Traps and Transfers

Edward Jee*
University of Chicago

Latest Version Here

27th December 2025

Abstract

Do poverty traps exist? This question is central to development economics as it implies initial conditions can have lasting effects on long-run prosperity. I estimate the existence and prevalence of poverty traps using harmonized microdata from 27 randomized controlled trials covering 75,000 households. By estimating how poor households' assets evolve over time, accounting for differences in their ability and consumption-savings decisions, I find that poverty traps are widespread: 60% of studies show signs of them. Yet within a study, on average, only 25% of households are actually trapped. This is driven by differences in productivity and forward-looking consumption responses which allow households to preemptively invest and escape low-asset states. A calibrated household growth model shows that optimal transfers target the poorest, not those near the trap threshold, in over 95% of contexts.

*The University of Chicago; email: edjee@uchicago.edu. I am grateful to my advisors, Michael Kremer, Stéphane Bonhomme, Sid Kankanala, and Anne Karing. I also thank Christina Brown, Fiona Burlig, Joshua Dean, Rachel Glennerster, Mikhail Golosov, Greg Kaplan, Erin Kelley, Gregory Lane, Rafe Meager, and Alex Torgovitsky, as well as Dylan Balla-Elliott, Thomas Bourany, Deniz Dutz, Sofia Gallo, Maryiam Haroon, Tom Hierons, Tyler Jacobson, Varun K, Andrew Kennedy, Zachary Kuloszewski, Hugo Lopez, Ian Pitman, Julius Ruschenpohler, Giulio Schinaia, Daniel Sonnenstuhl, Arnstein Vestre, and the participants of the Development Tea and Econometrics advising groups at the University of Chicago. I am immensely grateful to Sebastian Araya-Cardoso-Muent, Al-E-Iba Dahri, Juhyeon Eom, Bethany Gardner, Rob Lytle, Elliott Paschal, Jianing Shao, and Yishu Yu for excellent research assistance. This research was supported by The Weiss Fund for Research in Development Economics at the University of Chicago, the 'Structural Transformation and Economic Growth' (STEG) initiative funded by the Foreign, Commonwealth, & Development Office (FCDO) and the International Growth Centre (IGC). The views expressed are those of the author and are not necessarily those of the FCDO. All errors are my own.

1 Introduction

A large theoretical literature argues that poverty traps—self-reinforcing mechanisms that keep the poor poor—are a key feature of economic development (Rosenstein-Rodan 1943; Murphy et al. 1989; Banerjee and Newman 1993; Dasgupta 1997). Others, however, contend that the conditions required for poverty traps to exist are empirically implausible, and that markets, random shocks, or contracts should preclude their presence in practice (Cass 1965; Lucas 1988; Aiyagari 1994; Moll 2014). This theoretical divide is mirrored by mixed empirical evidence, which has traditionally focused on national, cross-country comparisons (Quah 1996; Sachs et al. 2004; Azariadis and Stachurski 2005) or, more recently, on micro-level evidence in specific contexts and interventions (McKenzie and Woodruff 2006; Santos and Barrett 2011; Banerjee et al. 2019; Balboni et al. 2022). The central question therefore remains: do poverty traps exist, and if so, how prevalent are they?

To answer this question, I harmonize microdata from randomized cash and asset transfer experiments. I then develop an empirical framework to identify and estimate poor households' asset accumulation equation, which describes how their assets change over time as a function of their production and consumption choices:

$$\dot{k}_t = f(k_t) - c(k_t) - \delta k_t$$

In this framework, a poverty trap exists if the accumulation equation has multiple steady states, that is, multiple levels of assets where $\dot{k}_t = 0$. This is because the accumulation equation determines whether a household's assets grow, shrink, or remain constant over time. Multiple steady states imply that some households are poor not because of innate differences in ability, but rather their initial endowment, and that a sufficiently large asset injection can move them to a higher-asset steady state.

I begin with the theoretical foundation, outlining a dynamic model of household production with endogenous consumption decisions. The model, originally developed in Skiba (1978); Buera (2009); Ikegami et al. (2016) has two key ingredients: (i) households produce output in the presence of fixed costs, such as large indivisible investments like a cow or machine, that create non-convexities in their production technology;[1] and (ii) households differ in ability. While these two features are common in poverty trap models, empirical work has largely ignored the implications of endogenous consumption decisions in this context.[2]

The theoretical model illustrates three challenges any empirical analysis must address.

[1] Fixed costs generate non-convexities in the production set which lead to non-concavities in the production function.

[2] Carter and Lybbert (2012) provide a detailed discussion of the empirical evidence on consumption smoothing in poverty trap models.

First, observed household assets reflect a sequence of optimizing consumption choices and are therefore endogenous. Second, fixed costs alone do not create poverty traps: forward-looking households may forgo current consumption to save and acquire indivisible assets, sorting into higher-asset steady states. Third, differences in ability create differences in who faces poverty traps. Some households may be so productive that, regardless of their starting assets, they accumulate wealth and reach a higher-asset steady state. Others are so unproductive that they lose assets until they settle at a low-asset state, no matter their initial assets.

Having reviewed theoretical conditions for indivisibility-based poverty traps, I next turn to identification. My empirical framework follows three main steps: First, to address endogenous asset formation, I use the randomized transfers to identify the household production function.[3] Second, I recover the consumption function from observed consumption data, allowing heterogeneous households to optimally adjust spending in response to potential production indivisibilities. Finally, combining these two functions, I construct the household asset accumulation equation and simulate its dynamics to determine the presence of poverty traps. Crucially, my approach explicitly accounts for household optimization and heterogeneity in ability, allowing the transition equation to vary across households and reflect their endogenous consumption responses.

To implement this framework empirically, I adopt a flexible, nonparametric dynamic panel approach that captures the joint evolution of income, consumption, and assets. This approach allows for, but does not impose, household optimization and enables me to flexibly estimate the asset accumulation equation without restrictive functional form or parametric assumptions. However, this generality comes at a cost: the underlying technology, preferences, and nature of the household optimization problem remain unspecified. I must recover the full production and consumption functions directly, not the finite dimensional parameters of some model.

This represents the next empirical challenge of the paper: nonparametric instrumental variable estimators can generate imprecise estimates, potentially producing spurious patterns that resemble production function non-concavities or poverty traps. I approach this problem by introducing regularization in the form of a prior, leveraging quasi-Bayesian methods (Kankanala 2025). Intuitively, I encode prior beliefs about the smoothness of the production function which stabilizes estimates, allowing me to recover economically plausible functions.

I then turn to estimation of the consumption function, which poses additional challenges due to its nonlinear dependence on unobserved household ability and productivity shocks. These difficulties are compounded by the short panels typical of randomized controlled trials, where direct estimates of ability and productivity shocks are prone to

[3] I supplement the experimental variation with internal instruments, using dynamic panel methods following Arellano and Bond (1991); Blundell and Bond (2000).

severe bias and high variance: with few observations per household, a sequence of "lucky" idiosyncratic shocks can be easily mistaken for high ability. To address this, I avoid estimating the unobserved ability and productivity directly and instead integrate over their distribution. Specifically, I implement a particle Markov Chain Monte Carlo (PMCMC) algorithm (Andrieu et al. 2010; Arellano et al. 2024) which samples latent productivity paths in proportion to their likelihood. By simulating and averaging across many such paths, I obtain estimates of the consumption function that integrate out latent states, obviating the need to estimate them explicitly.

The next step is to bring data to the dynamic panel model. I compile and harmonize microdata from 27 randomized cash and asset transfer experiments in low- and middle-income countries. This dataset is uniquely suited to the task of identifying poverty traps for three reasons. First, each study features a randomized transfer to households, totalling over $50 million in randomized assistance to the poor. These transfers generate a large, exogenous shock that can push households from low-asset to higher-asset steady states, allowing me to trace out the dynamics in between. Second, the studies feature a baseline survey and at least one follow-up, allowing me to observe households' asset dynamics over time. Third, detailed household surveys provide data on households' income and consumption decisions, which allow me to identify the production function, consumption function, and, in turn, the asset accumulation equation.

My first set of findings focus on study level average structural functions. I find that the household production function often exhibits non-concavities at low asset levels. Of the studies in my sample, 13 (48%) show evidence of increasing marginal returns over some range of assets, deviating from the standard assumption of diminishing marginal returns. In contrast, the remaining studies conform more closely to neoclassical predictions, with smooth, concave production functions and strictly diminishing returns. These results suggest that non-concavities are a genuine feature of poor households' production environment.

Second, I find the household consumption function demonstrates clear non-monotonicities in assets, consistent with forward-looking behavior aimed at reaching a higher steady state, in 16 (64%) studies. These non-monotonicities are strongly predicted by the shape of the production function: when the production function is non-concave, there is a 69% chance the consumption function is non-monotonic, whereas when the production function is concave, both monotonic and non-monotonic consumption are equally likely.

Third, the asset accumulation equation shows clear evidence of multiple steady states, and therefore poverty traps, in 15 (60%) of studies, with the remaining studies showing a single, stable steady state.

While the average asset accumulation function is informative, it masks substantial heterogeneity across households. Simulating each household's asset accumulation equation, I find that, on average, only 25% of households are in a poverty trap – even though

59% hold fewer assets than the threshold implied by the study level average structural functions. This surprising result stems from two sources: differences in ability and endogenous consumption responses. Some households are so productive that, regardless of their starting assets, they accumulate wealth and reach a high-asset steady state. Others so unproductive that they lose assets until they settle at a low-asset equilibrium, no matter their initial assets. A third group lies between these two extremes: for them, multiple steady states exist, and their long-run outcomes depend on initial assets. Similarly, endogenous consumption responses mean that while households typically face fixed costs, many reduce consumption as they near a critical threshold, allowing them to accumulate assets and escape the trap.

Next, I examine the mechanisms that could generate non-concave production functions and therefore poverty traps. Using data on livestock ownership and estimated proximity to the household poverty trap threshold, I find a sharp increase in the probability of ownership of large indivisible assets – specifically, a 23% increase in goat ownership and a 42% increase in cow ownership – at the threshold. These assets likely represent fixed costs that induce production non-concavities. In contrast, ownership of smaller animals, such as chickens, remains unchanged.

I also find qualitative support for fixed costs as a mechanism. Two studies in my sample discuss fixed costs as a barrier to investment: cows in Balboni et al. (2022) and a metal roof in Haushofer and Shapiro (2016). In both cases, my estimated poverty trap thresholds align closely with the price of these assets. In addition, I find a significant correlation between poverty trap status and household education and age, consistent with the idea that ability interacts with fixed costs to determine who faces a trap.

Finally, I use the estimated production and consumption functions to calibrate the guiding theoretical model. By imposing household optimization and a specific utility function, I can solve the model and evaluate household growth dynamics across a range of counterfactuals, using a finite-difference upwind scheme to account for non-convexities in the value function (Achdou et al. 2022). Households in this calibrated model endogenously adjust their consumption and savings decision in response to changes in the stochastic environment or promises of future income support, which I cannot infer from the dynamic panel estimates alone.

I use the calibrated model to study how optimal transfer policies should be designed when poverty traps may exist. A central question in the anti-poverty literature is whether to target impact or deprivation (Haushofer et al. 2025). Should transfers go to the very poorest, those with the highest marginal utility of consumption, or to households near the poverty-trap threshold, who have the highest marginal return? I find that optimal policies typically target deprivation, rarely diverting resources from the poorest towards those near the trap threshold.

The paper contributes to three literatures. First, it provides harmonized experimen-

tal evidence on poverty traps across diverse settings. Previous studies examined single contexts in isolation (Bleakley and Ferrie 2016; Barrett et al. 2018; Banerjee et al. 2019; Balboni et al. 2022; Bari et al. 2024), while broader reviews (Azariadis and Stachurski 2005; Kraay and McKenzie 2014) rely on cross-country or study-by-study comparisons. I assemble a large global panel of randomized transfer experiments, enabling systematic estimation of poverty traps from harmonized household microdata.

Second, I examine the empirical implications of dynamically optimizing households in settings where poverty traps may arise. Prior empirical work, such as Balboni et al. (2022), treats baseline assets as quasi-random and focuses on estimating a representative asset transition path. I instead incorporate household optimization and latent ability into a dynamic panel framework, yielding household-level asset transition equations. This richer structure makes it possible to separate households who are truly trapped, those who would reach a higher-asset steady state if shifted above a threshold, from those who are merely poor today but would naturally accumulate assets. To identify these dynamics, I use the randomized transfers and internal instruments that address endogenous investment under a weaker condition: households cannot systematically deviate from their long-run steady state, rather than requiring asset levels to be orthogonal to asset accumulation.

