low-level equilibrium, depends on whether the country begins above, or is able to boost itself over a critical minimum threshold level of capital or income.

In contrast to club convergence vision of intrinsic differences between nations, the multiple equilibria growth models posits the possibility of poverty traps related to thresholds at which returns are locally increasing (Azariadis and Drazen, 1990; Fiaschi and Lavezzi 2003; Murphy et al., 1989). These theories formalise earlier, informal models of economic ‘take off’ or ‘big bang’ (Young, 1928; Rosentstein-Rodan, 1943; Nurkse, 1953; Myrdal, 1957; Rostow, 1960), which likewise depended fundamentally

on locally increasing returns. The key insight of these models is that without a coordinated push that dramatically increases the scale of production and, or market size past some threshold, so that firms can tap into (locally) increasing returns to scale, the economy will get stuck at a low income equilibrium.¹¹

While both the club and threshold perspectives can account for patterns of divergence in the global economy, distinguishing between them empirically is clearly important from a policy perspective. While the econometric challenges are daunting, several recent papers have made significant progress on this front. Canova (2004) shows that countries can be divided into groups or clubs that gravitate toward distinct equilibrium income levels, while Hansen (2000) identifies critical threshold levels of (initial) per-capita GDP and literacy that divide low from high equilibrium countries. However, because Hansen's thresholds are cast in terms of time-invariant initial conditions (that is, initial conditions effectively define clubs), it is not yet clear whether the data support the notion that there are multiple equilibria available to any given country. Nonetheless, the notion that poverty traps can result from either unfavourable intrinsic characteristics, or from locally increasing returns processes that generate multiple equilibria is an important idea to carry forward to the analysis of household level poverty.

(b) Microeconomics of Poverty Traps and Asset Dynamics

As with nations, individuals may also have intrinsic characteristics (skills, savings propensities, discount rates, and geographic locations) that condition their desired level of accumulation and ultimate equilibrium level of well-being. However, there may also be analogues to the locally increasing returns to scale that generate multiple equilibria and thwart the ability of initially poor households to catch up and converge with their wealthier neighbours. This section focuses on forces that can create locally increasing returns at the individual level and draws out their implications for poverty traps and asset dynamics.

When returns are locally increasing, there will be a positive relation between wealth (level of assets) and the marginal returns to assets. At the microeconomic or household level, a positive relationship between wealth and marginal returns can exist for at least three reasons:

- (1) the underlying income generating process may itself directly exhibit increasing returns to scale, either because the primal technology exhibits locally increasing returns or because input (output) prices, or transactions costs are negatively (positively) related to scale over some significant range;
- (2) some high return production processes may require a minimum project size such that only wealthier households can afford to switch to and adopt the high return process; and
- (3) risk and financial market considerations may cause some lower wealth households to allocate their assets so as to reduce risk exposure, trading off expected gains for lower risk, thereby making expected marginal returns to wealth lower for lower wealth households.

For expository purposes, we will examine the second of these three reasons in detail.

Consider the case where a household can allocate its productive wealth to two distinct productive activities, L_1 and L_2 . Both activities exhibit diminishing returns to wealth, as under the canonical neoclassical growth model. However, activity L_2 has a minimum scale of operation due to sunk costs of operation or of switching into L_2 (that is, it generates no returns if the wealth dedicated to this activity is below this minimum level). Figure 3 graphs these two production technologies as well as the steady state asset values that a household would choose if it were exogenously restricted to one technology or the other. Note that the graph is drawn for a given set of intrinsic characteristics (individual time preferences, technical efficiency or skill, and so forth).

For an individual with these characteristics, the value A_1^* denotes the steady state value for a household restricted to livelihood activity L_1 , yielding income or material well-being level U_L^* . The value A_2^* denotes the same thing for L_2 , yielding the higher level steady state income, U_H^* .¹² For illustrative purposes, Figure 3 places the asset poverty line, \underline{A} , between A_1^* and A_2^* . Note that this implies that any individual who settles into equilibrium at A_1^* would be caught in a poverty trap even though in principle a higher, non-poor equilibrium exists.

So how would a household sort itself between activities and their implied equilibrium asset and well-being levels. Assuming that no risk or other constraints limit the adoption of the technologies, Figure 3 shows that the optimal livelihood choice for households is activity L_1 for households with asset stocks up to A_S , and L_2 for households with assets in excess of A_S . Although each of these livelihood functions exhibits diminishing returns, there are locally increasing returns in the neighbourhood of A_S , the threshold at which households optimally switch from L_1 to L_2 . There are plentiful empirical examples of such patterns, for example, households possessing more assets who adopt higher return crop varieties or agronomic practices, wealthier households who get skilled salaried employment rather than unskilled casual wage labour, or households who graduate from poultry or small ruminants to indigenous cattle to improved dairy cattle and advanced




Figure 3. Asset poverty with multiple livelihood options

animal husbandry practices (for example, artificial insemination, supplemental feeding, and so forth) as wealth grows and these methods become affordable.

