Poverty Traps, Convergence, and the Dynamics of Household Income

Raj Arunachalam*
Ajay Shenoy†

November 11, 2015
First Version: August 8, 2012

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 percent, 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.

JEL Codes: O15, E24, D31
Keywords: poverty trap, convergence, income mobility

*University of Michigan, Ann Arbor

†University of California, Santa Cruz; Corresponding author: email at azshenoy@ucsc.edu. Phone: (831) 359-3389. Postal Address: Rm. E2455, University of California, M/S Economics Department, 1156 High Street, Santa Cruz CA, 95064. 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 also gratefully acknowledge the support of NVIDIA Corporation, which donated a Tesla K40 GPU used in this research.

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, it is hard to find panel data that follows 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

[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).

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 served as the discipline's 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.

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. The data suggest that wealth and income have broadly increased. Most households had income growth of over 1.1 percent from 1969 to 1982, and this rate accelerated to 2.6 percent from 1982 to 1999. Income mobility is high; over 60 percent 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-degree line more than once. Bloom et al. (2003) use maximum likelihood to test whether geography traps some countries in a low output regime. Bianchi (1997) proposes a nonparametric test for two peaks in the distribution of national output.

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.

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:

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

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

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

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

$$y_t = Z_t y_t^*$$

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.

Figure 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 production function to cross the steady-state condition, represented by the dashed line, more than once. Each crossing is a steady state, though the middle one is unstable. A household in the region below the unstable steady state—either because it starts there or because a negative shock drops it there—will converge to the low steady state. Such households are in the poverty trap. A household with income above this region converges to the high steady state. Such households have escaped the poverty trap. The distance between the low steady state and the unstable steady state is a rough measure of the size of the poverty trap.

We use standard parameters for the elasticity of capital, the depreciation rate, and the investment rate: $\alpha = .3$, $\delta = .1$, $s = .2$. In the appendix we show that our results depend only on the size of the poverty trap, not the exact choice of parameters. For simplicity we assume zero population growth. We choose the level of initial technology A_0 to make (log) income in the low steady state one standard deviation below the true mean in our data. Throughout the main text we assume the rate of technological progress g is zero; we show in the appendix that a household poverty trap becomes almost irrelevant if there is technological progress, making the case against poverty traps even tighter.[2]

[2]The assumption that the investment rate is constant may seem strong, but letting the household choose its investment would make poverty traps even less relevant than our empirical results suggest. Moving from the low steady state to the high steady state permanently

The top left panel shows our baseline case, which sets the fixed cost at 75 percent above the low steady-state and makes the advanced technology five times more productive.[3] Income in the high steady state is roughly four times that in the low steady state—roughly the gap between the 80th and 20th percentile of income in the data at baseline. The top middle panel raises the fixed cost to 100 percent above the low steady-state and makes the advanced technology six times more productive. In this large trap, households in the high steady state earn 4.75 times as much as those in the low steady state.

For the top right panel we choose parameters that pull the steady states closer. The household can pay a fixed cost just 50 percent above the level of capital in the low steady state to use an advanced technology only 3.7 times more productive. Now income in the high steady state is only 2.5 times that in the low steady state, roughly the gap between the 70th and 30th percentile in 1969. Making the trap much smaller makes it harder to detect but also less meaningful. Given random shocks to income—an unusually good harvest, an unusually bad seling price—poverty traps become less meaningful as they get smaller. With ever higher probability, a household in the low steady state can earn a higher income than one in the high steady state.[4]

The bottom left panel has no poverty trap, but rather a "poverty morass." A household that starts with very little capital will be resigned to slow growth for several years until it passes the kink. Then its growth rate explodes until it draws near the steady state. In a dataset that follows households for just a few years this case may look like a poverty trap because many households cluster at low levels of income. But a decade later these households will have long since escaped the morass. This suggests that any careful search for poverty traps requires a long panel. Finally, the bottom right panel shows the case where $F = 0$, which is a simple Solow production function.

---

raises household income. A household in the low steady state will exploit any positive shock to invest a higher fraction of its income until it climbs out of the poverty trap

[3]Here, as in the empirical section, the units are 1960 Indian rupees.

[4]If households could choose their level of investment, a poverty trap this small would be meaningless even without shocks. In the low steady state a household has roughly 180 rupees of capital and saves roughly 18 rupees per year. By doubling its savings to 36 rupees for a single year it can afford the fixed investment of 216 rupees. This is not enough to get it out of the poverty trap in a single year because the fixed investment depreciates without adding to output. But by saving above the average for several years the household can climb out of poverty.

Figure 1
Steady State Diagrams

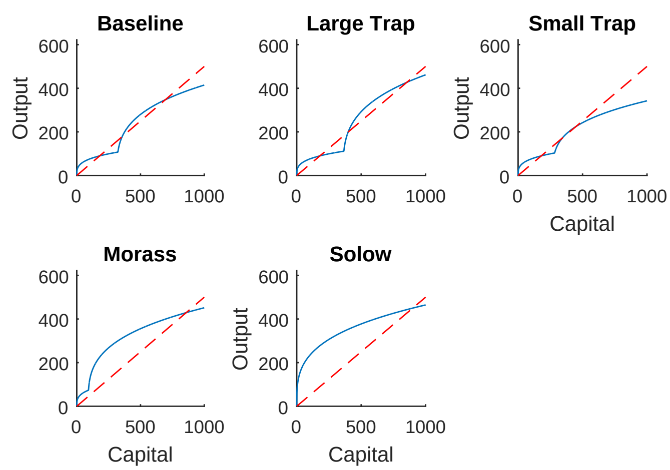

2.2 The Challenge of Detecting Poverty Traps

We use each of these production functions to create a simulated dataset that looks like our actual dataset. Each dataset contains 4000 households observed in year 1 (1969), in year 14 (1982), and year 31 (1999). We set the initial distribution of log income to be normally distributed with a mean and variance calibrated to match the actual distribution of household income in 1969, the first year in our dataset. Assuming there is no productivity shock in the first year, the initial distribution of income implies an initial distribution of capital. Since the initial distribution assumes no productivity shock, we set the following year to be 1969 (otherwise detecting the poverty trap is far easier than it would be in actual data).

Given that there are two equilibria the obvious sign of a poverty trap is an income distribution with two peaks. But the random shock Z_t may obscure this sign. Figure 2 illustrates this challenge using the baseline poverty trap. The plots on the left assume the standard deviation of the productivity shock is 0.1, roughly one-eighth the standard deviation of log income in 1969. The top left plot graphs growth from 1982 to 1999 against income in 1982. The two steady states are obvious; the observations cluster in two groups, each of which crosses

Figure 2
Phase Plots, Small and Big Shocks

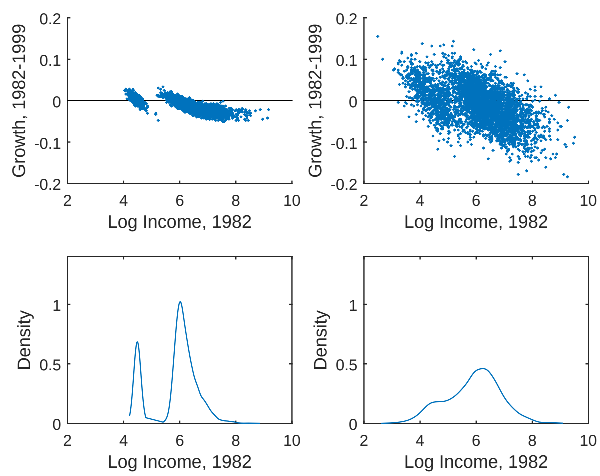

the horizontal axis (where growth is zero). The two peaks in the income distribution are visible in the kernel density estimate graphed in the bottom left plot.

The plots on the right run the same simulation assuming shocks have a standard deviation of 0.55, roughly two-thirds of the observed standard deviation in 1969 log income. It is harder to tell apart households in the low steady state from those in the high steady-state. It looks like the growth-income relationship crosses the horizontal axis only once, as it would if households were converging to a single steady state. Detecting the poverty trap in the kernel density estimates of the middle right plot is even harder, as the two peaks have merged into one.[5]

2.3 A Method to Detect Poverty Traps

Suppose there are a rich and a poor steady state, as in Panel A of Figure 3. The basin of attraction for the poor steady state ends at the orange dotted line while

[5]The bandwidth of the kernel density estimate is set exactly as Stata 13 sets its default bandwidth. That is, it computes the standard deviation and also the interquartile range divided by 1.349. Call the smaller of these two numbers M. If N is the number of observations, the bandwidth is set to $h = (0.9N^{-1/5})M$.

Figure 3

A Decrease in the Probability of Negative Income Growth Indicates a Poverty Trap

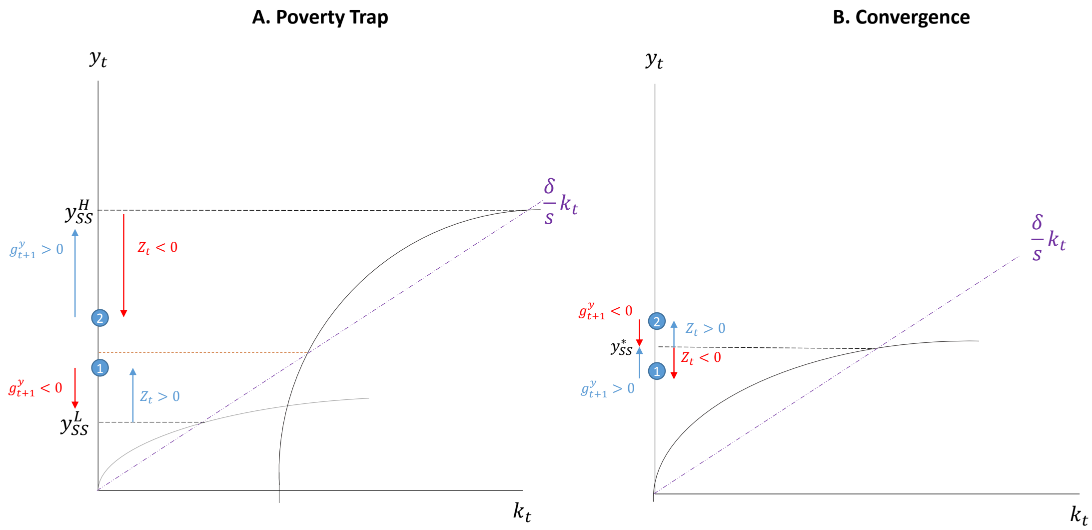

Note: The incomes of Household 1 and 2 are marked on the y-axis. In the case of the poverty trap (Panel A), Household 1 has been shocked above steady state (y_{SS}^L) and has a high probability of negative income growth between t and $t + 1$. Household 2 has been shocked below steady state and has a low probability of negative income growth. Thus the probability of negative growth as a function of current income decreases when income crosses the unstable steady-state (orange dotted line). By constrast, when there is no poverty trap (Panel B) the probability of negative growth is always increasing in current income.

that of the rich steady state begins. Suppose household 1 has the income given by the dot on the vertical axis—above the poor steady state but within its basin of attraction. The household likely landed above its steady state because it had a positive productivity shock Z_t. Since it has been shocked above steady state the household's income will likely have decreased when it is next observed at time $t + 1$. This is true as long as the household does not have an even larger positive shock, which becomes less likely the further it is above steady state. Thus, the probability any household has negative income growth is increasing as its income rises.

But this logic breaks down at levels of income above the orange line. Household 2 is likely in the high steady state, and thus has its current income because it suffered a negative shock. Since it is far below steady state its income is almost certain to rise between t and $t + 1$. That is, Household 2 has a lower probability of negative income growth than Household 1 even though Household 2 is richer. This is the direct effect of the poverty trap: the probability of negative income growth decreases.

There is no such decrease in the absence of a poverty trap. Panel B of Figure 3 shows the case of convergence. Since there is a unique steady state the probability of negative income growth is always increasing in current income. Household 1 is below steady state and thus expected to grow richer; Household 2 is above steady state and thus expected to grow poorer. Since all households have the same steady state, this logic is global: a richer households is always more likely to have negative income growth than a poor household. If at any point a richer household is less likely to have negative growth, it is evidence that there is another steady state and thus a poverty trap.

Making this intuition a formal test is simple:

1. For each household and each span of time (say, 1969 to 1982), define an indicator that equals one if income decreased and zero otherwise.

2. Discard outliers at the top and bottom of the distribution of initial log income (say, 1969 income). We discard the top and bottom 2.5 percent. Split initial income into J equally spaced bins. (We set $J = 10$.)

3. Compute the mean of the indicator for negative growth within each bin, and the standard error of the mean. This mean is a consistent estimator of the probability of negative income growth. (We compute the means and standard errors by regressing the indicator on a set of bin fixed effects.)

4. Compute the t-statistic for the difference between each mean and the mean in the next bin. That is, if within bin j we estimate the mean $\hat{\alpha}_j$ with variance $\hat{v}_j$, define the j^{th} statistic as

$$[D\hat{i}f]_j = \frac{\hat{\alpha}_j - \hat{\alpha}_{j+1}}{\sqrt{\hat{v}_j + \hat{v}_{j+1}}}$$

   Let $\hat{\lambda}$ be the largest (most positive) of these statistics.

