Stagnation or Transition? Poverty Traps and the Dynamics of Household Income

Raj Arunachalam
University of Michigan, Ann Arbor

Ajay Shenoy*
University of California, Santa Cruz

November 12, 2014
First Version: August 8, 2012

Abstract

We find no evidence of poverty traps in a thirty year panel of rural Indian households. We instead find strong evidence of income convergence, but also rising inequality. Simulating a model of Solow households, we reconcile convergence with inequality by relaxing the assumption that the variance of household productivity shocks remains fixed. The results suggest that rural Indian households faced an economic environment in rapid transition. (JEL Codes: O12)

(Keywords: poverty trap, Solow model, rural India)

*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.

1 Introduction

Development economics is in some sense the study of why things do not change—why a poor household stays poor, and why a stagnant economy stays stagnant. At the intellectual core of the discipline, 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). And as the earliest source of credible microdata, rural India has been the discipline's canonical example of an economy so ensnared (Bardhan, 1984).

This paper shows that in rural India neither the income of the poor nor the economic environment they face stays fixed. We use a nationally representative panel of rural households to construct a measure of income that is consistent across three survey rounds, letting us follow households over thirty years. Ours is the first study that can track the fortune of a dynasty over a period long enough to discern whether it is truly trapped in poverty. We find no evidence of household poverty traps. Simple tests show that the incomes of all households are converging to a single steady-state.

Our tests derive from a dynamic model of household income which assumes only that current income depends on current wealth (or any persistent asset), and that future wealth depends on current wealth and current income. The framework, which can mimic the dynamics of many more complex models (e.g. Banerjee and Newman, 1993; Galor and Zeira, 1993), implies a poverty trap if the relation between baseline income and income growth crosses the x-axis more than once; if the relation is monotonically decreasing, there is convergence to a single steady-state. Given a large dataset, this relation can be estimated nonparametrically.

We estimate the relation between growth from 1969 to 1982 and 1982 income and find no evidence of poverty traps. The period from 1982 to 1999 is similar. Our results do not seem to be driven by error in the measurement of income. Using a three-year average for baseline income lowers the rate of convergence, but only slightly. Furthermore, wealth and savings, arguably less prone to measurement error, converge even more rapidly than income. Nor do negative income shocks—which we proxy using rainfall—drop poor households into lower income trajectories, as a theory of asset-based poverty traps might predict. Given the length of our panel, the results reinforce a growing sense among development economists (e.g. Kraay and McKenzie, 2014) that we must reassess the empirical relevance of household poverty traps.

Yet all is not rosy in rural India. All else equal, convergence should imply a decline in inequality. While this holds in the earlier period, inequality rises sharply from 1982 to 1999. To resolve income convergence with rising inequality, we explore whether the economic environment has changed. To do so, we simulate the moments of our data using the simplest implementation of the general model: a model of Solow households. Each household produces income using a general form of capital, which represents anything from tractors to education to family members with jobs in the city. The rate of investment and of technology growth may vary between villages. Each household is subject to its own productivity shocks.

We simulate the model over 30 years for more than 50 thousand sets of parameters. For each set we create a simulated panel to match the actual data. We compute two simulated moments: the correlation between income in 1982 and growth from 1982 to 1999, and the standard deviation of log income in 1999. For each set of parameters these two moments measure the rate of convergence and the level of inequality.

The simulations show that a fixed economic environment—one in which parameters remain fixed over the period—is inconsistent with the data. Only by raising the volatility of the shock to household productivity between 1982 and 1999 can we reconcile convergence with inequality. Since the productivity shock is independent of household income, an increase in its volatility means that a greater portion of income growth is unexplained by current income. One way to interpret rising volatility in the productivity shock is therefore a rise in opportunity; indeed, we find that the increase in volatility suggested by the simulation coincides with a rise in educational mobility in the data.

Our results are inconsistent with models of poverty traps for households but do not rule out club convergence for countries, and cannot speak directly to broader debates about the global income distribution. While household income converges to a single steady state, nothing we find implies the steady state is as high as it should be. Indeed, the growth rate of income in rural India in this period is roughly half that of the country as a whole. Any model in which countries or sectors are trapped in poverty even as households converge would be consistent with our result. In particular, our rejection of a fixed economic environment is consistent with theories of regime transition, such as unified growth models (e.g. Galor and Weil, 2000; Hansen and Prescott, 2002; Galor et al., 2009) and models with two sectors and surplus labor (e.g. Lewis, 1954).

2 Data

The data we use are particularly suited to the inquiry: a nationally representative rural household panel from a developing country which spans three decades. We are aware of no other such resource; the closest approximations are small sample data 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.

In the late 1960s the National Council of Applied Economic Research (NCAER) initiated a panel study of rural households. Over 250 villages in 100 districts were identified so as to be representative of India's rural population in 17 major states. From these villages, just over 4500 households were surveyed in three rounds (crop years 1968-1969, 1969-1970, and 1970-1971). This Additional Rural Incomes Survey (ARIS) provides an array of information about income by source, wealth, and human capital attainment. In 1982, the Rural Economic Development Survey (REDS) located and resurveyed approximately 70 percent of the original sample—a large chunk of attrition was due to Assam being dropped due to political disturbances. 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 all original ARIS households, excluding eight in Jammu and Kashmir due to ongoing conflict, and again added a small additional random sample, bringing the sample to almost 7500 households.[1] As a relatively early series of large sample economic microdata from a developing country, individual as well as linked rounds of the ARIS-REDS have long represented a valuable resource for

[1] A new round was collected in 2006, but the data have not been publicly released.

researchers. Perhaps most prominent are a series of papers by Andrew Foster and Mark Rosenzweig and coauthors, e.g. Rosenzweig and Wolpin (1980); Foster and Rosenzweig (1995, 1996); Behrman et al. (1999); Foster and Rosenzweig (2002).

In this study, we focus on the long panel of income data. We generate a measure of income that maximizes consistency across rounds, aggregating reported income from specific sources (self-employment in and outside farming, agricultural and non-agricultural wages, salaries, interests and dividends on financial investment, rents from land and house property, pensions and transfers) as well as imputed income (rental income from owner occupied houses and imputed income of own-farm labor). Since most of our analysis focuses on income dynamics, we only include in our estimation samples households that were observed more than once. We also consider for robustness various specifications, including households that were observed in all three rounds, and alternatively households as originally defined from the earliest round, to avoid bias in our convergence estimates due to household splitting.[2]

For most of the analysis we consider two "rounds" of income growth: from 1969 to 1982 and from 1982 to 1999. We approximate average annual income growth as the change in the log of income from the baseline divided by the length of the round, where the baseline is 1969 for the first round and 1982 for the second.

3 Poverty Traps and the Dynamics of Income

[2]Split-off households are included in the ARIS-REDS data only in specific circumstances, as discussed in Foster and Rosenzweig (2002).

3.1 Tests for Poverty Traps and Convergence

Let y_t be household income at time t, where $t+1$ denotes either the same household in the future or the next generation of the dynasty. The household has some persistent asset—capital or wealth—that determines income today. Income today in turn partly determines the level of the asset tomorrow. Formally, let

$$y_t = f(k_t; \theta) \tag{1}$$

$$k_{t+1} = T(y_t, k_t; \gamma) \tag{2}$$

For the right choices of f and g this simple framework nests the dynamics of a more complex model such as Banerjee and Newman (1993) and Galor and Zeira (1993). The vectors θ and γ contain parameters that determine the economic environment.

Suppose f is strictly increasing in k_t, meaning more assets always produce more income. Rewrite assets as a function of income y_t, then rearrange:

$$\begin{aligned} y_t &= f(k_t;\theta) \\ \Rightarrow k_t &= f^{-1}(y_t;\theta) \end{aligned}$$

$$\begin{aligned} y_{t+1} &= f(k_{t+1};\theta) \\ &= f[T(y_t, k_t;\gamma);\theta] \\ &= f[T(y_t, f^{-1}(y_t;\theta);\gamma);\theta] \\ &= h(y_t;\theta,\gamma) \end{aligned}$$

Future income is now a function of current income, and the long run dynamics of income depend entirely on the shape of $h(\cdot)$. Figure 1 illustrates the outcomes when the transition function takes different shapes.

[Figure 1 about here.]

It is a new statement of an old result. If the transition function crosses the 45 degree line only once, as function h_1 does, all households converge to the same income. When the transition function crosses more than once, as function h_2 does, there is more than one steady state. Households in the lower steady state are in a poverty trap.