While it thus seems reasonable to postulate that poorer households might utilise technology L_I , the key dynamic question is whether or not the pattern of locally increasing returns would impede the ability of this household to accumulate, cross over asset level A_S , and catch up with wealthier households. Consider an individual with assets between A_1^* and A_S . Two features of this problem are relevant to the long-term accumulation choices of this individual. First, the individual will be earning relatively low rates of return on their modest asset holdings, a factor which further perpetuates their poverty because they earn less investible surplus, after meeting immediate consumption needs, than do richer households. Second, the marginal short-term, or myopic, incentives to save are depressed. If household accumulation decisions were driven by these depressed returns and liquidity constraints, then the household would indeed be expected to reach an equilibrium asset holding at the relatively low level, A_1^* .

The key question then becomes whether or not household savings and accumulation behaviour will be driven by these low marginal returns. A forward-looking household would know that while the marginal returns to further accumulation are low, increased accumulation has strategic value in moving the household closer to the asset level(s) where returns sharply increase. Clearly the household's first best option would be to borrow sufficient funds so that it could leap forward to a higher return asset level. Increasing returns would therefore not suffice by themselves to trap poor households at low asset levels.

If, however, poor households are rationed out of credit markets, as a now voluminous literature suggests, or if they lack socially mediated access to capital, as Mogues and Carter (forthcoming) suggest occurs in many polarised societies, then discrete jumps enabled by strategic borrowing may not be possible. In the face of exclusion from financial markets, a poor household's only option would be to move forward slowly with an autarchic savings strategy. This approach would require substantial short-term sacrifice (diminished consumption) with little return even in the medium term (as marginal returns to new assets are low until the household reaches A_S). If the poor household finds it desirable and feasible makes this sacrifice, then it will – with sufficient time – reach the asset level necessary to achieve the higher returns and will eventually converge toward the asset and income levels of initially wealthier households. But many very poor households cannot afford to reduce consumption further, or at least the opportunity cost of tightening their belts further – for example, in terms of foregone energy for work, withdrawing children from school, and so forth – make autarchic accumulation unattractive. If the poor household opts not to undertake extraordinary savings, it then settles into a poverty trap.

A somewhat complex theoretical literature explores the conditions under which each of these two outcomes is most likely to occur (for example, Loury, 1981; Banerjee and Newman, 1993; Galor and Zeira, 1993; Mookherjee and Ray, 2002). The basic intuition is, however, simple. It would seem likely that if a household was not 'too far', in some sense, from the asset level where increasing returns occur, then it would be likely to pursue the autarchic accumulation strategy. However, as the distance from that level increases, it seems less likely that households would find it

feasible and desirable to pursue the autarchic accumulation strategy. Zimmerman and Carter (2003) identify a Micawber threshold, the critical asset threshold below which it is no longer rational or feasible to pursue the autarchic accumulation strategy. If it exists, the Micawber threshold thus constitutes a dynamic asset poverty threshold, analogous to the static asset poverty line discussed in the previous section. Households whose assets place them above that threshold would be expected to escape poverty over time, while those below would not. One needs to identify this dynamic asset poverty threshold in order to disaggregate the structurally poor into those expected to escape poverty on their own over time through predictable asset accumulation and those expected to be trapped in poverty indefinitely.

As with the existence of multi-equilibria in macro growth models, the existence of the Micawber threshold has important policy implications (some of which we discuss in the conclusion to this paper). While the theoretical literature offers insights as to when such a threshold will occur, the really important question is the empirical one of whether such a threshold exists and, if so, where. As a first step in this direction, we now consider testable implications of such a threshold if it exists.

For illustrative purposes, denote $A^* < A_S$ as the critical dynamic asset poverty threshold. As discussed before, households with assets in excess of A^* will choose to save and accumulate (despite low marginal returns to accumulation) until they reach the point A_S where it becomes optimal to switch to livelihood strategy L_2 and to grow to a steady state level of capital, A_2^* . Households below this threshold will by definition not find it optimal to make the sacrifices needed to reach A_S . Absent access to intermediate capital, such households will thus revert to a steady state level of capital, A_1^* . Figure 4 portrays this scenario and its implication for asset dynamics. The top panel depicts the two distinct livelihood strategies of Figure 3, L_1 and L_2 . The bottom panel shows the asset dynamics that ultimately drive the system. Now we can better see how the critical threshold for poverty dynamics is neither \underline{A} , the static asset poverty line, nor A_S , the point at which households rationally switch from L_1 to L_2 in the static model, because while adoption of improved livelihood strategies is indisputably important, such choices are also reversible. Rather, the critical threshold is A^* , the unstable dynamic asset equilibrium, the threshold at which accumulation dynamics bifurcate. A household with initial wealth just above A^* will naturally accumulate assets, at some point pass A_S and switch from L_1 to L_2 , and ultimately settle at a long-term equilibrium asset stock of A_2^* , yielding steady state utility U_H^* above the income poverty line. By contrast, a household with initial wealth just below A^* will naturally shed assets down to A_1^* , never switch to the more remunerative livelihood strategy, and settle ultimately at an equilibrium welfare level of U_L^* , well below the income poverty line. Note that in this particular case illustrated in Figure 4 ($A_1^* < A^* < \underline{A}$), the structurally poor at any point time (those with assets below \underline{A}) can be divided into those who will be persistently poor ($A < A^*$) and those who will eventually surpass \underline{A} on their way to the high level equilibrium, A_2^* ($A^* \leq A < \underline{A}$).