5. Let P_{90} be the 90th percentile of the distribution of $\hat{\lambda}$ under the null hypothesis that all of the statistics $[Dif]_j$ are zero. There is evidence of a

poverty trap (at the 10 percent level) if $\hat{\lambda} > P_{90}$.[6]

Since $\hat{\lambda}$ is an order statistic its distribution is not normal. The following proposition gives the asymptotic distribution. The proof is in Appendix 1.1.

Proposition 1 *Suppose there are J bins and thus $J-1$ statistics. Define*

$$\mathcal{V}_j = \frac{v_{j+1}}{\sqrt{v_j + v_{j+1}} \cdot \sqrt{v_{j+1} + v_{j+2}}}$$

Let $\Phi_{J-1}(x_1, \ldots, x_{J-1})$ be the cumulative distribution function of a multivariate normal distribution with mean zero and variance

$$\Sigma = \begin{bmatrix} 1 & -\mathcal{V}_1 & 0 & \cdots & 0 \\ -\mathcal{V}_1 & 1 & -\mathcal{V}_2 & \cdots & 0 \\ 0 & -\mathcal{V}_2 & 1 & \ddots & 0 \\ \vdots & \vdots & \ddots & \ddots & -\mathcal{V}_{J-2} \\ 0 & 0 & 0 & -\mathcal{V}_{J-2} & 1 \end{bmatrix}$$

Then under the null hypothesis that $[Dif]_j = 0$ for $j = 1, \ldots, J-1$ the asymptotic distribution function of $\hat{\lambda}$ is

$$F(\lambda) = \Phi_{J-1}(\lambda, \lambda, \ldots, \lambda)$$

2.4 How Well Does the Method Perform?

Figure 4 shows an example of how the method works on data simulated from four of the five production functions illustrated in Figure 1. (We leave out the "large trap" case because it looks like the baseline case.) The small grey dots mark observations. The hollow circles mark the estimate of the probability of average growth (center circle) and the boundaries of the 95 percent confidence interval around the estimate (top and bottom circles). The estimates that give the largest t-statistic (those used to compute $\hat{\lambda}$) are marked in red. The p-values

[6]This is a conservative null hypothesis. Suppose there is a single steady state and household income is propelled away from it by shocks. Then the probability of negative growth is increasing with income, making each statistic negative. But exploiting this fact would require knowing how negative the statistics ought to be, which requires making an assumption about the relationship between growth and initial income in the absence of poverty traps.

for the test show that, at least in these realizations of the data, the method detects the poverty traps without mistaking the Solow production function or the poverty morass for poverty traps.

In Figure 5 we assess the method more systematically. We vary the standard deviation of the productivity shock from 0.4 to 0.9 (as compared to a total standard deviation of log income equal to 0.84 in 1969). For each value we produce 200 simulated datasets and record in what fraction of those datasets the method rejects the null of no poverty trap at the 10 percent level. We repeat the procedure for each of the five production functions illustrated in Figure 1.

The top panel shows the rejection rate for 1969 to 1982. In this earlier period, households have not yet converged to their steady state. The growth from convergence interferes with the mean reversion the method relies on, making it less powerful than it would be otherwise. The rejection rate for 1982 to 1999, shown in the top-middle panel, is higher at any standard deviation because households are close to their steady states.

The baseline and large traps are found with near certainty from 1969 to 1982 when the shock has a standard deviation less than 0.5. Given that the total standard deviation of income in 1969 is 0.84, this is a noisy shock indeed. The method does even better from 1982 to 1999, when at any standard deviation less than 0.6 both traps are found with near certainy.

Meanwhile, the rejection rate in the two cases where there is no poverty trap, the poverty morass and Solow convergence, show the probability of falsely rejecting the null hypothesis. Thus the rejection rate in these two cases traces out the size curve. The method almost never falsely detects a poverty trap.

The small trap is harder to detect, but is still found most of the time when the standard deviation is less than 0.5 in 1969 to 1982. The rejection rate is even higher from 1982 to 1999. At a standard deviation of 0.6, however, the small trap is found less than 20 percent of the time, and the rejection rate drops to zero when the standard deviation gets much bigger.

But though detecting such a poverty trap is more difficult, it is also less important. The bottom-middle panel computes the probability that, at least once over five years, a household in the low steady state will get an income shock large enough to let it earn as much income as it would in the high steady state. When the standard deviation has risen to 0.6 the small trap is crossed with a

Figure 4
Method Applied to Simulated Data

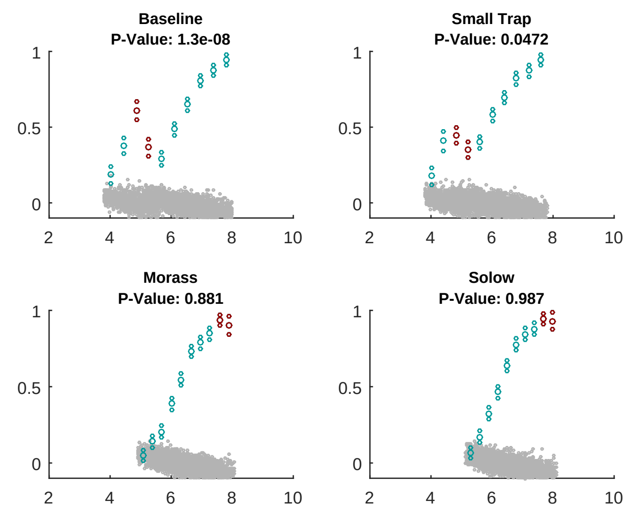

Note: Transition from 1982 to 1999.

probability of 30 percent. It is no surprise the method has trouble finding such a small trap when masked by such large fluctuations.

The bottom panel of Figure 1 shows what large shocks and a small trap imply for income mobility. The figure follows households that were in the bottom quartile of income in 1969, graphing the fraction that are still in the bottom quartile in 1999. A higher fraction implies a less income mobility. In the absence of a poverty trap—if there is Solow convergence or a poverty morass—the fraction is roughly 0.4 no matter how large the shock. If there is a poverty trap, this fraction is higher, though how much higher depends on the size of the trap and the size of the shocks. When the trap is small this fraction starts at just above 0.8 and falls to 0.6 by the time the trap becomes undetectable. A fraction of 0.6 implies 40 percent of those in the bottom quartile in 1969 have escaped it by 1999. Mobility in the presence of a larger trap—either the baseline case or the case of a large trap—is at a similar level when they become undetectable. As we show in Section 4, the true level of income mobility looks far more like Solow convergence than a poverty trap.

In summary, the simulations show three facts about using our method to detect poverty traps. First, studying a long horizon makes it more likely that households have converged to their steady states, and thus more likely the method will find the poverty trap. Second, large income shocks make the method less effective, especially when households have not yet converged to their steady state. And third, though large shocks reduce our power to detect a poverty trap, they also reduce its effect on income mobility.

Since the power of the method depends on the standard deviation of the shocks, we do a rough calculation of this statistic in our data (see the online appendix). Using income from 1969, 1970, and 1971 we compute it to be less than 0.4. Assuming the shocks do not get much larger in 1982, our power calculations suggest the test should easily detect a poverty trap.

Finally, in Appendix C we show that the method is robust to several of the simplifying assumptions made in this section. We show that assuming different values for the savings rate, depreciation rate, and production elasticity only affects the power of the method by widening or narrowing the gap between steady states. It is thus reasonable to focus only on the size of the gap, as we do here. We also show that allowing some heterogeneity in the location of the poverty trap—as would happen if households in the low steady state have different production functions—has only modest effects on the rejection rate.

2.5 Comparison to Other Methods

How does the negative growth test compare with other methods to detect poverty traps? The top panel of Figure 6 compares the negative growth test (solid lines) to the nonparametric multimodality test (dashed lines) of Bianchi (1997).[7] We apply his test to income in 1999 and compare its rejection rate to that of the negative growth test applied to growth from 1982 to 1999. We show only the cases of a small trap, which shows how often each method finds the most undetectable trap, and Solow convergence, which shows how often each method finds a trap that does not exist.[8]

[7] His test is an application of Silverman's (1981) test of multimodality to the distribution of income. The test finds the smallest bandwidth that makes the distribution of income unimodal, and rejects the null hypothesis of unimodaility if this bandwidth is large.

[8] Other work—for example, that of Lybbert et al. (2004), Quah (1996), and Lokshin and Ravallion (2004)—proposes ways to find evidence of poverty traps, but not formal tests. Bloom et al.

Figure 5
Power and Size Curves

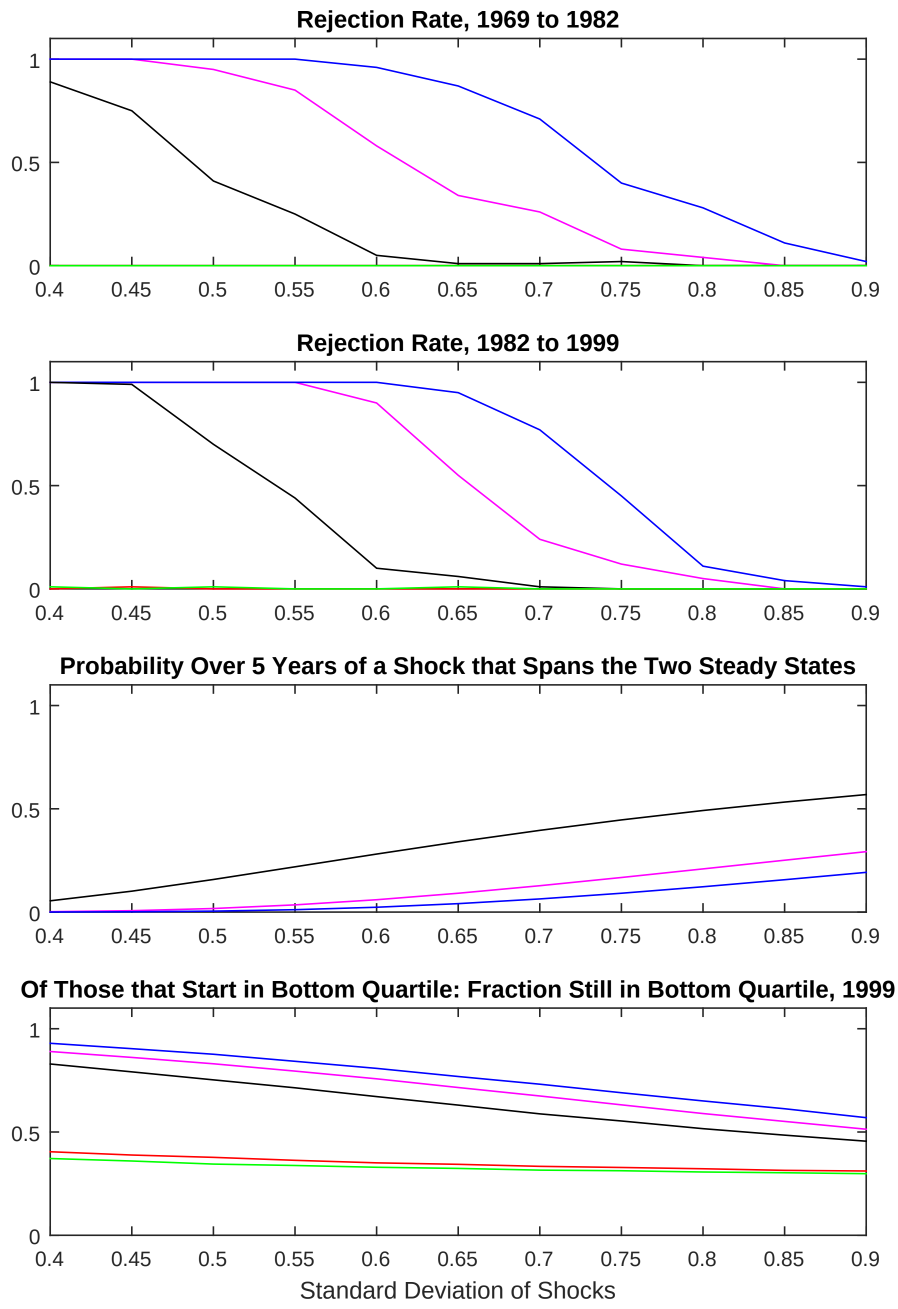

Legend: Baseline, Large Trap, Small Trap, **Morass**, **Solow**.

The top panel of Figure 6 shows that neither method finds a nonexistent trap more often than it should (the green lines), but the negative growth test is far more likely to find a real trap. This is not surprising, as Bianchi's test uses less information—income from only a single year. This feature is ideal when only a single year is available; otherwise the negative growth test is more powerful.[9]

The way we implement the negative growth test—splitting households into bins to compute average growth rates—is not the only way. The binning approach is a simple quasi-parametric test for monotonicity. How well does it compare with a nonparametric approach? We apply Chetverikov's (2013) test, which nests several other tests for monotonicity, to the probability of negative income growth from 1982 to 1999.[10] We graph the resulting power curve alongside that generated by our method in the bottom panel of Figure 6. As before, the black lines show the case of the small trap while the green lines show the case of Solow convergence. The binning approach (solid lines) and the nonparmetric approach (dashed lines) have nearly identical power and size. The nonparametric test, which takes roughly 20 minutes, does only marginally better than the binning approach, which takes less than one second.[11]

2.6 Conditional Convergence

An absence of poverty traps does not imply convergence—or rather, it does not imply unconditional convergence. If some households face disadvantages

(2003) do propose a test using maximum likelihood, but it is only consistent under certain assumptions about the error term, and requires much fine-tuning to ensure the maximum of likelihood function is found. The additional assumptions make it hard to compare their method with ours, and it is not clear whether the fine-tuning can be automated for a simulation. Given these issues, we compare our method only to that of Bianchi (1997).