One test for poverty traps is to estimate the relationship between future and current income. Suppose future income depends on a shock ε_{t+1} as well as current income—for example, a shock to the productivity of capital. Then

$$y_{t+1} = h(y_t;\theta,\gamma) + \varepsilon_{t+1}$$

We can estimate h by running a nonparametric regression of y_{t+1} on y_t. The simplest test for poverty traps is to count how often h crosses the 45 degree line. If it crosses only once we can reject the presence of a poverty trap.

Alternatively we can study how growth in income varies with baseline income. Define

$$\begin{aligned} g^y_{t+1} &= \frac{y_{t+1} - y_t}{y_t} \\ &= \frac{h(y_t; \theta, \gamma) - y_t}{y_t} + \frac{\varepsilon_{t+1}}{y_t} \\ &= g(y_t; \theta, \gamma) + \tilde{\varepsilon}_{t+1}. \end{aligned}$$

Figure 2 shows that $g(y_t; \theta, \gamma)$ crosses the income axis at each steady state. If it crosses with negative slope the steady state is stable; otherwise it is unstable. Poverty traps, which make $g(y_t; \theta, \gamma)$ non-monotonic, are easy to spot. One need but estimate $\mathbb{E}[g^y_{t+1} \mid y_t]$ and check if at any point it slopes upward.

This equation is similar to that estimated in Barro and Sala-i Martin (1992), the difference being that our test is nonparametric. It can detect if the growth rate crosses zero multiple times, which would imply there are poverty traps. Most importantly, any convergence we find is unconditional. If income is correlated with unobserved parameters (e.g. if some households have better health, have better institutions, or invest more than others), convergence implies that these differences do not prevent poor households from catching up with rich households.

[Figure 2 about here.]

3.2 Empirical Evidence of Convergence

Figure 3.A plots the distribution of log income in 1969, 1982, and 1999. Though most of the distribution does not move between 1969 and 1982 the lower tail rises; a household is less likely to be very poor in 1982 than in 1969. This decline in severe poverty is the first sign of income convergence. By contrast the most striking change from 1982 to 1999 is how the distribution shifts to the right. Though it does not dominate the old distribution, the new distribution still suggests a wealthier society. The reason is clear in Figure 3.B, which plots the density of the average annual growth in income for individual households. The median rate of income growth is only 1.1 percent from 1969 to 1982; many households grew poorer over this period. But from 1982 to 1999 income grew more than twice as quickly, and more than two-thirds of households grew richer. Wealth grew even more quickly; half of households saw their wealth grow at a yearly rate of over 3.8 percent.

[Figure 3 about here.]

[Figure 4 about here.]

But the distribution of income can rise even if poor households stay trapped in poverty. If it exists such a poverty trap would cause the income of a poorer household to grow more slowly. Figure 4.A plots what Section 3.1 calls $h(y; \theta, \gamma)$, the transition of log income from 1969 to 1982. A poverty trap would cause this function to cross the 45-degree line more than once. We find the opposite. The relationship between current and future income is almost linear and crosses only once. Figure 4.B tells a similar story about the transition from 1982 to 1999.

The relationship between current and future income crosses the 45-degree line at just one point.

Figures 4.C and 4.D show the relation between baseline income and income growth, what Section 3.1 calls $g(y; \theta, \gamma)$. Each point where the relation crosses the x-axis—that is, each level of income at which growth is zero—is a steady state. Poverty traps arise when there are multiple steady states. But in both periods a higher income predicts lower growth, meaning the poor on average catch up to the rich. The relation is monotonic, ruling out multiple steady states.

Could measurement error be the real cause of this supposed convergence? If each household's income is actually fixed but measured with error—because households forget their earnings or because income is measured inconsistently across rounds—it would look as though households are all converging to the same level of income. Though we take pains to make our measures consistent, income in poor countries is always measured with error. Wealth by contrast is easier to measure because it reflects the value of the household's assets. Figure 5.A shows a convergence diagram for wealth from 1982 to 1999 (the earliest round lacks the data needed to make a consistent measure of wealth). Just like the growth-income relation the growth-wealth relation is a downward-sloping line. Savings, shown in Figure 5.B, is easily measured though more volatile than income or wealth. It too converges.

[Figure 5 about here.]

If measurement error is not perfectly correlated across time—the shopkeeper who overestimates her revenue in 1969 does not overestimate it by exactly the same percentage in 1970—then averaging a household's income across time will shrink the error. Since income grows slowly from 1969 to 1982 we can use the average of income in 1969, 1970, and 1971 as a cleaner measure of baseline income without introducing too much bias. Figure 5.C compares the rate of convergence calculated from actual 1969 income to that calculated with the cleaned measure of income. Averaging across three rounds purges some random but real shocks to income, reducing the rate of convergence even when income is measured without error. Even so, income still converges.

One last test for poverty traps is to check whether negative shocks cause the poor to grow poorer. In a typical model of poverty traps, a household that suffers a bad shock must sell its capital. For example, a poor farmer who reaps a bad harvest must sell his bullocks to feed his family. Come the next season he cannot plow his fields, leaving him even poorer. If bad rainfall drops poor households into a spiral of negative income growth, it might be evidence of a poverty trap. We measure the effect of rainfall in the kharif season (wet season) and rabi season (dry season) on household income in 1982, and use the coefficients to predict the effect of rainfall in 1983 on income. If there is a poverty trap the households with the least income in 1982 should have the trajectory of their income growth knocked down by negative shocks in 1983. But Figure 5.D shows that bad shocks do not hurt the poorest decile any more than the general population. The graph plots the smoothed effect of the standardized 1983 rainfall shock for the entire sample alongside the effect on just the poorest 10 percent. The predicted growth for the poorest households is everywhere above the general population—a sign of convergence—but looks otherwise identical. It does not turn negative for negative shocks.

Income does converge, but how does it converge? Do the poor within each village catch up to the rich, or do poor villages simply catch up to rich villages?

Figure 6.A plots growth within a village—that is, growth after removing village fixed effects—against within-village baseline income. If anything, income converges more rapidly within villages. Linear regressions show that a log point increase in overall income predicts a decrease in growth of 3.4 percentage points. But a log point increase in within-village income predicts a decrease in within-village growth of 4.1 percentage point. Figure 6.B plots the growth of each village's average income against its average baseline income. Given that our sample has fewer than 250 villages, it is hard from this graph to be sure that income does converge between villages. If it does, it converges more slowly. After imposing linearity, we estimate that a log point increase in average income predicts a decrease in the growth of average income of 2.1 percentage points. Convergence happens mainly within rather than between villages.

[Figure 6 about here.]

Does income grow and converge with equal speed across all of India? Figure 7 highlights observations from several Indian states in the scatter plot of income in 1982 against growth from 1982 to 1999. Though income converges in all states, income does not converge as favorably in Karnataka as it does in Kerala. On average, households who start at a similar level of income get rich more quickly in Kerala than elsewhere. The opposite is true in Bihar. A poor Bihari household gets rich more slowly than an equally poor household elsewhere. Bihar's feeble growth should surprise no one; when observers of India need an example of poverty and misgovernance, they usually point to Bihar. Kerala's excellence is more surprising. Though praised for its investments in education and health, Kerala has struggled to convert those investments into higher aggregate output. Evidently it has converted them into higher income for rural households.

[Figure 7 about here.]

[Figure 8 about here.]

But even as the data yield clear evidence of convergence they also yield a paradox. Convergence should reduce inequality. According to Figure 8, which plots the standard deviation of log income in each round, that does seem to happen from 1969 to 1982. But between 1982 and 1999 income inequality rises sharply; the same pattern holds for the Gini or other measures of inequality. Although we find convergence—the poor catch up to the rich—we find no evidence of its most obvious consequence.

4 The Solow Household

Tracking the dynamics of income for thousands of households is monstrously complex. But to resolve the paradox in the data we must reconcile two parameters, the rate of convergence and the level of inequality, that are functions of the entire distribution of household income. The challenge is to find a computationally feasible way to simulate how this distribution evolves.

Our solution is to adapt the Solow model to describe household income. Depending on the parameters, the model can reproduce convergence or divergence, and inequality that rises or falls. Yet it remains simple enough to simulate for thousands of combinations of parameters. By trying different combinations we can check if any unchanging economic environment can reconcile rising inequality with rapid convergence. And when we ultimately show that no unchanging environment can, we can easily simulate and interpret an environment in transition.