While Figure 4 was drawn with $A_1^* < A^* < \underline{A}$, other configurations are possible. Adato et al. (this volume) estimate that $A_1^* < \underline{A} < A^*$ in South Africa. In this case, all the currently structurally poor, and a subset of the non-currently structurally poor would be expected to gravitate to the low-level equilibrium. These different cases

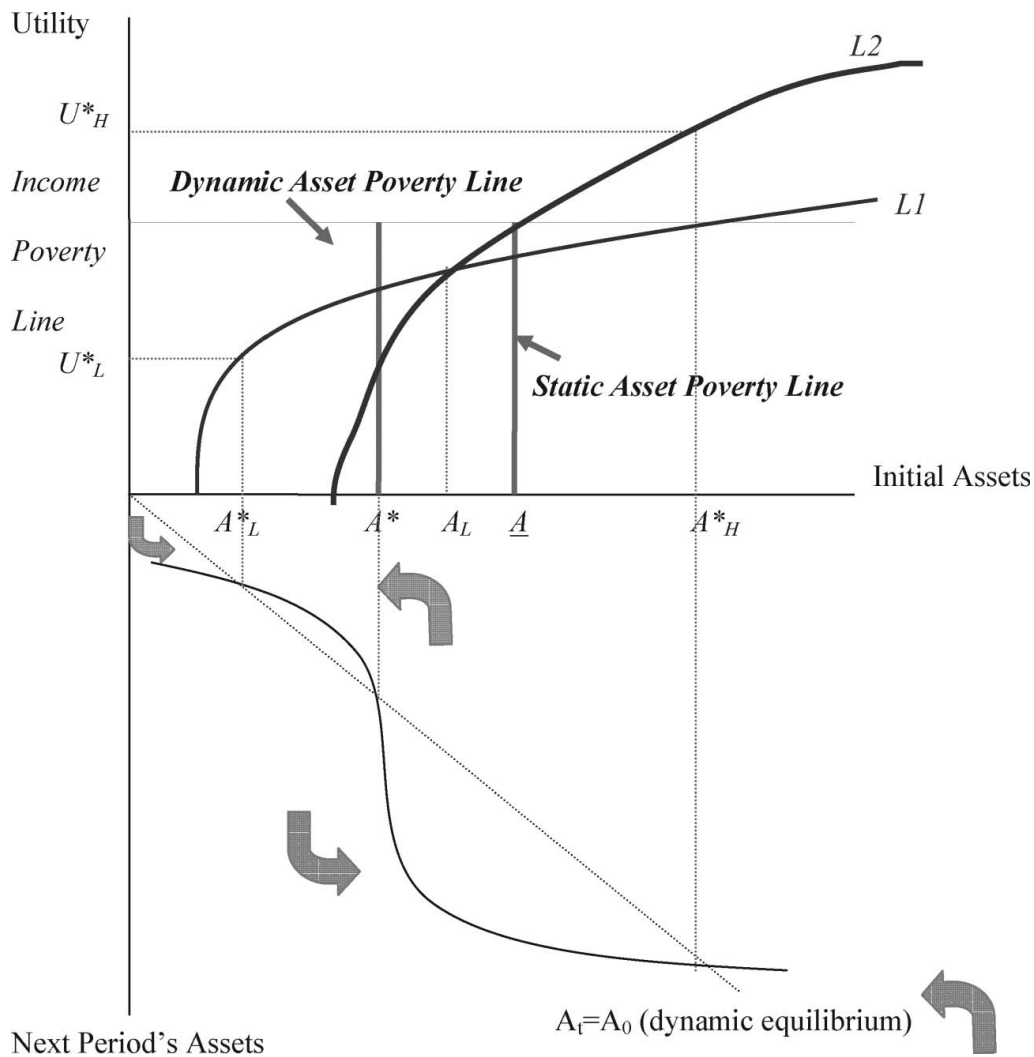


Figure 4. The dynamic asset poverty line

suggest the importance of poverty measures based on the Micawber threshold, something to which the next section now turns.

(c) Using the Dynamic Asset Poverty Threshold to Measure Chronic Poverty

The standard, money-metric poverty line is frequently criticised as an arbitrary construct which has no behavioural foundation.¹³ In contrast, the Micawber or dynamic asset poverty threshold is an empirical construct whose foundation is observed behaviour. Conceptually, the Micawber threshold can separate households expected to be persistently poor from those for whom time is an ally that promises better standards of living in the future. Poverty measures based on the Micawber threshold thus promise not only to distinguish stochastic from structural changes, but also to identify the long-run health of an economy as judged by its ability to facilitate growth in living standards amongst its least well-off members.