Finally, this paper contributes to a growing literature in economics that combines experimental evidence through meta-analysis (Meager 2019, 2022; Bandiera et al. 2022; DellaVigna and Linos 2022; Kremer et al. 2023; Crosta et al. 2024; Breza et al. 2025). Like Kremer et al. (2023), who study the link between water chlorination and child mortality, I revisit primary data to address a new question: the shape of household asset dynamics. My focus is not on treatment impacts, but on whether these dynamics imply multiple steady states consistent with poverty traps. By harmonizing microdata and embedding it in a unified structural framework, my analysis uncovers dynamic mechanisms that single trials struggle to reveal in isolation.

2 A Dynamic Model of Poverty Traps

This section develops a dynamic model of household growth to formalize the mechanisms that can generate poverty traps. The goal is to guide the empirical analysis that follows and understand why some households escape poverty while others remain stuck, even when facing similar environments. I review a standard model of poverty traps that integrates three features: (i) a production technology with a fixed cost to access high returns, which creates a non-convex opportunity set; (ii) two sources of unobserved heterogeneity—persistent shocks and time-invariant ability—that explain why different households follow different paths and; (iii) forward-looking households who optimize consumption and savings, allowing their expectations and patience to shape their own trajectory (Skiba

1978; Buera 2009; Ikegami et al. 2016). Together, these features yield predictions about how income, consumption and assets behave under poverty traps.

2.1 Consumption and Production with Fixed Costs

To build intuition, consider a model in which households can produce with two technologies: a low-return technology, $f_L(k)$, and a high-return technology, $f_H(k)$, available only after paying a fixed cost $F > 0$. The production function is the upper envelope of these two technologies, illustrated in orange in Figure 1:

$$y_{it} = f(k_{it}; \omega_{it}, \eta_i) = \exp(\omega_{it} + \eta_i) \max \left\{ \underbrace{A_L k_{it}^{\alpha}}_{f_L(k_{it})}, \underbrace{A_H((k_{it} - F)^+)^{\alpha}}_{f_H(k_{it})} \right\}, \quad A_H > A_L \tag{1}$$

Here, k_{it} denotes household i's assets at time t, and y_{it} is output. The fixed cost, F, creates a barrier to accessing the high-return technology – an indivisible investment such as a cow in Balboni et al. (2022), which costs a year's income, or a carpentry machine in Bassi et al. (2022), which requires savings equivalent to 18 months of average profit.

Two unobserved factors drive differences in production outcomes: a time-invariant component, η_i, and a persistent shock, ω_{it}. The fixed term η_i captures traits such as ability, motivation, or grit that remain stable over time. The shock ω_{it} represents serially correlated fluctuations, like health problems or weather shocks, that vary across periods but remain linked over time. Together, these components generate heterogeneous and uncertain production paths, even among households with the same asset levels.

Given their production possibilities, households choose a consumption path, c_{it}, to maximize lifetime expected utility:

$$V(k_{i0}, \omega_{i0}, \eta_i) = \max_{c_{it}} E_0 \int_0^{\infty} e^{-\gamma t} u(c_{it}) dt \tag{2}$$

subject to

$$dk_{it} = f(k_{it}; \omega_{it}, \eta_i) - c_{it} - \delta k_{it} \quad dt \tag{3}$$

where $u(\cdot)$ is the household's utility function, γ is the discount rate, and ω_{it} follows a stochastic process.

The solution is an optimal policy rule $c_{it} = \pi(k_{it}, \omega_{it}, \eta_i)$, linking consumption to assets, current shocks, and ability. Households are assumed to have no access to credit or rental markets–that is, they cannot borrow or rent capital. Although restrictive, this assumption isolates the pure dynamics of asset accumulation and plausibly reflects the empirical context: across studies that measure debt, 61% of households report none, and

the median positive debt is only \$88.02 compared with median assets of \$415.36.[4]

2.2 Model Dynamics and Testable Implications

Having defined household production and the budget constraint, I next examine the dynamics they imply–how fixed costs and forward-looking consumption decisions interact to generate, and potentially reveal, poverty traps.

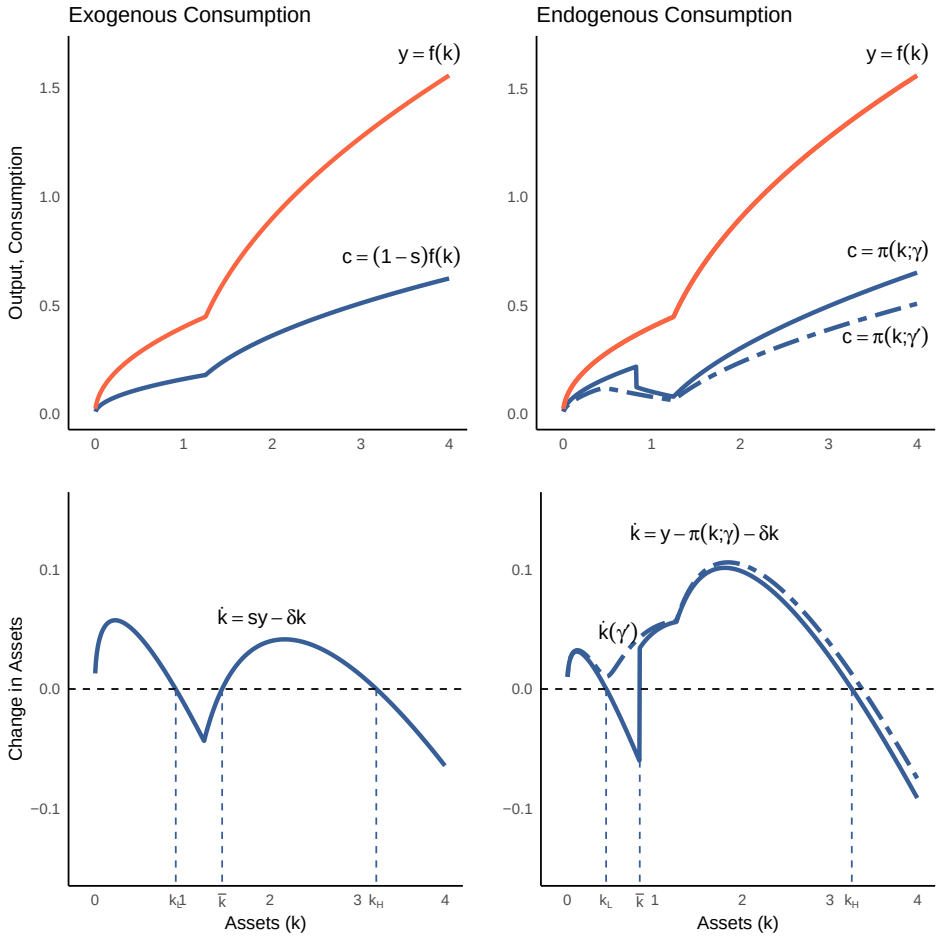


Figure 1: Poverty traps with fixed costs – exogenous and endogenous consumption

Notes: This figure shows the production and consumption functions (top row) and the resulting asset accumulation equations (bottom row) for three scenarios. Production functions are depicted in orange, whilst consumption and asset accumulation functions are in blue. The left column shows a Solow model with fixed costs and exogenous consumption, $c_{it} = (1-s)y_{it}$, generating a poverty trap. The solid lines in the right column show the same production function, but now with endogenous consumption, $c_{it} = \pi(k_{it};\gamma)$. The dot-dashed lines in the right column show a scenario where households value future consumption more, $c_{it} = \pi(k_{it};\gamma')$, eliminating the poverty trap. Consumption policies are generated using value function iteration with no stochastic shocks, $\omega_{it} = 0$, and fixed productivity, $\eta_i = 0$. Parameter values are: $A_L = 0.4, A_H = 0.9, \alpha = 0.5, F = 1.0, \delta = 0.25, \gamma = 0.05, \gamma' = 0.0425$.

[4]There is little evidence that low average debt reflects 'study-level' market clearing, where other sample households lend to indebted ones. Across studies, only 19% of households report any lending, with a median positive loan of \$33.33. However, only four studies provide detailed lending data, so this evidence is only suggestive.

2.2.1 Exogenous Consumption and Poverty Traps

The left column of Figure 1 presents a benchmark in which households consume a fixed share of income, $c_{it} = (1-s)y_{it}$, rather than making forward-looking choices. This setting resembles a Solow model with a constant savings rate s, but includes indivisibilities in production. The asset-accumulation equation crosses zero three times, yielding two stable steady states, k_L and k_H, separated by an unstable threshold $\bar{k}$. Households starting below $\bar{k}$ converge to poverty, while those above it converge to prosperity – a classic poverty trap.

2.2.2 Endogenous Consumption and Trap Thresholds

When households instead optimize consumption to maximize lifetime utility, the dynamics change. The right column of Figure 1 shows that the optimal policy, $c(k)$, departs from proportional consumption: households deliberately reduce consumption near the threshold to accumulate assets and reach the high-return technology. This endogenous savings response moves the threshold leftward, allowing households to escape poverty from lower initial asset levels than in the fixed-consumption benchmark.

In some cases, forward-looking behavior can eliminate the poverty trap entirely. The dot-dashed line in the right column of Figure 1 shows that when households place greater value on future consumption ($\gamma' < \gamma$), they accept very low consumption temporarily to save and accumulate assets from almost any starting point. The bottom-right panel shows that the asset-accumulation equation now has a single stable high-asset steady state, k_H.

The model implies three testable predictions. First, a fixed cost changes the shape of the production function: instead of smooth diminishing returns, it creates a kink with low returns below and high returns above. Second, facing such costs, forward-looking households reduce consumption near the fixed cost to self-finance the indivisible investment. Third, the interaction of these forces produces distinctive asset dynamics, households that heavily discount the future may face multiple steady states, while more patient households can save through the threshold and remove the trap. Together, these results suggest that identifying poverty traps requires observing the joint behavior of production, consumption, and assets rather than focusing solely on asset dynamics alone.

2.3 Implications of Stochastic Shocks and Heterogeneity

The deterministic model clarifies the core intuition, but real households face shocks and differ in productivity. Incorporating these features illustrates how poverty traps operate in practice and highlights why simple averages can be misleading.

Random productivity shocks, ω_{it}, complicate the definition of a poverty trap. In the

deterministic model, a household's future depends solely on its initial endowment. In contrast, with stochastic shocks, even households with substantial assets may fall into poverty due to repeated crop failures or chronic illness. Conversely, households with few initial assets may escape poverty after a series of favorable events. Many will lie in between these extremes, moving towards or away from steady states depending on the shocks they receive. As a result, the sharp threshold, $\bar{k}$, from the deterministic model no longer applies. Instead, there is a probability distribution over future asset levels that depends on current assets, shocks, and ability.

Fixed productivity differences, η_i, further interact with these shocks to shape households' economic trajectories. More productive households (high η_i) may accumulate assets despite occasional setbacks, while less productive ones (low η_i) may remain trapped even after favorable shocks. Households understand these dynamics, and their optimal consumption policy, $c_{it} = \pi(k_{it}, \omega_{it}, \eta_i)$, responds nonlinearly to assets, shocks, and ability. Thus, whether households in a study appear "trapped" depends on the joint distribution of initial ability, shocks, and assets. In some environments, most households remain trapped; in others, the same technology produces a unique, upward path for the majority. Averaging across these divergent trajectories can obscure the true mechanisms. It may even create the illusion of multiple steady states by "joining the dots" across households with differing abilities, when in fact, only one steady state exists for each household.

3 Data and Stylized Facts

I now describe the data and key variables I harmonize for use in my empirical analysis, before turning to stylized facts using the microdata.

3.1 Sample Construction

To identify poverty traps, I construct a novel panel dataset of poor households from 27 randomized transfer experiments across 17 countries, drawn from a wider harmonized sample of 58 RCTs. I assemble the sample through a three-step process. First, I identify a comprehensive list of potential studies by drawing from existing meta-analyses of transfer programs and conducting systematic keyword searches in academic databases and trial registries. Second, this list is screened for studies that meet two inclusion criteria: (1) the study must feature a randomized assignment to a transfer, and (2) it must measure at least one of the primary outcomes of interest (assets, income, or consumption) at the household level or finer detail. Third, I acquire the raw microdata for all qualifying studies, either from publicly available sources (such as the J-PAL or IPA dataverses) or by direct request to the principal investigators. Studies that meet the selection criteria but for which microdata are ultimately unobtainable are excluded from the final sample.

This leaves 58 studies with microdata that I subsequently harmonize. The main analysis sample in this paper focuses on the 27 studies which measure income and assets at least twice.[5][6]

The final sample is geographically diverse, covering 17 countries with broad representation across Sub-Saharan Africa (15 studies), South Asia (6 studies), and Latin America and the Caribbean (5 studies), with a single study drawn from East Asia and the Pacific, shown in Appendix Table A.1. Reflecting researcher focus on the poorest regions, the sample covers countries predominantly classified as low income (5 countries) or lower middle income (10 countries) by the World Bank, shown in Appendix Table A.2. In total, the sample comprises 74,811 households and 263,292 observations across all survey waves, with 15,416 households in low income countries and 40,358 households in lower middle income countries, approximately 75% of the full sample. Households in my sample are typically poorer than the average household in their respective countries, with all but two studies showing lower average daily consumption than the World Bank's country-level measure.