4.1 A Simple Model of Income Dynamics

First we prove that our adaptation of the Solow model can produce either convergence or divergence. Consider the per-capita production function

$$\begin{aligned} f(k_t;\theta) &= \frac{1}{L_i} K_i^{\alpha} (A_v L_i)^{1-\alpha} \\ &= k_i^{\alpha} (A_v)^{1-\alpha} \end{aligned}$$

where k is a generalized form of capital. A household's generalized capital includes everything that will boost its income this year and persist in some form into next year. Physical capital—the tractor used to grow wheat or the oven used to bake bread—is just one part. The education needed to run a business or the children who send remittances are also part of generalized capital.

Now let the capital transition function be

$$T(y_t, k_t; \gamma) = s_v y_i - \delta k_i$$

and suppose for now that time is continuous, so g represents the instantaneous change in capital. The subscript v is a village, and $i \in v$ is a household in the village. The important parameters—the investment rate and the level and rate of technological progress—can vary by village v. We force parameters to be the same within each village because, as shown in Figure 6 of Section 3.2, most of the convergence we find is within rather than between villages.

Since members of a village share parameters, income must converge within each village. But income only converges across the sample if parameters are not too different. In this way, the model may produce either convergence or divergence. It is this combination of simplicity and flexibility that makes a model of Solow households ideal to test whether the economic environment remains fixed.

To prove that the model may produce either convergence or divergence let g^X be the growth rate of variable X. Then income growth is

$$g_i^y = (1-\alpha)g_v^A + \alpha g_i^k$$

From the law of motion,

$$g_i^k = s_v \left(\frac{A_v}{y_i}\right)^{\frac{1-\alpha}{\alpha}} - \delta.$$

and back-substitution gives

$$g_i^y = (1-\alpha)g_v^A + \alpha s_v \left(\frac{A_v}{y_i}\right)^{\frac{1-\alpha}{\alpha}} - \alpha\delta \tag{3}$$

Linearize income growth (3) in terms of $\log y_i$ around the sample means $\bar{y}, \bar{A}$:

$$g_i^y \approx \omega_1 + \omega_2 g_v^A + \omega_3 s_v + \omega_5 s_v \log A_v^0 + \omega_4 s_v g_v^A - \omega_5 s_v \log y_i \tag{4}$$

where

$$\bar{X} = (\bar{A}/\bar{y})^{\frac{1-\alpha}{\alpha}}$$
$$\omega_1 = -\alpha\delta$$
$$\omega_2 = 1 - \alpha$$
$$\omega_3 = \alpha\bar{X}(1 - \log\bar{X})$$
$$\omega_4 = (1-\alpha)\bar{X}t$$
$$\omega_5 = (1-\alpha)\bar{X}$$

Define the aggregation error as

$$\begin{aligned}\varepsilon_i &= \omega_2(g_v^a - g^a) + \omega_3(s_v - s) + \omega_5(s_v \log A_v^0 - s\log A^0)\\ &+ \omega_4(s_v g_v^A - s\log g^A) - \omega_5(s_v - s)\log y_i \end{aligned} \tag{5}$$

where variables with no subscripted v are sample-wide averages. Then we can rewrite linearized income growth (4) as

$$g_i^y \approx \beta_0 + \beta_1 \log y_i + \varepsilon_i. \tag{6}$$

Equation 4.1 is exactly the regression used to estimate the rate of convergence in Section 3, and $\hat{\beta}_t < 0$ is the condition for sample-wide convergence. Applying the Omitted-Variable Bias formula shows that it is a sufficient condition for convergence if the variance of ε_i is small relative to the variance of

$\log y_i$.[3] As long as parameters vary little between villages, income across the sample converges.

It is the lack of variation in parameters that rules out an unchanging process for income growth. Income converges because villages have similar steady states, but if villages have similar steady states inequality must decrease. Only by changing the process, as we do in Section 4.3, can the model reconcile convergence and inequality.

4.2 Simulating the Model

We make time discrete for simplicity and simulate the model at the annual level. We simulate 242 villages (the total number that appear in the data) with 17 households per village, making a sample roughly the same size as the number of unique households with positive income in 1969. We assume away household splitting and attrition.[4] We set the depreciation rate to 0.1 and set the initial level of technology to match the median rate of income growth from 1969 to 1982.[5]

Figure 9 summarizes how we assign parameters to villages and households. We assume the investment rate s_v and the rate of technological progress g_v^A vary between but not within villages. We assume investment is independent of technological progress, and that both have normal distributions. Realizations of the

[3] See, for example, Section 3.2.2 of Angrist and Pischke (2009) for a derivation and interpretation of the formula.

[4] The data still show convergence when we study dynasties rather than households, which means household splitting does not drive the results. The tendency to attrit does not seem strongly correlated with baseline variables, suggesting attrition does not drive the results, either.

[5] We simulate the model for many combinations of the other parameters and a range of values for A_0. For each combination of parameters we keep the value of A_0 that produces a median rate of income growth closest to what we see in the data. We then regress these values of A_0 on a flexible polynomial of the other parameters of the simulation and use the coefficients to impute the best A_0 for the wider range of parameters used to run our later simulations.

investment rate less than zero are set to zero, and those greater than 1 are set to 1. We vary the mean and variance of the distribution for each parameter to find the rate of convergence closest to what we estimated in Section 3.2 (see below). The production return to capital α is the same for all households, and we vary it as well to match the rate of convergence. We assume the mean and standard deviation of log income within each village has a normal distribution, and the mean and standard deviation are themselves normally distributed across villages. We calibrate these distributions to match their analogs in the data from 1969.

[Figure 9 about here.]

We simulate the model from 1969 to 1999 and discard income for all years except 1969, 1982, and 1999. We calculate two statistics from Section 3 using the simulated data: the rate of convergence, measured as the coefficient β_1 from the regression of annual income growth from 1982 to 1999 on the log of 1982 income; and the standard deviation of log income in 1999. For each set of parameters we run 20 simulations and take the average of each statistic across simulations.

[Table 1 about here.]

Figure 10A shows the rate of convergence and the standard deviation of income from 65610 different combinations of parameter values, where each circle represents one combination. The red lines intersect at the rate of convergence and standard deviation of income we observe in the data (see Section 3). Table 1 shows the five simulations that get closest to the correct rate of convergence. The table suggests an investment rate of .31, a return to capital of .5, and an average rate of technological progress of .015 best fit the data. The model can only reproduce the rate of convergence in the data when the important village-level parameters—the investment rate and the rate of technological progress—do not vary much. Figure 11 holds the average investment rate, average rate of technological progress, and the production elasticity of capital fixed while varying the standard deviations of the investment rate and technological progress, plotting them on the x- and y-axis against the resulting rate of convergence on the z-axis. The basin of convergence—the set of parameters that cause poor households to grow faster than rich ones—are colored blue. The model only comes close to the convergence in the data when the investment rate and rate of technological progress do not vary at all.

[Figure 10 about here.]

[Figure 11 about here.]

4.3 An Unchanging World Versus A Rise in Opportunity

Figure 10A shows that the model can only get the standard deviation right if it gets the rate of convergence wrong. Any set of parameters that produces rapid convergence must also leave very little variation in the distribution of income by 1999.

Simply adding random shocks to income—in this case, random deviations in a household's productivity from the village mean A_v—does not solve the problem. Still using the parameters that produce the closest rate of convergence ($\mathbb{E}[s] = .31, \alpha = .5, \mathbb{E}[g^a] = .015, SD[s] = SD[g^A] = 0$), we vary the standard deviation of an i.i.d. log normal productivity shock. To be precise, we assume

$$A_t = (1 + g^a)A_{t-1} \cdot \exp(u_t)$$

where $u_t \sim N(0, \sigma^2)$. The parameter we vary is σ^2. Figure 10B plots the rate of convergence and the standard deviation of income when the standard deviation of the productivity shock varies from 0 to 1 in increments of .01. The diagram still shows that rapid convergence comes only with low inequality. Differences in productivity within villages cause divergence for exactly the same reason as differences between villages. To summarize, the standard Solow assumption of fixed parameters—which mirrors the assumption that poor economies are stagnant—cannot explain the data.

We now drop this assumption and let the standard deviation of the productivity shock rise halfway between 1982 and 1999. A positive shock to productivity lets households get more income from less capital. One way to interpret a more volatile shock is that a household's assets matter less for its prospects. A landless laborer is able to start a business or a poor widow can send her son to work in the city. Another interpretation is that greater variance in productivity represents greater ease in acquiring generalized capital. The child of an illiterate might get a college education and start sending remittances between two rounds of the survey. With either interpretation a more volatile shock makes a household's past less informative about its future. In this sense, bigger shocks represent a rise in social mobility.

