A note on the identification of poverty traps.*

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

October 28, 2026

1 Aim and scope

The existence of poverty traps has become almost inextricably linked to the estimation of an S-shaped transition equation for productive assets, despite compelling theoretical arguments that cast doubt on the reliability of this approach in the presence of heterogeneity (Alloush and Carter, 2024).

This note discusses the identification strategy employed in Balboni et al. (2022) that decouples the identification of the threshold from the dynamics around it. Balboni et al. (2022) make use of data on the entire distribution of assets to identify candidate threshold regions, and tests for divergence using responses off the equilibrium path generated by a large scale transfer. To compare this with the more standard use of non-parametric estimates of the transition equation as the sole test for poverty traps, we draw on the extensive analysis of our data by Karlan, Raswan, and Udry (2026). We show that the discrepancies between their evidence and the conclusions drawn are reconciled once we allow for heterogeneity. Therefore we conclude that, taken together, the evidence for the existence of a poverty traps in Balboni et al. (2022) is robust to the issues raised by Karlan, Raswan, and Udry (2026). Finally, we outline a path for future research in this field.

2 Identification off the equilibrium path

Poverty traps exist if there is an unstable equilibrium, that is a threshold value of income or productive assets such that growth is positive above and negative below. Testing for a poverty trap then requires simultaneously locating the threshold and showing that it is an unstable equilibrium, that is, that its dynamics are repelling. This is what the estimate of the transition equation tries to achieve but instability itself creates a catch-22: because repelling dynamics push households away from the threshold, few households are observed near it, so

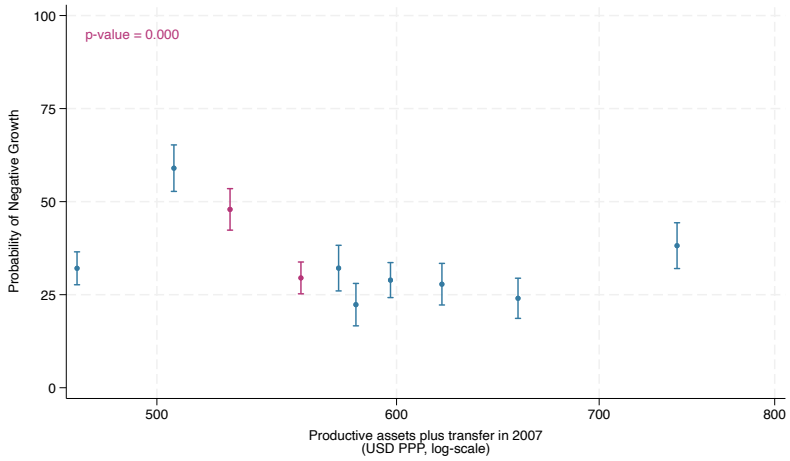
*We thank Kate Banner and Jacopo Olivieri for outstanding research assistance

the test for bifurcation dynamics has very low power. It is not surprising, then, that evidence for or against poverty traps has proved elusive. The main innovation in Balboni et al. (2022) is to separate the identification of the threshold from the test of repelling dynamics. Two features of the setting make progress possible. First, we draw a sample that is representative of the entire wealth distribution and hence we can find regions with low population density. This is not evidence of poverty traps *per se* but it is informative whether there are low density troughs because, if a threshold exists, it must lie in a trough. In our context the distribution is bimodal with a single trough so we know where to look. This is helpful but not sufficient, because the very same fact that makes it a valid candidate also makes testing impossible in equilibrium. The second feature of our setting is that we observe this candidate region off the equilibrium path because BRAC's ultra-poor graduation programme shifts over 3,000 beneficiaries from the low mode into the candidate region. This gives us the power to detect bifurcation dynamics as the economy returns to its equilibrium configuration.

These features are rarely available in other settings where people have looked. Tests with observational data rely solely on naturally occurring shocks both to find the threshold and study the dynamics. Randomized transfers create exogenous shocks but this is moot unless researchers have access to auxiliary information that allows them to gauge whether the magnitude of the transfer is large enough to bring people close to the threshold but not so large as to overshoot it. These studies rely on transition equations to jointly test for the existence of a threshold and asset dynamics. These tests are typically very low powered and their findings inconclusive as they cannot separate "absence of evidence" from "evidence of absence". By contrast, we already know where the potential unstable equilibrium is, and we have the power to test the null of convergence against the alternative of poverty traps.

While transition equations are the most visually impactful illustration of poverty traps, a more robust test for short run divergence is the negative growth rate test proposed by Arunachalam and Shenoy (2017). The intuition is simple: under the null of no poverty traps, the share of households experiencing negative growth must be increasing in baseline assets. Rejecting the null, that is, finding a significant drop in the share of households with negative growth, indicates a poverty trap. Figure 1 reports the results for our sample. In line with the transition equation estimates in Balboni et al. (2022), we find an interior threshold such that, to its left, nearly 50% of beneficiaries lose assets over the four-year period, while to its right the share is about half as large (27%). The test rejects the null of convergence to a unique steady state with p-value < 0.001.