Analogous to the discussion in Section III, consider the following restatement of the FGT family of poverty measures around the Micawber threshold:

$$P_{\alpha}^{A^*} = \frac{1}{N} \sum_{i=1}^N I_i^{A^*} \left(\frac{A^* - A_i}{A^*} \right)^{\alpha} \quad (3)$$

where A_i is asset stock of household i and the binary indicator variable $I_i^{A^*} = 1$ if $A_i < A^*$ and reflects whether the household i 's asset stock falls below the dynamic asset poverty threshold. The dynamic asset poverty gap ($A^* - A_i$) indicates the asset transfer that is necessary to place a household in a position from which they can grow and sustain a non-poor standard living in the future. This class of dynamic asset poverty measures allows for the very real possibility that some of those who are presently structurally non-poor would be expected to decumulate assets over time and fall into poverty. This is precisely the sort of forward-looking measure that policymakers need in order to gauge the health of an economy and to target poverty reduction interventions appropriately. Analogous to other FGT-based measures, $P_{\alpha}^{A^*}$ will give a persistent poverty headcount measure for $\alpha=0$, a measure of the average transfers needed to eliminate persistent poverty for $\alpha=1$, and a distributionally sensitive measure of the severity of persistent poverty for $\alpha \geq 1$. In conjunction with the standard poverty measures, and the measures based on the asset poverty line discussed earlier, which describe conditions in the recent past and the near future, respectively, $P_{\alpha}^{A^*}$ measures offer a longer-term perspective on the likely evolution of well-being, and thereby flesh out the dynamics of well-being at the lower tail of the wealth distribution.

V. Empirical Strategies to Identify Poverty Dynamics and Critical Asset Thresholds

A very recent empirical literature has begun to test statistically for the existence of poverty traps.¹⁴ Unfortunately, much of this literature has taken its cue from the macroeconomic growth literature on convergence and often uses parametric methods that assume globally decreasing returns to scale to explore the dynamics of household income or expenditure. However, as the discussion here has made clear, poverty traps are defined by a threshold in *asset* space around which accumulation dynamics bifurcate and are defined by the existence of some range over which increasing returns might prevail. A household that suffered a temporary income shock that pushed it below the poverty line, but which did not degrade its asset base, would be expected to recover to its pre-shock level of well-being.¹⁵ That is, in the language of this paper, households that suffer stochastic income poverty transitions should not be expected to fall into poverty traps.

In contrast, a household that suffered a loss of productive assets (for example, a loss of assets that pushed it below the dynamic asset poverty threshold A^* in Figure 4) might indeed fall into a poverty trap. In short, without a firm grounding in an asset-based approach to poverty – which permits us to distinguish the dynamics of households that experience stochastic from structural transitions – we cannot test empirically for the existence of poverty traps.¹⁶ Furthermore, imposing assumptions of strict concavity (that is, globally decreasing returns), assumes away one of the key features for which one ostensibly wishes to test.

To date there has been no systematic development of empirical strategies to identify poverty dynamics and critical asset thresholds. This section briefly maps out key elements of the extant tool kit for exploring this exciting topic.

(a) Flexible Methods for Estimating Dynamic Asset Poverty Thresholds

Estimation of the sort of asset dynamics displayed in Figure 4 in order to test for the existence of a dynamic asset poverty threshold confronts two basic problems. First, not only is the relationship potentially highly non-linear, but also the dynamic asset poverty threshold is an unstable equilibrium, away from which households move over time. This means that we would expect few observations in the neighbourhood of the threshold itself in any data set and an unstable equilibrium can easily be mistaken for heteroskedastic errors (Barrett, 2005). The second problem is that most households possess a portfolio comprised of multiple assets. Estimation of asset dynamics must somehow deal with this dimensionality problem.

Lybbert et al. (2004) examine a pastoral population whose major productive asset is livestock. This feature of the economy they study makes the second basic issue, asset aggregation, relatively easy to solve.¹⁷ To solve the first problem, Lybbert et al. estimate livestock dynamics using a non-parametric kernel estimator. This estimator is sufficiently flexible to capture high-order non-linearities. In addition, because it is non-parametric, local curvature is estimated using nearby points, meaning that a local twist in the asset dynamics relationship is not overwhelmed by the weight of distant points, as might happen using parametric regression methods. Lybbert et al. find strong evidence of a dynamic asset poverty threshold as well as evidence that, as predicted, recovery from shocks depends fundamentally on whether or not the shock casts the household below that threshold. These results corroborate qualitative ethnographic research among the same subject population.

The asset aggregation issue is less easily solved in the case of more complex economies. Barrett et al. (this volume) and Adato, Carter and May (this volume) flexibly estimate asset aggregation weights using factor analysis or by regressing expenditure or other well-being measures on households' productive assets. As detailed in those papers, these approaches permit the creation of asset indices in which the weights can both vary over time and depend themselves on the presence or absence of complementary assets in the household's wealth portfolio. While the properties of these asset indices have yet to be fully worked out, they permit the authors to test for the presence of dynamic asset poverty thresholds in more complex economies, again using relatively simple non-parametric kernel or nearest neighbour estimators.¹⁸ Both papers find evidence of such thresholds.

(b) Directions for Future Analysis of Asset Dynamics

The bivariate non-parametric methods employed by Lybbert et al. (2004), Barrett et al. (this volume) and Adato, Carter and May (this volume) depend on two non-trivial econometric assumptions. First, these studies presume that all households in the same structural position all lie within the same accumulation regime. However, as the theory of poverty traps makes clear, households facing otherwise identical initial conditions may follow different accumulation trajectories if one enjoys better

capital or insurance access than the other. And households with equal access to finance may face quite different accumulation trajectories when they have different livelihood functions due to spatio-temporal variation in agro-ecological or policy conditions. The challenge is to find ways of separating households into distinct capital access and accumulation regimes conditional on underlying livelihood mappings, either via *ex ante* measurement or through the development of sufficiently flexible econometric methods. In addition, a secondary problem is to control for other factors may influence accumulation (for example, life cycle household savings patterns) that could be spuriously correlated with initial asset holdings.