[9]The method of Quah (1996) may be regarded as a multimodality test that uses two years of income. But Quah does not propose a formal test, and it is not obvious how to extend Silverman's method to a bivariate density. For example, it is not clear what the testing statistic would be, as a bivariate density has two bandwidths.

[10]Chetverikov's method computes a statistic that is negative when applied to an increasing function but positive when applied to a decreasing function. The algorithm searches over the domain of the independent variable for a region and a weighting function that maximizes the testing statistic (a region where the function is decreasing), then applies a wild bootstrap to compute the distribution of the testing statistic. We apply his one-step method, which seems an acceptable compromise between time-to-compute and power.

[11]This is how long it took Matlab to compute each measure on Ajay's computer, which has a 3.4 gigahertz processor and 32 gigabytes of memory.

Figure 6
Comparison to Other Approaches

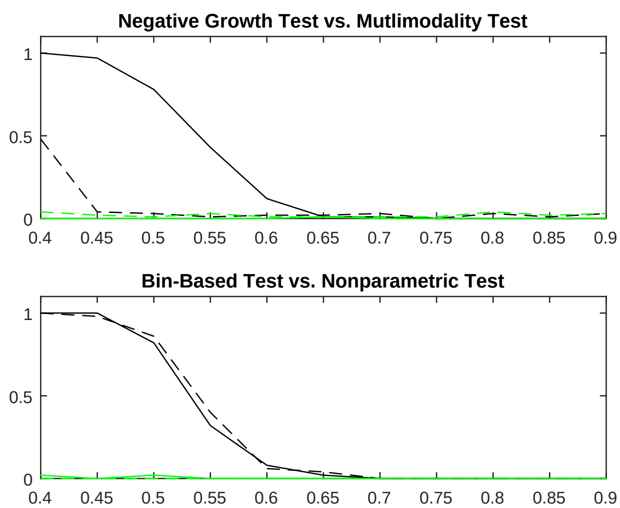

Note: **Top panel—** The solid lines show the rejection rate of the negative growth test applied to growth from 1982 to 1999. The dashed lines show the rejection rate of Bianchi's (1997) multimodality test applied to income in 1999. The black lines show rejection rates for the small poverty trap and the green lines for the case of Solow convergence. **Bottom panel—** As in the top panel, black lines show rejection rates for the small trap and green lines the rate for Solow convergence. The solid lines show rejection rates for the negative growth test using our binning method; the dashed lines show the negative growth test using the nonparametric test of Chetverikov (2013).

Figure 7
Conditional Convergence Implies
Multiple Growth-Capital Relationships

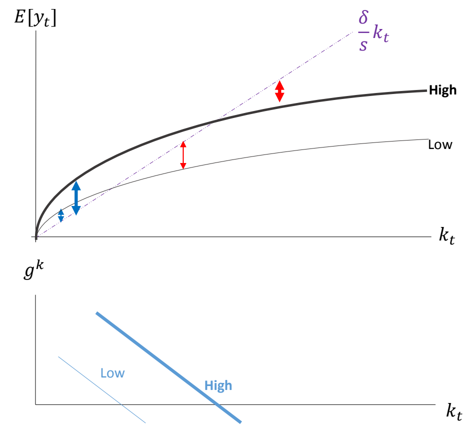

beyond their current income—if society discriminates against their caste or creed—they will converge to a lower steady state. To adapt the model of Section 2.1 for conditional convergence, suppose there is no fixed cost for the advanced technology (and no poverty trap), but some fraction of households have a low level of initial technology A_0^L. The favored group would have a higher level $A_0^H > A_0^L$. Then the favored group would converge to a steady state income y_{SS}^H higher than that of the disfavored group y_{SS}^L.

Figure 7 shows why the negative growth test of Section 2.3 would have trouble detecting such conditional convergence. The negative growth test relied on the current income of the household as being informative about its steady state. But as Figure 7 shows, two households can have the same level of income and still converge to different steady states. Unless the two steady states are far apart, there is no point at which the probability of negative growth abruptly reverses.

But Figure 7 suggests that if the sample can be split into the favored and disfavored group it is easy to identify their steady states. The favored group with the high technology will grow faster at any level of capital, meaning the growth-capital relationship is higher at every level of capital. Since the point at which the growth-capital relationship crosses the horizontal axis—the point at which growth equals zero—is the steady state, knowing the relationship is the same as knowing the steady state.

This is the key to our test for conditional convergence. Since we have data on income and not "capital," whatever its definition in this context, we study the growth-income relationship. The point where this relationship crosses the horizontal axis gives not steady state capital but steady state income. Given any two groups we test for whether they converge to the same steady state:

1. For each group (say, high caste versus low caste) estimate a linear regression of growth over a span of time (say, 1969 to 1982) on initial income (1969). Let $\hat{\beta}^H, \hat{\beta}^L$ be the coefficient vectors of the two regressions and $\hat{V}^H, \hat{V}^L$ their estimated variance matrices.

2. Compute the steady state income for the favored group as $\hat{y}^H_{SS} = -\hat{\beta}^H_0/\hat{\beta}^H_1$, where $\hat{\beta}^H_0$ is the intercept and $\hat{\beta}^H_1$ the slope of the regression. Follow analogous steps to compute the steady state $\hat{y}^L_{SS}$ for the disfavored group.

3. Let $\hat{J}^H = [-1/\hat{\beta}^H_1, \hat{\beta}^H_0/(\hat{\beta}^H_1)^2]$ be the Jacobian of the steady state. By the Delta method, $\hat{v}^H_{SS} = \hat{J}^H\hat{V}^H(\hat{J}^H)'$ is a consistent estimator for the variance of the estimated steady state $\hat{y}^H_{SS}$.

4. Form the testing statistic

$$\hat{\kappa} = \frac{\hat{y}^H_{SS} - \hat{y}^L_{SS}}{\sqrt{\hat{v}^H_{SS} + \hat{v}^L_{SS}}}$$

5. The null hypothesis is that the steady state of the favored group is no higher than that of the disfavored group. Let Φ_1^{-1} be the inverse distribution function for a (univariate) standard normal random variable. Reject the null at the 10 percent level if $\hat{\kappa} > \Phi_1^{-1}(0.9)$.

Figures 8 and 9 apply the negative growth test and the multi-line test to a simulated set of data generated with $A^H_0 = 2$ and $A^L_0 = 1$. Figure 8 confirms that the negative growth test cannot detect the two steady states. But Figure 9 shows that the multi-line test has no trouble finding them. For income growth from both 1969 to 1982 and from 1982 to 1999 the test overwhelmingly rejects that both groups converge to the same steady state.

Figure 8
Negative Growth Test Does Not Detect Conditional Convergence

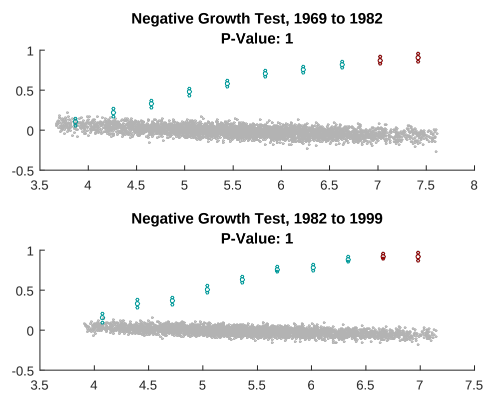

Figure 9
Multi-Line Test Does Detect Conditional Convergence

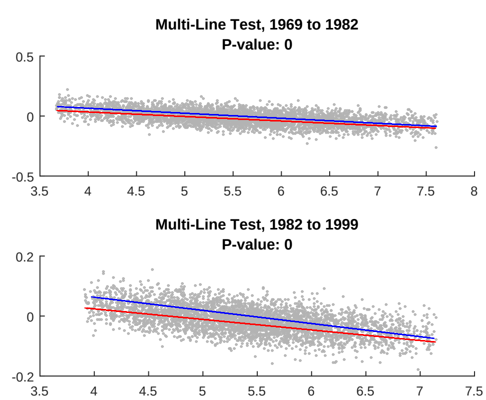

3 Data

3.1 Description of the Survey

The data we use are particularly suited to the inquiry: a nation-wide panel that follows rural households in a developing country over three decades. We are aware of no other such resource. The closest alternatives are a small sample from six ICRISAT villages beginning in the mid-1970s (Naschold, 2009; Dercon and Outes, 2009), and the long-term study of the village of Palanpur since the 1950s (Himanshu and Stern, 2011), both from India as well. Neither covers as many households over such a wide region as our data.

In the late 1960s the National Council of Applied Economic Research (NCAER) began a panel study of rural households. Roughly 250 villages in over 100 districts were sampled to be representative of India's rural population in 17 major states. From these villages an initial sample of 4500 households were surveyed. Of these, 4111 were found and surveyed across the crop years 1968-1969, 1969-1970, and 1970-1971. This Additional Rural Incomes Survey (ARIS) provides an array of information about income and its sources.

In 1982, the Rural Economic Development Survey (REDS) found and resurveyed roughly 70 percent of the original sample. The splitting of some original households and the inclusion of a small additional random sample raised the 1982 sample to just under 5000 households. In 1999, a second round of REDS revisited the households surveyed in 1982, excluding those in Jammu and Kashmir due to ongoing conflict, and again added a random sample of new households, bringing the sample to almost 7500 households. As one of the first large panels of household data from a developing country, ARIS-REDS has long been a valuable resource for researchers (for example: Rosenzweig and Wolpin, 1980; Foster and Rosenzweig, 1995, 1996; Behrman et al., 1999; Foster and Rosenzweig, 2002).

Taken together, the three rounds track households over thirty years. By contrast, the longest panel considered in McKay and Perge (2013) is a panel of Ugandan households from 1992 to 1999. Lybbert et al. (2004) use retrospective data on the herds of 55 pastoral households, who recalled the sizes of their herds over 17 years. Most studies of household poverty traps use panels of similar length or shorter, and sometimes rely on households to remember their income many years in the past rather than measuring their current income at different points in time.

Likewise, the coverage of our dataset—thousands of households and hundreds of villages—makes it larger than any similar panel from a developing country. This is important because, as Barrett and Carter (2013) note, households within a village may be receive common shocks to their income, making inference difficult. Our panel is wide enough to avoid this problem.

3.2 Our Measure of Income and Wealth

Unlike in rich countries, where most people get their income from a single paycheck, measuring income in an Indian village is not straightforward. Households earn income from several sources, and their main source is usually a farm or business. Since different sources pay out at different times the household may not ever compute its annual income. Any self-report of annual income cannot be trusted.

Instead, we define our own measure of income. Each round of the survey asks about the revenue and cost of each crop grown, each business run, each herd raised, and more. Given the precision of these questions, households are more likely to answer them accurately. The household may not know its total earnings for the year, but it probably knows the value of its rice harvest.

Defining our own measure also helps ensure the components of income stay fixed across rounds of the survey. If one round includes income from beekeeping in its measure of income while the next round does not, households that keep no bees—likely poorer households—would falsely appear to catch up with households that do. The omission would make it seem rich households had lost part of their income. Making income consistent takes many steps. Several forms of imputed income from family labor must be added or subtracted from 1999 income to make it consistent with the earlier rounds. We list the components of income in more detail in the data appendix. Given the complexity of a poor household's balance sheet, it is not clear what the ideal measure of income is, let alone whether our definition matches it. But since our aim is to follow household income across many decades, what matters most is consistency.

Nevertheless, these precautions may not remove all measurement error. To address this problem we confirm that our results hold not only for income but wealth. We define wealth as the value of buildings, land, farm equipment, animals, non-farm business assets, farm and non-farm inventory, consumer durables, cash and non-cash savings, and the value of loans owed to the household minus loans owed by the household. We compute wealth only for 1982 and 1999 because the 1969-1971 data lack the information we need to measure wealth consistently. Aside from being more acurately measured—a household is unlikely to forget how much land it owns—wealth is also subject to smaller shocks. This makes wealth a valuable check on the results, as Section 2 shows that large shocks reduce our power to detect a poverty trap. It is reassuring that the results using both income and wealth are similar.

In all cases, our measures are per capita—that is, we divide by the number of people in the household.

3.3 What is a Household?

In a thirty-year panel the answer is not obvious. One definition, as defined in the survey, is a group of people who live and eat together under a single head of household. Under this definition, a single household in 1969 may become three households in 1982 if two children grow up and move out with their families. These three households may further divide (or combine) before the next round of the survey. We assign each descendent household the income of its antecedent—all three 1982 households are assigned the income of the 1969 household from which they split.

Another definition imposes that all of these descendent households are part of the same household, which we call a dynasty. A dynasty is defined as all the members of all the households descended from a particular 1969 household. We compute the dynasty's per capita income in later rounds by summing the income of all descendent households and dividing by the total number of people in these households. The education of the head is defined as the highest education attained by any descendent head.

