Why do people stay poor? A research agenda for poverty traps

Clare Balboni, Oriana Bandiera, Robin Burgess, Maitreesh Ghatak, & Anton Heil

May 15, 2026

1 Introduction

Since the publication of "Why do people stay poor?" (Balboni et al., 2022), several scholars have used our data and additional data from other settings, to replicate and build on our results (Alloush and Carter, 2024; Jee, 2025; Karlan, Raswan, and Udry, 2026).[1] A focus of this work has been the importance of heterogeneity and understanding which individuals may be trapped in poverty within and across settings. This note discusses the identification of poverty traps with an explicit focus on heterogeneity and sets out a way forward. It draws on selected material from Bandiera et al. (2026).

There are three broad reasons why an individual may be in poverty at any given point in time. First, transitory misfortune: a negative shock – drought, illness, asset loss – has pushed them below the poverty line, but they have the means to recover. Second, persistent disadvantage: low skill, poor health, or other characteristics permanently limit their earning potential and trap them at a low steady state. Third, a poverty trap: they have the capacity to escape poverty, but their wealth falls below a critical threshold from which self-financed escape is not possible. Importantly, heterogeneity in the traits of individuals and of the economy they live in implies that all three explanations might coexist. Across contexts, there is likely to be significant heterogeneity in which, if any, individuals are trapped, as well as in our ability to test for this based on available data. As we show below, understanding the relative importance of the three channels – and the data required to do so – depends critically on the nature of unobservable heterogeneity.

[1] We also thank, without implicating them, Antonio Ciccone, Alexander Göppert, Ellen Munroe, Alexander Newton, and Meet Shah for helpful discussions.

2 Heterogeneity

Balboni et al. (2022) included evidence that asset dynamics may vary with characteristics at the village level (such as distance to markets) or individual level (such as entrepreneurial ability or saving rates). Here we expand on the differing implications of heterogeneity at the individual and village level for identifying poverty traps. This explains how unobservable heterogeneity affects the interpretation of the position of the transition equation, highlighting that neither the shape nor the position of the transition equation are necessary or sufficient conditions for identifying poverty traps or the analysis in Balboni et al. (2022).

Case I: identical individuals in identical villages If all individuals were identical, or if we knew and were able to measure all the traits that matter, panel data over two periods would be sufficient to test for poverty traps against the alternative of convergence by estimating the transition equation (k_{t+1} on k_t). Under conditional convergence, this function is concave and crosses the 45-degree line exactly once, from above. Wherever an individual starts, the dynamics pull them toward the unique stable equilibrium. If there are poverty traps, the function is S-shaped, crossing the 45-degree line three times: at a low-wealth stable equilibrium, an intermediate unstable threshold, and a high-wealth stable equilibrium. Individuals starting below the threshold lose wealth and converge to the low equilibrium; those starting above it accumulate and converge to the high one. If the assumption of identical traits holds, an S-shaped transition equation is both necessary and sufficient for the existence of a poverty trap.

Case II: different individuals, identical villages If individuals have different traits that affect their earnings potential, the transition equation becomes individual-specific, shifted up or down by unobserved productivity. What we estimate from the data is an average across types. This mixture can produce an S-shape even when no individual faces a poverty trap. Highly productive individuals have transition equations everywhere above the 45-degree line; low-productivity individuals have transition equations everywhere below it. If low-productivity individuals are over-represented at low wealth levels – which selection into poverty makes likely – the average transition equation will appear S-shaped purely as a statistical artifact of mixing types. The problem runs in the other direction too. If high-productivity individuals are sufficiently numerous or their transition equations sufficiently high, they can dominate the average and produce an estimated transition equation that lies everywhere above the 45-degree line, masking genuine poverty traps faced by low-productivity households. The S-shape is therefore neither necessary nor sufficient for poverty traps. More fundamentally, the aggregate transition equation is not a structural object: it is a statistical average that conflates the production technology with the distribution of unobservable productivity.

Transition equations estimated with panel data are not informative because both transitory shocks and unobservable traits generate variation in wealth, and we are unable to separate them. The solution is to use, or create, an exogenous shock to wealth that allows us to observe how trajectories respond to a push, independently of who received it.

The randomized evaluations of transfer programs stand out as a natural candidate because, by virtue of the randomization protocol, individual traits do not affect their assignment.[2] The logic is straightforward: give a randomly selected group of poor households a one-time wealth transfer and observe their subsequent trajectories. Under conditional convergence, the transfer moves recipients temporarily away from their steady state, but the effect fades. Treated and untreated households eventually converge back to similar trajectories. Under the poverty trap model, something different happens for households pushed past the threshold: the transfer triggers self-sustaining accumulation, and their trajectories permanently diverge from those of untreated households. Those who received the transfer but remained below the threshold behave just as the convergence model predicts and drift back.

