Journal of Development Economics 126 (2017) 215–230

Contents lists available at ScienceDirect

Journal of Development Economics

journal homepage: www.elsevier.com/locate/jdeveco

Poverty traps, convergence, and the dynamics of household income[☆]

Raj Arunachalam[a], Ajay Shenoy[b,*]

[a] Bates White, L.L.C., Washington, D.C., United States
[b] University of California, Santa Cruz, M/S Economics Department, Rm. E2—455, 1156 High Street, Santa Cruz, CA 95064, United States

ARTICLE INFO

JEL classification:
O15
E24
D31

Keywords:
Poverty trap
Convergence
Income mobility

ABSTRACT

We design a new method to detect household poverty traps. We apply the method to a unique panel that follows rural Indian households over thirty years. We find no evidence of poverty traps. Most households had annual income growth of over 2%, and income mobility is high. We then design and apply a method to detect conditional convergence. We find that upper castes are converging to a level of wealth 3 times as high as disadvantaged castes.

1. Introduction

At the intellectual core of development economics, offered as metaphor in the age of "high development theory" (Krugman, 1994) and formalized ever since, is the unifying concept of the poverty trap: a self-reinforcing mechanism that causes poverty to persist (Azariadis and Stachurski, 2005). The neoclassical model of growth promises that all countries and all households, no matter how poor in the beginning, will be equally rich in the end. Models of poverty traps make no such promise. Even when equally productive and equally thrifty the poor may not catch up to the rich.

The best-known theories of poverty traps focus on entire economies. Theories of geography (Krugman, 1991), imperfect credit (Matsuyama, 2004; Quah, 1996), and coordination failure (Murphy et al., 1989) all try to explain global inequality—why India, for example, is poorer than the U.S. But another set of theories focuses on households. Theories of occupational choice (Banerjee and Newman, 1993), human capital (Galor and Zeira, 1993), and nutrition (Dasgupta and Ray, 1986) try to explain local inequality—why one family is poorer than another. Given that inequality within countries explains a large part of the global distribution of income (Bourguignon and Morrisson, 2002), the household poverty trap—if it exists—is no less important than the economy-wide poverty trap. But compared to the aggregate poverty trap, the household poverty trap has received less attention in empirical work.[1]

That may be because detecting a household poverty trap is hard. When household income is subject to large shocks—illness, failed monsoons, and sudden movements in crop prices—it is hard to tell whether poverty persists. Moreover, few panel surveys follow households for more than a few years, whereas a true poverty trap immiserates households for decades. Simple parametric tests for convergence, especially when run on short panels, may give misleading results.

This paper develops a method to detect household poverty traps and applies it to a unique set of household data. The method exploits a simple fact. A household just inside the threshold of a poverty trap is likely to suffer negative income growth; the trap pulls income back towards the low steady state. But a slightly wealthier household—one that has just escaped the trap—is propelled to a higher steady state. Thus at the threshold of the poverty trap, the probability a household suffers negative income growth decreases. By contrast, if households are converging to a single steady state the probability of negative income growth is always rising. By running simulations we show that the method finds poverty traps even when income is subject to shocks larger than those in our data. The method is not sensitive to the parameters of the simulation and can tolerate heterogeneity between households.

We apply the method to a unique panel that follows rural Indian households over thirty years. As the earliest source of credible microdata, rural India has at least historically served as the discipline's

☆ We are grateful to Charles Coop, Hugo Le Du, Nana Kikuchi, Martha Johnson, Jia De Lim, and Xuequan Peng for excellent research assistance on this paper. We want to thank Andrew Foster for giving us access to the data, and to the NCAER for helping us understand it. We would also like to thank Carlos Dobkin, Jon Robinson, David Weil, Oded Galor, and seminar participants at U.C. Santa Cruz, Michigan, and Brown for helpful suggestions. We are grateful for the attention and suggestions of the editor Chris Udry and two anonymous referees. We also gratefully acknowledge the support of NVIDIA Corporation, which donated a Tesla K40 GPU used in this research.

* Corresponding author.
E-mail address: azshenoy@ucsc.edu (A. Shenoy).

[1] Aside from the papers we discuss in detail below, some notable exceptions are Estudillo et al. (2013) Quisumbing and Baulch (2013), Krishna (2013), Kwak and Smith (2013), Michelson et al. (2013).

http://dx.doi.org/10.1016/j.jdeveco.2017.02.001
Received 27 April 2016; Received in revised form 21 December 2016; Accepted 3 February 2017
Available online 15 February 2017
0304-3878/

canonical example of an economy caught in a poverty trap (Bardhan, 1984). This dubious honor, together with India's sheer size, make it the perfect place to search for poverty traps. The length of the panel lets us test whether households stay trapped in poverty over decades and across generations, and whether poverty traps existed during the period of India's stagnation in the 1970s as well as the period of its rapid expansion afterwards.

We find no evidence that they do. At no level of income does the chance of negative growth significantly decrease. The result holds whether we apply the method to the period from 1969 to 1982, the period from 1982 to 1999, or the combined period from 1969 to 1999. It is of special note that we find no poverty trap in either of the two periods despite that the earlier was one of relative stagnation while the latter one of rapid growth. Our results suggest that neither stagnation nor growth in the overall economy left a subset of households languishing in a lower steady state. Instead the data suggest that wealth and income have broadly increased. Most households had income growth of over 1.1% from 1969 to 1982, and this rate accelerated to 2.6% from 1982 to 1999. Income mobility is high; over 60% of households in the bottom quartile of income in 1969 rise to a higher quartile by 1982. There is no evidence that the poor are more likely to suffer persistent negative income growth.