Figure 1: A-S negative growth test: $k + \min$ district k



Notes: The sample is restricted to ultra-poor households in treatment villages. Productive assets are measured as the total value, in 1,000 Bangladeshi taka, of all livestock, poultry, business assets, and land owned by the households. Post-transfer assets are imputed by adding to each household's baseline assets the median value of a cow in the catchment area of a household's BRAC branch. The figure plots the estimated probability of negative growth within each decile ($k + \min$ district k). The p-value is for the null hypothesis of convergence, based on the A-S test.

The strongest test for the existence of poverty traps comes from the long-run data. Our five-wave panel follows the same households for 11 years, and we show that the bifurcation observed after the transfer does indeed set households on fundamentally different paths. Households above the threshold have higher consumption and holdings of productive assets and the gaps grow over time.

Our analysis also sheds light on the underlying mechanism: long-run occupational change. We observe beneficiaries' earnings in two periods in which they had no choice of occupation: at baseline, when none could afford livestock and so all worked as casual laborers, and in the first two years of the program when, by design, all were committed to running a livestock business. Using these earnings, we calibrate a model of occupational choice and estimate the optimal allocation of hours between the two activities as a function of productive asset holdings, and project it to later periods. The fact that actual allocations match closely for 9 years up to 2018, indicates that occupational change is at the heart of the escape from the trap. In summary, Balboni et al. (2022) show that:

1. There is a threshold value of productive assets such that, in the short run (two years after the end of the program and four years after the transfer), households whose initial assets are lower than this value are more likely to lose assets, while those who have more gain. The threshold (around 9,500 BDT) corresponds to:

    - the trough of the estimated asset distribution;

- the point where the estimated transition equation crosses the 45 degree line from below; and
- the mid-value of the decile where the share of households experiencing negative growth drops in the A-S test shown in Figure 1.

2. The bifurcation that emerges after four years grows after seven years and continues to grow further after eleven years.

3. The estimates of a structural model of occupational choice support the mechanism that motivated the design of the intervention in the first place.

3 Identification using the transition equation

The appeal of the S-shaped transition equation is unquestionable as it can potentially identify both the value of the threshold and the dynamics around it. However, this holds only under restrictive assumptions, most importantly, it fails when individuals are heterogeneous along any relevant dimension. In what follows we draw on Karlan, Raswan, and Udry (2026)to illustrate the difficulties in interpreting non-parametric estimates of the transition equation and how apparently innocuous choices can have disproportionate consequences depending on the nature of the heterogeneity.

3.1 Log scales and the zero trap

Balboni et al. (2022) transform productive assets x according to: $\ln(x/Q + 1)$, where Q equals 1000. Each step has a simple rationale: we divided by a constant to simplify and ease interpretation of figures and tables, took logarithms to limit the influence of outliers (see next section), and added one to retain the households whose assets fall to zero. The first puzzle is that the estimated position, but not the shape, of the transition equation appears sensitive to the choice of the scaling factor Q.

Figure 2 replicates the estimation of the transition equation in Balboni et al. (2022) using different scaling factors from $Q = 1000$ to $Q = 1$. For $Q = 500$ and $Q = 1000$, the estimated transition equation intersects the 45 degree line three times, consistent with two stable steady states separated by an unstable threshold. For $Q = 1, 10$, and 100, by contrast, the estimated transition equation lies below the 45 degree line throughout the plotted asset distribution (i.e., the average transformed wealth of households declined for all levels of baseline wealth).[1]

Why does this happen? The solution to the puzzle has two pieces. The first part of the explanation is that, because of the addition of one, changing the scaling factor Q does more

[1]We choose these values to illustrate the issue concisely, Karlan, Raswan, and Udry (2026) report an extensive range of values that place the transition equation below the 45 degree line