The key predictions therefore center on persistence and non-linearity. If a person is trapped, their long-run response to a shock should depend discontinuously on whether it was large enough to push the household past the threshold, not merely on its size. If instead the effect of a shock is proportional to its size and fades over time for all recipients, this points toward conditional convergence. The negative growth test by Arunachalam and Shenoy (2017) formalizes this intuition: without a trap, the probability of negative growth increases monotonically in current assets. When there is a trap, there can be discontinuous decreases in the probability of negative growth. As the test focuses on the extensive margin of asset loss, it is robust to scaling issues and outliers. Alloush and Carter (2024) extend this idea to account for a mixed population of households that are trapped and those that are conditionally converging to their unique steady state. The test they develop requires knowledge of shocks and is thus well suited to the study of randomized transfer programs. They implement the test using the data in Balboni et al. (2022) and reach similar conclusions.

Case III: different individuals, different villages The same logic applies to heterogeneity at higher levels of aggregation. Here we use the village as an example, but the logic is the same at any level. Village characteristics – market access, prices, infrastructure, agro-climatic conditions – shift the transition equation up or down for all households in a village simultaneously. Households in well-connected villages face lower costs of selling goods and higher returns to assets, such that their transition equations are uniformly higher than those

[2]That said, randomization does not guarantee identification because transfers might be too small or too big to create sufficient variation around the threshold, or because take-up is too low.

in remote villages.

When we pool across villages, we again estimate a combination of transition equations from different environments. If well-connected villages are more numerous, or if the transition equations in those villages are substantially higher than in remote ones, their observations dominate the pooled estimate. A genuine S-shape in the transition equations of remote villages, reflecting real poverty trap dynamics, is swamped by the mass of observations from favorable environments, and the pooled estimate appears to lie everywhere above the 45-degree line. Conversely, if remote villages dominate the sample, an apparent S-shape can emerge in the pooled estimate even though no household in any village faces a poverty trap, simply because low- and high-transition equation villages contribute observations at different points in the wealth distribution.

The solution follows the same logic as the previous case; that is, randomization of program assignment at the village level instead of (or in addition to) the individual level, as is the case in the program studied in Balboni et al. (2022). Control villages, assigned randomly, face the same distribution of environmental conditions as treated villages and provide a direct estimate of where the transition equation sits in each environment in the absence of the program. By observing the wealth trajectories of households in control villages, we measure the baseline dynamics – how much households accumulate or lose given their local environment – with no intervention. Differences in long-run trajectories between treated and control villages can then be attributed to the shock itself rather than to differences in market access, prices, agroclimatic conditions, or any other village-level shifter. Without control villages, any observed trajectory in treated villages could reflect the local environment rather than the program, and the two could not be separated.

3 Functional forms and sample splits

The fact that the estimated transition equation is a combination of different individual transition equations explains why, as noted by several readers, the position of the transition equation relative to the 45-degree line in Balboni et al. (2022) depends on the transformation used. In particular, the estimates in levels are everywhere above the 45-degree line whereas the log transformation used in the paper has three crossings. This is due to the fact that the two do not estimate the same object; rather, levels estimate the equation for the average individual whereas logs estimate it for the median, under a symmetry assumption on the distribution of log productivity (see Appendix for a formal proof).[3] Using the same data,

[3]In Balboni et al. (2022), we estimate the equation in logs to avoid the influence of outliers that would shift the curve upwards. However, this raises a by now well known issue of log transformations with zeros (Chen and Roth, 2024). Karlan, Raswan, and Udry (2026) propose a solution based on correcting measurement challenges that can result in zero measured assets even when asset-holdings are non-zero. In line with this, assuming that all households own at least the lowest value asset owned by poor households in the same sub-

Karlan, Raswan, and Udry (2026) estimate the transition equation under a wide variety of transformations and obtain an equally wide variety of vertically shifted transition equations.

This is a clear illustration of the fact that, if there is unobserved heterogeneity, the shape or position of the transition equation alone is not informative about poverty traps. Instead, identification requires observing long run dynamics for treated and control households and, ideally, control areas. When there are no data on control areas to account for aggregate heterogeneity, as in the other datasets analyzed by Karlan, Raswan, and Udry (2026), an alternative is the negative growth test developed by Arunachalam and Shenoy (2017). This test is robust to moderate levels of heterogeneity and is not sensitive to the unit of measure or the transformation used.

Heterogeneity also implies that, within the same sample, there can be individuals trapped in poverty and others that are not, and the existence of the latter does not rule out the former. For instance, Karlan, Raswan, and Udry (2026) argue that in areas with high cow prices there cannot be a trap and from this it follows that there is no evidence for a trap in areas with low cow prices either, but rather the two areas converge to different equilibria. However, this runs contrary to the evidence from control villages in the same areas presented in Balboni et al. (2022), and the tests proposed by Arunachalam and Shenoy (2017) that the authors employ, which reject the hypothesis of convergence in both areas.