But the absence of poverty traps need not imply convergence. Some households, whatever their initial income, may hold a privileged place in society that lets them converge to a higher steady state. In other words, there may be conditional rather than unconditional convergence. We derive another simple test that detects whether households in one social group converge to a higher steady state than those of another.

In India the natural division in society is caste. We apply our method to three groups: members of the heavily disadvantaged Scheduled Castes and Tribes, members of what India calls the “Other Backwards Castes,” and members of upper castes. The test shows that upper castes converge to a higher steady state than backwards castes, who in turn converge to a higher steady state than scheduled castes. Compared to a household of a scheduled caste, a household of an upper caste can in the long run expect wealth nearly three times higher.

We make two contributions, one methodological and one empirical. Ours is hardly the first method proposed to detect a poverty trap. Quah (1996) looks at the bivariate density of national output and its fifteen-year-lag, taking density with two peaks as evidence of a poverty trap. Lybbert et al. (2004) trace out the relationship between past and current wealth to see whether this transition function crosses the 45° line more than once. Bloom et al. (2003) use maximum likelihood to test whether geography traps some countries in a low output regime. Carter and May (2001) and Carter and Barrett (2006) use deviations in consumption from that predicted by asset holdings to distinguish temporary from structural poverty. Bianchi (1997) proposes a nonparametric test for two peaks in the distribution of national output, while Vollmer et al. (2013) proposes a parametric test for mixtures of single-peaked distributions.[2]

We extend this literature in three ways. First, our method is simpler and less computationally intensive than previous methods, yet gives a formal test for poverty traps. Second, our method balances the flexibility of a nonparametric approach against the computational ease of a parametric approach. Such balance is ideal for detecting household poverty traps, which might be smaller than national poverty traps but can be sought in larger datasets. Finally, to our knowledge we are the first to not only propose a method but test its properties. Our simulations are grounded in theory and let us measure the power and size of our test.

Our second contribution is empirical. To our knowledge we are the first to look for poverty traps in a large household dataset that spans several decades. We construct a consistent measure of income from three waves of a national survey that was conducted in a country home to one-quarter of the world's poor. A growing literature has sought and failed to find much evidence of conditions that might cause a poverty trap—for example, high fixed costs or low returns to capital. But in the words of Kraay and McKenzie (2014), a direct test for the household poverty trap is impossible “until improved data becomes available.” Our panel is precisely the improved data needed for a direct test. Our results suggest the traditional theory of household poverty traps does not explain inequality in India.

The poverty trap, though central to development economics, has implications far beyond the field. Inequality in rich countries has recently seized the attention of economists from all fields of the profession (e.g. Chetty et al., 2014; Clark and Cummins, 2015; Piketty and Saez, 2003). By keeping the poor in poverty, a poverty trap perpetuates inequality and shuts down social mobility. In the U.S. and Europe, lawmakers and protesters alike worry that this is exactly what has happened in their countries.

The poverty trap in our model is phrased as a fixed cost that must be paid before a household (say, a farmer) can produce using a more advanced technology. But it could just as easily describe the up-front cost of tuition for a college degree. This poverty trap is familiar to economists who study social mobility in the U.S. Also familiar are the arguments we make about conditional convergence by caste, as they could apply just as easily to race or ethnicity in the U.S. As a result, the methods we develop could be applied to detect household poverty traps or conditional convergence in any country, be it rich or poor.

2. Defining and detecting a poverty trap

2.1. Setup

Consider the simplest of poverty traps: the need for a fixed capital investment (Quah, 1996; Banerjee and Duflo, 2011). The household can use either of two technologies, basic and advanced, both of which are Cobb-Douglas in capital and labor. The basic technology gives total income $Y_t = K_t^{\alpha}(A_t L_t)^{1-\alpha}$ or per capita income $y_t = k_t^{\alpha}(A_t)^{1-\alpha}$. The advanced technology is identical except the level of technology is scaled up by $\Omega > 1$. But in any year the household can only use the advanced technology if it makes a fixed investment F. For simplicity we assume the capital is not lost but tied up. For example, the household pays F to buy a power generator, which produces nothing but lets the household irrigate its farm with electric rather than hand pumps.

Given these options the household picks whichever earns higher income:

$$y_t^* = \max[k_t^{\alpha}(A_t)^{1-\alpha}, (k_t - F)^{\alpha}(\Omega \cdot A_t)^{1-\alpha}]$$

Aside from the fixed investment, all else is as in the Solow model. The law-of-motion is

$$k_{t+1} = sy_t + (1 - \delta)k_t$$

and the level of technology is

$$A_t = A_0(1 + g)^t.$$

Finally, output is subject to a Hicks-neutral productivity shock Z_t that is independent and identically distributed across time. Actual output is

$$y_t = e^{Z_t} y_t^*$$

The shock Z_t represents bad weather, illness, and other random events that cause household income to be higher or lower than implied by its level of capital.

Fig. 1 shows the steady state diagram for each of several combinations of the fixed cost F and the technology scalar Ω. The max operator in the production function creates a kink. This kink makes it possible for the

[2] There is a distinct but related literature that tests not for poverty traps, but for state dependence in the probability someone transitions into or out of poverty (see, for example, Cappellari and Jenkins, 2002, 2004). Though a poverty trap implies state dependence, the converse is not necessarily true. A case of a “poverty morass,” where poor households grow slowly at first but eventually catch up to the rich, would imply state dependence (at least in the short run) but is not a poverty trap. We consider this case in Section 2.