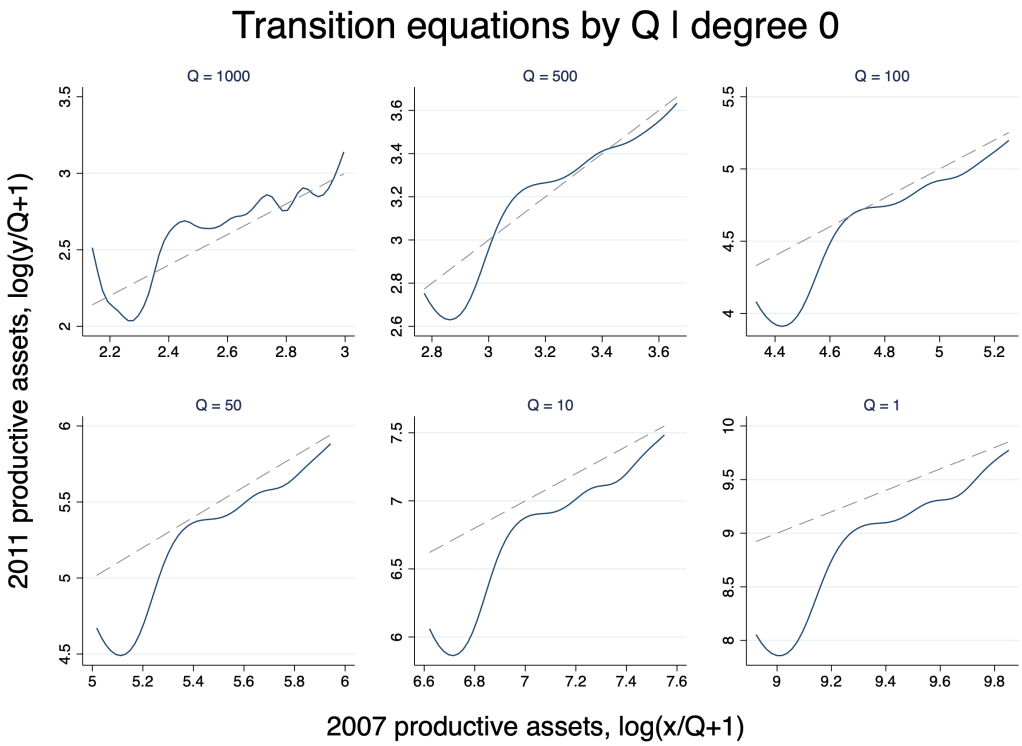


Figure 2: Transition equation estimates using the default polynomial degree in Stata

than simply change the units in which the assets are measured: it determines the extent to which observations with zeros are outliers. To see this, consider first what would happen without the additive constant. In that case, $\ln(x/Q) = \ln(x) - \ln(Q)$. Changing Q would therefore subtract the same constant, $\ln(Q)$, from every observation. The entire distribution would shift, but the distances between observations, and hence its shape, would remain unchanged. This is no longer true once one is added to accommodate households with zero assets. In particular, $\ln(x/Q + 1) = \ln(x + Q) - \ln(Q)$, so that the scaling factor is now an additive constant on raw assets inside the logarithm. The distance between two observations becomes $\ln\left(\frac{x_1+Q}{x_2+Q}\right)$, which depends on Q. As a result, changing Q changes relative distances between observations and therefore the shape of the transformed asset distribution.

This distinction is particularly consequential in the presence of a zero mass. As highlighted by a recent literature on log transformations (Chen and Roth, 2024) which post-dates Balboni et al. (2022), the choice of units can affect treatment effects and their statistical significance when outcomes include zero. This is because zero is invariant to the transformation: $\ln(0/Q+1) = 0 \quad \forall Q > 0$. Positive asset holdings, by contrast, change with Q. For all $x > 0$, increasing Q reduces x/Q and thus reduces $\ln(x/Q + 1)$. Moreover, the magnitude of this change depends on x; households with different positive asset holdings move by different amounts on the transformed scale, while households with zero assets do not move at all. This is illustrated in Figure 3.

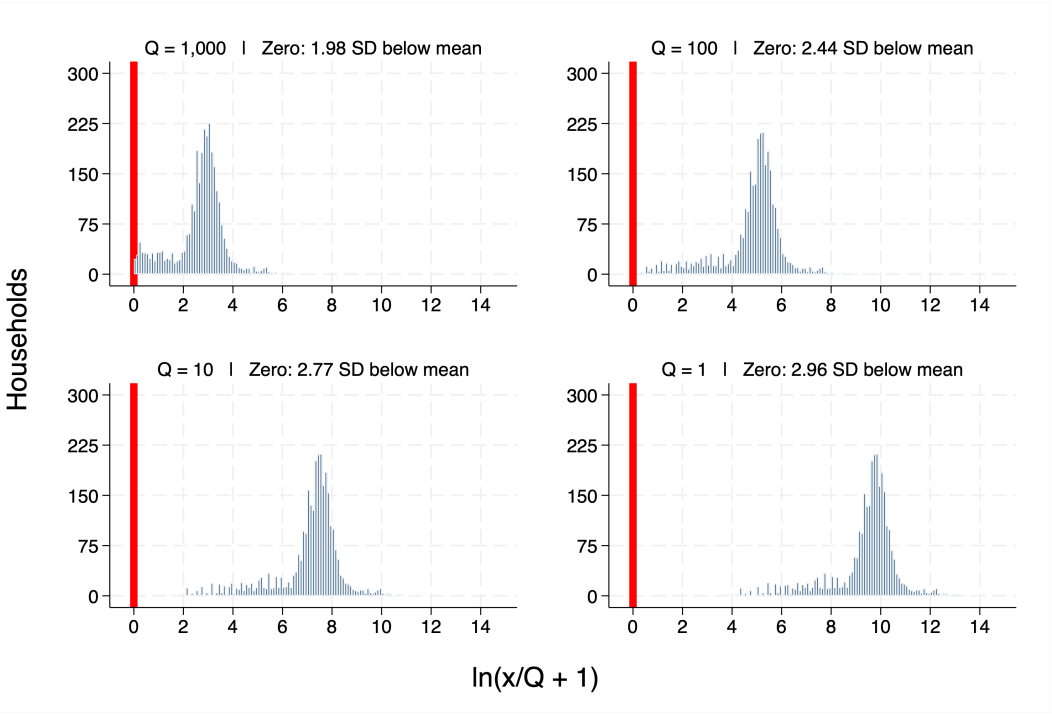


Figure 3: Distribution of $ln(x/Q+1)$ across Q

Notes: Productive assets in 2011 (x) are measured in BDT. Histograms exclude households with zero assets; red lines indicate the position of zero. Distances are calculated using means and standard deviations that include zero-asset households.