In addition to these regime identification and prospective omitted relevant variables problems, these existing studies implicitly assume that unobserved household characteristics are uncorrelated with initial asset condition. If this assumption is incorrect, then these analyses are likely to confound the true structural state dependence of multiple equilibria models with what might be termed a spurious state dependence. The latter would result if households with initially low levels of assets remain in a structurally poor state, not because of barriers to accumulation, but because they share intrinsic characteristics (for example, low work ethic or a high discount rate) that place them in a low-level equilibrium ‘club.’

To resolve this identification problem, one would conceptually like to observe an experiment in which higher and lower asset households experienced fundamental changes in their asset positions (for example, through weather or other exogenous shocks that fundamentally altered asset holdings). True structural dependence would suggest that households with these exogenously altered asset positions would shift between low and high equilibrium positions. Under spurious dependence, households would be expected to return to their original equilibrium attractor point. In practical econometric terms, resolution of this problem is likely to require at least three periods of observations with significant random perturbations in the asset positions of households. As with the macro-growth literature reviewed above, significant additional work needs to be done before we can fully distinguish between club convergence and multiple equilibrium theories.

Short of resolving these econometric challenges, the study of poverty dynamics and the identification of critical asset thresholds are among the tasks best suited to mixing qualitative and quantitative methods (Hulme and Shepherd, 2003). Panel data can be used to stratify households for qualitative study via oral histories (for example, Barrett et al. this volume, Adato, Carter and May, this volume), participatory methods can be used to define poverty transitions then studied quantitatively using survey methods (Krishna, 2003; Krishna et al., 2004; Kristjanson et al., 2004; Krishna et al., this volume), or other means of sequential or simultaneous mixing of qualitative and quantitative methods can be effectively employed. Qualitative analysis can be especially valuable in identifying historical causes of structural transitions that predate initial surveys.

VI. Toward Persistent Poverty Reduction Strategies

Reformulating poverty analysis explicitly on an asset basis offers important advantages. Identification of the asset poverty line makes it possible to distinguish structural from stochastic poverty transitions. Identification of the dynamic asset

poverty threshold permits a further refinement of poverty measurement, making it possible to distinguish households likely to escape poverty over the longer term from those apparently mired in a poverty trap. Application of these structural or asset-based approaches to poverty should ultimately underwrite a more satisfying analysis of the contentious question of the impact of market-oriented liberalisation policies on long-term poverty dynamics.

While these measurement issues are important, perhaps the deeper value of an asset-based approach is that it offers three important policy insights. First, it permits us to determine whether there exists a minimum configuration of assets or economic conditions required for households to ultimately engineer their own escape from poverty. From this perspective, the asset-based approach we advance adds specificity to John Williamson's (2003) call for minimum asset bundle.

Second, the asset based approach promises insights into the enabling conditions necessary to assure that time is an ally of poor households. As discussed above, the existence of the threshold depends on the degree to which the household is excluded from intertemporal exchange through credit, insurance or savings, whether formally or through social networks. A household with perfect access to capital over time and across states of nature would not face a critical threshold. Such a household would always be able to access the funds needed to build assets so as to move onto a natural growth trajectory. Similarly, such a household could use (formal or informal) insurance relationships to protect its assets from shocks that might otherwise threaten its ability to generate a high rate of return and reach or recover to a non-poor equilibrium.

Third and finally, this asset approach has implications for the design and positioning of safety net policies. The arguments put forward here indicate that the long-term implications of shocks depend not on the absolute magnitude of the shock, but where a shock leaves the household *ex post*. Households that do not fall below the Micawber threshold would be expected to use time and markets to engineer an eventual recovery to a higher, non-poor equilibrium living standard. Households that fall below that threshold would not be expected to recover but, instead, to suffer a permanent deterioration in their position (Carter et al., 2005). Thus households' need for a safety net depends less on the magnitude of the shock they experienced – as it is usually conceptualised based on the standard economics of insurance – and more on their asset position *ex post* of a shock.

As policymakers and development practitioners increasingly turn their attention to the problem of persistent poverty and its correlates of hopelessness and polarisation, we researchers must adapt our analytical toolkits. In this paper, we have made a case for availing ourselves of emerging longitudinal data on households and individuals and framing poverty analysis using an asset-based approach that pays particular attention to the underlying, systemic dynamics of critical assets. This is the only way to distinguish deep-rooted, persistent structural poverty from poverty that passes naturally with time due to economywide growth processes. While empirical research that takes asset dynamics and the possibility of poverty traps seriously remains in its infancy, the promise of such research is considerable in its capacity to inform targeting of interventions, the identification of key enabling conditions to open up pathways out of poverty for the structurally poor, and the design of safety net programming.

Acknowledgement

The authors thank Harold Alderman, Felix Naschold, participants in the November 2004 USAID BASIS CRSP policy conference on Combating Persistent Poverty in Africa and an anonymous referee for helpful comments on an earlier draft. This work has been made possible by support from the United States Agency for International Development (USAID), through grant LAG-A-00-96-90016-00 to the BASIS CRSP, and the Strategies and Analyses for Growth and Access (SAGA) cooperative agreement, number HFM-A-00-01-00132-00. The views expressed here and any remaining errors are the authors' and do not represent any official agency.