We use both definitions in our analysis to ensure the results are robust. Thus we run negative growth tests on income from 1969 to 1982 and from 1982 to

Table 1
Sample Sizes for Negative Growth Tests

Households	
Income, 1982 to 1999:	5930
Income, 1969 to 1982:	2849
Wealth, 1982 to 1999:	5985
Dynasties	
Income, 1969 to 1982:	2240
Income, 1982 to 1999:	2240
Wealth, 1982 to 1999:	2231

1999; we also run the test on wealth from 1982 to 1999. We run each of these three tests on both households and on dynasties. Table 1 gives the sample size for each test.

3.4 Attrition

Attrition is inevitable in a thirty-year panel. The overall rate of attrition in our panel of dyansties is roughly 46 percent. This may seem high, but it implies an average attrition of just 2 percent per year compounded over many years.

Aside from its usual causes, such as migration, some attrition is caused by the survey design. First, political violence after 1971 made it too dangerous to resurvey villages in the states of Assam and Jammu and Kashmir. Second, the NCAER took an unusual approach to following households in the 1982 round. If the original 1969-1971 household remained intact—regardless of whether the original head of household were still alive—it was found and resurveyed. Likewise, if the household had split but the original head of household were still alive, each of the descendent households were resurveyed. But if the original head had died and the household had split, the descendent households were not resurveyed. Instead, they were randomly replaced from the pool of all households in the village that had split after the 1969 head had died. This practice was only followed in 1982—in 1999 all original 1982 households were sought and, if found, resurveyed.

Together with migration, these patterns could cause differential attrition. Table 2 looks for whether income quartile or the level of education predicts attrition. We regress an indicator for leaving the panel between, say, 1969 and

Table 2
Attrition of Households and Dynasties

	Households 1969-1982	Households 1982-1999	Dynasties Any Round
Income Quartile:			
-Lowest	-0.019	**0.056**	**-0.080**
	(0.021)	**(0.017)**	**(0.023)**
-Mid-Low	**-0.047**	0.022	-0.005
	(0.021)	(0.017)	(0.023)
-Mid-High	-0.020	-0.010	0.015
	(0.021)	(0.017)	(0.022)
Head's Schooling:			
-Illiterate	-0.007	-0.055	0.018
	(0.061)	(0.042)	(0.066)
-Primary or Below	0.011	-0.054	0.015
	(0.062)	(0.044)	(0.067)
-Pre-Matric	-0.024	-0.065	-0.016
	(0.065)	(0.044)	(0.070)
-Matric to University	0.016	-0.020	0.017
	(0.068)	(0.046)	(0.073)
Reference Group	0.331	0.254	0.459
	(0.061)	(0.041)	(0.065)
Observations:	4110	4740	4110

Note: The columns report the coefficients from linear regressions of a dummy for attrition on dummies for income quartile and level of schooling. Statistically significant coefficeints are **bold**. The reference group is households in the highest income quartile whose head has had some tertiary schooling.

1982 on 1969 schooling and income quartile. (Since all of our regressors are dummy variables, a linear regression will predict a probability of attrition that lies between zero and one.) The first two columns show attrition in our dataset of households. The first column shows that households in the second income quartile in 1969 were less likely to be dropped from the 1982 survey (relative to top quartile). But more troubling is the second column, which shows that from 1982 to 1999 the poorest households were nearly 6 percentage points more likely to attrit than the richest.

However, Column 3 suggests a solution. Column 3 uses the panel of dynasties to compute the probability a dynasty does not appear in either 1982 or 1999. This regression shows that the poorest dynasties are less likely to attrit than the richest. The difference between the panel of households and dynasties is likely selective migration. When a poor household splits, it may not have enough land to divide among all the descendants. Instead, some parts of the household may migrate in search of work. However, some branch of the dynasty is more likely to stay behind compared to the richer quartiles.

Since the household and dynasty panels show opposite patterns of attrition, they serve as useful robustness checks on one another. Finding similar results in both would suggest the results are not driven by differential attrition. We also show in the appendix that reweighting observations to account for both the design of the survey and the predicted probability of attrition does not change the results.

4 Wealth and Income, 1969 to 1999

Before applying the tests derived in Section 2 we present simple graphs to show how wealth and income have changed. These graphs illustrate a fact that the formal tests confirm: income has risen for nearly everyone, and the poor do seem to catch up to the rich. Caste, however, continues to divide society.

Figure 10 shows the cumulative distribution of (log) income in 1969, 1982, and 1999. From 1969 to 1982 the bottom of the distribution shifted up even as the top stayed the same. The poorest members of society grew richer even as the rich made no progress. Moving from 1982 to 1999, however, the distri-

Figure 10
The 1999 Distribution of Income
Nearly Dominates the Earlier Distributions

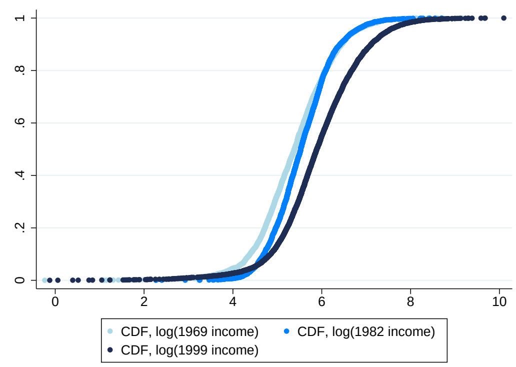

bution shifts outward nearly everywhere; indeed, the 1999 distribution nearly nominates the 1982 distribution. Households were slightly more likely to be very poor in 1999, but the mass shifted to the lower tail is small compared to the overall shift to higher levels of income.

Figure 11, which plots the density of log income in each year, makes this even clearer. The density in 1999 has shifted well to the right of the density in the previous years. Figure 11 also shows little evidence of two peaks in the distribution of income, as might be expected if there were a poverty trap.

Figure 12 plots kernel density estimates of annual income growth from 1969 to 1982 and 1982 to 1999 in the panel of households. (The densities look nearly identical for the panel of dynasties.) The numbers indicate the median rate of growth for each transition.

The median growth rate from 1969 to 1982 is only 1.1 percent; nevertheless, it shows that most households did earn at least 15 percent more in 1982 than they did in 1969. The median growth rate more than doubles in the period from 1982 to 1999. Most households grew by at least 2.6 percent per year, making their income at least 56 percent higher in 1999 than in 1982. In the panel of dynasties, the median household's income grew by 81 percent between 1969 and 1999. The figure also plots the density of growth in wealth from 1982 to 1999, the only years for which wealth can be measured. Wealth grew even more

Figure 11
No Sign of Two Peaks in the Density of Income

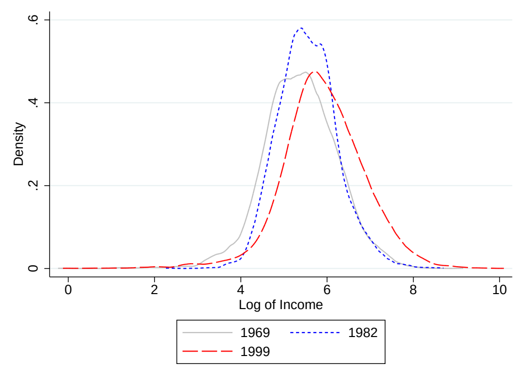

rapidly than income; most households had 89 percent more wealth in 1999 than in 1982.

The consequences of growth are clear in Figure 13. We draw a line at the median income in 1969 and measure what fraction of dynasties below that line in 1969 are still below it in later years. The bar marked "Below" represents all dynasties with income below the line in 1969 (the top left figure) and 1982 (the top right figure). The bottom part of the bar shows what fraction of those households were still below the line in 1982 (top left) and 1999 (top right).

Clearly, it is easier for households that start above the line to stay above it than for households that start below the line to cross it. Neither overall growth nor convergence can completely erase the disadvantage of poverty. But during the 1982 to 1999 transition over 60 percent of households that started below the line did cross it. The bottom figure shows that of the households that started below the line 1969, 65 percent had crossed it by 1999. It is hard to find evidence that a large part of the population has been left with stagnant income.

Yet not all households had positive income growth during both transitions, and roughly 20 percent had negative growth during both. Is it possible that such persistent negative income growth is concentrated among the poor? Figure 14 compares the income distribution of households with persistent negative growth, defined as negative growth from both 1969 to 1982 and from 1982 to 1999, to that of the other households. If anything it is richer households that

Figure 12
Income and Wealth Grow for Most Households

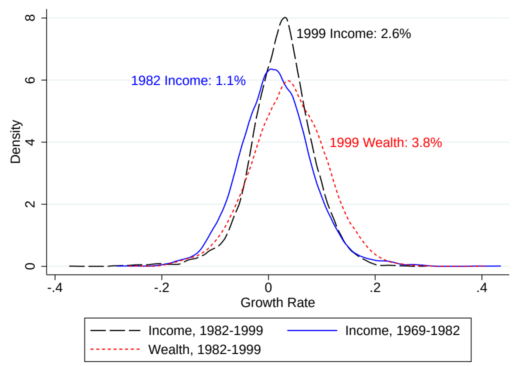

Note: The numbers give the median of each density.

Figure 13
Few Households Below the 1969 Median are Still Below it in 1999

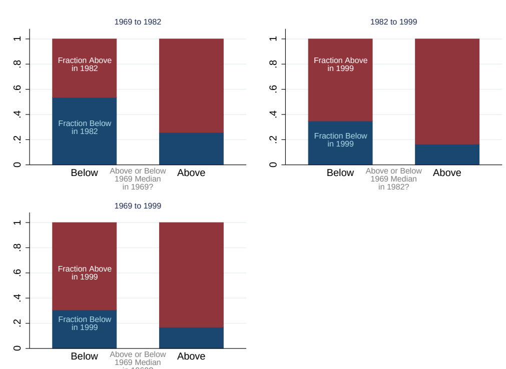

Figure 14
Persistent Negative Growth is Not Concentrated Among the Poor

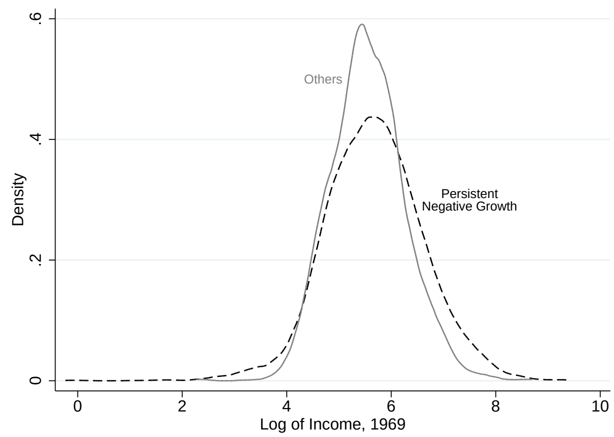

Note: This graph uses the dynasty dataset. We say a household has persistent negative growth if income growth is negative both from 1969 to 1982 and from 1982 to 1999.

suffer persistent negative growth.

The figures shown thus far make it clear that nearly every household and dynasty has grown richer. But are the poor catching up to the rich? Figure 15 suggests that they are. We use a kernel-weighted moving average to show the nonparametric relationship between initial income and income growth. For both periods the relationship is negative and almost linear; the income of the poor grows faster than that of the rich. The result applies equally to growth in wealth from 1982 to 1999. Figure 15 also shows that the result holds for dynasties. It estimates the nonparametric relationship between income in 1969 and income growth over the entire thirty-year sample. The poorest dynasties in 1969 grew the fastest from 1969 to 1999.

Following Quah (1996), Figure 16 graphs the bivariate density of log income at the beginning and end of each period. Unlike Quah's sample of countries, our sample of rural Indian households does not seem to have two peaks in these bivariate distributions. The same holds for wealth. The presence of two peaks might be evidence that one group of households is stuck in a poverty trap. Its absence makes it less likely any such trap exists. Though a rich household is of course more likely to stay rich, most of the distribution lies above the 45-degree

Figure 15
The Poor Have the Fastest
Income and Wealth Growth

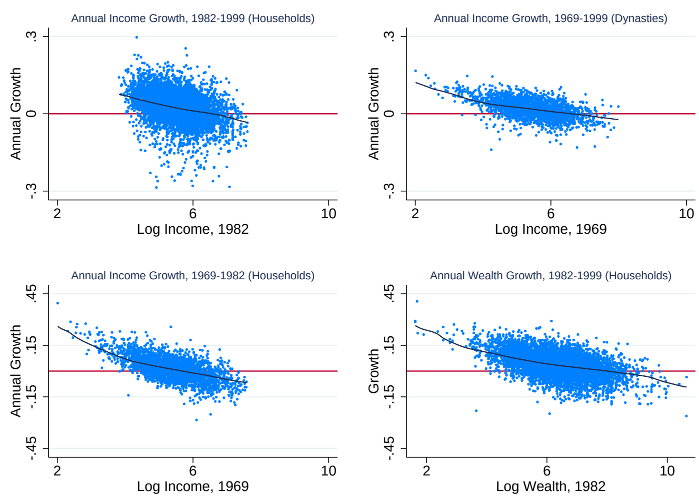

Note: The graphs give kernel-weighted moving averages of income growth as a function of initial income, and wealth growth as a function of initial wealth. These graphs exclude a few outliers whose presence does not change the overall relationship but does make the figure hard to read. The bandwidth is set to one-half of the standard deviation of initial income.