3.2 Variable Measurement and Harmonization

To create a consistent panel dataset, I harmonize the primary outcomes and treatment variables across all studies. My general principle is to defer to the contextual knowledge of the original authors, using their variable definitions wherever possible before attempting reconstruction or imputation.

3.2.1 Treatment

The key treatment variable is the randomized transfer, which provides a source of exogenous variation in household assets. For each household, I harmonize data on the size of the transfer (converted to 2020 USD), its timing, and its type (cash or in-kind). Harmonizing this information allows me to construct a household-month level panel, indicating precisely when each household received the treatment relative to each survey wave. Table A.3 shows study level treatment summary statistics. On average, studies randomly assign 44% of households to treatment, with a median transfer of $448, corresponding to approximately 10 months of baseline consumption. Transfer installments are typically one-off or monthly, with households receiving on average two installments of cash or assets. The transfers, summed across all studies, constitute approximately 50 million USD in randomized assistance to the poor.

[5] In future work I hope to expand the analysis in this paper to include cases where income or assets are not measured.

[6] I discuss differences between the wider sample and analysis sample in more detail in Section 3.2.4, in short I find little evidence of systematic differences or selection.

3.2.2 Outcome Variables

The primary outcome variables–household assets, income, and consumption–are harmonized to the household month level to allow for direct estimation of household's production function and consumption choices.

To measure asset dynamics, I extract the authors' provided measures of total, productive, and livestock assets. If these aggregates are not directly available, I reconstruct them using reported quantities and prices.[7] I do not redefine the authors' original asset definitions (for instance, by adding the value of land to total assets if the paper originally excludes it) in order to preserve the authors' intents given the studies' specific contexts.

Income and consumption outcomes are harmonized to the household-month level. I collect the authors' measures of total income, wage income, microenterprise profits, and microenterprise revenue, as well as total consumption. The harmonized sample also includes data on savings and debt, but these are measured less frequently across studies and therefore do not form a central part of the analysis. Finally, the sample includes measures on the number and price of any livestock the household owns since these are key productive assets in many low-income contexts and are frequently recorded across studies.

Households in the sample are very poor, with the median household owning a total of just $595 assets, shown in Appendix Table A.4 – the randomized transfer represents a near doubling of assets for the median household. Households have very low levels of consumption, typically $196 per month, corresponding to an average daily value of $1.41 per person, below the World Bank's $2.15 per day poverty line.

3.2.3 Baseline Covariates

In addition to the primary treatment and outcome variables, I harmonize a set of baseline household covariates from the pre-treatment surveys. These demographic and economic characteristics are used to describe the sample population and explore mechanisms and correlates of traps. The harmonized covariates include the age and years of education of the household head; the total household size; the percentage of household members who are female; and a common proxy for housing quality and baseline wealth, an indicator for the material of the household's roof. The median household has five members, with an average (adult) age of 37 years – most households have very little education, with the median household completing on average 6 years of education, shown in Appendix Table A.6.

[7] When only quantities are available, I impute the price of a given good using the median price from the most recent, price-reporting study in the same country (within a five-year window). This only occurs in a single study, Ethiopia, Blattman (2018), and results are robust to excluding this study.

3.2.4 Sample Selection and Panel Construction

The paper's main analysis requires observations of households' income, consumption, and assets at multiple time periods, naturally not all studies measure these outcomes at every survey wave, or at all. Therefore, I restrict the sample to households that report at least two observations of income and assets, leaving a final sample of 74,811 households and 263,292 observations across all survey waves.[8] I find little evidence of systematic differences between this sample and the wider set of studies with harmonized microdata. Appendix Table A.7 shows evidence of only two imbalances at the ten percent level between the two samples: households in the main analysis sample own slightly more pigs (+0.4, $p = 0.036$) and receive a higher transfer (+\$77.8, $p = 0.052$) however across twenty comparisons one would expect two to be significant at the ten percent level by chance.

Finally, my preferred specification pools multiple income, consumption, and asset measures to generate proxies for each outcome. For instance, income is measured using total income, but if this is not available it is substituted with total reported profits, revenue, or wage income in that order. Similarly, assets are measured by taking total assets, productive assets, or livestock assets respectively. Outcome statistics using these proxy measures are reported in Appendix Table A.5.

3.3 Stylized Facts

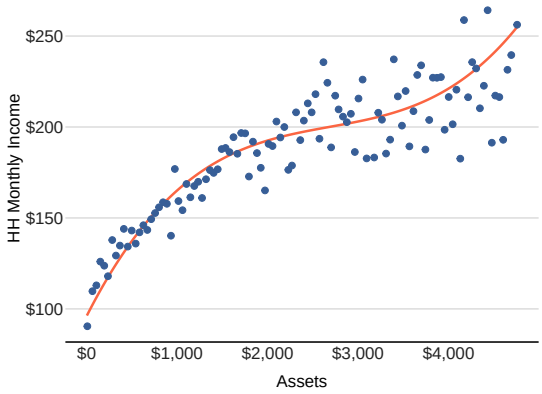
A first look at the data reveals a puzzle: on aggregate, the data shows patterns that are remarkably consistent with the Solow model but when considering individual studies, there are patterns suggestive of poverty traps.

In Figure 2, I pool all the microdata across studies and find incomes smoothly rise with assets (Figure 2a), consumption similarly increases with assets (Figure 2b), and asset dynamics reflect this, with a single stable steady state (Figure 2c). These patterns all match the predictions of a simple Solow model, suggesting the poor are not necessarily trapped, but rather differences in human capital, institutions, or other unobservables may explain persistent poverty.

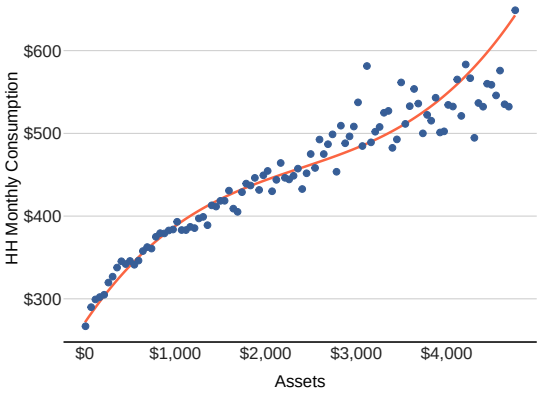
However, turning to individual studies, I find cases that match the predictions of the endogenous consumption model presented in Section 2, shown in Figure 3. For instance, in Tanzania (Baird et al. 2024) incomes initially slowly rise with assets, but only after assets approach the approximate price of a key lumpy investment, a cow, do incomes increase rapidly (Figure 3a).[9] Similarly, in Bangladesh (Balboni et al. 2022), incomes plateau around the price of a cow, before rising sharply thereafter (Figure 3b).

[8] Consumption is also required for the full model, however it is possible to estimate the production function using only observations of income and assets.

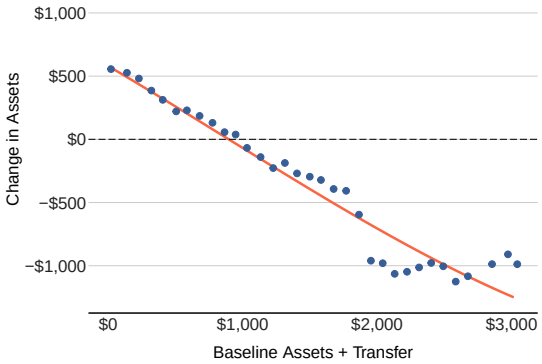
[9] There is no direct price data for cattle in Baird et al. (2024) so this figure is calculated by extrapolating from 2010 prices reported in Pica-Ciamarra et al. (2011), approximately 250-450 thousand Tanzanian Shillings.



(a) Household income by assets



(b) Household consumption by assets

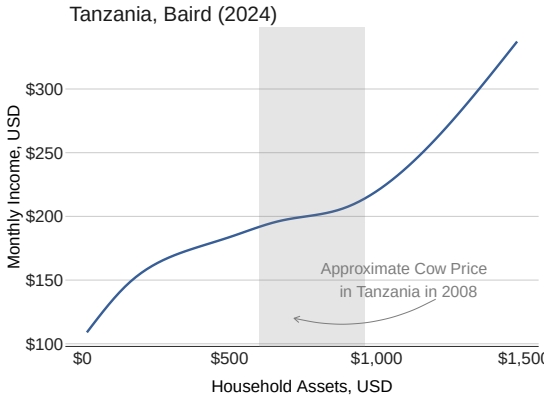


(c) Change in assets by assets + transfer

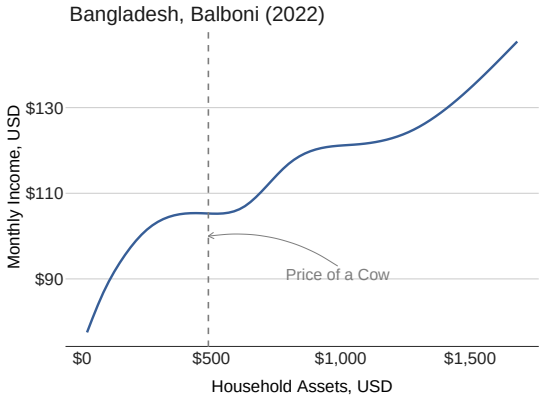
Figure 2: Pooled Microdata Stylized Facts

Notes: This figure shows three stylized facts about income, consumption, and assets amongst the poor. Income, consumption, and assets are winsorized at the 5th and 95th percentiles to reduce the influence of outliers. The solid line is a global third order polynomial regression line of best fit. The binscatter points are estimated using Cattaneo et al. (2024)'s binscatter method, with dummies for a common intercept and each study, the structural function is plotted at the common intercept level.

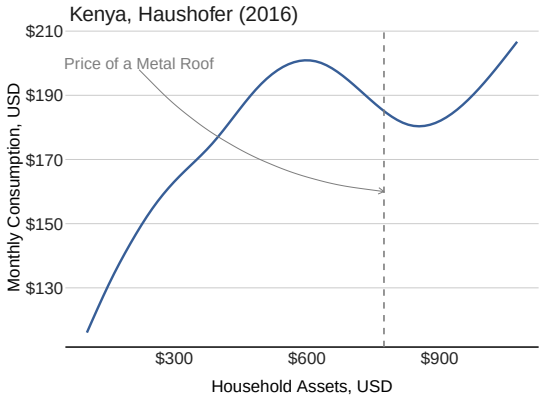
Consumption patterns also reflect non-linearities consistent with endogenous consumption behavior around lumpy investments. In Kenya, (Haushofer and Shapiro 2016), households appear to reduce their consumption *before* their asset holdings approach the price of a large indivisible asset discussed in the original paper, a metal roof (Figure 3c). Similarly, in Bangladesh, consumption dips sharply as households approach the price of a cow (Figure 3d). In the case of Bangladesh, the production and consumption patterns align and repeat – plausibly around multiples of the cow price – suggesting that households may be making lumpy investments in livestock as they accumulate assets.



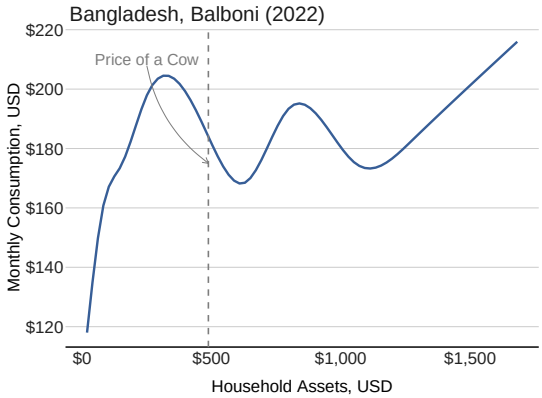
(a) Income by assets in Tanzania (Baird et al. 2024)



(b) Income by assets in Bangladesh (Balboni et al. 2022)



(c) Consumption by assets in Kenya (Haushofer and Shapiro 2016)



(d) Consumption by assets in Bangladesh (Balboni et al. 2022)

Figure 3: Study-Level Stylized Facts

Notes: This figure shows two stylized facts about income and consumption responses to assets in individual studies. The first panel fits a flexible regression function to income and assets in Tanzania and Bangladesh. The second panel fits a flexible regression function to consumption and assets in Kenya and Bangladesh. The sample is trimmed to remove outliers, keeping only observations with assets, consumption, or income between the 5th and 95th percentiles. The approximate cow price in Tanzania is calculated by taking the upper and lower end of cattle prices reported in Pica-Ciamarra et al. (2011) and adjusting for inflation to 2020 USD, I find a range of $598 to $958.