Zero follow-up outcomes therefore lie increasingly far below positive follow-up outcomes, relative to the dispersion of the latter, on the transformed scale as Q becomes smaller. Because the transition equation is estimated as the conditional mean of *transformed* follow-up assets given values of the baseline assets, these extreme zero outcomes pull the transition equation downward.

The second piece of the puzzle lies in making the estimates with different Q directly comparable, which requires a shared criterion for the choice of the degree of the polynomial smoother. We estimate the transition equation non-parametrically using Stata's *lpoly* command, which approximates the conditional mean of $\ln(y/Q+1)$, where y is 2011 productive assets, at each level of baseline assets $\ln(x/Q+1)$, where x is 2007 productive assets. It does so by fitting a polynomial to nearby observations, weighting closer observations more heavily. All smoothing, including the choice of the degree, is carried out on the transformed variables, which change with Q. The polynomial degree determines the flexibility of this approximation: degree zero is a locally weighted mean, degree one fits a local linear relationship, and higher degrees allow progressively more curvature.

lpoly keeps the degree of the polynomial at 0 by default. So using the default imposes the same degree at every Q. Because the transformation $\ln(x/Q+1)$ changes the shape of the data as Q changes, the same degree does not approximate the conditional mean equally well at different values of Q. Differences between estimated transition equations then reflect how

well the same smoother fits the transformed data, rather than differences in the underlying relationship.

To address this challenge, we follow Pei et al. (2022) and choose the polynomial degree that minimizes the mean squared error. Specifically, for each Q we consider degrees 0 to 5 and select the degree that minimizes the leave-one-out cross-validation mean squared error of the transformed outcome, using a Gaussian kernel and the bandwidth from *lpoly*.

Figure 4 reports the estimates of the transition equation for different values of Q using the degree of the polynomial that minimizes the mean squared error for that specific Q. With the degree chosen using the same criterion at every Q, the estimated transition equations are similar across values of Q, and our conclusions are not sensitive to the choice of Q. Moreover, for every value of Q, the estimated thresholds, once transformed back to original units, approximately recover the poverty trap threshold of Balboni et al. (2022).

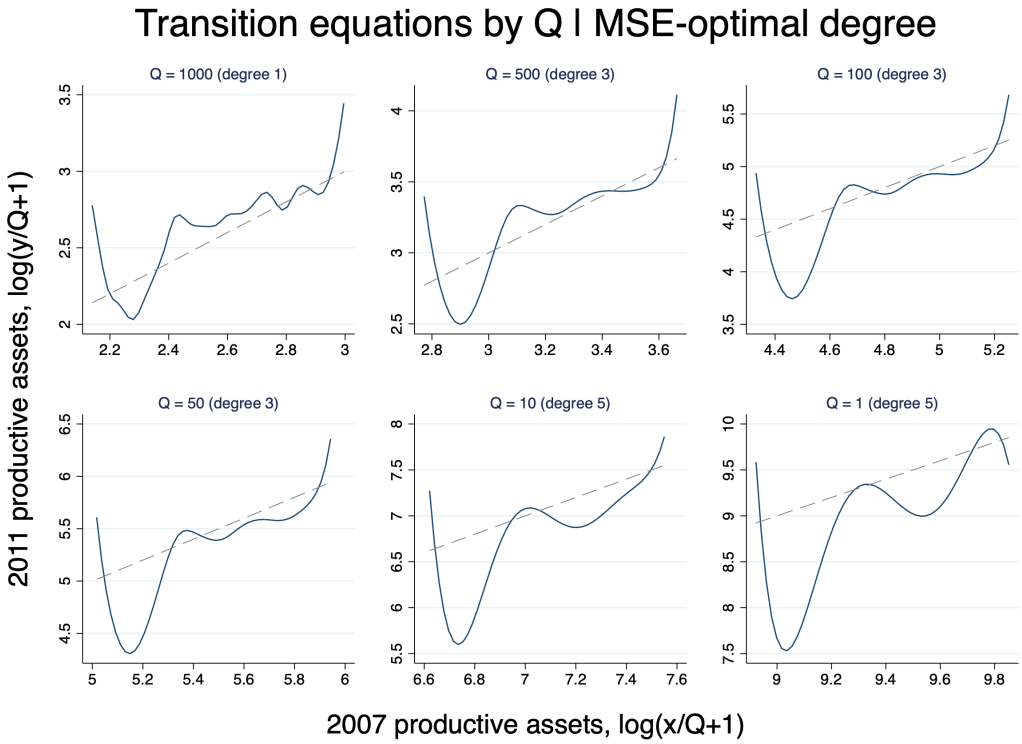


Figure 4: Transition equation estimates using the same criterion to choose the polynomial degree in Stata