Figure 16
No Sign of Two Peaks in the Bivariate Density of Income

line; even among the very poor, most households get richer.

Figure 17 casts doubt on the existence of poverty traps using a different metric: mobility across the income distribution. The figure graphs what fraction of dynasties that start in, say, the bottom quartile of income in 1969 will still be in the bottom quartile versus one of the other quartiles. In other words, the figure graphs a transition matrix, where each bar is a row of the matrix and the shaded area within each bar gives the transition probability into the new quartile conditional on starting in the old quartile. The figure shows that only 40 percent of households that started in the bottom quartile in 1969 were still in the bottom quartile in 1982. This percentage is exactly what we computed in Section 2.4 in our simulations of Solow convergence. The numbers are similar for the period from 1982 to 1999. Likewise, more than half of households that began in the top quartile dropped down to a lower quartile. Though only 10 percent of households managed a rags-to-riches transition from the bottom to the top, it is hard to imagine many households make that transition in any society. In short, there is far more income mobility than might be expected in the presence of a poverty trap.

Figure 17
Income Mobility Among Dynasties is High

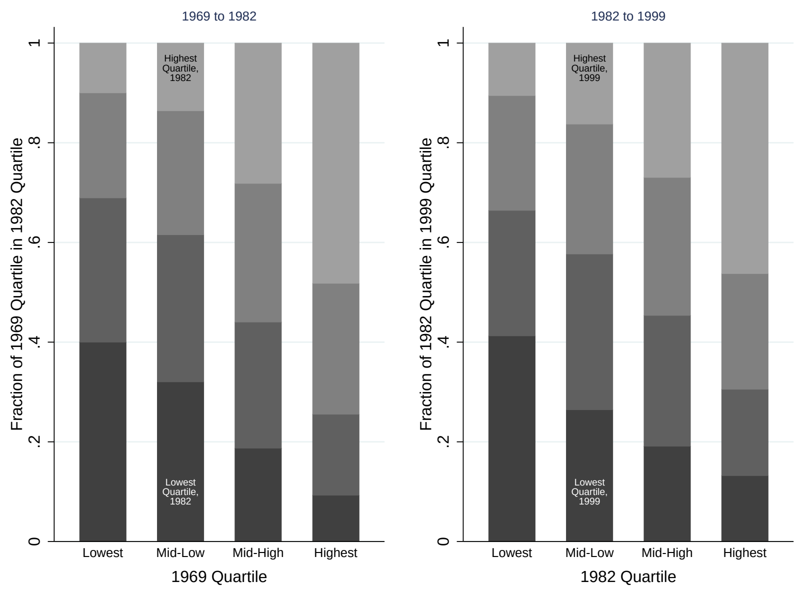

The evidence so far suggests that income has risen, that the poor are catching up to the rich, and that a dynasty's place in the income distribution is not fixed. But is this rosy image equally clear for everyone? In particular, have the gains accrued equally to those of high and low caste?

The Indian constitution, as detailed in The Constitution (Scheduled Castes) Order, recognizes a list of castes in its first schedule. These Scheduled Castes, together with disadvantaged tribal people called the Scheduled Tribes, suffered grave discrimination throughout history. These groups have been granted benefits to compensate for this legacy but remain disadvantaged. The government later recognized another group of castes, called the Other Backwards Castes, that have also suffered discrimination. Though less disadvantaged than the Scheduled Castes and Tribes, the Other Backwards Castes still suffered for the benefit of the upper castes.

Figure 18, which graphs the density of log income for each of the three groups, confirms that the disparities between them have not vanished. In 1969 it is clear that income is higher among the upper castes than among the Other Backwards

Figure 18
Disadvantaged Castes Have Lower Income and Wealth

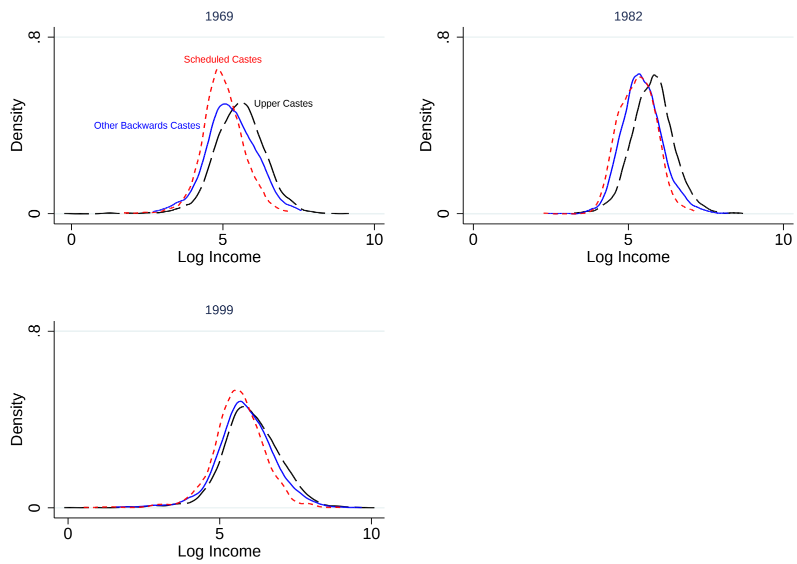

Castes, who in turn have higher income than the Scheduled Castes and Tribes. Though the distributions overlap—not all scheduled castes are poor and not all upper castes are rich—a household of an advantaged caste is more likely to be rich than one of a disadvantaged caste. There is some evidence that the distributions are converging in later years. By 1982, the scheduled castes seem almost to have caught up with the backwards castes, and by 1999 both have moved much closer to the upper castes. Yet the gaps remain.

These gaps do not appear in the growth rate of income. Figure 19 plots the density of income growth for each group. From both 1969 to 1982 and from 1982 to 1999, the density of income growth for each group lies atop that of the others. On first glance it seems that the upper castes enjoy no advantages; everyone is growing at the same rate. But as we show in Section 5, in fact this graph suggests the opposite. Given that disadvantaged castes are poorer, their income should be growing faster. Since it does not, the relationship between income and growth must be less favorable for the disadvantaged castes. They are converging to a lower steady state.

Figure 19
Income Growth Appears the Same Across All Castes

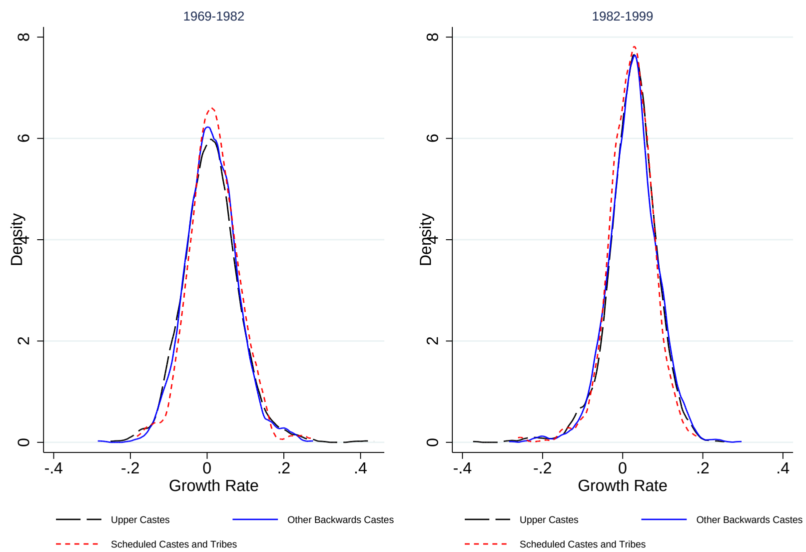

5 Results of the Tests

The simple graphs of Section 4 look inconsistent with a poverty trap. But as Figure 2 shows, looks can deceive; income shocks may hide a poverty trap. The negative growth test proposed in Section 2, however, can detect what the naked eye cannot. Does this test find evidence of a poverty trap in rural India?

It does not. Figure 20 is made much like Figure 4, except it applies the test to real rather than simulated data. It splits households into 10 equal bins based on their (log) income in 1969, estimates the probability of negative growth in income from 1969 to 1982, and tests for a decrease in the probability between each bin and the bin above it. We do the same for income growth from 1982 to 1999, using log income in 1982 to define the bins; and for growth in wealth from 1982 to 1999, using log wealth in 1982 to define the bins.

In none of the three cases do we reject the null of no poverty trap. In two of the three cases the p-value is almost 1. Figure 21 applies the same tests to the dynasty data; not only are the outcomes of the test similar, but the figures look almost identical. Using dynasties we confirm that the test for the entire period from 1969 to 1999 also finds no poverty trap. In Appendix D we show that reweighting the dynasty estimates to adjust for sample selection and the probability of attrition does not change the result of the test. Finally, in unreported results we find that running these tests within the caste groups defined in Section 4 still yields no evidence of poverty traps.

Are the shocks to income are too big for us to detect a trap? It is impossible to get a precise answer, but under some assumptions we can bound the standard deviation of the shock. In addition to the 1969 round used in the test, the Additional Rural Income Survey had two more rounds in 1970 and 1971. We regress income in each of these years on a household fixed effect and a common time trend (see Appendix 4.1). The deviations around the fixed effect and the trend are an estimate of income shocks—likely an over estimate, as each dynasty probably has its own trend. By imposing a common trend we leave more variation to the residual. Nevertheless we find that the residual has a standard deviation of only 0.4. According to the simulations the test should have no trouble detecting even a small trap (see Figure 5). It is possible the shock became far more variable in 1982, reducing the power to detect a trap from 1982 to 1999 (it is shocks in the initial year that most affect the power of the test). But given that the standard deviation of income actually falls from 1969 to 1982, this seems unlikely.

To summarize, there no evidence of a decrease in the probability of negative income growth. If anything, the probability seems to always be increasing, much like the panel in Figure 4 that shows Solow convergence. In short, the dynamics of household income in rural India look more like Solow convergence than any model of poverty traps.

But an absence of poverty traps need not imply convergence, as different households may converge to different steady states for reasons unrelated to their income. Some households may enjoy privileges that let their income grow faster at any level of income. In India, the main source of historical privilege is caste. As shown in Section 4, the upper castes enjoy income growth as rapid as the disadvantaged castes despite being richer. Is this a sign that they are converging to a higher steady state?

We apply the multi-line test proposed in Section 2 to test whether conver-

Figure 20
No Evidence of Poverty Traps Among Households

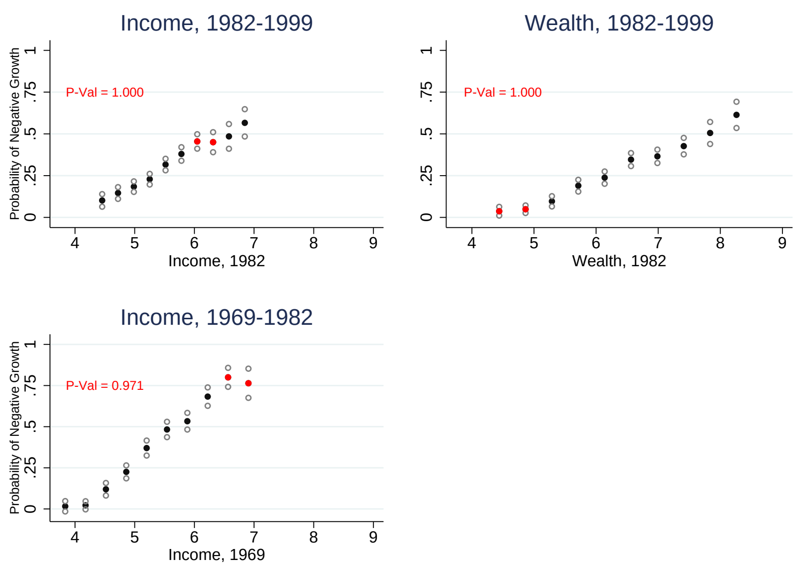

Note: These graphs show the estimated probability of negative growth within each bin. The p-value is for the null hypothesis of no poverty trap from the negative growth test derived in Section 2.3. The variance matrix of these estimates, which was fed into the test and used to compute the confidence intervals in each figure, was clustered by antecedent household. That is, if one 1982 household split into three 1999 households the errors of these three households were allowed to be correlated.

Figure 21

No Evidence of Poverty Traps Among Dynasties

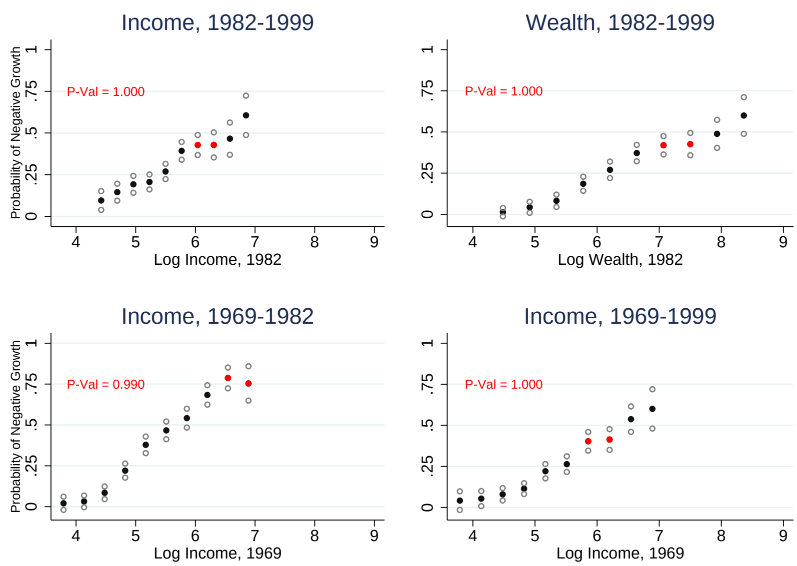

gence is conditional on caste. We compare upper castes to the so-called Other Backwards Castes and the heavily disadvantaged Scheduled Castes and Tribes. Figure 22 graphs the result of applying the test to growth in wealth from 1982 to 1999. Each panel highlights a different group and shows the fitted regression line of growth on initial wealth for that group. The level of wealth at which this line crosses the horizontal axis—the level of wealth at which growth is zero—is the steady state. As the graph shows, there are big gaps between these three steady states. The upper castes are converging to a higher steady state than the backwards castes, who in turn are converging to a higher steady state than the scheduled castes.