The discrepancy between the aggregate and individual study patterns motivates why considering heterogeneity is so important. Averaging over diverse contexts can blur sharp,

non-linear patterns in the data that are key to identifying poverty traps. Furthermore, such heterogeneity can appear across studies, motivating an approach that estimates the asset transition equation separately for each study but also within studies, as households may differ in their individual productivity or consumption preferences.

In the next sections, I develop a systematic approach to identifying and estimating poverty traps that accounts for both types of heterogeneity.

4 A Dynamic Panel Model of Income, Consumption, and Assets

To formally examine the existence of poverty traps, I develop and estimate a dynamic model of joint household income, consumption, and asset accumulation. The framework is deliberately flexible, designed to capture the complex, non-linear dynamics that the theoretical poverty trap model, discussed in Section 2, predicts without imposing restrictive functional form assumptions ex-ante. Lowercase variables of income, consumption, and assets, $\{y_{ijt}, c_{ijt}, k_{ijt}\}$, denote natural logarithms throughout this section and the next.

4.1 Production Function

I model the household's income generating process to capture three sources of heterogeneity that are critical for understanding poverty dynamics. First, the production function, $f_j(\cdot)$, is study-specific and nonparametric, allowing the relationship between assets and income to vary freely across different economic contexts. Second, the framework accounts for ex-ante heterogeneity through a latent, time-invariant household productivity term, η_{ij}, which may represent fixed differences in ability or "grit". Finally, it incorporates ex-post heterogeneity via a persistent productivity shock, ω_{ijt}, to capture serially correlated events such as weather shocks or lingering health issues. These features are formalized in the standard log-production function:

$$y_{ijt} = f_j(k_{ijt}) + \eta_{ij} + \omega_{ijt} \tag{4}$$

$$\omega_{ijt} = \rho_j \omega_{ijt-1} + \xi_{ijt} \tag{5}$$

where y_{ijt} is the log-income generated by household i in study j at time t using log-assets k_{ijt}.

The persistent productivity shock, ω_{ijt}, follows a study-specific AR(1) process with i.i.d. innovations ξ_{ijt}. Households choose their asset holdings, k_{ijt}, based on knowledge of their fixed productivity, η_{ij}, and the previous period's shock, ω_{ijt-1}, but before realizing the current innovation, ξ_{ijt}. Thus, assets are predetermined when current income is

realized.

4.2 Consumption and Asset Process

Within this production environment, households make endogenous consumption and investment decisions based on current state variables. Following the theoretical framework laid out in Section 2, these state variables include the current asset stock, k_{ijt}, fixed productivity, η_{ij}, and the current productivity shock, ω_{ijt}.

To capture the complex and nonlinear relationship between these variables and household behavior, I treat the consumption function as a flexible, nonparametric function of the state variables and their interactions. This allows, for instance, households to reduce their consumption sharply when approaching a production function non-concavity and for this response to vary as a function of ability and the current shock. Combined with the standard intertemporal budget constraint, this consumption decision governs the evolution of assets:

$$c_{ijt} = g_j(k_{ijt}, \eta_{ij}, \omega_{ijt}) + u_{ijt} \tag{6}$$

$$K_{ijt+1} = \underbrace{\exp\left(f_j(k_{ijt}) + \eta_{ij} + \omega_{ijt}\right)}_{Y_{ijt}} - \underbrace{\exp\left(g_j(k_{ijt}, \eta_{ij}, \omega_{ijt})\right)}_{C_{ijt}} + \delta(K_{ijt}) + \nu_{ijt} \tag{7}$$

The error term u_{ijt} captures unobserved preference shocks or measurement error in consumption, whilst ν_{ijt} reflects idiosyncratic disturbances in the asset accumulation process. The budget constraint is therefore the difference between a household's production and consumption, net of depreciation $\delta(K_{ijt})$, with an additional i.i.d. error term, ν_{ijt}.

4.3 Model Assumptions

The model makes several simplifying assumptions to maintain tractability. First, the production function has a log-additive structure, so that ability and productivity shocks enter multiplicatively in levels. While this is a standard assumption, it does assume higher ability households produce proportionally more output at all asset levels. This may be a reasonable approximation in rural settings, where households use similar technologies and face comparable markets, but could be more restrictive in urban settings with greater occupational and technological heterogeneity.[10]

Second, the model assumes that a household's state variables are limited to current assets, ability, and the current productivity shock. While this nests models of endogenous consumption, discussed in Section 2, this simplification rules out richer dynamics that could arise if, for instance, ability evolves over time or depends on past consumption or investments in human capital. It would also be violated if households interact through

[10] I thank David McKenzie for raising this point.

local markets with aggregate shocks, insure each other, or benefit from peer effects, in which case the relevant state would include aggregate or peer variables.

5 Identification

The complete model of production, consumption, and asset accumulation forms a nonlinear dynamic panel model with multiple latent state variables, η_{ij} and ω_{ijt}. Identifying such models is a well-known challenge, the strategy in this paper combines standard techniques for nonparametric panel models with exogenous variation from a randomized cash/asset transfer. This section outlines the identification strategy in two parts: first, the production function; and second, the consumption function. A more detailed discussion of the identifying assumptions and identification strategy is provided in Appendix E.

5.1 Production Function Identification

The production function is identified from two complementary sources of variation through a set of conditional moment restrictions. The primary source is an external instrument: a randomized asset transfer that exogenously shifts a household's stock of productive assets, which traces out the overall shape of the production function. However, the transfer alone may not provide sufficient variation to identify local nonlinearities in the production function due to its semi-discrete nature – most RCTs only include a few discrete transfer amounts. Therefore, I supplement this external variation by a set of internal instruments, derived from the model's dynamic panel structure, which provide continuous variation in assets, tracing out finer details of the production function around the randomized transfer.

Substituting the AR(1) process for the productivity shock in Equation 5 into the production function in Equation 4 generates the following, standard, expression for income:

$$y_{ijt} = \rho_j y_{ijt-1} + f_j(k_{ijt}) - \rho_j f_j(k_{ijt-1}) + \overbrace{(1-\rho_j)\eta_{ij} + \xi_{ijt}}^{w_{ijt}} \tag{8}$$

$$\Delta w_{ijt} = \xi_{ijt} - \xi_{ijt-1} \tag{9}$$

where w_{ijt} is the composite error term in the income equation, our goal will be to find valid instruments that are correlated with income and assets but uncorrelated with this error term, which is a composite of fixed household ability and idiosyncratic productivity innovations.

5.1.1 Randomized Transfer as an External Instrument

The randomized transfer provides a core source of identifying variation for the production function. Its key advantage is that it solves the classic endogeneity problem where unobserved factors, such as households' innate ability, η_{ij}, are correlated with both their income and accumulated assets. By design, the random allocation of the transfer provides an exogenous shock to a household's asset path that is independent of this latent ability, η_{ij}, and any productivity innovations, ξ_{ijt} and therefore w_{ijt}. This is formalized by the moment condition:

$$E\left[\overbrace{(1-\rho_j)\eta_{ij}+\xi_{ijt}}^{w_{ijt}}|F_{ij}^T\right]=0 \tag{10}$$

where F_{ij}^T is the history of transfers for the household. While F_{ij}^T provides exogenous shifts in assets, randomized transfers typically only take a few discrete values, limiting their ability to trace out the full local shape of the production function. Therefore, I supplement this external variation with a set of internal instruments derived from the panel structure of the data described next.

5.1.2 Lagged Levels as Internal Instruments

To complement the randomized transfer instrument, I introduce internal instruments that provide a more continuous source of identifying variation. These instruments help trace how output responds to incremental changes in assets, beyond the discrete shifts induced by the transfer. Specifically, I implement the System GMM (SYS-GMM) estimator developed in Arellano and Bover (1995) and Blundell and Bond (1998, 2000), which introduces two sets of instruments, in differences and levels, to exploit these internal moments efficiently.

The first set of instruments, following the standard Arellano-Bond procedure, uses lagged levels of assets, k_{it-1}, and income, y_{it-2}, as instruments for the differenced production equation. Here, households make decisions about investment at time t with knowledge of their ability, η_{ij}, and the current productivity innovation, ξ_{ijt}, but changes in assets are only realized next period, at $t+1$. Therefore, previous assets are valid instruments: they are correlated with current assets and uncorrelated with the change in productivity innovations, $\Delta w_{ijt}=\xi_{ijt}-\xi_{ijt-1}$. A similar argument works for income – however, income is realized instantaneously at time t so we must go back one additional period to ensure the instrument is uncorrelated with last period's innovation, ξ_{ijt-1}, thus income two periods removed, y_{it-2}, is a valid instrument.[11] The first set of instruments

[11]In the presence of measurement error in the production process, it is necessary to lag instruments an additional time period. I show similar results that incorporate measurement error using studies with

generate the moment conditions:

$$E\left[\Delta w_{ijt} | k_{ijt-1}\right] = 0 \tag{11}$$

$$E\left[\Delta w_{ijt} | y_{ijt-2}\right] = 0 \tag{12}$$

5.1.3 Lagged Differences as Internal Instruments

The second set of instruments use lagged first-differences of assets and income as instruments for the equation in levels. In this case, Δy_{ijt-1} and Δk_{ijt-1} are used as instruments which are orthogonal to $w_{ijt} = (1-\rho_j)\eta_{ij} + \xi_{ijt}$. The identifying assumptions underpinning this approach, $E\left[(1-\rho_j)\eta_{ij} | \Delta k_{ijt}\right] = 0$, require additional restrictions on the initial conditions, $E\left[(1-\rho_j)\eta_{ij} | \Delta y_{ij2}\right] = 0$, and are satisfied by assuming joint stationarity of income and assets.

Intuitively this requires that expected deviations in income and assets from a household's long run mean–in this case the steady state–are not systematically related to the household's fixed productivity term, η_{ij}. These generate the following moment conditions:

$$E\left[w_{ijt} | \Delta k_{ijt-1}\right] = 0 \tag{13}$$

$$E\left[w_{ijt} | \Delta y_{ijt-1}\right] = 0 \tag{14}$$

This strategy is a common way of identifying production functions in the presence of unobserved heterogeneity and persistent shocks, as laid out in Blundell and Bond (2000). However, one concern may be that the households facing poverty traps do not satisfy the standard moment conditions and their dependence on initial condition restrictions – particularly given models of poverty traps often feature path dependence based on initial asset conditions. To alleviate this concern, I simulate from the dynamic model with endogenous household consumption presented in Section 2, calibrated to match the empirical moments of a representative study in the sample. I then show the population moments of the simulated data satisfy the SYS-GMM moment conditions. Table A.10 shows the results of this exercise and a placebo test using endogenous instruments – all moments are satisfied when using the correct instruments. Finally, I report robustness results re-estimating the model without the level moments in Appendix D, which show similar production function estimates and poverty trap results (Appendix Figure D.12, Appendix Table D.28).

5.1.4 Instrument Relevance and Completeness

Nonparametric instrumental variables (NPIV) identification of a structural function, such as a production function, requires a *completeness* condition analogous to the relevance

longer panels in the appendix.

condition in standard IV settings (Newey and Powell 2003). Completeness ensures that the instruments vary sufficiently across the support of the endogenous variable, household assets, so that no other function can generate the same conditional expectation. A randomized transfer alone is unlikely to satisfy this condition, as most RCTs offer only a single or a few discrete transfer amounts, leaving the function unidentified outside those points. Internal instruments from panel data provide continuous variation and can potentially meet the condition over some asset range, but in the presence of a poverty trap, shocks may rarely shift households out of low-asset states, limiting their identifying power.

However, identification relies on the combined strength of all instruments. My strategy exploits both the discrete shift provided by the randomized transfer and the continuous variation from internal instruments. The transfer helps move households across the poverty trap threshold, while the internal instruments trace the shape of the production function on either side. Together, they are more likely to satisfy the completeness condition and enable credible identification of the production function.

5.1.5 Pooling Information Across Studies

A final identification challenge arises in studies with short panels, where within study variation is insufficient to identify the persistence parameter ρ_j. To address this, I impose a hierarchical prior, assuming each study's ρ_j is drawn from a common normal distribution:

$$\rho_j \sim N(\mu_\rho, \sigma_\rho^2) \tag{15}$$

This partial pooling borrows strength from data rich studies to share information with data poor ones, enabling identification of productivity dynamics without assuming identical persistence across studies.

In practice, longer panels identify μ_ρ and inform cross-study variation σ_ρ, while shorter panels draw from this distribution and average over uncertainty in ρ_j. As a robustness check, I report results excluding short panels (Appendix D).

By combining external randomization, internal panel variation, and the multi-study nature of my sample to pool information, I can identify the nonparametric production function, $f_j(\cdot)$, and the persistence of productivity shocks, ρ_j, across diverse study contexts.