Notes

1. Vulnerability-based approaches to poverty, which ask who is at risk of being poor in the future, share this forward-looking approach to poverty (for example, Dercon, 2004).
2. The sense that the development liberalism of the 1980s and 1990s came up short on reducing poverty has led to a similar search to understand the minimum asset positioning needed for households to constructively participate in the opportunities afforded by a market economy. John Williamson, who coined the term Washington Consensus that is often used to label the suite of liberalisation policies popular in the 1980s and 1990s, argued that governments must ensure that citizens have the minimum asset base and market access required to save, accumulate and succeed in a market economy (Williamson, 2003). Williamson suggests that without such assurance, some households will be trapped in poverty, unable to use time and markets to fundamentally improve their well-being or that of their children.
3. We owe this label to Michael Lipton (1993), who wrote of a 'Micawber threshold', below which it is difficult for agents ever to accumulate assets. The image echoes the Dickensian travails of Wilkins Micawber, the perpetually insolvent debtor with whom David Copperfield took up residence, who moves in and out of different jobs and debtor's prison, unable to advance until he encounters and ultimately exposes the evil Uriah Heep.
4. The magnitude of measured transitory expenditure or income poverty may also reflect the measurement error to which flow-based welfare measures are especially prone. Transitory poverty would be purely a statistical artefact of imprecise measurement when, for example, non-poor households are mis-measured as poor in one period, but correctly measured as non-poor in another period when nothing fundamentally changed between survey periods. Barrett et al. (this volume) show that measurement error and stochastic components to income data generating processes can completely mask structural patterns of income change over time.
5. Slightly more formally, the second generation approaches to poverty measurement cannot differentiate between stationary and non-stationary shocks to individuals' welfare.
6. None of these observations are meant to imply that the stochastically poor are somehow not 'really' poor. Instead, the message is that the transition of a stochastically poor person to a non-poor status conveys fundamentally different information about the economy than does the identical transition made by a structurally poor person.
7. The curvature of the livelihood mapping is itself interesting, as Carter and May (1999) and Finan, Sadoulet and de Janvry (forthcoming) discuss in detail.
8. In general, we would expect the livelihood function, and therefore the asset poverty line, to move as rates of return change due, for example, to price changes or to technological change that affects productivity. We address this possibility shortly, as illustrated by the dashed livelihood function, $\tilde{u}(A)$ in Figure 2.
9. See the account given in Romer (1994).
10. Empirical work has been cast almost exclusively in terms of income, not in terms of assets (capital stocks of various sorts). The primary exception has been the literature on 'green national accounts', which worries about depreciation of the stock of natural capital (that is, environmental resources) and the resulting sustainability of income levels as measured in the standard national accounts. Given that

- asset transactions overwhelmingly occur within rather than between countries, a large part of the asset changes that matter considerably at more micro (for example, household) levels of analysis do not matter at the macro level of nation states.
11. Note that private firms will also be in this low-level equilibrium trap. A big push based on public coordination is thus typically seen as necessary to cross the threshold and move toward a higher equilibrium.
 12. Note that the household restricted to *L1* would choose a lower steady state level because the marginal returns to further accumulation (given by the slope of the production function) do not warrant additional savings. As illustrated in Figure 3, households restricted to either technology accumulate assets only up to the point where marginal returns are equalised.
 13. Similar criticisms would also apply to the static asset poverty line, introduced earlier in this paper.
 14. See for example, Adato, Carter and May (this volume), Barrett et al. (this volume), Dercon (2004), Elbers et al. (2002), Jalan and Ravallion (2002, 2004), Lokshin and Ravallion (2002), Lybbert et al. (2004), Ravallion and Jalan (1996).
 15. The experience of graduate students who leave professional employment to go back to school offers an intuitive example from a very different context. The student's income typically falls sharply, often dropping the student and her family below the income poverty line, but her asset stock is preserved, even built up, enabling predictable, subsequent recovery to a non-poor equilibrium income level.
 16. In principle, the same comment could be made about the macroeconomic literature. Note that growth models ultimately concern the steady state levels of productive assets (capital), with steady state growth equal to the rate of technological change. However, in the case of nations, national output or income is a relatively stable index of the underlying level of productive assets as national income deviates relatively little from its expected value. At the micro-economic level, household income can depart far more significantly from its expected income and thus offers a far less reliable index of underlying assets.
 17. Following common practice in the study of livestock, Lybbert et al. (2004) aggregate heterogeneous livestock into 'tropical livestock units' using a generally accepted weighting system that permits sheep and goats to be aggregated with larger animals such as cattle and camels.
 18. Barrett et al. (this volume) also use parametric methods that yield qualitatively identical estimates of the critical asset threshold, but which fit the data far less well in the tails of the wealth distribution.