Table 3 tests for whether the gaps between these steady states are statistically significant. From both 1969 to 1982 and from 1982 to 1999; for both income and wealth; and among both households and dynasties, the gaps are highly significant. The first row in each panel tests whether the steady state of upper castes is higher than that of scheduled castes. From 1969 to 1982, upper caste households are converging to a level of income 70 percent higher than that of

Figure 22
Privileged Castes Converge to a Higher Steady State

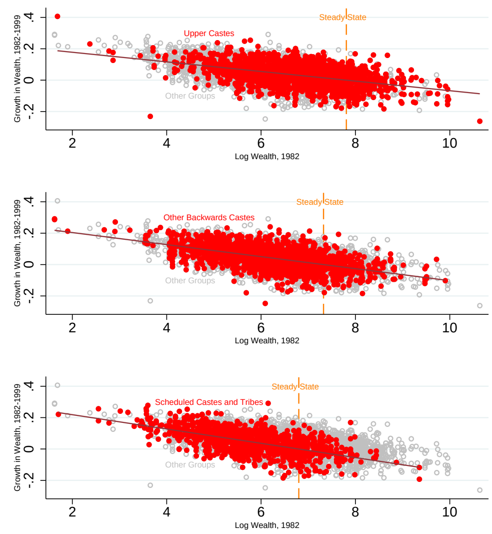

scheduled castes. From 1982 to 1999, they are converging to a level of income 100 percent higher, and a level of wealth nearly 200 percent higher. The second row confirms a similar pattern (but smaller gaps) between upper castes and backwards castes. The third row shows that the backwards castes in turn are converging to higher steady states than the scheduled castes.

The sheer size of these gaps is stunning. In the long run, a household from an upper caste can expect nearly three times the wealth of a household from a scheduled caste. More troubling still is that over time the gaps get bigger rather than smaller. India's many forms of affirmative action may not have been as successful as policymakers had hoped.

6 Discussion

Our results are not surprising given the existing literature on household poverty traps. Since a thirty-year panel like ours was previously unavailable, the prior literature has used cross-sections and shorter panels to test for the mechanisms

Table 3
Conditional Convergence: Percent Gap Between Steady States

Households:

	1969-1982	1982-1999	
	Income	Income	Wealth
High Caste – Scheduled Caste/Tribe	71***	101***	174***
High Caste – Other Backwards Castes	40***	54***	62***
Other Backwards Castes – Scheduled Caste/Tribe	22***	30***	69***

Dynasties:

	1969-1982	1982-1999	
	Income	Income	Wealth
High Caste – Scheduled Caste/Tribe	73***	131***	186***
High Caste – Other Backwards Castes	39***	65***	65***
Other Backwards Castes – Scheduled Caste/Tribe	24***	40***	73***

that might cause poverty traps (Kraay and McKenzie, 2014). Studies of herders in Ethiopia (Lybbert et al., 2004; Santos and Barrett, 2011) and Kenya and Madagascar (Barrett et al., 2006) have found evidence that herds below a critical size cannot migrate to fresh pastures and will remain small, creating a poverty trap.

But there is little consistent evidence of a similar mechanism in other contexts. De Mel et al. (2008) and Fafchamps et al. (2014) find in randomized controlled trials that small firms in Sri Lanka and Ghana could reap large additional profit if given extra funds. If these firms were in the low steady state of an asset-based poverty trap, such small investments would earn no additional profit. Though some studies have found that combining capital with management training can raise income (e.g. Bandiera et al., 2013; Banerjee et al., 2015), others have found that similar programs have no net impact (Morduch et al., 2012). Meanwhile, there has been little evidence for a nutritional poverty trap (see for example Subramanian and Deaton, 1996; Banerjee and Duflo, 2011).

More recent models have blamed behavioral poverty traps—traps that arise because poverty saps a person's self-control or attention (e.g Banerjee and Mullainathan, 2010; Shah et al., 2012). Much work in behavioral economics suggests the poor or less educated have trouble saving, have inaccurate expectations, or manage their businesses suboptimally.[12] But economists and psychologists have found similar behavior in subjects from all backgrounds. It is not

[12]For a tiny sample see Beaman et al. (2014); Duflo et al. (2009); Dupas and Robinson (2014); Dizon-Ross (2014); Shenoy (2015).

clear that such suboptimal behavior creates a poverty trap. Indeed, Kraay and McKenzie (2014) argue that suboptimal behavior does not prevent the income of the poor from rising in tandem with national income.

It is possible that there is some mechanism for a poverty trap not yet modeled by theorists or tested for by empiricists. But our results, by testing directly for the income dynamics implied by a poverty trap, suggest otherwise. Both the summary statistics and the formal test suggest the income of the poor grows and that it grows faster than the income of the rich.

Our other result—that the incomes of disadvantaged castes are converging to a lower steady state—is also supported by prior literature. For example, Pande (2003) and Besley et al. (2004) show that politicians of high castes are less likely to support policies that help low castes. Given that most politicians are of high caste, such discrimination could persist and preserve the second-class status of low castes. We show that the resulting gaps between steady states are large and have only widened over time.

7 Conclusion

The household poverty trap has long eluded empirical work in economic development. We show that it remains elusive. We derive a new method to detect poverty traps; in simulations the method performs well. We apply the method to a panel dataset that follows thousands of rural Indian households over thirty years. Nevertheless we find no evidence of a poverty trap. The income of most households grew, and the income of the poorest grew the fastest. But we also find that income does not converge—or rather, it converges only conditionally. Households of high castes are converging to a higher steady state than those of low caste. The gaps are large and highly significant.

Taken together our results suggest that inequality within India is not caused by a poverty trap, though some of it is caused by caste. Our results cannot rule out that inequality across the globe is caused by a poverty trap. The type of poverty trap that might prevent a poor Indian from catching up to his richer countryman may differ from the type that prevents India from catching up to the U.S. A household may be stuck in a poverty trap because it lacks financial or human capital, whereas a country may be stuck in a poverty trap because it lacks good institutions. Weak institutions would hurt both rich and poor, making everyone converge to a lower steady state than they might otherwise reach. Unlike a poverty trap based on financial or human capital, an institutional poverty trap cannot be found in a household panel and must be sought elsewhere. We leave the search to future research.

References

AZARIADIS, C. AND J. STACHURSKI (2005): "Poverty Traps," in *Handbook of Economic Growth*, ed. by P. Aghion and S. Durlauf, North-Holland, vol. 1A, 295–384.

BANDIERA, O., R. BURGESS, N. DAS, S. GULESCI, I. RASUL, AND M. SULAIMAN (2013): "Can Basic Entrepreneurship Transform the Economic Lives of the Poor?" .

BANERJEE, A., E. DUFLO, N. GOLDBERG, D. KARLAN, R. OSEI, W. PARIENTÉ, J. SHAPIRO, B. THUYSBAERT, AND C. UDRY (2015): "A Multifaceted Program Causes Lasting Progress for the Very Poor: Evidence from Six Countries," *Science*, 348, 1260799.

BANERJEE, A. AND S. MULLAINATHAN (2010): "The Shape of Temptation: Implications for the Economic Lives of the Poor," Tech. rep., National Bureau of Economic Research.

BANERJEE, A. V. AND E. DUFLO (2011): *Poor Economics: A Radical Rethinking of the Way to Fight Global Poverty*, Public Affairs.

BANERJEE, A. V. AND A. F. NEWMAN (1993): "Occupational Choice and the Process of Development," *Journal of Political Economy*, 101, 274–298.

BARDHAN, P. K. (1984): *Land, labor, and rural poverty: Essays in development economics*, Columbia University Press.

BARRETT, C. B. AND M. R. CARTER (2013): "The Economics of Poverty Traps and Persistent Poverty: Empirical and Policy Implications," *The Journal of Development Studies*, 49, 976–990.

BARRETT, C. B., P. P. MARENYA, J. MCPEAK, B. MINTEN, F. MURITHI, W. OLUOCH-KOSURA, F. PLACE, J. C. RANDRIANARISOA, J. RASAMBAINARIVO, AND J. WANGILA (2006): "Welfare Dynamics in Rural Kenya and Madagascar," *The Journal of Development Studies*, 42, 248–277.

BEAMAN, L., J. MAGRUDER, AND J. ROBINSON (2014): "Minding Small Change among Small Firms in Kenya," *Journal of Development Economics*, 108, 69–86.

BEHRMAN, J. R., A. D. FOSTER, M. R. ROSENZWEIG, AND P. VASHISHTHA (1999): "Womens Schooling, Home Teaching, and Economic Growth," *Journal of Political Economy*, 107, 41–57.

BESLEY, T., R. PANDE, L. RAHMAN, AND V. RAO (2004): "The Politics of Public Good Provision: Evidence from Indian Local governments," *Journal of the European Economic Association*, 2, 416–426.

BIANCHI, M. (1997): "Testing for Convergence: Evidence from Non-parametric Multimodality Tests," *Journal of Applied Econometrics*, 12, 393–409.

BLOOM, D. E., D. CANNING, AND J. SEVILLA (2003): "Geography and Poverty Traps," *Journal of Economic Growth*, 8, 355–378.

BOURGUIGNON, F. AND C. MORRISSON (2002): "Inequality among World Citizens: 1820-1992," *American economic review*, 727–744.

CHETTY, R., N. HENDREN, P. KLINE, E. SAEZ, AND N. TURNER (2014): "Is the United States Still a Land of Opportunity? Recent Trends in Intergenerational Mobility," *The American Economic Review*, 104, 141–147.

CHETVERIKOV, D. (2013): "Testing Regression Monotonicity in Econometric Models," .

CLARK, G. AND N. CUMMINS (2015): "Intergenerational Wealth Mobility in England, 1858–2012: Surnames and Social Mobility," *The Economic Journal*, 125, 61–85.

DASGUPTA, P. AND D. RAY (1986): "Inequality as a Determinant of Malnutrition and Unemployment: Theory," *The Economic Journal*, 1011–1034.

DE MEL, S., D. MCKENZIE, AND C. WOODRUFF (2008): "Returns to Capital in Microenterprises: Evidence From a Field Experiment," *The Quarterly Journal of Economics*, 123, 1329–1372.

DERCON, S. AND I. OUTES (2009): "Income Dynamics in Rural India: Testing for Poverty Traps and Multiple Equilibria," Working Paper.

DIZON-ROSS, R. (2014): "Parents' Perceptions and Children's Education: Experimental Evidence from Malawi," *Unpublished Manuscript. Massachusetts Institute of Technology. http://web. mit. edu/rdr/www/perceptions. pdf.*

DUFLO, E., M. KREMER, AND J. ROBINSON (2009): "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya," Tech. rep., National Bureau of Economic Research.

DUPAS, P. AND J. ROBINSON (2014): "The Daily Grind: Cash Needs, Labor Supply and Self-Control," *Unpublished paper*.

ESTUDILLO, J. P., Y. MANO, AND S. SENG-ARLOUN (2013): "Job Choice of Three Generations in Rural Laos," *The Journal of Development Studies*, 49, 991–1009.

FAFCHAMPS, M., D. MCKENZIE, S. QUINN, AND C. WOODRUFF (2014): "Microenterprise Growth and the Flypaper Effect: Evidence from a Randomized Experiment in Ghana," *Journal of development Economics*, 106, 211–226.

FOSTER, A. D. AND M. R. ROSENZWEIG (1995): "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal of Political Economy*, 103, 1176–1209.

——— (1996): "Technical Change and Human Capital Returns and Investments: Evidence from the Green Revolution," *American Economic Review*.

——— (2002): "Household Division and Rural Economic Growth," *Review of Economic Studies*, 69, 839–869.

GALOR, O. AND J. ZEIRA (1993): "Income Distribution and Macroeconomics," *Review of Economic Studies*, 60, 35–52.

HIMANSHU AND N. STERN (2011): "India and an Indian village: 50 years of economic development inPalanpur," Working Paper.

KRAAY, A. AND D. MCKENZIE (2014): "Do Poverty Traps Exist? Assessing the Evidence," *Journal of Economic Perspectives*, 28, 127–48.

KRISHNA, A. (2013): "Stuck in Place: Investigating Social Mobility in 14 Bangalore Slums," *The Journal of Development Studies*, 49, 1010–1028.