5.2 Consumption Function Identification

In this section, I outline the identification of the nonparametric consumption function, $g_j(\cdot)$, the second component of the asset transition equation.

In a setting where the panel length grows, identification is straightforward because the fixed ability component can be consistently estimated from long-run averages of the income residual. However, in short panels with fixed T, identification is more challenging as the fixed effect cannot be consistently estimated. To address this, I follow the approach in Arellano et al. (2017, 2024) and leverage a nonparametric instrumental variables strategy that exploits the latent structure of the productivity shocks. More detailed discussion of the identification strategy and identifying assumptions are provided in Sections E.2 and E.3 of the Appendix.

5.2.1 Consumption and Asset Equations

Referring back to the production equations in (4) the residual income net of the returns to assets, is:

$$e_{ijt} = y_{ijt} - f_j(k_{ijt}) = \eta_{ij} + \omega_{ijt} \tag{16}$$

The consumption and asset accumulation equations are:

$$c_{ijt} = g_j(k_{ijt}, \eta_{ij}, \omega_{ijt}) + u_{ijt} \tag{17}$$

$$k_{ijt+1} = y_{ijt} - c_{ijt} + \delta(k_{ijt}) + \nu_{ijt} \tag{18}$$

where u_{ijt}, ν_{ijt} are i.i.d. mean zero measurement error or preference shocks independent of $k_{ijt}, \eta_{ij}, \omega_{ijt}$.

5.2.2 Identification with Long Panels

To build intuition, I first consider identification as the panel length grows, $T \to \infty$. With $T \to \infty$, identification is straightforward as I can consistently recover each household's fixed productivity component, η_{ij}, by averaging residual income, e_{ijt}, over time because ω_{ijt} is mean zero. Once η_{ij} and the persistent shock ω_{ijt} are identified from the income process, they can be treated as observed regressors in the consumption equation. Identification of $g_j(\cdot)$ then reduces to a standard nonparametric regression problem: with $(k_{ijt}, \eta_{ij}, \omega_{ijt})$ observed and u_{ijt} mean independent, it is possible to recover the functional form of g_j directly from the joint distribution of observables.

5.2.3 Identification in Short Transfer RCT Panels

With short panels, this direct approach fails because η_{ij} cannot be consistently estimated when T is fixed.[12] Averaging over time introduces a non-vanishing estimation error, the

[12] I show using simulations in Appendix Table G.31 that the direct fixed-effects approach performs poorly in this setting.

"incidental parameters" problem, so η_{ij} and ω_{ijt} must be treated as latent variables. Identification then requires additional structure. I therefore follow the arguments in Arellano et al. (2017) (ABB) with some additional parametric assumptions to simplify the identification argument.

Following this NPIV approach, I identify the consumption function by treating the latent productivity components, η_{ij} and ω_{ijt}, as unobserved regressors in the initial period consumption function, $g(\cdot)$. I use future residual income, $E_i = (e_{i2}, \ldots, e_{iT})$, as instruments since they are correlated with the latent productivities (due to the structure of the income process) but are independent of the initial period consumption error.[13] This yields the integral equation:

$$E[c_{i1}|k_{i1}, E_i] = \iint g(k_{i1}, \eta, \omega) \cdot f(\eta, \omega|k_{i1}, E_i)\, d\eta\, d\omega \tag{19}$$

In this equation, the left-hand side is observed from the data, and $g(\cdot)$ is the unknown consumption function to be estimated. The main simplification in my approach is the assumption that the latent variables are Gaussian.[14][15] This means the conditional density $f(\eta, \omega|k_{i1}, E_i)$ is identified and has a known, analytical (multivariate normal) form, as its underlying parameters $(\rho, \sigma^2_\eta, \sigma^2_\xi)$ are already identified from the production function above. Given a standard completeness condition to ensure this integral equation is invertible, the consumption function $g(\cdot)$ is identified. The argument then proceeds sequentially for the next time period, using the identified $g(\cdot)$ to recover the asset accumulation dynamics, allowing the full system to be identified even with short panels. The necessary completeness conditions and additional details are provided in Appendix E.3. With the production, consumption, and asset accumulation function identified, I turn next to estimation.

6 Estimating the Panel Model

Understanding the dynamics of poverty traps requires flexibly estimating both the production function and household consumption behavior. This poses two econometric challenges: estimating a non-parametric production function using the derived moment conditions, and recovering a consumption function dependent on latent productivity. This section details a two-part estimation strategy, which addresses these challenges by pairing a novel quasi-Bayesian approach for the production function with a stochastic expectation-

[13] This is a common strategy to identify consumption responses in the literature (Hall and Mishkin 1982; Blundell et al. 2008) as noted in Arellano et al. (2017).

[14] ABB show how to relax Gaussianity by nonparametrically identifying the joint distribution of latent components under additional completeness conditions.

[15] This Gaussian assumption is specific to the consumption function, the identification and estimation of the production function do not rely on it.

maximization (EM) algorithm for the consumption policy. Implementation details are provided in Appendix F.

6.1 Estimating Production: A Quasi-Bayesian Approach

The objective is to estimate the production function $f_j(\cdot)$ flexibly in each site j, without imposing a specific functional form. This requires nonparametric estimation under instrumental-variable moment conditions (10)–(14), a setting known to be difficult: standard nonparametric instrumental variables methods and dynamic panel SYS-GMM estimators often perform poorly in finite samples. The challenge is therefore to recover the shape of $f_j(\cdot)$ in a way that permits nonlinearities while still imposing enough structure to deliver stable, precise estimates.

A key difficulty is the ill-posed inverse problem (Blundell et al. 2007): the data offer limited information about the full shape of $f_j(\cdot)$, so small sampling noise can lead to unstable, wiggly estimates. In the context of poverty traps, this is especially problematic, as spurious non-convexities may be misread as evidence of multiple steady states.

6.1.1 Quasi-Bayesian Framework

To address these challenges, I adopt a quasi-Bayesian estimator (Chernozhukov and Hong 2003; Kankanala 2025) that combines the GMM moment conditions with a prior over the shape of the production function. The prior encodes beliefs about plausible shapes for the production function, such as smoothness or monotonicity, providing additional structure in data sparse settings, yet becomes dominated by the data when enough information is present. By introducing a prior over the object of interest, the production function itself, while leaving the likelihood and nuisance parameters unrestricted, quasi-Bayes offers a powerful tool to estimate structural functions in data demanding environments with minimal additional assumptions.

6.1.2 Gaussian Process Prior and Kernel Specification

Therefore, following Kankanala (2025), I implement a quasi-Bayesian approach that places a Gaussian Process (GP) prior over the production function, $f_j(\cdot)$, shown in Equation (20). A Gaussian Process can be thought of as a probability distribution over *functions*, much like a multivariate normal distribution is a probability distribution over *vectors*. Intuitively, instead of specifying a single functional form for production, the GP prior describes a whole family of plausible smooth curves that could represent the production technology.

The "smoothness" and shape of these curves are governed by the covariance kernel, $\Sigma(k_{jt}, k'_{jt})$, which determines how similar the outputs (income, y) should be when the inputs (assets, k) are close together. The kernel, shown in Equation (21), thus controls

how quickly correlation decays as the distance between households' assets increase.

$$f_j(k_{jt}) \sim \mathcal{GP}\left(0, \Sigma(k_{jt}, k'_{jt})\right) \tag{20}$$

$$\Sigma(k_{jt}, k'_{jt}) = \left(1 + \frac{\sqrt{3}|k_{jt} - k'_{jt}|}{\ell_j}\right) \exp\left(-\frac{|k_{jt} - k'_{jt}|}{\ell_j}\right) \tag{21}$$

The kernel has a single hyperparameter: the length scale, ℓ_j. The length scale governs how quickly correlation decays as asset levels diverge: a large ℓ_j forces incomes at different assets to move together, yielding a smoother function; a small ℓ_j allows sharper local changes. When estimating poverty trap models, too large a length scale can obscure true nonlinearities, while too small a length scale can overfit noise, generating spurious "wiggles" that look like non-convexities.

6.1.3 Implementation

To construct the quasi-posterior, I form the sample analog of the conditional mean of the moment conditions, $\hat{m}(Z_n, f_j)$, where Z_n represents the corresponding set of instruments.[16] In practice, this involves drawing a function $f_j(\cdot)$ from the GP prior, computing the residuals, $w_{ijt}(f_j), \Delta w_{ijt}(f_j)$, and regressing these residuals on a flexible set of instrument basis functions.[17] I then construct the quasi-likelihood:

$$L_n(f) = E_n\left[\hat{m}(Z_n, f)' \hat{W}(Z_n) \hat{m}(Z_n, f)\right] \tag{22}$$

and combine it with the GP prior, $\pi_0(f)$, to obtain the quasi-posterior:

$$\pi\left(f | X_n, Y_n, Z_n\right) = \frac{\exp\left(-\frac{n}{2} L_n\left(f\right)\right) \pi_0\left(f\right)}{\int \exp\left(-\frac{n}{2} L_n\left(f\right)\right) \pi_0\left(f\right) df} \tag{23}$$

Sampling from Equation (23) produces a distribution over the production function, $f_j(\cdot)$, which under weak regularity conditions contracts around the true structural function (Kankanala 2025). After obtaining these draws, I enforce monotonicity in assets by projecting the GP posterior onto the space of monotone functions, following Lin and Dunson (2013).

I implement the model in Stan, a probabilistic programming language that uses Hamil-

[16] For example, the conditional mean for the moment described in Equation (10) is:

$$\begin{aligned} m(Z, f_j) &= E\left[w_{ijt}(f_j) | Z\right] = E\left[y_{ijt} - \rho_j y_{ijt-1} - (f_j(k_{ijt}) - \rho_j f_j(k_{ijt-1})) | F_{ij}^T\right] = 0 \\ &= E\left[(1 - \rho_j)\eta_{ij} + \xi_{ijt} | F_{ij}^T\right] = 0 \end{aligned}$$

[17] I use a 10th-order Bernstein polynomial basis for each level/difference instrument and their interaction. The transfers enter directly.

tonian Monte Carlo to sample from a Markov chain that approximates the posterior distribution. I estimate the length scale ℓ_j simultaneously with the production function using a hyperprior prior across studies, shrinking estimates of the length scale towards a common global mean. To ensure results are not driven by hyperparameter estimation choices, I show results are robust to two alternative estimation approaches in Appendix D: i) fixing the length scale after a burn-in period; and ii) removing the hierarchical prior across studies.

6.2 Estimating Consumption: A Stochastic EM Algorithm

The primary challenge in estimating households' consumption behavior is that the consumption function depends on unobserved productivity–both the fixed household effect and the persistent productivity process:

$$c_{ijt} = g_j(k_{ijt}, \eta_{ij}, \omega_{ijt}) + u_{ijt} \tag{24}$$

Estimating these latent variables directly is difficult in short panels – intuitively, it is hard to distinguish lucky farmers from good farmers when you only observe a few seasons of harvests. My estimation strategy therefore mirrors my identification approach: I recover the consumption function integrating over the distribution of latent productivity rather than attempting to estimate each household's productivity directly. Intuitively, it is easier to estimate how many good farmers there are in a study, than identify which one is which.

6.2.1 Why Direct Estimation Fails

In long panels, one could estimate the consumption function by first recovering each household's latent productivity. This involves decomposing the production residual $e_{ijt} = y_{ijt} - f_j(k_{ijt})$ into a fixed component $\hat{\eta}_{ij}$ and shocks $\hat{\omega}_{ijt}$, and then regressing consumption on assets and these estimated latent variables:

$$c_{ijt} = \sum_{b=0}^{B} \beta_{bj}\, \varphi_b(k_{ijt}, \hat{\eta}_{ij}, \hat{\omega}_{ijt}) + u_{ijt}.$$

This direct approach is infeasible in my setting due to the high uncertainty in the latent variable estimates driven by short panel lengths. In short panels (as in transfer RCTs), any direct estimates of the latent productivity terms would suffer from the incidental parameter problem, propagating large uncertainty through to the final estimates of the consumption function.

6.2.2 Estimation via Latent Variable Integration

Therefore, rather than attempting to estimate the latent productivity components directly, I recover the consumption policy rule by integrating over their posterior distribution:

$$E\left[c_{ijt} \mid k_i^t, y_i^t, c_i^{t-1}\right] = \int g_j\left(k_{ijt}, \eta_{ij}, \omega_{ijt}\right)\, p\left(\eta_{ij}, \omega_{ijt} \mid k_i^t, y_i^t, c_i^{t-1}\right) d\eta_{ij}\, d\omega_{ijt}, \tag{25}$$

where $k_{ij}^t, y_{ij}^t, c_{ij}^{t-1}$ collect the observed histories of assets, income, and consumption respectively.