References

- Acemoglu, D., Johnson, S. and Robinson, J. A. (2001) The colonial origins of comparative development: an empirical investigation, *American Economic Review*, 91, pp. 1369–1401.
- Adato, M., Carter, M. and May, J. (this issue), Exploring poverty traps and social exclusion in South Africa using qualitative and quantitative data, *Journal of Development Studies*.
- Azariadis, C. and Drazen, A. (1990) Threshold externalities and economic development, *Quarterly Journal of Economics*, 105, pp. 501–26.
- Banerjee, A. V. and Newman, A. F. (1993) Occupational choice and the process of development, *Journal of Political Economy*, 101, pp. 274–98.
- Barrett, C. B. (2005) Rural poverty dynamics: development policy implications, in D. Colman and N. Vink (Eds) *Reshaping Agriculture's Contributions to Society* (Oxford: Blackwell).
- Barrett, C. B. and Carter, M. R. (2001–2) Can't get ahead for falling behind: new directions for development policy to escape poverty and relief traps, *Choices*, 16, pp. 35–8.
- Barrett, C. B., Marenja, P. P., McPeak, J. G., Minten, B., Murithi, F. M., Oluoch-Kosura, W., Place, F., Randrianarisoa, J. C., Rasambainarivo, J. and Wangila, J. (this issue), Welfare Dynamics in Rural Kenya and Madagascar, *Journal of Development Studies*.
- Baulch, B. and Hoddinott, J. (2000) *Economic Mobility and Poverty Dynamics in Developing Countries* (London: Frank Cass).
- Baumol, W. J. (1986) Productivity growth, convergence, and welfare: what the long-run data show, *American Economic Review*, 76, pp. 1072–85.
- Binswanger, H. and Rosenzweig, M. R. (1993) Wealth, weather risk and the composition and profitability of agricultural investments, *Economic Journal*, 103, pp. 56–78.

- Bloom, D. E., Canning, D. and Sevilla, J. (2003) Geography and Poverty Traps, *Journal of Economic Growth*, 8, pp. 355–78.
- Bloom, D. E. and Sachs, J. D. (1998) Geography, demography and growth in Africa, *Brookings Papers on Economic Activity*, 2, pp. 207–295.
- Canova, Fabio (2004) Testing for convergence clubs in income per-capita: a predictive density approach, *International Economic Review*, 45, pp. 49–77.
- Carter, M. R. (1997) Environment, technology and the social articulation of risk in West African agriculture, *Economic Development and Cultural Change*, 45, pp. 557–90.
- Carter, M. R., Little, P. D., Mogues, T. and Negatu, W. (2005) Poverty traps and environmental shocks: tracking the impacts of natural disasters on assets in Ethiopia and Honduras, working paper, University of Wisconsin.
- Carter, M. R. and May, J. (1999) Poverty, livelihood and class in rural South Africa, *World Development*, 27, pp. 1–20.
- Carter, M. R. and May, J. (2001) One kind of freedom: the dynamics of poverty in post-apartheid South Africa, *World Development*, 29, pp. 1987–2006.
- DeLong, J. B. (1988) Productivity growth, convergence, and welfare: comment, *American Economic Review*, 78(5).
- Dercon, S. (1998) Wealth, risk and activity choice: cattle in Western Tanzania, *Journal of Development Economics*, 55, pp. 1–42.
- Dercon, S. (2004) Growth and shocks: evidence from rural Ethiopia, *Journal of Development Economics*, 74, pp. 309–29.
- Dercon, S. and Krishnan, P. (1996) Income portfolios in rural Ethiopia and Tanzania: choices and constraints, *Journal of Development Studies*, 32, pp. 850–75.
- Easterly, W. and Levine, R. (1997) Africa's growth tragedy: policies and ethnic divisions, *Quarterly Journal of Economics*, 112, pp. 1203–50.
- Elbers, C., Gunning, J. W. and Kinsey, B. H. (2002) Convergence, shocks and poverty, Tinbergen Institute Working Paper 2002–035/2.
- Esteban, J. and Ray, D. (1994) On the measurement of polarization, *Econometrica*, 62, pp. 819–52.
- Feder, G. and Nishio, A. (1998). The benefits of land titling and registration: economic and social perspectives, *Land Policy Studies*, 15, pp. 25–43.
- Fiaschi, D. and Lavezzi, A. M. (2003) Distribution dynamics and nonlinear growth, *Journal of Economic Growth*, 8, pp. 379–401.
- Field, E. (2003) Entitled to Work: Urban Property Rights and Labor Supply in Peru Harvard University working paper.
- Finan, F., Sadoulet, E. and de Janvry, A. (forthcoming). Measuring the poverty reduction potential of land in Mexico, *Journal of Development Economics*.
- Foster, J., Greer, J. and Thorbecke, E. (1984) A class of decomposable poverty measures, *Econometrica*, 52(3), pp. 761–66.
- Galor, O. and Zeira, J. (1993) Income distribution and macroeconomics, *Review of Economic Studies*, 60(1), pp. 35–52.
- Gonzalez-Vega, Claudio (1977) Interest rate restrictions and income distribution, *American Journal of Agricultural Economics*, 59(5), pp. 973–6.
- Grootaert, C. and Kanbur, R. (1995) The lucky few amidst economic decline: distributional change in Cote d'Ivoire as seen through panel datasets, 1985–1988. *Journal of Development Studies*, 31, pp. 603–19.
- Hansen, Bruce (2000) Sample splitting and threshold estimation, *Econometrica*, 68, pp. 575–603.
- Hoddinott, J. (2003) Pathways from poverty in sub-Saharan Africa, International Food Policy Research Institute working paper.
- Hulme, D. and Shepherd, A. (2003) Conceptualizing Chronic Poverty, *World Development*, 31(3), pp. 403–23.
- Jalan, J. and Ravallion, M. (2002) Geographic poverty traps? A micro model of consumption growth in rural China, *Journal of Applied Econometrics*, 17, pp. 329–46.
- Jalan, J. and Ravallion, M. (2004) Household income dynamics in rural China, in S. Dercon (Ed.) *Insurance Against Poverty* (Oxford: Oxford University Press).
- de Janvry, A. and Sadoulet, E. (1993) Relinking agrarian growth with poverty reduction, in M. Lipton and J. van der Gaag (Eds) *Including the Poor* (Washington, DC, World Bank).