KRUGMAN, P. (1991): "Increasing Returns and Economic Geography," *Journal of Political Economy*, 99, 483–499.

——— (1994): "The fall and rise of development economics," in *Rethinking the Development Experience; Essays Provoked by the Work of Albert O. Hirschman*, Brookings Institution Press, ed. by Lloyd Rodwin and Donald A. Schon.

KWAK, S. AND S. C. SMITH (2013): "Regional Agricultural Endowments and Shifts of Poverty Trap Equilibria: Evidence from Ethiopian Panel Data," *The Journal of Development Studies*, 49, 955–975.

LOKSHIN, M. AND M. RAVALLION (2004): "Household Income Dynamics in Two Transition Economies," *Studies in Nonlinear Dynamics & Econometrics*, 8.

LYBBERT, T. J., C. B. BARRETT, S. DESTA, AND D. LAYNE COPPOCK (2004): "Stochastic Wealth Dynamics and Risk Management among a Poor Population*," *The Economic Journal*, 114, 750–777.
MATSUYAMA, K. (2004): "Financial Market Globalization, Symmetry-Breaking, and Endogenous Inequality of Nations," *Econometrica*, 72, 853–884.
MCKAY, A. AND E. PERGE (2013): "How Strong is the Evidence for the Existence of Poverty Traps? A Multicountry Assessment," *The Journal of Development Studies*, 49, 877–897.
MICHELSON, H., M. MUÑIZ, AND K. DEROSA (2013): "Measuring Socio-Economic Status in the Millennium Villages: The Role of Asset Index Choice," *The Journal of Development Studies*, 49, 917–935.
MORDUCH, J., S. RAVI, AND J. BAUCHET (2012): "Failure vs. Displacement: Why an Innovative Anti-Poverty Program Showed No Net Impact," *Institute of Economic Research, Hitotsubashi University, PRIMCED Discussion Paper*, 32.
MURPHY, K. M., A. SHLEIFER, AND R. W. VISHNY (1989): "Industrialization and the Big Push," *Journal of Political Economy*, 97, 1003–1026.
NASCHOLD, F. (2009): ""Poor stays poor": Household asset poverty traps in rural semi-arid India," Working Paper.
PANDE, R. (2003): "Can Mandated Political Representation Increase Policy Influence for Disadvantaged Minorities? Theory and Evidence from India," *American Economic Review*, 93, 1132–1151.
PIKETTY, T. AND E. SAEZ (2003): "Income Inequality in the United States, 1913–1998," *The Quarterly Journal of Economics*, 118, 1–41.
QUAH, D. T. (1996): "Convergence Empirics across Economies with (Some) Capital Mobility," *Journal of Economic Growth*, 1, 95–124.
QUISUMBING, A. R. AND B. BAULCH (2013): "Assets and Poverty Traps in Rural Bangladesh," *The Journal of Development Studies*, 49, 898–916.
ROSENZWEIG, M. AND K. I. WOLPIN (1980): "Testing the Quantity-Quality Model of Fertility: The Use of Twins as a Natural Experiment," *Econometrica*.
SANTOS, P. AND C. B. BARRETT (2011): "Persistent Poverty and Informal Credit," *Journal of Development Economics*, 96, 337–347.
SHAH, A. K., S. MULLAINATHAN, AND E. SHAFIR (2012): "Some Consequences of Having Too Little," *Science*, 338, 682–685.
SHENOY, A. (2015): "Does Schooling Make Our Expectations Accurate?" .
SILVERMAN, B. W. (1981): "Using Kernel Density Estimates to Investigate Multimodality," *Journal of the Royal Statistical Society. Series B (Methodological)*, 97–99.
SUBRAMANIAN, S. AND A. DEATON (1996): "The Demand for Food and Calories," *Journal of political economy*, 133–162.

A Technical Appendix (For Online Publication)

1.1 Proof of Proposition 1

Since each coefficient $\hat{\alpha}_j$ is asymptotically normal the difference $[\hat{Dif}]_j$ is asymptotically normal, and thus the vector of differences $([Dif]_1, \ldots, [Dif]_{J-1})$ are jointly asymptotically normal.[13] It is easy to show that since each coefficient $\hat{\alpha}_j$ is estimated using a different set of observations that $Cov(\hat{\alpha}_j, \hat{\alpha}_k) = 0$ for all $k \neq j$. Then $Cov([Dif]_j, [Dif]_k) = 0$ for all $k \neq j-1, j, j+1$.

Adjacent differences have a single coefficient in common and thus will be correlated. Suppose for a moment that the variance of $\hat{\alpha}_j$ (call it v_j) is known (in practice we replace the true variance with a consistent estimator). Then

$$\begin{aligned} Cov([\hat{Dif}]_j, [\hat{Dif}]_{j+1}) &= Cov\left(\frac{\hat{\alpha}_j - \hat{\alpha}_{j+1}}{\sqrt{v_j + v_{j+1}}}, \frac{\hat{\alpha}_{j+1} - \hat{\alpha}_{j+2}}{\sqrt{v_{j+1} + v_{j+2}}}\right) \\ &= \mathbb{E}\left[\frac{\hat{\alpha}_j - \hat{\alpha}_{j+1}}{\sqrt{v_j + v_{j+1}}} \frac{\hat{\alpha}_{j+1} - \hat{\alpha}_{j+2}}{\sqrt{v_{j+1} + v_{j+2}}}\right] \\ &= \frac{1}{\sqrt{v_j + v_{j+1}}\sqrt{v_{j+1} + v_{j+2}}} \mathbb{E}\left[(\hat{\alpha}_j - \hat{\alpha}_{j+1})(\hat{\alpha}_{j+1} - \hat{\alpha}_{j+2})\right] \quad (1) \end{aligned}$$

[13] This follows from the Cramér-Wold Theorem.

where the second equality follows because under the null $[Dif]_j = [Dif]_{j+1} = 0$. The expectation equals

$$\begin{aligned}\mathbb{E}\left[(\hat{\alpha}_j - \hat{\alpha}_{j+1})(\hat{\alpha}_{j+1} - \hat{\alpha}_{j+2})\right] &= \mathbb{E}[\hat{\alpha}_j\hat{\alpha}_{j+1}] - \mathbb{E}[\hat{\alpha}_j\hat{\alpha}_{j+2}] - \mathbb{E}[\hat{\alpha}_{j+1}^2] + \mathbb{E}[\hat{\alpha}_{j+1}\hat{\alpha}_{j+2}] \\ &= \mathbb{E}[\hat{\alpha}_j]\mathbb{E}[\hat{\alpha}_{j+1}] - \mathbb{E}[\hat{\alpha}_j]\mathbb{E}[\hat{\alpha}_{j+2}] - (\mathbb{E}[\hat{\alpha}_{j+1}^2] - \mathbb{E}[\hat{\alpha}_{j+1}]\mathbb{E}[\hat{\alpha}_{j+2}])\end{aligned}$$

Under the null hypothesis, $[Dif]_1 = [Dif]_2 = \cdots = [Dif]_{J-1} = 0$ which implies $\mathbb{E}[\hat{\alpha}_1] = \mathbb{E}[\hat{\alpha}_2] = \cdots = \mathbb{E}[\hat{\alpha}_J]$. We can then replace the expectations above with $\mathbb{E}[\hat{\alpha}_{j+1}]$, collapsing the expression to

$$\begin{aligned}\mathbb{E}\left[(\hat{\alpha}_j - \hat{\alpha}_{j+1})(\hat{\alpha}_{j+1} - \hat{\alpha}_{j+2})\right] &= -(\mathbb{E}[\hat{\alpha}_{j+1}^2] - \mathbb{E}[\hat{\alpha}_{j+1}]^2) \\ &= -v_{j+1} \qquad (2)\end{aligned}$$

Subbing (2) into (1) give the covariance

$$Cov([\hat{Dif}]_j, [\hat{Dif}]_{j+1}) = \frac{-v_{j+1}}{\sqrt{v_j + v_{j+1}}\sqrt{v_{j+1} + v_{j+2}}} \qquad (3)$$

Since the asymptotic variance of $[\hat{Dif}]_j = 1$, the variance matrix of $([\hat{Dif}]_1, \ldots, [\hat{Dif}]_{J-1})$ has ones along the diagonal, (3) in each position $(j, j+1)$ and $(j+1, j)$, and zeros everywhere else.

Since $\hat{\lambda}$ is the largest order statistic among $([\hat{Dif}]_1, \ldots, [\hat{Dif}]_{J-1})$, the probability that $\hat{\lambda} < \lambda$ is simply the probability that all of the estimated differences are less than λ:

$$\begin{aligned}Pr(\hat{\lambda} < \lambda) &= Pr([\hat{Dif}]_1 < \lambda, \ldots, [\hat{Dif}]_1 < \lambda) \\ &= \Phi_{J-1}(\lambda, \lambda, \ldots, \lambda)\end{aligned}$$

□

B Data Appendix (For Online Publication)

- **Household Income (1969-1971):** From the merged data compiled from the 1969-1971 household economic survey. Household income is computed from the sum of income from various sources and is defined as receipts net of expenditure. These sources of income are income from agriculture, plantations and orchards, income from self-employment in farm activities (livestock and allied activities (allied activities consist of bee-keeping, fishery, sericulture, forestry and other activities)), income from self-employment in non-farm activities (business, craft and professional activities), income from salaries (longer-term employment) and wages, income from house property, income from interest and dividends, and income from current transfers. Imputed value of family labor for investments are included as well.

- **Household Income (1982):** From various decks in the 1982 household economic survey. Definition is identical to that of 1969-1971

- **Household Income (1999):** From various decks in the 1999 household economic survey. Definition is identical to that of 1969-1971

- **Household Wealth (1982):** From the 1982 ARIS-REDS household economic survey. Household wealth is computed as the owners equity of the household. I.e. the value of all assets owned at the beginning of the RP net of the value of all outstanding liabilities at the beginning of the response period (RP). Value of assets owned is the sum of the following variables:

  - From Deck 16. Real value of buildings owned and real value of non-house land owned
  - From Deck 8. Real value of irrigation assets owned at BRP
  - From Deck 9. Real value of farm equipment owned at BRP
  - From Deck 10. Real value of other farm assets owned at BRP
  - From Deck 11. Real value of animals owned at BRP
  - From Deck 12. Real value of animal-related assets owned at BRP

  - From Deck 13. Real value of non-farm business assets and inventory owned at BRP
  - From Deck 18. Real value of consumer durables owned at BRP
  - From Deck 19. Real value of savings. For each household, real value of savings is measured as the sum of deposits with commercial banks, cooperative banks, post office savings banks and companies, shares and securities, small savings instruments, gold and jewellery and currency all at BRP
  - From Deck 21. Real value of outstanding loans made by household at BRP
  - From Deck 20. Value of outstanding liabilities is measured as the real value of outstanding liabilities at BRP.

- **Household Wealth (1999):** From the 1999 ARIS-REDS household economic survey. Household wealth is computed as the owners equity of the household. I.e. the value of all assets owned at the beginning of the RP net of the value of all outstanding liabilities at the beginning of the response period (RP). Value of assets owned is the sum of the following variables:
  - From Deck 102 and Deck 103. Real value of buildings owned and real value of non-house land owned
  - From Deck 52. Real value of irrigation assets owned at BRP.
  - From Deck 59. Real value of farm equipment owned at BRP
  - From Deck 62. Real value of other farm assets owned at BRP
  - From Deck 47. Real value of inventory of farm output at BRP
  - From Deck 69 and Deck 70. Real value of animals owned at BRP
  - From Deck 79. Real value of animal-related assets owned at BRP
  - From Deck 89 and Deck 93. Real value of non-farm business assets and inventory owned at BRP
  - From Deck 114. Real value of consumer durables owned at BRP
  - From Deck 121. Real value of savings. For each household, real value of savings is constructed from the sum of deposits with commercial

banks, cooperative banks, post office savings banks and companies, shares and securities, small savings instruments, gold and jewellery and currency all at BRP

- From Deck 126. Real value of outstanding loans made by household at BRP
- From Deck 125. Value of outstanding liabilities is measured as the real value of outstanding liabilities at BRP.

C More Simulation Results (For Online Publication)

3.1 Technological Progress Makes Poverty Traps Short-Lived

In the simulations reported in the main text we assume zero technological progress, as though the economy were stagnant. Figure 23 shows how the steady state diagram changes when we relax that assumption. We return to the parameters that generate a large poverty trap. The top left diagram shows the (technology-augmented) production function in 1969, 1982, and 1999 assuming labor-augmenting productivity grows at 2 percent per year. The bottom left diagram assumes growth of 4 percent per year. Both diagrams show why in the presence of overall growth the poverty traps are short-lived. Assuming the fixed investment remains constant, over time it becomes easier to afford.