This reframes the estimation problem: the goal is now to approximate the joint distribution of latent productivity conditional on the observed data, $p(\eta_{ij}, \omega_{ijt} \mid k_i^t, y_i^t, c_i^{t-1})$, well enough to compute the integral and recover $g_j(\cdot)$.

To do so, I begin by specifying Gaussian priors for the latent components:

$$\eta_{ij} \sim N\left(0, \sigma_{\eta j}^2\right) \tag{26}$$

$$\xi_{ijt} \sim N\left(0, \sigma_{\xi j}^2\right) \tag{27}$$

where the variances $\sigma_{\eta j}^2, \sigma_{\xi j}^2$ are estimated from moments of the production function problem (Equations (53)-(54)).

Although these priors are Gaussian, the posterior $p(\eta_i, \omega_{it} \mid k_i^t, y_i^t, c_i^{t-1})$ is not. Conditioning on the observed paths of income, assets, and consumption introduces nonlinear feedback through both the consumption function $g_j(\cdot)$ and the asset accumulation equation. If these functions were linear, the posterior would remain Gaussian and could be computed with a Kalman filter. But the nonlinear structure of $g_j(\cdot)$ breaks Gaussianity, ruling out closed-form solutions and shifting the empirical problem to approximating this intractable posterior through simulation methods.

6.2.3 Sampling Latent Productivity with PMCMC

To recover the posterior distribution of latent productivity, I use a procedure that alternates between two steps: proposing a household's underlying fixed productivity, and then simulating sequences of shocks that would make its observed income and asset path likely. This produces draws from the joint distribution of unobserved ability and productivity shocks implied by the data.

Formally, I implement Particle Markov Chain Monte Carlo (PMCMC), a class of methods designed for sampling from latent-variable posteriors in nonlinear, sequential settings (Andrieu et al. 2010; Creal 2012; Arellano et al. 2024). Specifically, I use a particle marginal Metropolis-Hastings (PMMH) algorithm which has two nested components. The outer layer is a standard Metropolis-Hastings sampler that proposes values for the

fixed household-specific productivity term η^*_{ij}. Conditional on each proposed η^*_{ij}, the inner layer runs a particle filter to simulate the sequence of productivity shocks and evaluate the likelihood contribution of the observed data.[18] As the filter processes the time series, it updates beliefs about the latent state and returns two key objects: (i) a simulated path of shocks, $\{\omega^*_{ijt}\}_{t=0}^T$, and (ii) an unbiased estimate of the marginal likelihood conditional on the proposed fixed effect, $\widehat{p}(y^T_{ij}, c^T_{ij}, k^T_{ij} \mid \eta^*_{ij})$.[19]

The unbiased likelihood estimate produced by the particle filter is then used in the Metropolis-Hastings acceptance step. This ensures that the accepted draws $\{\eta^*_{ij}, \{\omega^*_{ijt}\}_{t=0}^T\}$ jointly follow the posterior distribution of the latent productivity process. With a method to sample from the latent productivity posterior in hand, I can now turn to estimation of the consumption function.

6.2.4 Embedding in a Stochastic EM

Because the latent productivity draws produced by PMMH depend on the consumption function, and the consumption function itself depends on those same latent states, the two objects must be estimated together. A natural way to do this is to alternate between (i) updating beliefs about the unobserved productivity process and (ii) updating the consumption function given those beliefs. This leads to a stochastic EM algorithm, where each iteration refines estimates of both the latent states and the consumption function.

Formally, the PMMH sampler provides the E-step by integrating over the unobserved productivity, and the M-step updates the estimate of $g_j(\cdot)$ conditional on the simulated latent paths. The algorithm proceeds by repeating:

1. *E-Step*: Draw latent productivity paths, $(\eta^*_{ij}, \omega^*_{ij0}, \ldots, \omega^*_{ijT})$, from the joint posterior using PMCMC with Metropolis-Hastings.

2. *M-Step*: Regress observed consumption, c_{ijt}, on the simulated latent paths using a fourth-order tensor basis of ordinary polynomials, $\varphi_b(\cdot)$, to update the coefficients β_{bj} of the consumption function approximation.

Running the algorithm for S iterations and averaging over the implied functions pro-

[18] Specific implementation details can be found in Appendix Section F.2.2.

[19] Without measurement error in the production function, the particles collapse to a point mass because $e_{ijt} = \eta_{ij} + \omega_{ijt}$ exactly pins down the latent productivity path given η_{ij}. In that case, the algorithm simplifies to a standard Metropolis–Hastings sampler. In practice, however, allowing for a small amount of measurement noise in e_{ijt} improves mixing and helps avoid local modes in the stochastic-EM estimation. Incorporating measurement error in this way is not unusual in particle filters applied to problems in macroeconomics (Creal 2012). If the production function does in fact include measurement error, filtering is necessary because latent productivity must be inferred rather than directly recovered.

duces an estimate of $g_j(\cdot)$ that fully integrates over uncertainty in latent productivity:

$$\hat{\beta}_{bj} = \frac{1}{S}\sum_{s=1}^{S}\beta_{bj}^{s} \tag{28}$$

$$\hat{g}_j(k,\eta,\omega) = \sum_{b=0}^{B}\hat{\beta}_{bj}\varphi_b(k,\eta,\omega) \tag{29}$$

where β_{bj}^{s} are the M-step basis function coefficients estimated from iteration s.

6.2.5 Summary

Taken together, this procedure delivers a flexible and tractable way to estimate households' consumption behavior when productivity is only imperfectly observed. By integrating over latent productivity rather than attempting to estimate it directly, the stochastic EM algorithm avoids the incidental-parameter problems that arise in short panels, while allowing the consumption function to be estimated in a fully nonlinear and data-driven manner. With the estimation framework in place, I now turn to a series of Monte Carlo exercises that illustrate the production and consumption estimation performance across different environments.

6.3 Monte-Carlos: Validating the Estimation Strategy

I validate the estimation strategy through Monte-Carlo simulations. First, I generate data from the theoretical model, discussed in Section 2, and compare my estimator's ability to detect poverty traps to existing methods. Second, I simulate from the theoretical model using parameters calibrated to match empirical dynamics, and assess production and consumption estimator accuracy across sample sizes. Appendix G provides full details.

6.3.1 Comparing Poverty Trap Detection

I compare my dynamic panel framework with the nonparametric method of Balboni et al. (2022), which estimates asset transitions separately for treatment and control groups by regressing future on current assets. Poverty traps are identified by intersections with the 45-degree line. I simulate households with heterogeneous fixed productivity, η_{ij}, under two conditions: (i) no fixed costs and no traps, and (ii) fixed costs that generate traps for medium-ability households. Production and utility parameters remain constant; only initial conditions, shocks, and treatment status vary. I assess both methods' accuracy across 100 simulations.

Results (Appendix Tables A.8 and A.9) show that my method outperforms the direct regression approach. In the absence of poverty traps, it correctly identifies a unique steady state in 97% of simulations, while the regression approach does so in only 18%,

often detecting spurious traps due to inability to account for unobserved ability. By averaging over abilities with different transition slopes, the regression method produces misleading S-shaped functions.

When traps exist for medium-ability households, my estimator identifies the correct number of steady states in 86% of simulations, compared to 41% using the nonparametric method. The latter again over-detects traps due to ability pooling. By recovering ability-specific production and consumption dynamics, my framework distinguishes genuine traps from those caused by unobserved heterogeneity. Errors arise only near regime boundaries, where households transition across ability thresholds. These results underscore the importance of accounting for unobserved heterogeneity when identifying poverty traps.

6.3.2 Production, Consumption and Asset Transition Simulations

The simulations measuring production, consumption, and the joint asset transition estimator performance show that both components of the estimation strategy perform well. For production, the quasi-Bayes estimator accurately recovers the underlying function across all sample sizes, with errors declining sharply as the number of households increases. Credible intervals also achieve high coverage, even in small samples (Appendix Table G.29).

For consumption, the stochastic EM algorithm consistently outperforms the two-step fixed effects approach in recovering both the consumption function and latent productivity (Appendix Tables G.30 and G.31). Coverage rates are somewhat lower in the calibrated simulations because I fix the number of particles for computational reasons. However, the non-calibrated simulations show that increasing the number of particles restores coverage to near-nominal levels, and the empirical application uses greater computational resources, such as more stochastic EM iterations, to mitigate this issue.

Finally, the joint asset transition function, combining production and consumption estimates, is also recovered with reasonable accuracy (Appendix Table G.32). Errors decline with sample size, and while the calibrated simulations show some imprecision in small samples, this is consistent with the added complexity of estimating two nonlinear components jointly.

6.3.3 Summary

Overall, the Monte-Carlo simulations show that my estimation strategy reliably recovers the existence of poverty traps compared to existing methods and that as sample sizes increase, both the production, consumption, and asset transition functions are estimated with greater accuracy. The quasi-Bayesian production estimator performs especially well in small samples. While consumption is harder to estimate due to latent heterogeneity,

the stochastic EM algorithm consistently offers the best performance among available methods considered. Together, both components of the estimation framework provide increasingly precise estimates of the key structural functions governing household dynamics as the number of households increases.

7 Empirical Results

This section presents the core empirical findings. I first describe the estimated average structural functions and their implied asset dynamics, then quantify the prevalence of poverty trap thresholds, and finally explore fixed costs as a potential mechanism.

7.1 Average Structural Functions

The estimated functions depend on assets, latent ability, and idiosyncratic shocks. To summarize the results, I present average structural functions, which integrate over the distribution of ability and shocks. While these averages provide a useful summary, they may not reflect the dynamics faced by any individual household. These average functions are calculated as:

$$y_j(k) = E_{\eta,\omega}[f_j(k,\eta,\omega)] \tag{30}$$

$$c_j(k) = E_{\eta,\omega}[g_j(k,\eta,\omega)] \tag{31}$$

$$\dot{k}_j(k) = E_{\eta,\omega}[f_j(k,\eta,\omega) - g_j(k,\eta,\omega) - \delta_j] \tag{32}$$

where expectations are taken over the joint distribution of latent ability and shocks generated by the particle Markov chain Monte Carlo algorithm. To aid visualization, the main text plots standardized production and consumption functions for six highlighted studies.[20] I present results in the original units and for each study separately in Appendix Figures B.2-B.4.[21]

7.1.1 Production Function

The estimated production functions suggest that non-convexities—a key theoretical driver of poverty traps—are empirically widespread. The sites split into two groups (Figure 4): the first exhibits classical concave functions (Figure 4a), while the second shows that 13 studies (48%) feature increasing marginal returns over some asset range (Figure 4b).

[20] Specifically, $s(x) = (x-\mu_x)/\sigma_x$, where μ_x and σ_x are the mean and standard deviation of the average function evaluated over a uniform prediction grid. The prediction grid covers the support of the data for each site from the 1st to 99th percentile of assets.

[21] Fortuitously, the asset transition results are all approximately on the same scale, so I present these in levels without standardization.

This second group is characterized by households at very low asset levels (<$500) facing nearly flat production functions before output increases sharply. This pattern suggests fixed costs that households must overcome before engaging in productive activities, a mechanism I explore in Section 7.3.

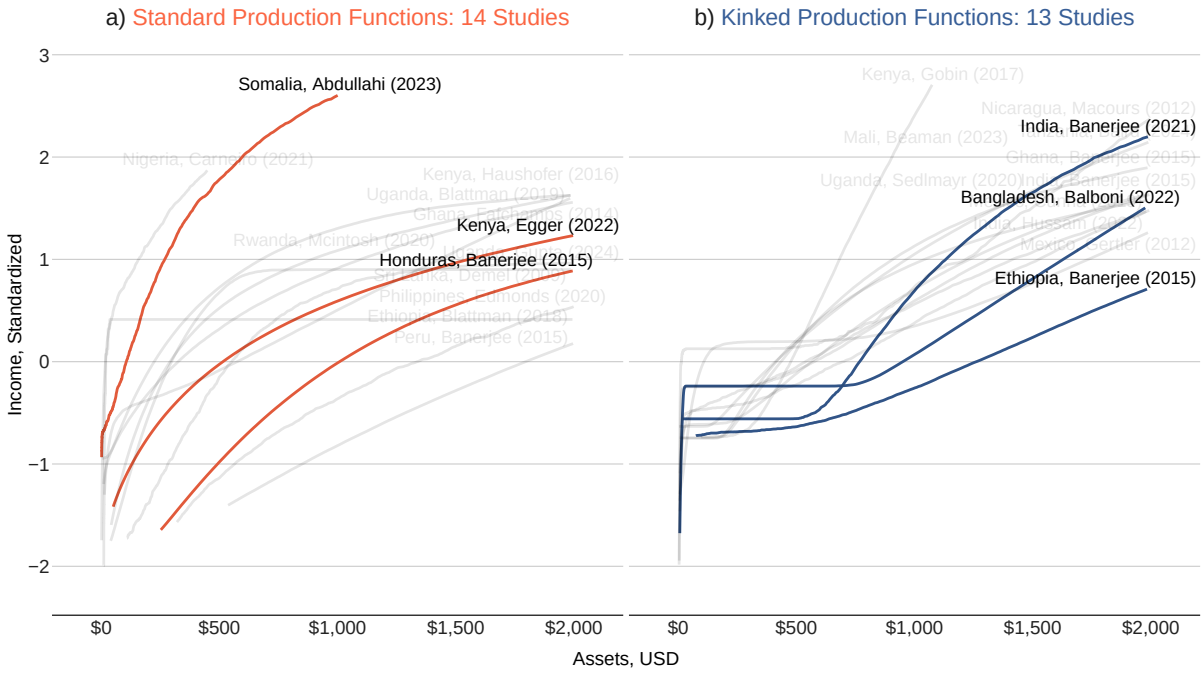


Figure 4: Household Production Function Estimates

Notes: Each panel shows estimated production functions for each study site. Estimates are posterior medians from the quasi-Bayesian procedure described in Section 6. Each line shows the predicted structural function in rescaled levels, $s(\exp(\hat{f}_j(k)))$, with η and ω set to zero. To aid visualization, the rescaling function, $s(\cdot)$, standardizes the predicted functions to have a mean of zero and a standard deviation of one across the support of the prediction grid.