Keeping the other parameters fixed, we let the standard deviation of the productivity shock change in 1990. Panels C and D of Figure 10 show contour plots of the rate of convergence and standard deviation of log income in 1990. Each point is a different combination of pre-1990 and post-1990 values for the standard deviation of the shock. Panel C shows that the post-1990 shock is almost irrelevant to the rate of convergence. The height of the contour rises as the pre-1990 shock gets more volatile but remains flat as the post-1990 shock changes. In contrast, Panel D shows that the standard deviation of log income in 1999 does rise when the post-1990 shock becomes more variable. Taken together Panels C and D suggest an initial regime with very small shocks—a standard deviation close to 0—and a later regime with much larger shocks—a standard deviation between .5 and .7—can reconcile convergence with rising inequality. The change represents a sharp rise in opportunity. Under the first regime the probability of a shock that at least doubles or halves expected income is nearly zero, whereas under the second regime the probability is between 20 and 30 percent.

Do the data confirm the model's story of rising opportunity? In ongoing work we compute for each period the intergenerational mobility between income quartiles, levels of education, and other measures of household well-being. Figure 12, taken from that paper, shows father-son educational mobility in 1971, 1982, and 1999. We group each set of bars by the father's level of education, and the colored areas represent the fraction of sons over 22 who report having a particular level of education (the lowest area is illiterate, the next lowest primary or below, and so on in ascending order). Since the sons in 1971 are likely the household heads in 1982, the bars marked ”71” describe the transition from 1971 to 1982 (which can stand in for the period 1969 to 1982) while the bars marked ”82” do the same for 1982 to 1999. The change between the two periods is stark. The son of an illiterate in 1971 was almost certain to be illiterate himself, suggesting almost zero educational mobility for the least educated group from 1971 to 1982. In 1982, however, the son of an illiterate was more likely to be literate than not. The odds that the child of an illiterate has tertiary education in 1982 is bigger than the odds he is even literate in 1971. The cause may be India's periodic bouts of school construction. Public schools gave children a chance to be more educated than their fathers. The model interprets the change as a rise in opportunity.

[Figure 12 about here.]

5 Conclusion

Our results suggest that the latter part of the 20th century was a period of dramatic transformation in rural India. We find no evidence that the rural Indian households in our sample were trapped in poverty between 1969 and 1999. Yet while households were unshackled from their dynastic history, income inequality climbed dramatically. We simulate a model of Solow households to show that rapid convergence and rising inequality are inconsistent when the parameters governing income growth remain fixed. We resolve the inconsistency by allowing the variance of productivity shocks to rise between 1982 and 1999. In contrast to models in which rural economies remain mired in stagnation, our findings support theories that emphasize regime transition and structural change.

References

Angrist, J. and Pischke, J. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton Univ Pr.

Azariadis, C. and Stachurski, J. (2005). Poverty traps. In Aghion, P. and Durlauf, S., editors, *Handbook of Economic Growth*, volume 1A, pages 295–384. North-Holland.

Banerjee, A. V. and Newman, A. F. (1993). Occupational choice and the process of development. *Journal of Political Economy*, 101(2):274–298.

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

Barro, R. J. and Sala-i Martin, X. (1992). Convergence. *Journal of Political Economy*, pages 223–251.

Behrman, J. R., Foster, A. D., Rosenzweig, M. R., and Vashishtha, P. (1999). Womens schooling, home teaching, and economic growth. *Journal of Political Economy*, 107(4):41–57.

Dercon, S. and Outes, I. (2009). Income dynamics in rural india: Testing for poverty traps and multiple equilibria. Working Paper.

Foster, A. D. and Rosenzweig, M. R. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*, 103:1176–1209.

Foster, A. D. and Rosenzweig, M. R. (1996). Technical change and human capital returns and investments: Evidence from the green revolution. *American Economic Review*.

Foster, A. D. and Rosenzweig, M. R. (2002). Household division and rural economic growth. *Review of Economic Studies*, 69:839–869.

Galor, O., Moav, O., and Vollrath, D. (2009). Inequality in land ownership, the emergence of human capital promoting institutions, and the great divergence. *Review of Economic Studies*, 76:143–179.

Galor, O. and Weil, D. N. (2000). Population, technology, and growth: From malthusian stagnation to the demographic transition and beyond. *American economic review*, 90(4):806–828.

Galor, O. and Zeira, J. (1993). Income distribution and macroeconomics. *Review of Economic Studies*, 60(1):35–52.

Hansen, G. D. and Prescott, E. C. (2002). Malthus to Solow. *American Economic Review*, pages 1205–1217.

Himanshu and Stern, N. (2011). India and an Indian village: 50 years of economic development inPalanpur. Working Paper.

Kraay, A. and McKenzie, D. (2014). Do poverty traps exist? assessing the evidence. *Journal of Economic Perspectives*, 28:127–48.

Krugman, P. (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.

Lewis, W. (1954). Economic development with unlimited supplies of labour. *Manchester School of Economic and Social Studies*, 22(2):139–91.

Naschold, F. (2009). "Poor stays poor": Household asset poverty traps in rural semi-arid India. Working Paper.

Rosenzweig, M. and Wolpin, K. I. (1980). Testing the quantity-quality model of fertility: The use of twins as a natural experiment. *Econometrica*.

A Data Appendix (For Online Publication)

Measures of household income and household wealth are constructed for each response period of the ARIS-REDS. The definition and computation of household income is identical across the 1969-71, 1982, and 1999 rounds of ARIS-REDS. Household income is defined as receipts net of expenditures and is computed as the sum of income from various sources. These sources of income that sum to equal total household income are: income from agriculture, plantations, and orchards; income from self-employment in farm activities, (livestock and allied activities, which include 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. In addition, the imputed value of family labor for investment is included in this measure.

For each round of ARIS-REDS, household wealth is defined as the owners equity of the household. Thus household wealth is computed as the value of all assets owned at the beginning of the response period net of all outstanding liabilities at the beginning of the response period. However the computation of total assets varies in different rounds.

For 1982 and 1999, the total value of assets is the sum of the following variables (all at the beginning of the response period): real value of buildings owned and non-house land owned; real value of irrigation assets owned; real value of farm equipment owned; real value of animals owned; real value of animal-related assets owned; real value of non-farm business assets and inventory owned; real value of consumer durables owned; real value of savings, which is measured as the sum of deposits with commercial banks, post office saving banks and companies, shares and securities, small savings instruments, and gold, jewelry, and currency; and real value of outstanding loans made by the household. The value of outstanding liabilities is measured as the real value of outstanding liabilities in the form of loans borrowed at the beginning of the response period.

Other variables in this constructed cross-panel dataset of the ARIS-REDS include household size, education, and savings. Household size is a single variable in both the 1982 and 1999 rounds, coded as the number of family members living in the household (this number does not include servants or permanent labours staying with the family). For the 1969-1971 round, there are three variables corresponding to household size, a single variable for each year, again coded as the number of family members living in the household.

The education variable is defined as the highest education of the household. This variable is coded as a categorical variable (0-4). The following details the definition of the education variable for each round.

In 1971 and 1982, the categorical breakdown for level of education is: 0 = Illiterate; 1 = Literate but no formal education to Primary or below; 2 = above

Primary but below Matriculation; 3 = Matriculation, Higher Secondary, Intermediate, or Pre-University; and 4 = Graduate in Arts and Sciences, including commerce (B.A., B.Sc., or B.Com.) and Above. In 1999, the categorical breakdown for level of education is: 0 = Illiterate; 1 = Literate but no formal education to Primary; 2 = above Primary but below Middle to above Middle but below Secondary School Certificate; 3 = Secondary School Certificate to Pre-University; and 4 = anything above Pre-University.

For 1982 and 1999, savings is measured as the sum of deposits with commercial banks, post office savings banks and companies; shares and securities; small savings instruments; and gold, jewelry, and currency. Note that this is the same savings measure that is used in the construction of total household assets.