Even households in the low steady state are able to produce more income despite having a level of capital that keeps them within the poverty trap. To illustrate this we run simulations for the cases showin in both the top left and bottom left. We assume there are (true) shocks to productivity with standard devation 0.1. The top and bottom diagrams on the right show the fraction of households that began with income and capital below the level of the fixed investment who are still below it in each year after. When growth is 2 percent per year, the poverty trap has not vanished by 1999. Thus no one has gotten enough capital to exit the poverty trap by 1999. But the top right panel shows that the poverty trap is less meaningful because by 1999 everybody earns enough income to exceed its limits. When growth is 4 percent per year, as in the bottom left diagram, the poverty trap has vanished by 1999. The bottom right diagram

Figure 23
Technological Progress Eliminates Poverty Traps

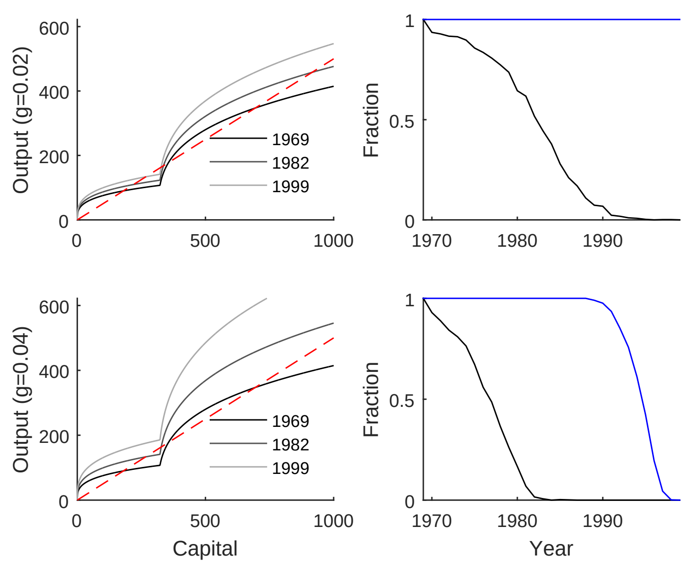

shows that everyone earns income beyond the limits of the original poverty trap by the mid-1980s and everyone has acquired enough capital to escape the low steady-state by 1999.

In both cases a poverty trap for households is steadily rendered moot by improvements in technology. For the households to actually escape the poverty trap—that is, move to the high steady-state—it must be that the fixed investment does not rise as quickly as the level of technology. This assumption might fail if the fixed investment is the cost of buying land, as land prices might rise in tandem with its productivity. But more likely the fixed investment is the cost of buying a generator, building an irrigation canal, or sending a child to work in the city. These costs will probably become more affordable as overall income rises.

Does this point apply to rural India? The incomes of most households in our sample grow by at least 2 percent per year between 1982 and 1999. If the source of that growth is "technological progress"—better farming, better education, better access to markets, better chances to send family members to work in cities—any poverty trap that existing in 1969 may cease to matter by 1999. In short, technological progress makes it even less likely there is a poverty trap.

3.2 Changing the Parameters Does Not Matter; Changing the Steady State Does

We check how well the negative growth test detects poverty traps when the parameters α and s are tweaked. Figure 24 shows the phase diagram for each new set of parameters (alongside the baseline case). The "High alpha" case sets $\alpha = 0.5$, leaving all else as in the baseline. The "Low alpha" case sets $\alpha = 0.15$, and the "Low alpha, low s" case sets $\alpha = 0.15, s = 0.15$. Finally, the "High s, High F" case sets $s = 0.25$ and uses the fixed cost from the "Large Trap" scenario of Section 2.1. (Using the fixed cost from the Baseline case would give only a single steady state.)

Figure 25 graphs the rejection rates for each set of parameters at different levels of variance for the shock. In all cases the test continues to detect poverty traps in the span from 1969 to 1982. The difference with Figure 5 arises in the span from 1982 to 1999. By comparing the rejection rate of each specification to its phase diagram in Figure 24 it is clear that the rejection rate is predicted by the size of the poverty trap. In both cases when α is low—when capital has sharply decreasing returns—the trap is large. A household must get many positive shocks to escape the trap, and if it falls just short it will be quickly dragged back to the low steady state. By contrast, when α is high

3.3 The Method Still Works When There is Heterogeneity in the Location of the Traps

A model of poverty traps typically assumes households differ only in their initial wealth. But in reality they may differ in the technologies they use and thus the locations of their steady states. Each household would have its own poverty trap, creating a continuum of poverty traps. How does the method fare when there is a continuum of traps?

Suppose the productivity of the inferior technology varies by household. This might represent the case in which some are more productive farmers (in

Figure 24
Steady State Diagrams for the New Parameters

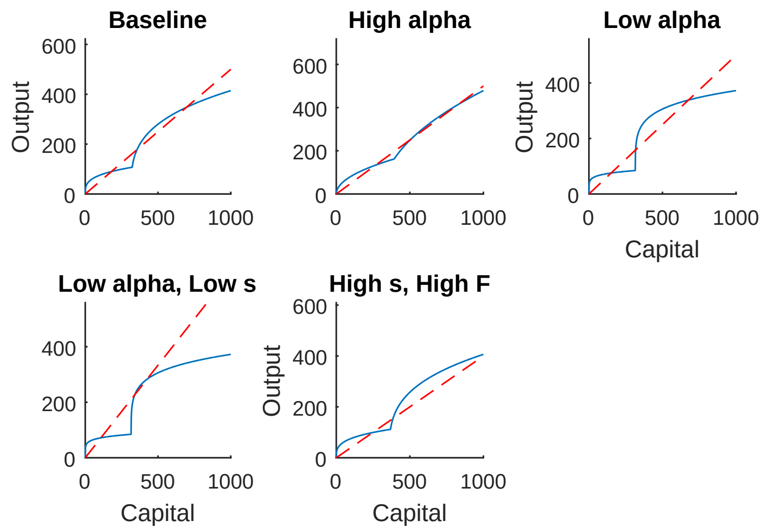

the case of a land-based poverty trap) or have more efficient metabolisms (in the case of a nutrition-based poverty trap). Then $A_0 = \tilde{A} \cdot A_0^*$, where A_0^* is calibrated as described in Section 2.1 and $\tilde{A}$ is a draw from a uniform distribution with support $[1 - h^A, 1 + h^A]$. The higher the bandwidth h^A the greater the heterogeneity. A household with a high productivity A_0 will have a low steady state that is higher than a household with a low A_0. There is now a range of low steady states and thus a range of thresholds for the poverty trap.

Figure 26 shows the rejection rate in the case of the baseline trap assuming several bandwidths. I consider $h^A \in \{0, .05, 0.1, \ldots, 0.25\}$. Lighter colored curves show the rejection rate with higher values of h^A. At the highest bandwidth the method does have less power than in the baseline case ($h^A = 0$). But the loss in power is not catastrophic. Though productivity varies over a range half the size of the mean, the power curve looks similar to the case with no variation.

This is not to say that the method is robust to arbitrarily large heterogeneity. No method for detecting poverty traps can adapt to large deviations from the theory. But the figure shows that the negative growth test does not simply collapse in the presence of heterogeneity.

Figure 25
Results Using the New Parameters

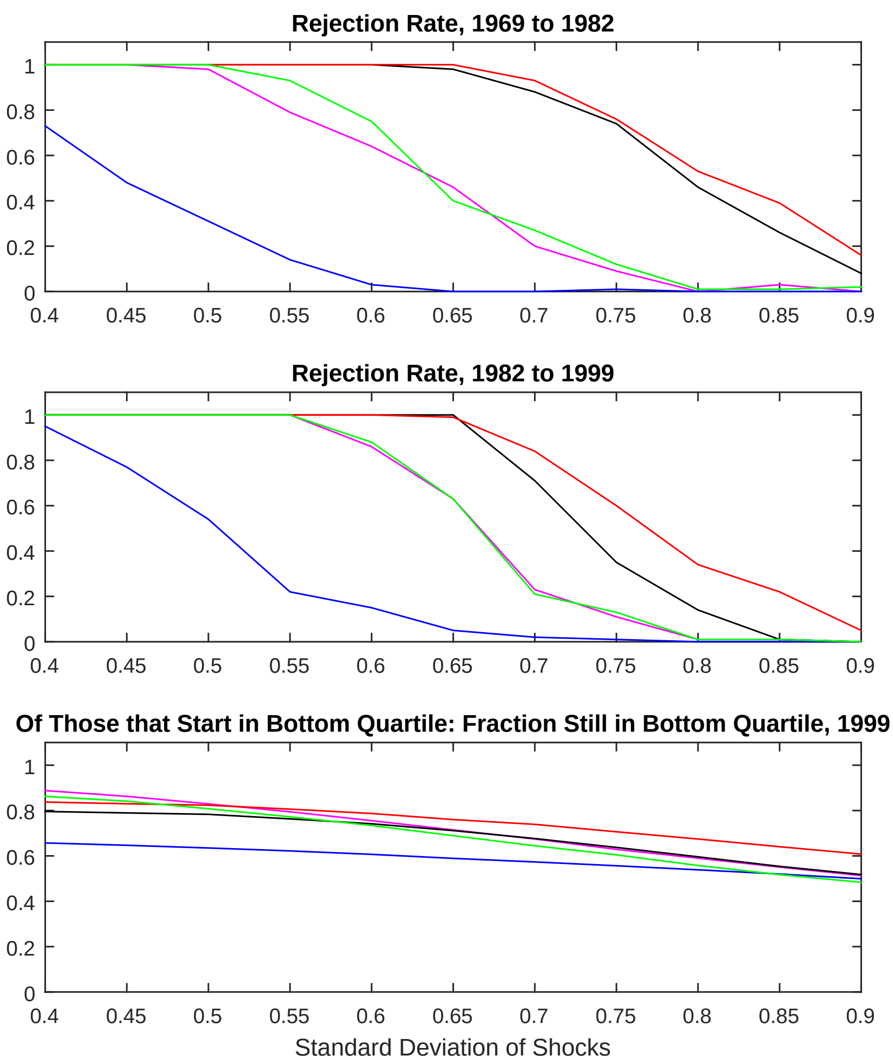

Legend: Baseline, High alpha, Low Alpha, **Low alpha, low s**, **High s, high F**.

Figure 26
A Continuum of Poverty Traps

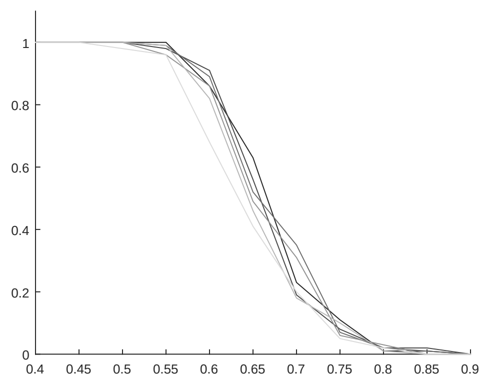

Legend: The base line case is in black; lighter colors show the power curve at higher levels of heterogeneity (higher bandwidths h^A), with the lightest shade giving a bandwidth of $h^A = 0.25$.

D Additional Empirical Tests (For Online Publication)

4.1 How Big are Shocks in the Data?

Our power to detect a poverty trap hinges on the size of the income shocks. If the standard deviation of the shocks are large relative to the total standard deviation of income the negative growth test may be too weak to find a small poverty trap.

Under some assumptions we can bound the size of the shock using the data from the Additional Rural Income Survey, which records household income in 1969, 1970, and 1971. In the model

$$\begin{aligned}\log y_{it} &= \log Z_{it} + \log F(k_{it}, A_{it}) \\ &= \log Z_{it} + \overline{\log F}_t^{kA} + (\log F(k_{it}, A_{it}) - \overline{\log F}_t^{kA}))\end{aligned}$$

where $\overline{\log F}_t^{kA}$ gives average contribution of capital and productivity to income. If we assume the deviation from the average is roughly constant—not im-

Figure 27
Distribution of Residuals (Shocks)

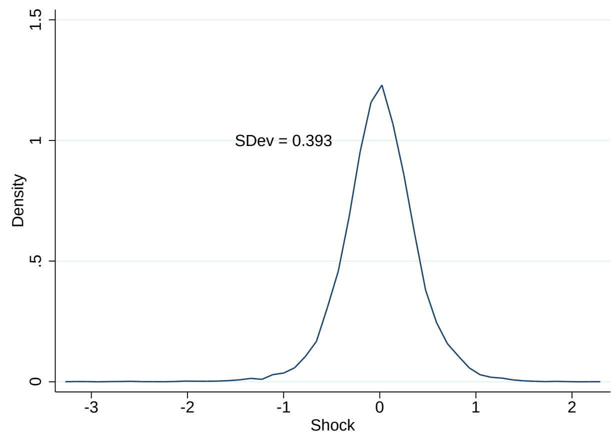

plausible given that we focus on just three years—then $\log F(k_{it}, A_{it}) - \overline{\log F}_t^{kA} \approx c_i$, a household fixed-effect. Using the household panel we regress

$$\log y_{it} = c_i + \beta t + \varepsilon_{it}$$

To be conservative we assume $\overline{\log F}_t^{kA}$ grows at a constant rate. (Year dummies would likely absorb some part of income that is actually caused by shocks.) We take the residual of the regression as a rough measure of $\log Z_t$. Figure 27 shows the distribution of the residual. The standard deviation is less than 0.4, well within the range in which the negative growth test detects a poverty trap. Assuming shocks do not become too much more variable in 1982, our tests should be able to detect a poverty trap.

4.2 Reweighting for Sample Selection and Attrition

Figure 28 reweights our tests for poverty traps using both sampling weights and the predicted probability of attrition, as computed using the coefficients in Table 2. The weights raise the variance of the estimates at several points of the range. We still find no evidence of a poverty trap.

Figure 28
Test for Poverty Traps: Reweighted