An alternative approach to resolving the "zero problem" is to cast it as a measurement issue. Households reporting no productive assets are in fact likely to own items that could be used in income generation, but these may not be reported because they are not used for that purpose.[2] Take for instance two individuals who own identical spades but use them for different purposes: one grows vegetables for sale, and the other grows flowers to decorate

[2]We thank Chris Udry for suggesting this approach.

their house. The former would list the spade as a productive asset, the latter would not. In principle one could review all assets owned by each household and make a list of "potentially productive" assets. As this entails subjective judgment that can add more noise of unknown nature, we assume that every household that reports owning assets, albeit no business assets, owns at least one asset that can be used to generate income. For each control household, we identify the lowest positive unit value among its reported assets. We then calculate a weighted average of these household-level minimum values within each subdistrict and survey wave, giving greater weight to households who report unit values for more asset categories. We add this amount to every household's productive assets within the corresponding sub-district and survey wave. As seen in Figure 5, this adjustment produces practically identical results to Balboni et al. (2022) which are not sensitive to the scaling factor.

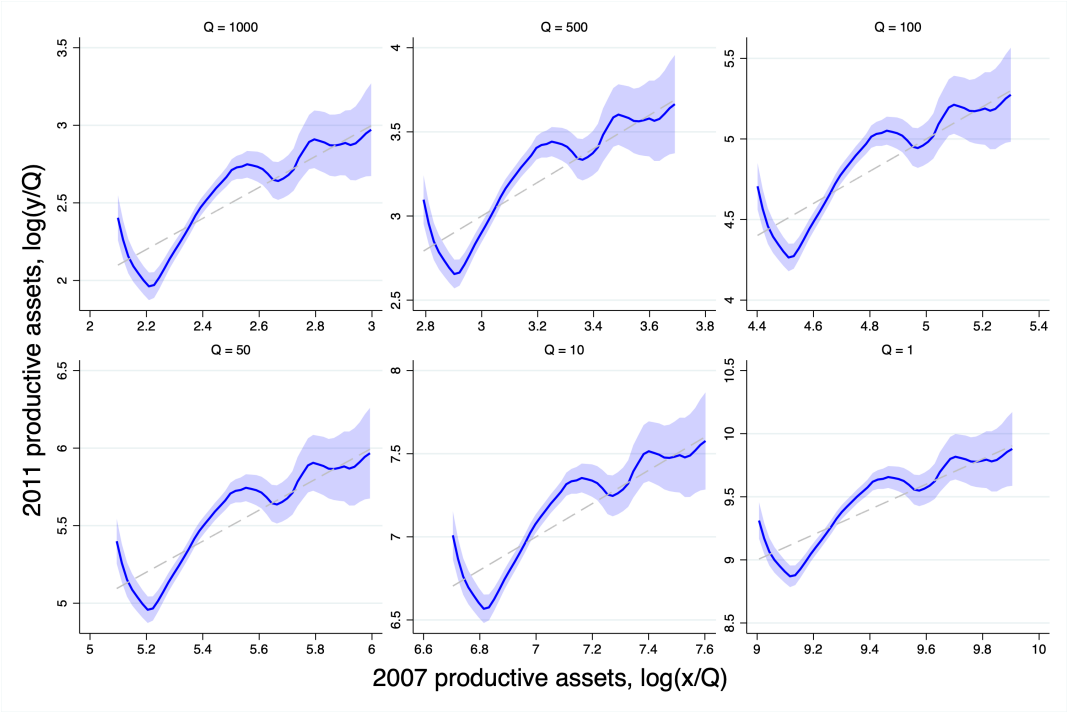


Figure 5: Transition equation estimates - k + min district k

3.2 Logs or Levels?

The second puzzle raised by Karlan, Raswan, and Udry (2026) is that estimating the transition equation in levels also shifts it vertically, albeit upwards rather than downwards. As shown in Balboni et al. (2026) these two specifications estimate fundamentally different objects: estimation in levels recovers the conditional arithmetic mean of future assets, whereas estimation in logs recovers the conditional mean of transformed future assets. Because the logarithmic transformation compresses large asset values, households with extremely large

future asset holdings have substantially less influence on the conditional mean of the transformed outcome than on the conditional mean in levels. The two transition equations can therefore have different shapes in the presence of a large upper tail. As a result, the log specification sits closer to the conditional median, and will recover a quantity closer to the outcome of the typical households than does the conditional arithmetic mean.

This distinction can reconcile a transition equation that lies everywhere above the 45 degree line in levels with one which exhibits a crossing in logs. Suppose that, conditional on a given level of baseline assets, most households experience relatively weak asset accumulation, while a small number of households with favorable unobserved characteristics experience very large gains. These large outcomes can raise the conditional arithmetic mean sufficiently for the transition equation in levels to lie above the 45 degree line. Taking logs compresses these large gains, reducing their influence on the conditional mean of the transformed outcome. The resulting transition equation can therefore lie below the 45 degree line over part of the asset distribution and thus exhibit a poverty trap threshold. The two estimates are both correct but they identify different objects and the researcher must specify which they want to identify. The appropriate specification depends on which feature of the conditional distribution is the object of interest.