4 Looking ahead

The iconic S-shape that is commonly associated with poverty traps serves as evidence only in the absence of unobservable heterogeneity. When individuals or local economic conditions vary, finding an S-shaped transition equation does not convey clean information on the poverty trap, regardless of whether it crosses the 45-degree line. To test for poverty traps, we need to combine individuals' long run wealth trajectories and investment choices with knowledge of the (large) shocks they experience – for example from the treatment and control group of a transfer program – while accounting for the heterogeneity in individual trajectories.

In Balboni et al. (2022), we combine randomization across areas with long-run follow-up data to show that the divergence in growth paths observed four years after the program persists and widens through eleven years, suggesting that at least a subset of households in our sample were trapped in poverty.

What we cannot say – and what is essential for policy – is which households were trapped and where they live. In recent work addressing this point, Jee (2025) constructs and es-

district negates the need for an arbitrary additive constant inside the log function. The resultant transition equation, invariant to the units in which assets are measured, is very similar to that in Balboni et al. (2022), crossing the 45 degree line from below.

timates a unified model incorporating the production function, intertemporal consumption choices, and the asset transition equation. The framework features a fixed cost generating non-convexities in production, a wealth threshold, and forward-looking households who may reduce consumption to accumulate assets and cross it. Estimation draws on harmonized microdata from 27 randomized cash and asset-transfer experiments covering nearly 75,000 households in 17 countries. The paper finds widespread evidence of poverty traps. In roughly half of the studies, including ours, household production functions exhibit increasing returns to scale, generating non-monotonic consumption and multiple steady states in asset accumulation. Accounting for heterogeneous ability and endogenous consumption allows Jee (2025) to distinguish households merely below the threshold from those genuinely trapped, and yields the conclusion that even in places where a trap exists, only a share of the population (with low productivity) is indeed trapped.

The most promising direction for future research lies in moving beyond the binary question of whether poverty traps exist, toward embracing heterogeneity among the poor. In any given setting, some households may be trapped, some may face persistent poverty at a unique low equilibrium, and some may have suffered a temporary shock from which they will escape unaided. Identifying which households fall into each group would allow the right policy to be targeted to each. We are optimistic that a collective effort in gathering long-run data from the large number of randomized evaluations over the past twenty years will help us to answer this important question.

References

Alloush, Mo and Michael R. Carter (2024). "Empirical Detection of Poverty Traps and Unnecessary Deprivation". Working Paper.

Arunachalam, Raj and Ajay Shenoy (2017). "Poverty Traps, Convergence, and the Dynamics of Household Income". *Journal of Development Economics* 126, pp. 215–230.

Balboni, Clare, Oriana Bandiera, Robin Burgess, Maitreesh Ghatak, and Anton Heil (2022). "Why do people stay poor?" *The Quarterly Journal of Economics* 137.2, pp. 785–844.

Bandiera, Oriana, Robin Burgess, Anton Heil, and Munshi Sulaiman (2026). *Poverty Traps.* Working Paper 73. BRAC Institute of Governance and Development (BIGD).

Chen, Jiafeng and Jonathan Roth (2024). "Logs with zeros? Some problems and solutions". *The Quarterly Journal of Economics* 139.2, pp. 891–936.

Jee, Edward (2025). "Traps and Transfers". Job Market Paper, University of Chicago.

Karlan, Dean S., Amol Singh Raswan, and Christopher Udry (2026). *The Sisyphean Pursuit of Evidence for Poverty Traps*. SSRN Working Paper 6675998.

Appendix

This appendix proves that the level specification recovers the transition equation for the mean productivity type, while the log specification recovers it for the median.

Suppose individual capital evolves according to

$$k_{i,t+1} = A_i f(k_{it}) + (1-\delta)k_{it}, \tag{1}$$

where $A_i > 0$ differs across individuals. Define

$$y_{it} \equiv k_{i,t+1} - (1-\delta)k_{it} = A_i f(k_{it}). \tag{2}$$

Assume that A_i is independent of k_{it} and that $\log A_i$ is symmetrically distributed. The symmetry assumption is needed only for the log result below.