7.1.2 Consumption Function

Consistent with households attempting to overcome production non-convexities, the estimated consumption functions are often non-monotonic. In 16 studies (64%), households appear to temporarily reduce consumption to accumulate assets–behavior predicted by the production non-convexities identified above. Production indivisibilities are a strong predictor of this pattern: when present, there is a 69% chance of observing non-monotonic consumption, compared to only an even chance when production functions are concave (Appendix Table A.12).[22]

As with the production results, the estimated consumption functions fall into two groups (Figure 5). In nine studies, consumption rises monotonically with assets (Figure 5a). In the remaining 16, it first increases at very low asset levels, then dips around

[22]Whilst I estimate production functions for 27 studies, two lack consumption data, and I exclude the flat production function in Blattman and Dercon (2018), where consumption predictions are unclear.

$500–$1,000 before rising again (Figure 5b)–a pattern consistent with households cutting consumption to save past a fixed-cost threshold.

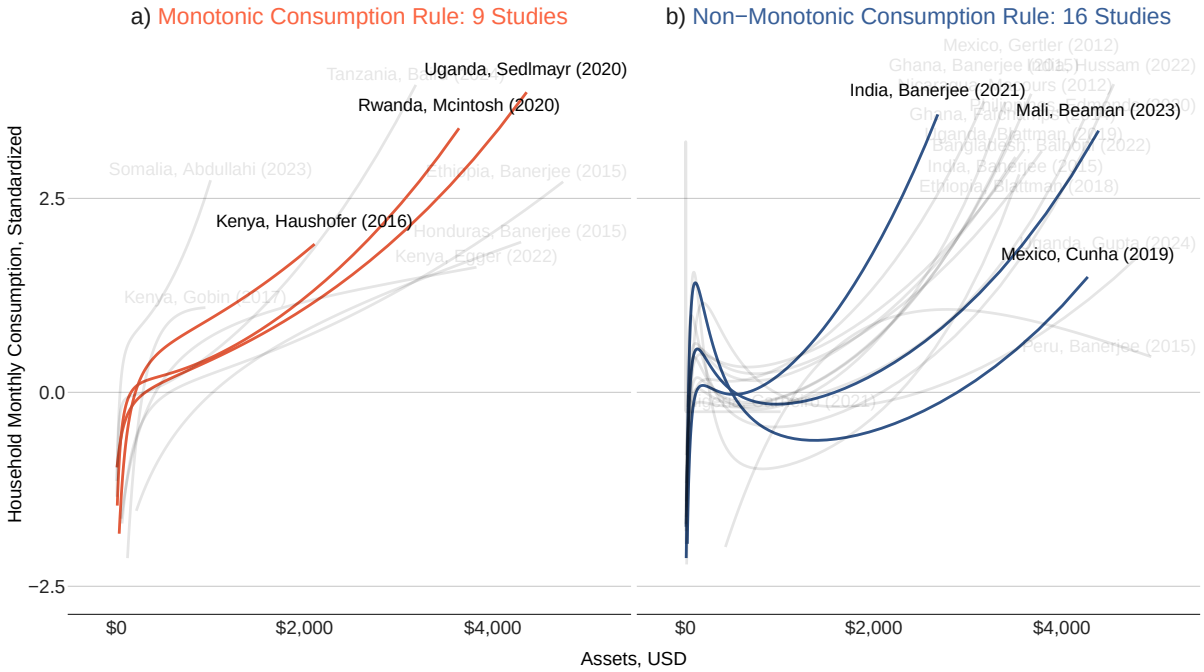


Figure 5: Average Consumption Function Estimates

Notes: Each panel shows estimated consumption functions for each study site. Estimates are from the stochastic EM procedure described in Section 6. Each line shows the predicted average structural function in rescaled levels, $s\left(E_{\eta,\omega}\left[\hat{g}_j(k,\omega,\eta)\right]\right)$, where expectations over η, ω are calculated by sampling from the posterior generated by the particle Markov chain Monte Carlo method. To aid visualization, the rescaling function, $s(\cdot)$, standardizes the predicted functions to have a mean of zero and a standard deviation of one across the support of the prediction grid.

7.1.3 Asset Transition Equation

Combining the production and consumption estimates shows that average asset dynamics in many sites feature multiple steady states. The estimated transition functions in Figure 6 illustrate the contrast between neoclassical convergence and poverty-trap dynamics. In the first group of studies (Figure 6a), households with low assets accumulate while those with high assets decumulate, leading to convergence toward a single steady state. In the second group (Figure 6b), the change in assets is positive at very low levels, negative at intermediate levels, and positive again beyond a threshold—implying multiple stable steady states and a classic poverty trap.

Overall, 15 of the 25 studies (60%) exhibit multiple steady states in their average asset dynamics—a large share given the mixed evidence for poverty traps in the literature. One concern is that these patterns might reflect averaging over heterogeneous households rather than true multiple equilibria. I address this next by estimating poverty-trap thresholds at the household level.

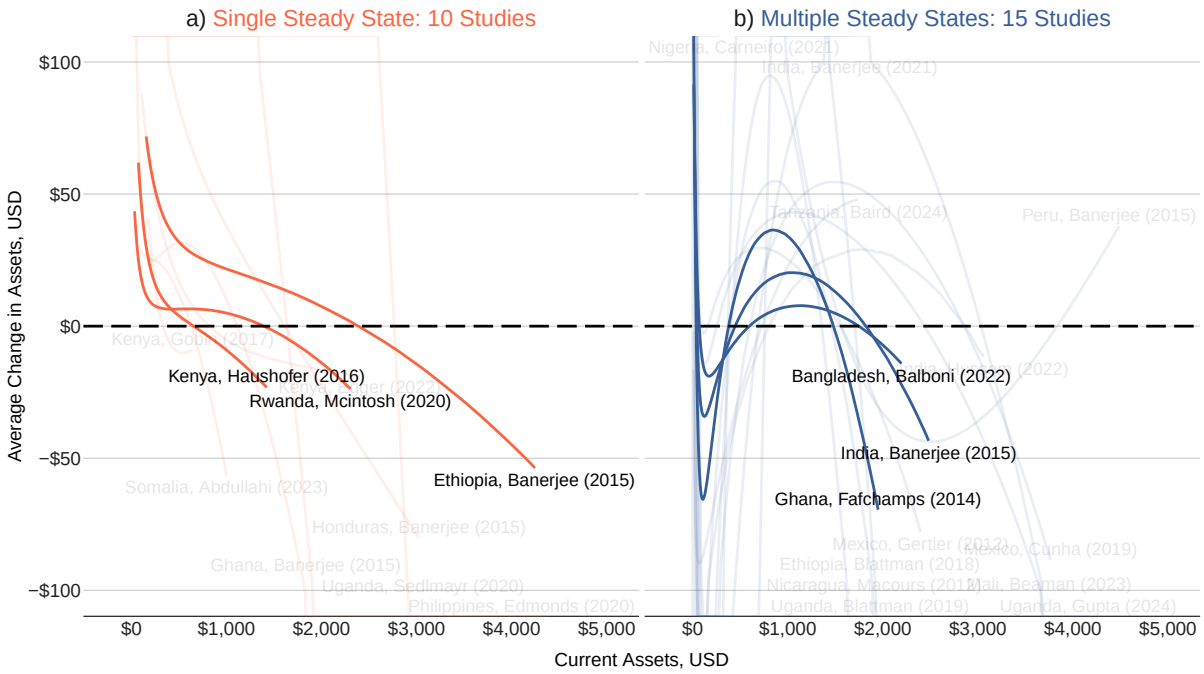


Figure 6: Average Asset Transition Function Estimates

Notes: This figure shows estimated average asset transition functions for each study site. Estimates are drawn from combining the estimated production and consumption functions, alongside the estimated depreciation function, as described in Section 6. Each line shows the predicted average structural function in levels.

7.2 Poverty Traps

While the average dynamics suggest widespread poverty traps, incorporating household specific dynamics reveal a more nuanced picture.

7.2.1 Thresholds

To quantify how many households face a poverty trap, I characterize the steady states and thresholds implied by each household's estimated asset transition equation. Because both asset accumulation and steady state locations depend on household ability, I define both these objects at the ability-specific (i.e. household) level and then aggregate.

I begin by defining the expected number of steady states, $C_j(\eta)$, for a household with ability η in study j:

$$C_j(\eta) = E_{\omega|\eta}\left[\#\{k : \dot{k}(k, \eta, \omega) = 0\} \mid J = j\right] \tag{33}$$

This function counts the number of points at which the household's asset dynamics satisfy $\dot{k} = 0$ – that is, the steady states where assets neither fall nor rise. Productivity shocks influence this equation, so the number of solutions to $\dot{k}(k, \eta, \omega) = 0$ may vary with ω; therefore, I average this count over the conditional distribution of shocks to obtain the

expected number of steady states for a household with ability η.[23]

Next, I define the expected unstable threshold, $T_j(\eta)$, associated with these steady states:

$$T_j(\eta) = E_{\omega|\eta}\left[\inf\{k : \dot{k}(k,\eta,\omega) = 0 \text{ and } \partial_k \dot{k}(k,\eta,\omega) > 0\} \mid J = j\right] \tag{34}$$

Among the set of steady states, the threshold is the *unstable* one: the point at which a small increase in assets causes the household to accumulate, while a small decrease sends it back towards a low-asset steady state. Mathematically, this corresponds to the steady state with a positive slope, $\partial_k \dot{k}(k,\eta,\omega) > 0$. I take the infimum to ensure that when multiple unstable steady states exist, I take the lowest one. As before, I average over productivity shocks to obtain an ability-specific expected threshold.

Finally, I aggregate these ability-specific thresholds to obtain a study-level threshold:

$$T_j = E_\eta[T_j(\eta) \mid J = j] \tag{35}$$

This expectation averages thresholds across the empirical distribution of household abilities in study j. The scalar T_j summarizes, for the study as a whole, where the poverty trap lies for households that face one.

The study-level thresholds, T_j, are reported in Column 1 of Table 1, with a median threshold of \$547. Column 2 shows that households lie, on average, \$425 below this level, implying that crossing the threshold would require roughly ten months of foregone consumption (Column 3). Most thresholds cluster between \$260 and \$1,000, with a few outliers such as Peru near \$4,000. These values align closely with prior work: Balboni et al. (2022) estimate a threshold of \$504 in Bangladesh-the same study that yields the median threshold in my sample, differing by only \$43.

7.2.2 Traps

Having defined steady states and thresholds, I now classify households as *trapped* based on their current asset level and the number of steady states they face. Intuitively, I define a household as trapped only if (i) it faces multiple steady states, and (ii) the household lies below the unstable threshold. The trap indicator for a household with assets k and

[23] These expectations are 'reduced-form' in the sense that I do not iterate the household policy rule forward to check for convergence. Instead, I average the estimated structural functions as-is. In Section 8.1, I use the calibrated model to simulate household paths and compute the long-run probability of being trapped.