To illustrate this, we simulate an economy with a poverty trap generated by barriers to enter a more profitable occupation as in Bandiera et al. (2017). The data generation process is as follows: we draw 1,000 values of baseline productive assets, k_t, from a uniform distribution on $[0, 25000]$, and create 100 individuals at each value. The production function is linear in k_t and jumps at the threshold, which we set equal to 12,000 BDT. Specifically,

$$k_{t+1} = \begin{cases} 500 + 0.8k_t, & k_t < 12{,}000, \\ 10{,}000 + 0.8k_t, & k_t > 12{,}000. \end{cases}$$

This function has two stable fixed points, at $k = 2,500$ and $k = 50{,}000$, separated by a threshold at $k_t = 12{,}000$. Crossing the threshold at $k_t = 12{,}000$ therefore gives access to a discrete increase in returns. Below the threshold, assets converge to a low stable equilibrium at $k = 2{,}500$; above it, they converge to a high stable equilibrium at $k = 50{,}000$. The ratio between k in the high and low equilibria is of the same order of magnitude as the ratio of productive assets held by middle income and very poor women in the Bangladeshi villages studied by Bandiera et al. (2017) and Balboni et al. (2022).

We then introduce unobserved heterogeneity in future asset outcomes by assigning a share θ of individuals an asset level at $t + 1$ equal to the high equilibrium, that is $k_\theta = 50{,}000$, irrespective of their baseline assets. This creates a small upper tail of individuals with unusually high future asset holdings. The exercise is designed to show how these observations affect transition equations estimated as conditional means. In levels, the large values receive

substantial weight in the arithmetic mean and can therefore pull the estimated transition equation sharply upward, potentially obscuring the underlying poverty trap dynamics. In logs, the same large outcomes are compressed and consequently exert much less influence on the conditional mean. As θ increases, the transition equation in levels can therefore fail to display the poverty trap even when the transition equation in logs continues to reveal it.

For simplicity, we assume unobserved heterogeneity is uncorrelated with baseline assets, but the general point can be made by assuming positive or negative correlation.

Table 1: Monte Carlo Poverty Trap Detection Rates

Share at high steady state (%)	Levels (%)	Logs (%)
0–1%	100.0	100.0
1–2%	100.0	100.0
2–3%	100.0	100.0
3–4%	100.0	100.0
4–5%	77.7	100.0
5–6%	9.8	100.0
6–7%	1.0	100.0
7–8%	0.0	100.0
8–9%	0.0	96.5
9–10%	0.0	59.6

Notes: The table reports the percentage of Monte Carlo replications in which the estimated transition equation detects a poverty trap. In each replication, a proportion θ of individuals reaches the high steady state of 50,000, regardless of initial capital. Detection requires an upward crossing of the 45-degree line. Intervals include the lower bound and exclude the upper bound. Rates are calculated directly from replication-level indicators without smoothing.

Table 1 shows detection rates, that is the share of 1,000 draws that detect poverty traps given that they exist, for both logs and levels. We see that both logs and levels are accurate as long as the high ability share is below 4%, after that the detection rate for levels falls drastically, such that when the population share of high ability individuals reaches 5% it does not even have the accuracy of a coin toss, and by 7% it fails to detect the threshold altogether. By comparison, the log estimates correctly detect the existence of a threshold in most cases. The simulation makes clear that a small fraction of outliers can create large differences between logs and levels.

Three patterns emerge. First, even modest differences in ability can place the transition equation above the 45 degree line, especially when assets are measured in levels. In levels, the estimated low-asset crossing moves rightward as θ increases, and once θ exceeds roughly 5% it would lie beyond the 12,000 BDT threshold, so no crossing remains and the poverty trap is no longer detected. Second, level estimates are systematically more sensitive to heterogeneity in ability than estimates in logs, because the log transformation compresses large next period

asset values contributed by high ability individuals i.e., the mean is more sensitive to outliers than the median. Third, this difference can determine whether a poverty trap is detected at all. Heterogeneity can shift the estimated transition equation entirely above the 45 degree line, eliminating the crossing even though a threshold is present.