Figure 1
The Transition Dynamics of Income

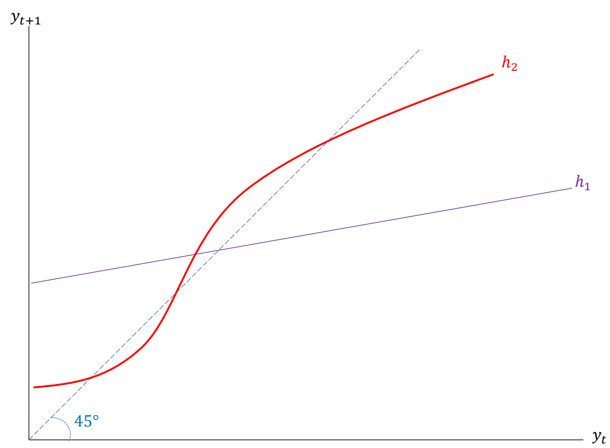

Note: The transition diagram shows the relation between current and future income. The relation h_1 would create convergence to a unique steady-state whereas relation h_2 would create a high-income steady-state and a low-income steady state (a "poverty trap).

Figure 2
Phase Diagram versus Transition Diagram

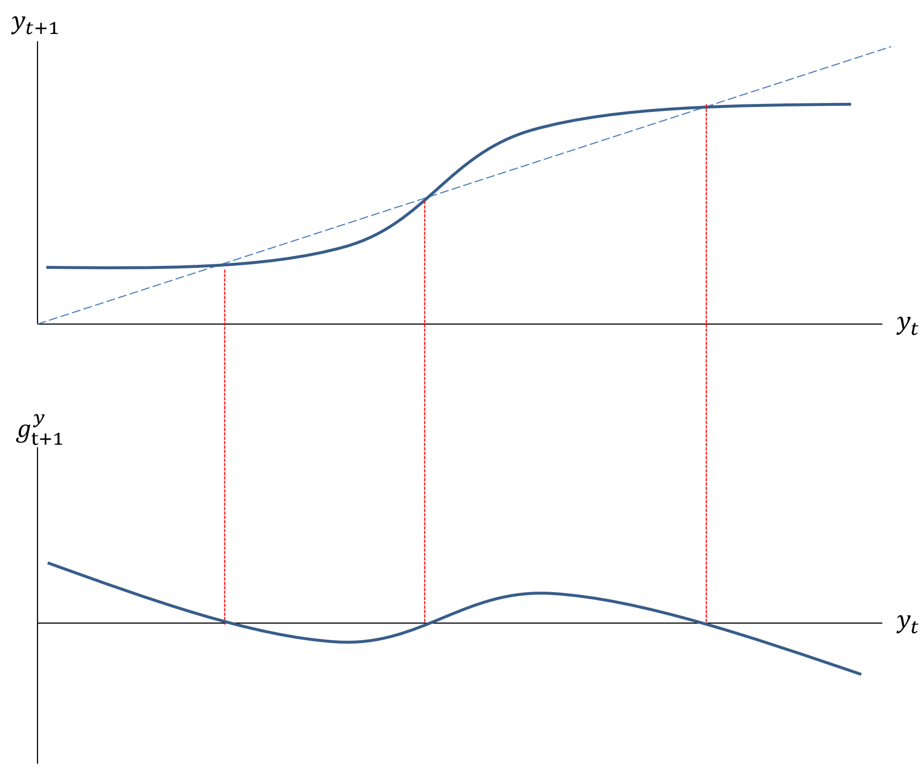

Note: The transition diagram in the top panel has a corresponding phase diagram, which plots the relation between current income and growth in income. There is a steady-state at every point where the relation crosses the income-axis.

Figure 3
Distribution of Income and Growth

A. Cumulative Distribution of Income **B. Density of Income Growth**

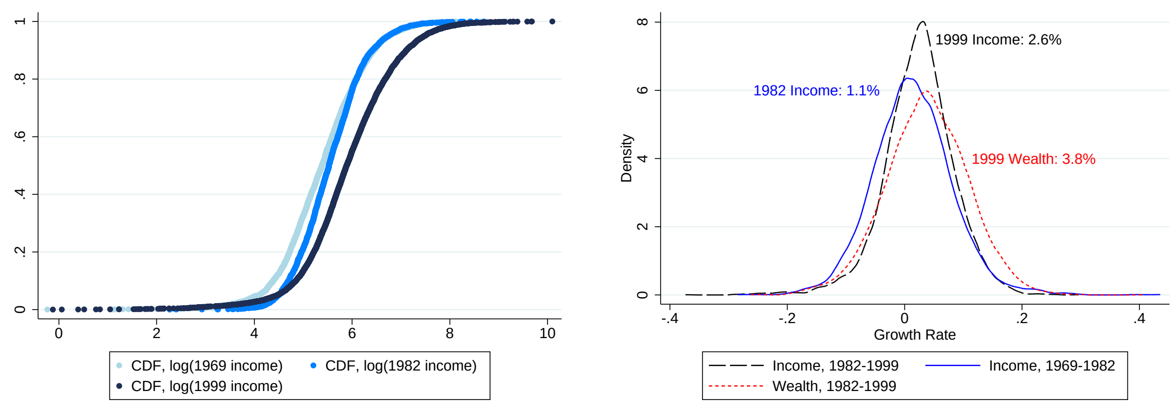

Note: The numbers above the kernel density estimates in Panel B give the mediate rate of growth for either income or wealth between the survey year given and the previous round of the survey.

Figure 4
Income Transition and Convergence Diagrams

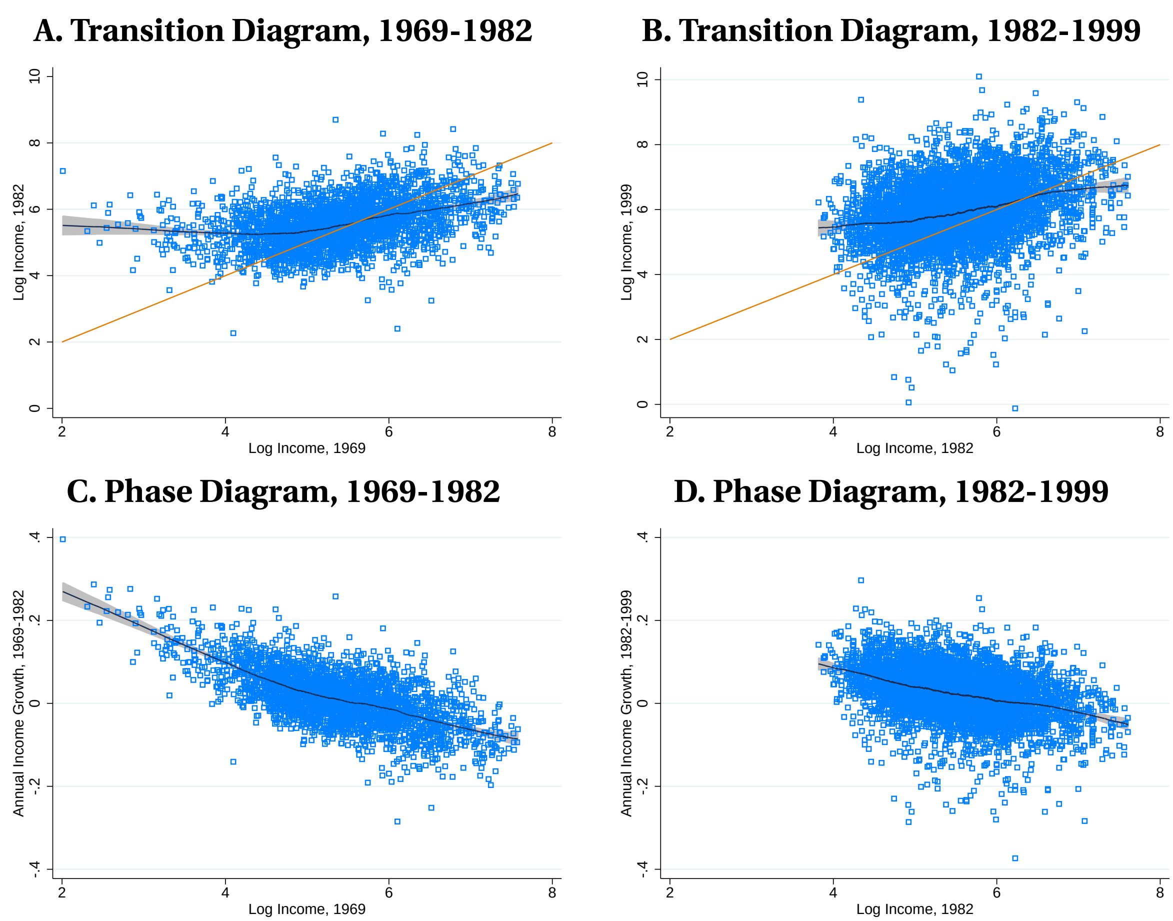

Note: The diagrams show nonparametric regressions (running averages calculated using the "running" command in Stata). The shaded region covers the 95 percent confidence interval at each point.