Then a local linear kernel regression of y_{it} on k_{it} identifies

$$m(k) \equiv \mathbb{E}[y_{it} \mid k_{it} = k] = \mathbb{E}[A_i]f(k), \tag{3}$$

so that the implied transition equation is

$$k_{t+1} = \mathbb{E}[A_i]f(k_t) + (1-\delta)k_t. \tag{4}$$

By contrast, a local linear kernel regression of $\log y_{it}$ on $\log k_{it}$ identifies

$$\mu(x) \equiv \mathbb{E}[\log y_{it} \mid \log k_{it} = x] = \mathbb{E}[\log A_i] + \log f(e^x), \tag{5}$$

and therefore

$$\exp\big(\mu(\log k)\big) = \mathrm{Med}(A_i)f(k). \tag{6}$$

Hence the level specification recovers the transition equation for the mean of A_i, whereas the log specification recovers the transition equation for the median of A_i.

See below for a formal proof:

Let

$$m(k) \equiv \mathbb{E}[y_{it} \mid k_{it} = k]. \tag{7}$$

The local linear estimator of $m(k)$ is obtained by choosing $(\hat{a}, \hat{b})$ to minimize

$$\sum_{i,t} K\left(\frac{k_{it}-k}{h}\right)\left(y_{it} - a - b(k_{it}-k)\right)^2, \tag{8}$$

and setting

$$\hat{m}_{LL}(k) = \hat{a}. \tag{9}$$

Under standard regularity conditions,

$$\hat{m}_{LL}(k) \xrightarrow{p} \mathbb{E}[y_{it} \mid k_{it} = k]. \tag{10}$$

Since

$$y_{it} = A_i f(k_{it}), \tag{11}$$

it follows that

$$\mathbb{E}[y_{it} \mid k_{it} = k] = \mathbb{E}[A_i f(k_{it}) \mid k_{it} = k] \tag{12}$$
$$= \mathbb{E}[A_i \mid k_{it} = k] f(k). \tag{13}$$

By independence of A_i and k_{it},

$$\mathbb{E}[A_i \mid k_{it} = k] = \mathbb{E}[A_i], \tag{14}$$

so

$$m(k) = \mathbb{E}[A_i] f(k). \tag{15}$$

Thus the local linear estimator in levels recovers

$$\hat{m}_{LL}(k) \xrightarrow{p} \mathbb{E}[A_i] f(k). \tag{16}$$

Adding back depreciation yields the transition equation

$$\widehat{k}_{t+1}(k) = \hat{m}_{LL}(k) + (1 - \delta)k \ \xrightarrow{p} \ \mathbb{E}[A_i] f(k) + (1 - \delta)k. \tag{17}$$

Now define

$$x_{it} \equiv \log k_{it}, \qquad z_{it} \equiv \log y_{it}. \tag{18}$$

Then

$$z_{it} = \log A_i + \log f(k_{it}) = \log A_i + \log f(e^{x_{it}}). \tag{19}$$

Let

$$\mu(x) \equiv \mathbb{E}[z_{it} \mid x_{it} = x]. \tag{20}$$

The local linear estimator of $\mu(x)$ chooses $(\hat{c}, \hat{d})$ to minimize

$$\sum_{i,t} K\left(\frac{x_{it} - x}{h}\right) \left(z_{it} - c - d(x_{it} - x)\right)^2, \tag{21}$$

and sets

$$\hat{\mu}_{LL}(x) = \hat{c}. \tag{22}$$

Again, under standard conditions,

$$\hat{\mu}_{LL}(x) \xrightarrow{p} \mathbb{E}[z_{it} \mid x_{it} = x]. \tag{23}$$

Therefore,

$$\mu(x) = \mathbb{E}[\log A_i + \log f(e^{x_{it}}) \mid x_{it} = x] \tag{24}$$

$$= \mathbb{E}[\log A_i \mid x_{it} = x] + \log f(e^{x}). \tag{25}$$

Since $x_{it} = \log k_{it}$ and A_i is independent of k_{it},

$$\mathbb{E}[\log A_i \mid x_{it} = x] = \mathbb{E}[\log A_i]. \tag{26}$$

Hence

$$\mu(x) = \mathbb{E}[\log A_i] + \log f(e^{x}). \tag{27}$$

Exponentiating gives

$$\exp\big(\mu(\log k)\big) = \exp\big(\mathbb{E}[\log A_i]\big) f(k). \tag{28}$$

Because $\log A_i$ is symmetrically distributed, its mean equals its median:

$$\mathbb{E}[\log A_i] = \mathrm{Med}(\log A_i). \tag{29}$$

Since the exponential function is strictly increasing, the median is preserved under transformation:

$$\exp\big(\,\mathrm{Med}(\log A_i)\big) = \mathrm{Med}(A_i). \tag{30}$$

Therefore

$$\exp\big(\mu(\log k)\big) = \mathrm{Med}(A_i) f(k), \tag{31}$$

so the log local linear estimator recovers the transition equation for the median productivity type:

$$\widehat{k}_{t+1}(k) = \mathrm{Med}(A_i) f(k) + (1 - \delta)k. \tag{32}$$