3.3 One trap, one threshold?

Balboni et al. (2022) show that heterogeneity in individual traits or at the level of the local economy might result in different thresholds without changing the interpretation of the findings. Heterogeneous thresholds, however, can be easily mistaken for different equilibria different people convergence to. To illustrate, Karlan, Raswan, and Udry (2026) put forward the hypothesis that, since in nearly half of the sample the branch-level median cow price alone exceeds the estimated threshold, the estimates in Balboni et al. (2022) do not identify a poverty trap but rather a composite of two stable equilibria in two different geographies. In other words, bifurcation around the threshold reflects differences in local economic conditions proxied by high and low cow price villages rather than differences in post-transfer capital. To tell these apart we need to go beyond the transition equation and make use of auxiliary data for which the predictions differ under the two hypotheses. First, if high cow prices simply identify locations with better underlying growth prospects, control households who would have been classified above the threshold may also be expected to accumulate faster. Table 2, from Balboni et al. (2022), provides a direct placebo test: among control households, those who would have been classified above the threshold had they received the imputed transfer do not accumulate faster than those who would have been classified below it. The estimated above-threshold difference is -0.020 (s.e. 0.057). By contrast, this above/below difference is 0.318 larger among treated households (s.e. 0.070), implying an above-threshold difference of approximately 0.298 for treated households. Thus, the divergence around the threshold appears specifically among households that actually received the transfer, rather than among households living in high cow price locations whose imputed transfer values would place them above the threshold.

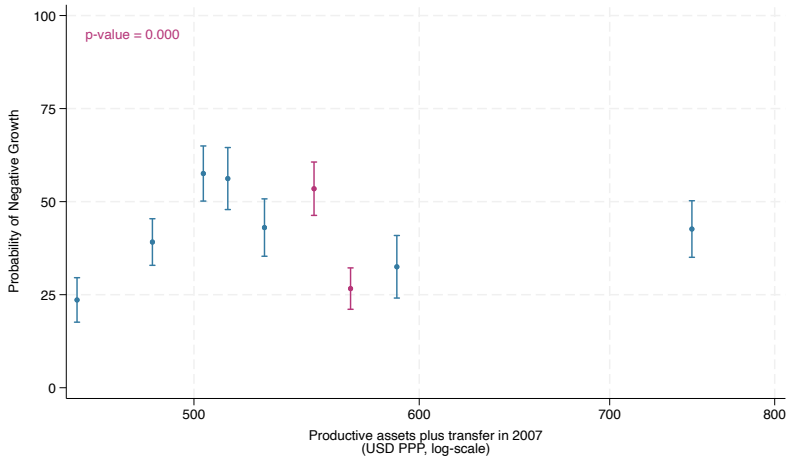
Table 2: Short term responses to the asset transfer

	Dependent variable: log change of productive assets 2007–2011				
	Panel A		**Panel B**		
	(1) Treatment	(2) Treatment	(3) Control	(4) Control	(5) Both
above $\hat{k}$	0.297***	0.475***	-0.020	-0.097	-0.020
	(0.043)	(0.070)	(0.052)	(0.598)	(0.057)
Baseline assets		-2.199***		-0.463*	
		(0.698)		(0.266)	
Above $\hat{k}$ × Baseline assets		1.969***		-0.097	
		(0.729)		(0.269)	
Treatment					-0.483***
					(0.059)
Above $\hat{k}$ × Treatment					0.318***
					(0.070)
constant	-0.138***	-0.282***	0.345***	-0.680	0.345***
	(0.033)	(0.057)	(0.046)	(0.592)	(0.050)
N	3292	3292	2450	2450	5742

Notes: *: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$. Standard errors in brackets. Sample: ultra-poor households in treatment and control villages with log baseline productive assets below 3 (observations from control households are excluded if their baseline productive assets were above 3 had they received the transfer). The dependent variable is the difference between log productive assets in 2011 and log productive assets in 2007, where productive assets are defined as the total value of livestock, poultry, business assets (e.g. tools, vehicles and structures), and land. Above $\hat{k}$ equals 1 if the baseline asset stock plus the imputed transfer is larger than 2.333, and 0 otherwise. In treatment, this represents households' actual post-transfer asset stock. In control, where no transfer was received, above $\hat{k}$ indicates whether the household would be above 2.333 if it had received a transfer. Baseline assets always refers to the actual asset stock, i.e. in control without the imputed transfer. Baseline assets are centered at 2.333, such that the coefficient on above $\hat{k}$ reflects the log change at the threshold. Treatment was assigned at the village level.

Second, there is clear evidence of poverty traps even when we restrict the sample to low cow price villages thus not making use of the identifying variation from cow prices. Figure 6 restricts the sample to branches with median cow prices below 10,000 BDT, thereby removing the high-cow-price branches that generate much of the concern. Within this restricted sample, the Arunachalam and Shenoy (2017) test continues to reject convergence ($p < 0.001$). Thus, the evidence of a threshold is not simply generated by comparing households in high and low cow price branches. This suggests that the threshold pattern is not driven solely by the mechanical role of local cow prices in determining households' post-transfer asset positions.

Figure 6: A-S negative growth test: median cow price < 10,000 BDT



Notes: The sample is restricted to ultra-poor households in treatment villages in BRAC branches with median cow price < 10,000 BDT. Productive assets are measured as the total value, in 1,000 Bangladeshi taka, of all livestock, poultry, business assets, and land owned by the households. Post-transfer assets are imputed by adding to each household's baseline assets the median value of a cow in the catchment area of its BRAC branch. The figure plots the estimated probability of negative growth within each decile (k + min district k). The p-value is for the null hypothesis of convergence based on the A-S test.