Figure 5
Verifying Convergence

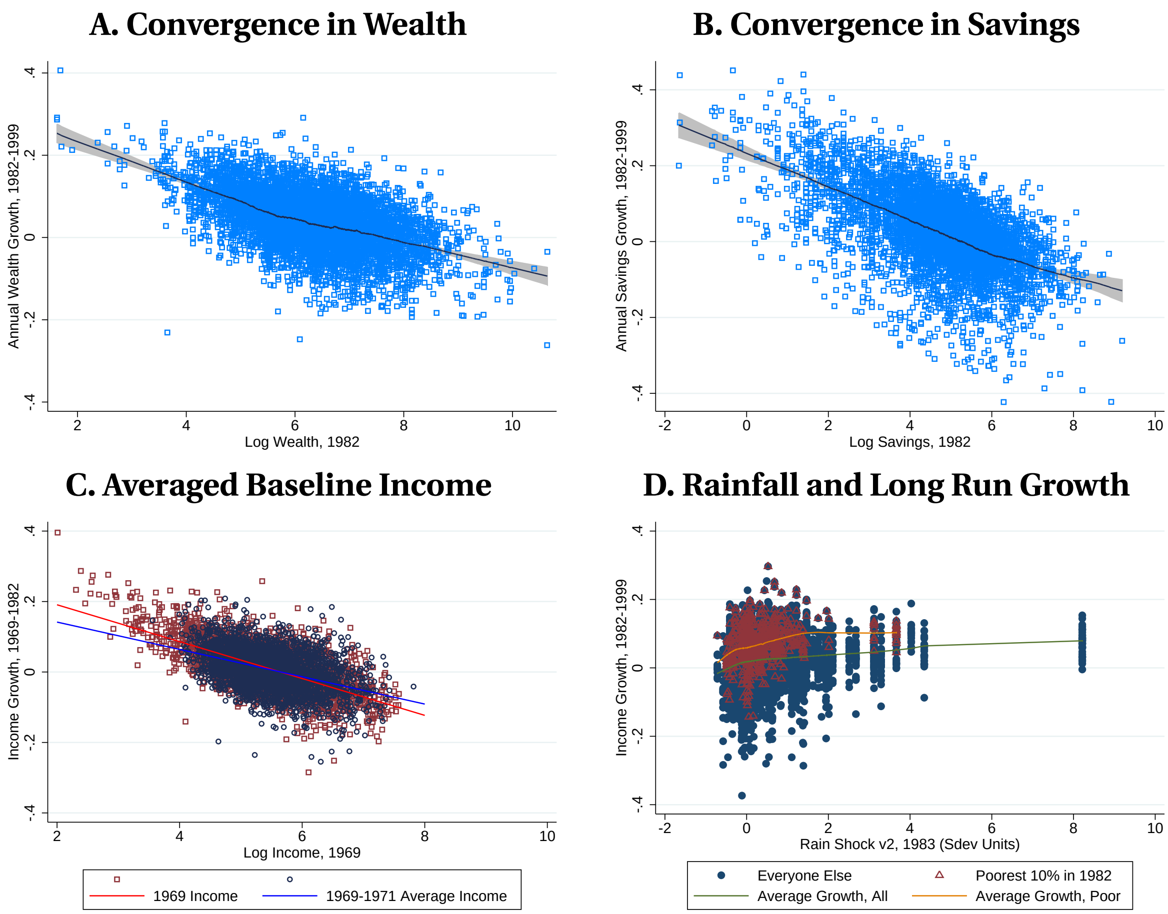

Note: Panels A and B show the convergence in wealth and savings. Panel C compares the rate of convergence calculated using as baseline income the log of income in 1969 with convergence calculated using the average of log income in 1969, 1970, and 1971. Panel D compares the growth rate of households in the lowest decile in 1982 to everyone else in response to a rainfall shock in 1983.

Figure 6

Within- Versus Between-Village Convergence

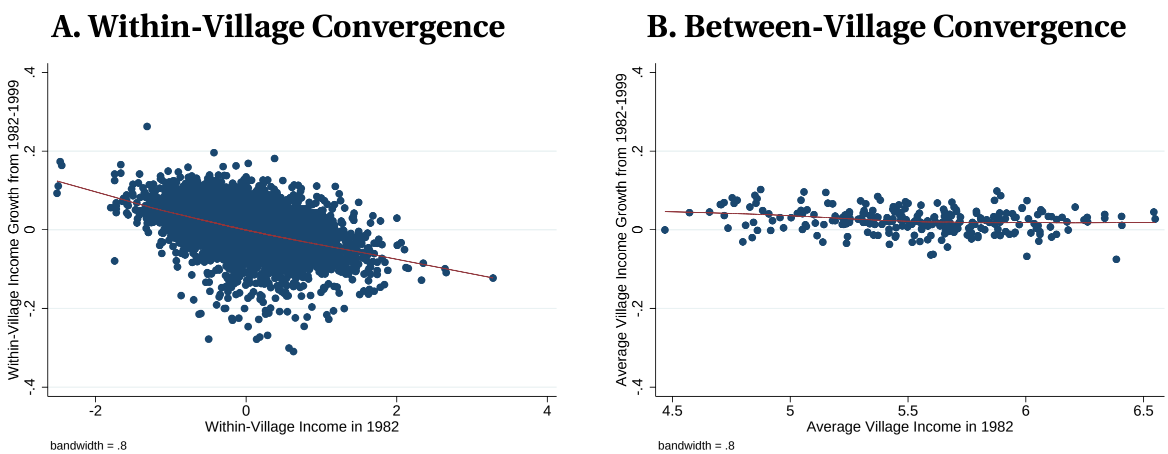

Note: Panel A removes the variation predicted by village fixed-effects in income growth and 1982 income and graphs the nonparametric relation between the residuals. Panel B graphs a nonparametric regression of each villages mean rate of income growth on its mean (log) income in 1982.

Figure 7
Income Converges Within Each State

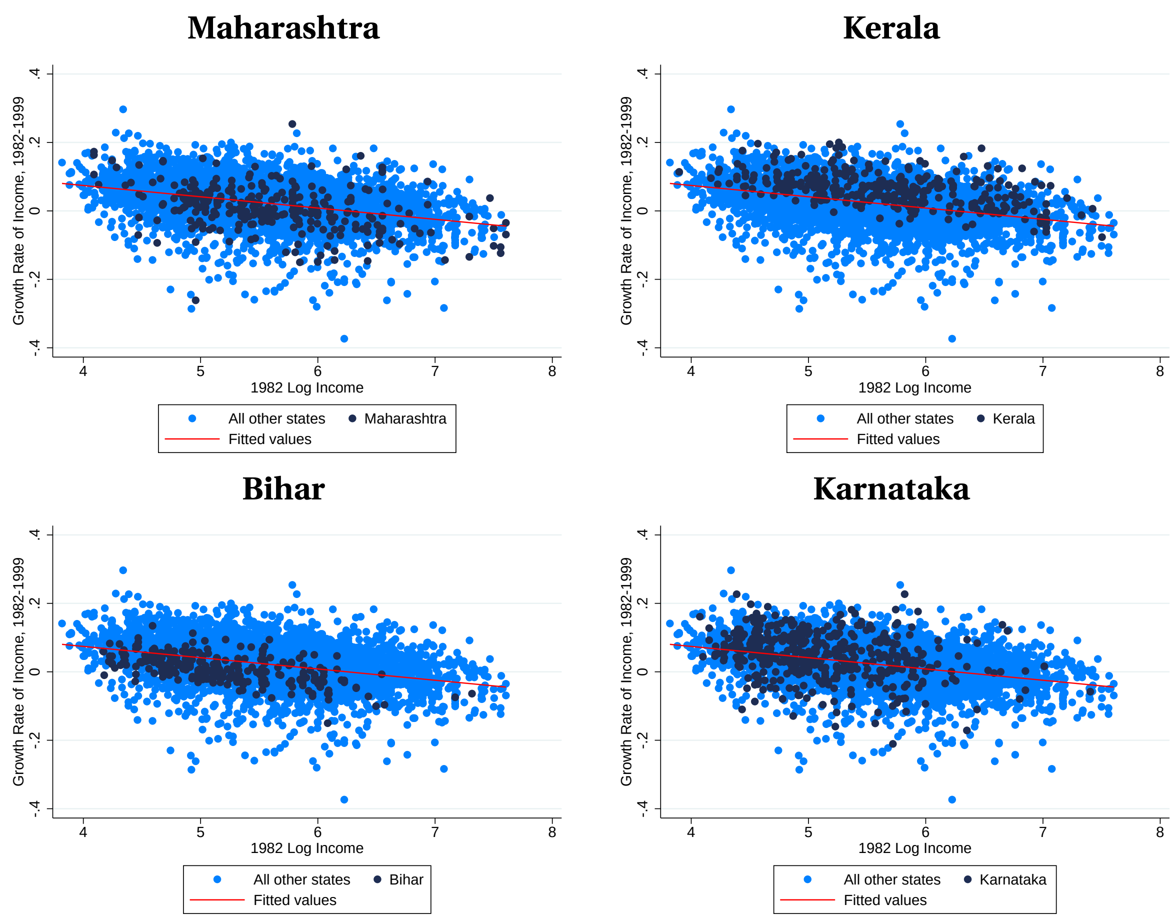

Note: Each panel shows the income-to-growth relationship for the whole sample from 1982 to 1999. Observations in the indicated state are darkened.