- Krishna, A. (2003) Escaping poverty and becoming poor: Who gains, who loses, and why? Accounting for stability and change in 35 North Indian villages, *World Development*, 32, pp. 121–36.
- Krishna, A., Kristjanson, P., Radeny, M. and Nindo, W., 2004, Escaping poverty and becoming poor in twenty Kenyan villages, *Journal of Human Development*, 5, pp. 211–26.
- Krishna, A., Lumonya, D., Markiewicz, M., Mugumya, F., Kafuko, A. and Wegoye, J. (this issue), Escaping poverty and becoming poor in 36 villages of Central and Western Uganda, *Journal of Development Studies*.
- Kristjanson, P., Krishna, A., Radeny, M. and Nindo, W. (2004) Pathways out of poverty in Western Kenya and the role of livestock, Working paper for the FAO/Pro-Poor Livestock Policy Initiative.
- Lipton, M. (1993) Growing points in poverty research: labour issues, ILO Symposium: Poverty: New Approaches to Analysis and Policy. Geneva: International Labour Organisation.
- Lokshin, M. and Ravallion, M. (2002) Household income dynamics in two transition economies, World Bank working paper.
- Loury, G. C. (1981) Intergenerational transfers and the distribution of earnings, *Econometrica*, 49, pp. 843–67.
- Lybbert, T. J., Barrett, C. B., Desta, S. and Coppock, D. L. (2004) Stochastic wealth dynamics and risk management among a poor population, *Economic Journal*, 114, pp. 750–77.
- Masters, W. A. and Macmillan, M. S. (2001) Climate and scale in economic growth, *Journal of Economic Growth*, 6, pp. 167–86.
- Mookherjee, D. and Ray, D. (2002) Contractual structure and wealth accumulation, *American Economic Review*, 92(4), pp. 818–49.
- Mogues, T. and Carter, M. R. (forthcoming), Social capital and the reproduction of inequality in socially polarized economies, *Journal of Economic Inequality*.
- Murphy, K. M., Schleifer, A. and Vishny, R. W. (1989) Industrialization and the Big Push, *Journal of Political Economy*, 97, pp. 1003–20.
- Myrdal, G. (1957) *Economic Theory and Under-developed Regions* (London: Duckworth).
- Nurkse, R. (1953) *Problems of Capital Formation in Underdeveloped Regions* (Oxford: Oxford University Press).
- Pritchett, L. (1997) Divergence, big time, *Journal of Economic Perspectives*, 11(3), pp. 3–17.
- Quah, D. (1993) Empirical cross-section dynamics for economic growth, *European Economic Review*, 37, pp. 426–34.
- Quah, D. (1996) Twin peaks: growth and convergence in models of distribution dynamics, *Economic Journal*, 106, pp. 1045–55.
- Quah, D. (1997) Empirics for growth and distribution: stratification, polarization and convergence clubs, *Journal of Economic Growth*, 2, pp. 27–59.
- Ravallion, M. (2001). Growth, inequality and poverty: looking beyond averages, *World Development*, 29(11), pp. 1803–15.
- Ravallion, M. and Jalan, J. (1996) Growth divergence due to spatial externalities, *Economics Letters*, 53, pp. 227–32.
- Ray, D., Duclos, J. Y. and Esteban, J. (2004) Polarization: concepts, measurement, estimation, *Econometrica*, 72(6), pp. 1737–72.
- Romer, P. M. (1994) The origins of endogenous growth, *Journal of Economic Perspectives*, 8(1), pp. 3–22.
- Rosenstein-Rodan, P. (1943) Problems of industrialization of Eastern and South-Eastern Europe, *Economic Journal*, 53, pp. 202–11.
- Rostow, W. W. (1960) *The Stages of Economic Growth* (Oxford: Oxford University Press).
- Sachs, J. D. and Warner, A. (1997) Fundamental sources of economic growth, *American Economic Review Papers and Proceedings*, 87, pp. 184–88.
- De Soto, H. (2000) *The Mystery of Capital: Why Capitalism Triumphs in the West and Fails Everywhere Else* (New York: Basic Books).
- Williamson, J. (2003) From reform agenda: a short history of the Washington Consensus and suggestions for what to do next, *Finance & Development*, (Sept.), pp. 10–3.
- Young, A. A. (1928) Increasing returns and economic progress, *Economic Journal*, 38, pp. 527–42.
- Zimmerman, F. J. and Carter, M. R. (2003) Asset smoothing, consumption smoothing and the reproduction of inequality under risk and subsistence constraints, *Journal of Development Economics*, 71, pp. 233–60.