The geographic confound explanation cannot be reconciled with this evidence. Furthermore, this apparent contradiction relies on interpreting the crossing of the estimated transition equation as a common structural threshold faced by all individuals. This interpretation need not hold in the presence of heterogeneity in productive ability, as discussed in Section IV.A of Balboni et al. (2022) and in Balboni et al. (2026). When we do not correct for differences in ability we effectively pool individuals with very different thresholds into a single estimate. This distinction matters for the apparent geography puzzle. A cow worth 10,000 BDT exceeds the crossing of the pooled transition equation, but this does not imply that it places every recipient above the unstable steady state of their own transition equation. A lower-ability individual could, for example, face an unstable steady state at 12,000 BDT and therefore remain below their poverty-trap threshold despite receiving a cow worth more than 9,309 BDT.

The apparent contradiction therefore arises from comparing the value of the transfer with a pooled crossing as though that crossing were a common steady state. With heterogeneous accumulation dynamics, the pooled crossing and the poverty-trap threshold relevant for an individual are distinct objects. The fact that cow values exceed 9,309 BDT in many branches therefore does not, by itself, imply that the program was sufficiently generous to move every recipient out of the poverty trap.

Conclusion

The research on poverty traps is at a crossroads. Until recently evidence has been elusive because it requires simultaneously locating the threshold and showing that it is an unstable equilibrium, which is challenging because by its very nature few households are observed near it, so statistical power is lowest precisely where it is most needed. The many identification issues that make the search for poverty traps using transition equations in observational data difficult are well articulated in Karlan, Raswan, and Udry (2026). In the presence of heterogeneity, the average nonparametric relationship is difficult to interpret: it is an average over types whose shape depends on the composition of the sample and on the scale of estimation, and it need not resemble the transition function faced by any individual household. The presence of an S-shape is accordingly neither necessary nor sufficient evidence of a poverty trap. It is not necessary, because aggregation across types can erase a threshold that every type faces; it is not sufficient, because aggregation across types can generate one where none exists (Alloush and Carter, 2024; Bandiera et al., 2026). In Balboni et al. (2022), we make progress by exploiting temporary disequilibrium created by a large transfer that placed beneficiaries in the low density region which is a good candidate for the threshold. Because of this feature we do not have to rely on estimates of the transition equation to provide evidence for the existence of poverty traps.

In the past twenty years, economists have evaluated a plethora of cash and other forms of transfers that are now yielding long run outcomes. Many of these transfers will be too small or too large to study off-equilibrium dynamics in the right place, but some will be of the right magnitude to help us make progress on this important question. In pioneering work exploiting this data Jee (2025) constructs and estimates 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.

Future research should move beyond asking whether poverty traps exist and embrace heterogeneity among the poor. Within any setting, some households may be trapped, some persistently poor at a unique low equilibrium, and some only temporarily poor after a shock

they will overcome unaided. Distinguishing these groups would allow policy to be targeted accordingly. We believe that a collective effort to gather long-run data from the many randomized evaluations over the past two decades will shift the equilibrium from a frustrating search for an elusive phenomenon to a thriving field of research and policy, which will ultimately make progress in long-term poverty reduction.

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. ISSN: 0304-3878. DOI: 10.1016/j.jdeveco.2017.02.001. (Visited on 07/28/2025).

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.

— (2026). *Why do people stay poor? A research agenda for poverty traps*. Response hosted by the National Bureau of Economic Research. URL: https://data.nber.org/data-appendix/w35253/povertyTraps_reply.pdf.

Bandiera, Oriana, Robin Burgess, Narayan Das, Selim Gulesci, Imran Rasul, and Munshi Sulaiman (2017). "Labor Markets and Poverty in Village Economies". *Quarterly Journal of Economics* 132.2, pp. 811–870.

Bandiera, Oriana, Robin Burgess, Anton Heil, and Munshi Sulaiman (2026). *Poverty Traps*. Working Paper 73. BRAC Institute of Governance and Development (BIGD). URL: https://bigd.bracu.ac.bd/publications/poverty-traps/.

Chen, Jiafeng and Jonathan Roth (2024). "Logs with Zeros? Some Problems and Solutions*". *The Quarterly Journal of Economics* 139.2, pp. 891–936. ISSN: 0033-5533. DOI: 10.1093/qje/qjad054. (Visited on 10/04/2024).

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

Karlan, Dean, Amol Singh Raswan, and Christopher R Udry (2026). *The Sisyphean Pursuit of Evidence for Poverty Traps*. Working Paper 35253. National Bureau of Economic Research. DOI: 10.3386/w35253. URL: http://www.nber.org/papers/w35253.

Pei, Zhuan, David S. Lee, David Card, and Andrea Weber (2022). "Local Polynomial Order in Regression Discontinuity Designs". *Journal of Business & Economic Statistics* 40.3, pp. 1017–1030. DOI: 10.1080/07350015.2021.1920961. URL: https://www.tandfonline.com/doi/abs/10.1080/07350015.2021.1920961.