Figure 8
Inequality Rises in 1999

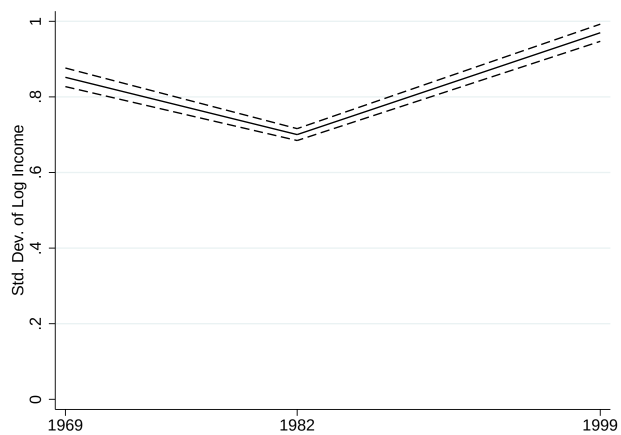

Note: Dashed lines show 95 percent confidence intervals for the estimate of the standard deviation.

Figure 9
Simulation Parameters

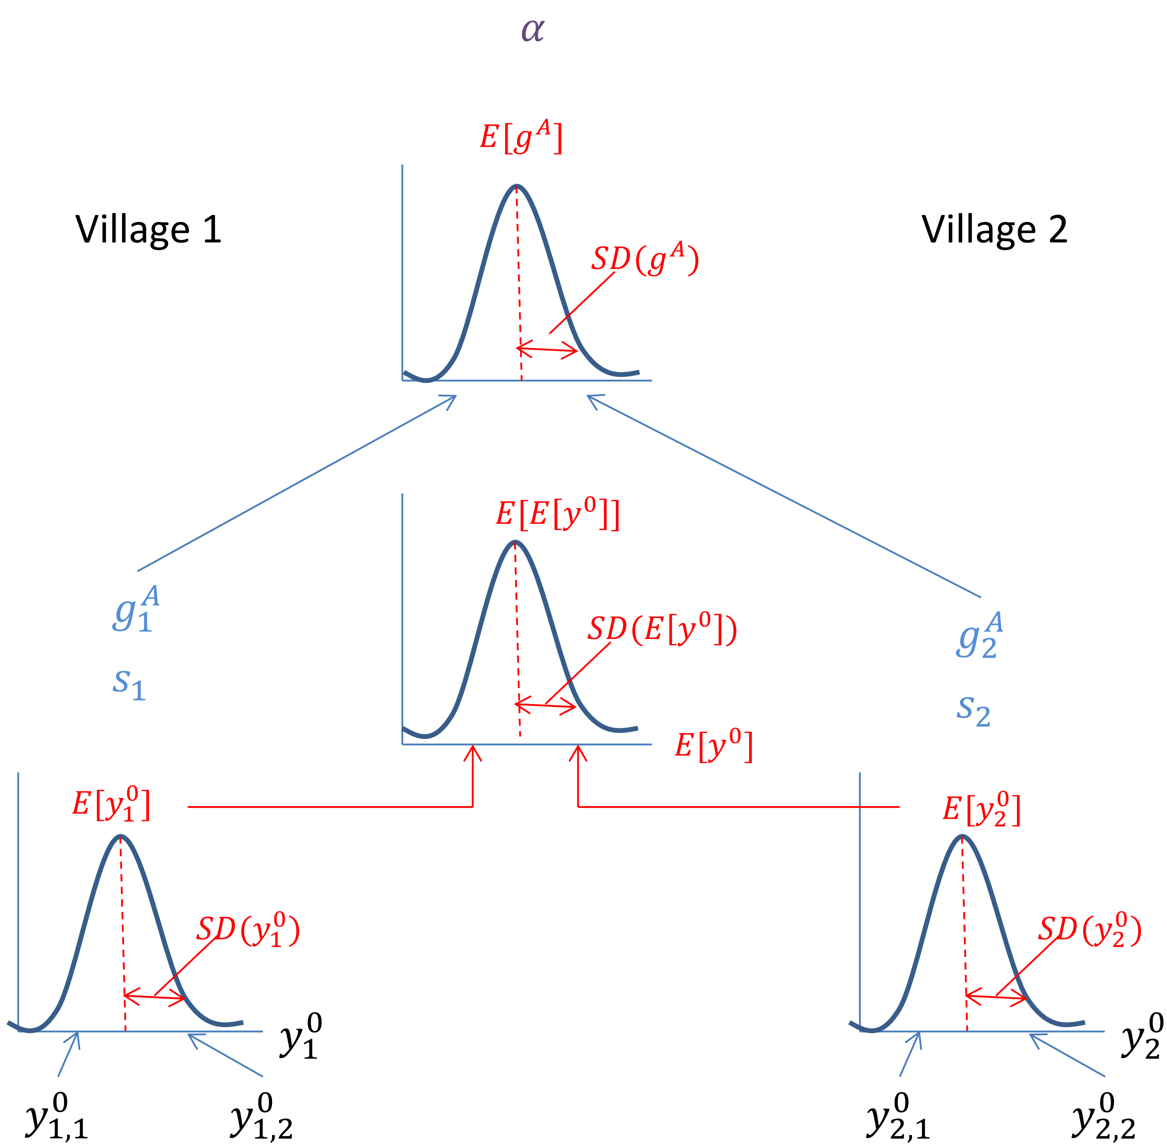

Note: All households within a village have the same savings rate s and rate of technological progress g^A, and across villages each rate is normally distributed. The distribution of baseline (1969) income within each village is normally distributed, and the parameters of each village's income distribution ($\mathbb{E}[y^0], SD(y^0)$) are in turn normally distributed across villages. The production return to capital α is constant across all households in all villages.

Figure 10
Simulation Results

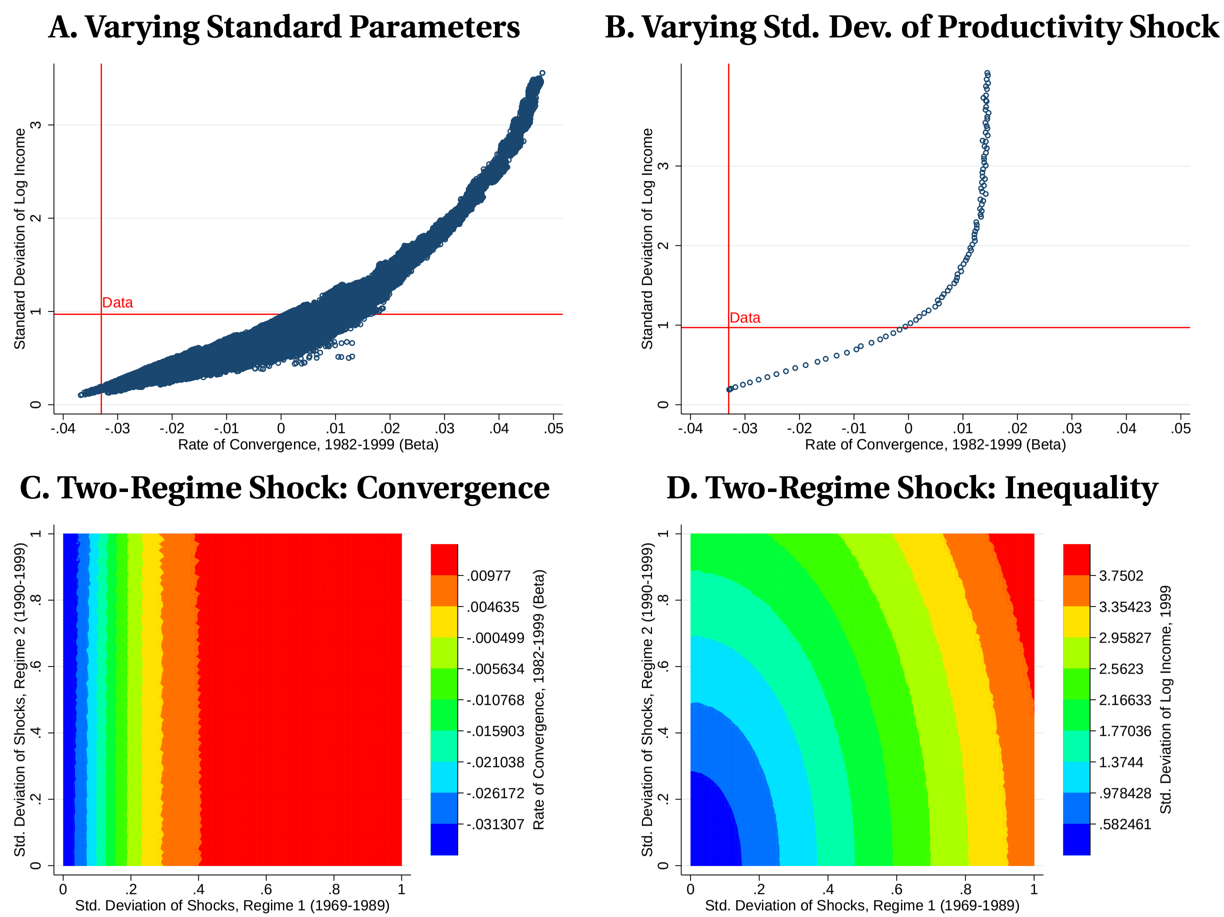

Note: We simulate the model with 65610 sets of values for each parameter. Panel A plots the rate of convergence against the standard deviation of log income in 1999. The red lines intersect at the rate of convergence and standard deviation found in the data. We then take the set of parameters that best replicates the rate of convergence and run simulations that vary the standard deviation of a household productivity shock. Panel B plots the rate of convergence against the standard deviation of income. Finally we run simulations that let the standard deviation of the shock change in 1990. For each combination of pre- and post-1990 values for the standard deviation, Panel C plots a heat map of the rate of convergence. Panel D does the same for the standard deviation of income.

Figure 11
Basin of Convergence

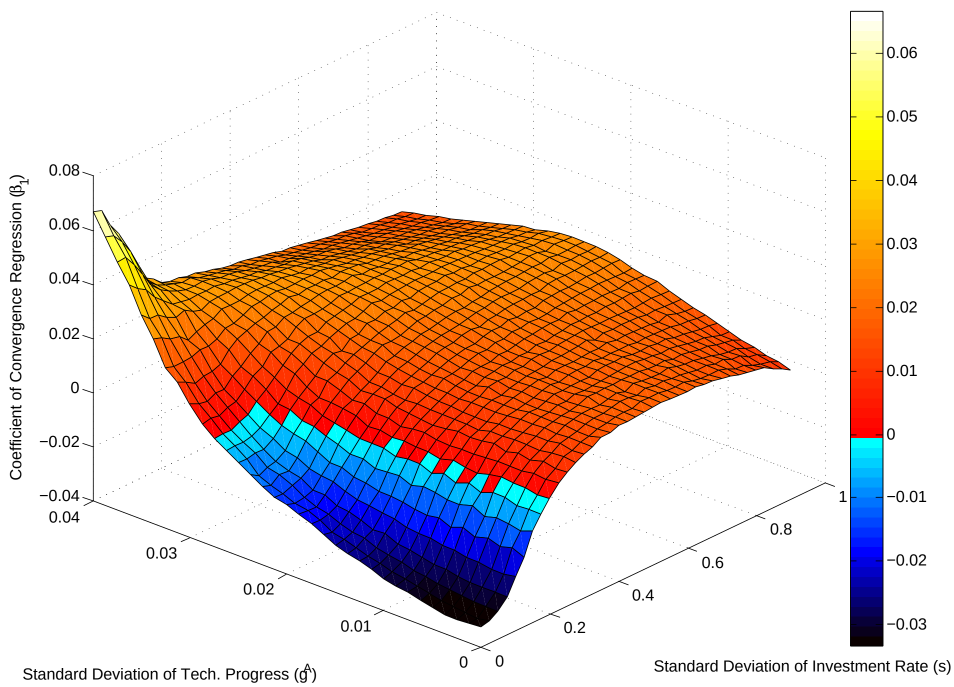

Note: We simulate the model of Solow households varying the standard deviations of technological progress and the investment rate (all other parameters are as shown in the top row of Table 1). The z-axis graphs the rate of convergence estimated by regressing household income growth from 1982 to 1999 on log income in 1982 using data from the simulations. Parameters that produce convergence are colored blue; those that produce divergence are colored red.

Figure 12
Father-Son Educational Mobility

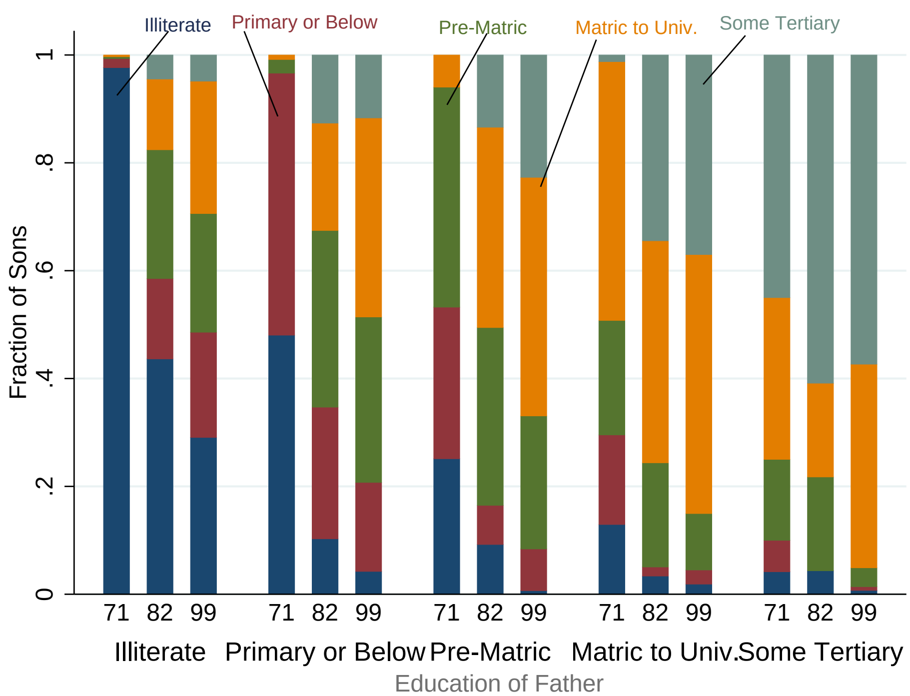

Note: We split sons aged 22 or older into five categories based on the schooling of their fathers. Each bar represents all the sons in a survey year. The bar is split into regions that show what fraction of sons have the indicated level of schooling. The top region is the fraction with some tertiary education, and each region below matches a region on the horizontal axis (so the bottom region is those who are illiterate with no schooling). For example, of the men whose fathers had tertiary education in 1999, roughly 60 percent had tertiary education themselves.

Table 1
Simulation Parameters with Closest Rate of Convergence

$SD[s]$	$SD[g^A]$	$\mathbb{E}[s]$	α	$\mathbb{E}[g^A]$	$\beta_{1,1999}$	$SD[\log y_{1999}]$
0.00	0.00	0.31	0.50	0.015	-0.033	0.20
0.00	0.00	0.91	0.40	0.007	-0.033	0.15
0.00	0.00	0.31	0.50	0.018	-0.033	0.20
0.00	0.00	0.21	0.40	0.015	-0.033	0.17
0.00	0.00	0.21	0.50	0.020	-0.033	0.21

Note: $SD[s]$ and $SD[g^A]$ are the standard deviations of village investment rates and rates of technological progress. $\mathbb{E}[s]$ and $\mathbb{E}[g^A]$ are likewise the mean rates. α is the production return to capital. See Figure 10 for more details on how these moments are used to assign parameters to each village in the simulation. $\beta_{1,1999}$ and $SD[\log y_{1999}]$ are the rate of convergence the standard deviation of log 1999 income estimated from the resulting simulation.