Table 1: Poverty Trap Threshold Results

Program	Threshold (USD PPP) (1)	Distance Below Threshold (USD PPP) (2)	$\frac{\text{Threshold}}{\text{Consumption}}$ (3)	Transfer (USD PPP) (4)	HHs Below Threshold (%) (5)	HHs Trapped Below Threshold (%) (6)
Mali, Beaman (2023)	263	214	12	140	24.75	24.75
Ethiopia, Blattman (2018)	298	287	6	1,030	26.92	25.00
Mexico, Gertler (2012)	306	250	12	379	59.25	35.75
Ghana, Fafchamps (2014)	343	218	6	120	59.25	19.75
Nicaragua, Macours (2012)	387	341	7	485	76.50	69.00
Uganda, Blattman (2019)	395	316	2	400	68.00	67.00
India, Banerjee (2015)	409	337	9	788	62.50	45.50
India, Hussam (2022)	533	367	11	100	71.48	34.90
Bangladesh, Balboni (2022)	547	425	11	542	54.00	14.75
Mexico, Cunha (2019)	621	532	9	21	75.68	75.68
Tanzania, Baird (2024)	633	515	42	350	76.25	31.00
Uganda, Gupta (2024)	721	639	8	2,967	50.00	50.00
Kenya, Gobin (2017)	815	648	24	238	98.75	36.00
Uganda, Sedlmayr (2020)	828	570	8	155	63.00	0.25
India, Banerjee (2021)	997	865	73	547	86.00	31.00
Kenya, Egger (2022)	2,323	1,674	47	1,871	95.25	1.25
Peru, Banerjee (2015)	3,906	1,996	34	1,433	78.25	76.25
Somalia, Abdullahi (2023)	-	-	-	500	-	-
Ethiopia, Banerjee (2015)	-	-	-	338	-	-
Ghana, Banerjee (2015)	-	-	-	406	-	-
Honduras, Banerjee (2015)	-	-	-	754	-	-
Nigeria, Carneiro (2021)	-	-	-	528	-	-
Philippines, Edmonds (2020)	-	-	-	518	-	-
Kenya, Haushofer (2016)	-	-	-	404	-	-
Rwanda, Mcintosh (2020)	-	-	-	328	-	-
Overall	547	425	10	338	59.25	25.00

Notes: This table reports poverty trap statistics, estimated using the structural functions for each study site. The 'Threshold' column reports the average asset threshold a household must cross to move from a low- to high-steady state in 2020 USD, conditional on households facing a threshold. The 'Distance Below Threshold' column reports the average amount of assets a household would need to accumulate to reach the threshold, conditional on having less assets than the threshold. The 'Transfer' column shows the median cash/asset transfer administered to households in each study. 'HHs Below Threshold' describes the proportion of households with assets below this threshold at baseline. The final column, 'HHs Trapped Below Threshold', shows the proportion of households below the threshold that are 'trapped' – that is, the proportion of households that are productive enough such that if they were given a transfer to cross the threshold, they would continue to accumulate assets, but are not so productive that they are predicted to cross this threshold without a transfer. All monetary values are in 2020 USD. Studies where a threshold does not exist for any household are shown with dashes.

ability η in study j is given by:

$$\text{Trapped}_j(k,\eta) = \mathbf{1}\left\{ \overbrace{C_j(\eta) > 1}^{\text{Multiple steady states}} \quad \text{and} \quad \overbrace{k < T(\eta)}^{\text{Below threshold}} \right\} \tag{36}$$

This indicator equals one when the household faces multiple steady states and has assets below the threshold. In these cases, the estimated dynamics predict convergence to the low-asset steady state unless the household receives a sufficiently large shock or transfer. This definition reflects the idea that, counterfactually, the household could escape the trap if it started above the threshold. By contrast, households with only a single, low, steady state are not classified as trapped, even though they may remain persistently poor.

I then compute the study-level proportion of trapped households by averaging this indicator over the empirical joint distribution of assets and abilities:

$$P_{\text{Trapped},j} = E_{k,\eta}\left[\text{Trapped}_j(k,\eta) \mid J = j\right] \tag{37}$$

This summarizes, for each study, how many households are trapped – not just whether traps exist on average, but what share of households are predicted to fall into them given their observed asset levels and abilities.

The probability of being trapped can be contrasted with the probability of simply being below the study-average threshold, $P_{\text{Below Threshold},j}$:

$$P_{\text{Below Threshold},j} = E_{k,\eta}\left[\mathbf{1}\{k < T_j(\eta)\} \mid J = j\right] \tag{38}$$

Distinguishing between these two probabilities is important. Many households may lie below the study-level threshold because of low productivity or adverse shocks, but this does not necessarily imply that they are trapped. Only by estimating the household-specific asset transition equation from the dynamic model can we determine whether a household faces multiple steady states and would start accumulating assets after a sufficient transfer, thereby separating those who are merely below the threshold from those who are genuinely trapped.

Figure 7 and Columns 5 and 6 of Table 1 show these proportions. Many studies have a large share of households below the threshold (59% on average), but only about 25% are truly trapped. This gap reflects two key sources of heterogeneity. First, some households are sufficiently productive that their individual asset transition functions exhibit only a single high steady state, while others are so unproductive that even owning assets greater than the study-level average threshold does not ensure escape – these households only have a single, low steady state. Second, optimal consumption responses to production non-

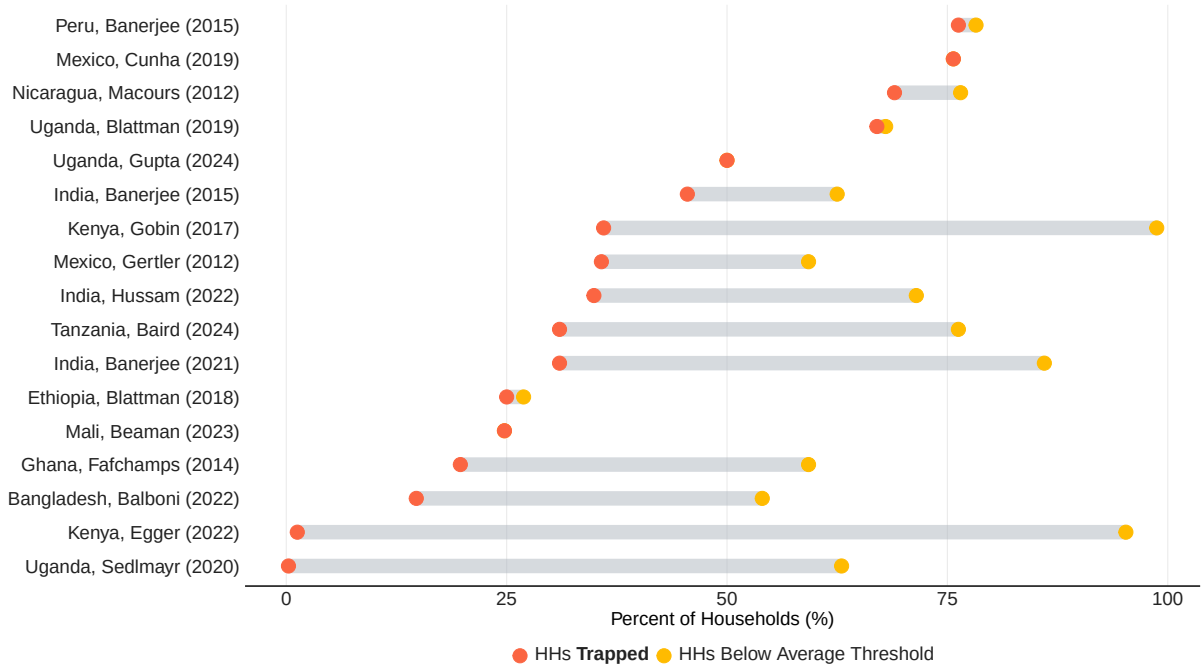


Figure 7: Proportion of Households Below Trap Thresholds and Proportion Trapped

Notes: This figure shows the proportion of households with assets below the study-average threshold, $P_{\text{Below Threshold},j}$, in yellow, and the proportion of households classified as 'trapped', $P_{\text{Trapped},j}$, in orange. Studies where no threshold exists for any household are omitted.

convexities often mitigate their impact on asset dynamics: households can temporarily reduce consumption, accumulate assets, and escape the trap region.

Together, these findings suggest that while poverty-trap dynamics are a feature of many environments, they do not determine the long-run fate of most households below the threshold.

7.3 Evidence For Fixed Costs

One theoretical mechanism underlying poverty traps is the presence of fixed costs required to engage in high-return production. The estimated production functions, which exhibit flat marginal returns at low asset levels, are consistent with this mechanism. I next provide additional suggestive evidence that indivisible, high-value assets do indeed generate such fixed costs.

To test this, I plot the probability that a household owns a livestock asset (chicken, goat, or cow) against its asset level relative to the household's estimated poverty-trap threshold, focusing on households for whom a threshold exists.[24] Figure 8 shows no relationship between asset levels and chicken ownership, a small and inexpensive asset. In contrast, there is a 23 percentage-point jump in goat ownership and a 41.9 point jump in cow ownership above the threshold. This pattern is consistent with fixed costs:

[24] I exclude a $250 "donut" around each household's threshold to account for estimation uncertainty.

households appear to acquire large, indivisible assets as they cross the threshold, enabling participation in higher-return production. Appendix Figures B.5 and B.6 confirm that these discontinuities are unlikely to arise by chance, using randomization-inference and placebo thresholds I find the probability of observing a jump at least as large as those observed to be below 0.1% for goats and cows, but 86% and 69.7% for chickens.

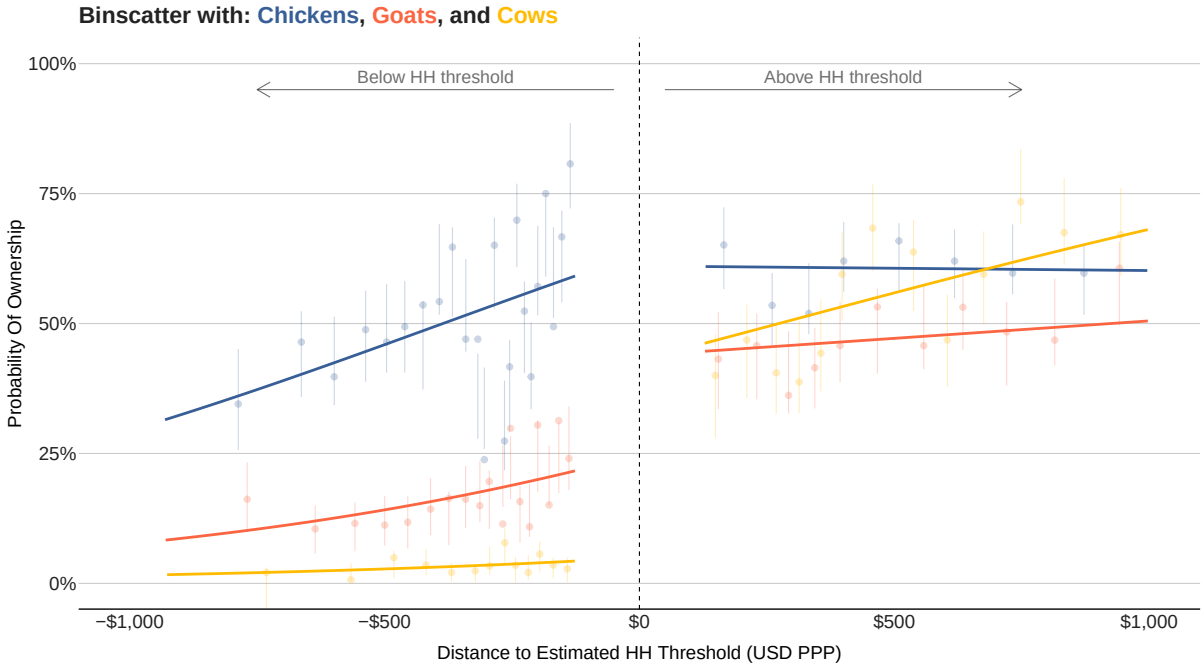


Figure 8: The Probability of Livestock Ownership Around the Poverty Trap Threshold

Notes: This figure shows the probability of livestock ownership (chickens, goats, and cows) against assets relative to the household poverty trap threshold, excluding a $250 donut around the threshold to account for estimation uncertainty. Each point is generated via Cattaneo et al. (2024)'s binscatter method and the lines show linear fits on either side of the threshold. I exclude households more than $2,000 from the threshold.

In Appendix Figure B.8, I return to two original studies–Balboni et al. (2022) in Bangladesh and Haushofer and Shapiro (2016) in Kenya–which explicitly describe households saving to acquire large, indivisible assets. In both cases, nonparametric regressions of consumption on assets show households reducing consumption as they approach the estimated threshold, which coincides with the price of the fixed-cost asset: a cow in Balboni et al. (2022) and a metal roof in Haushofer and Shapiro (2016).[25] This provides further evidence that fixed costs are a key mechanism generating poverty traps and that the estimated thresholds capture a salient feature of household behavior.

[25]Using the dynamic panel model, I do not find any households with multiple steady states in Haushofer and Shapiro (2016) using my preferred specification. The structural model, discussed in the next section, does predict a small fraction of households with multiple steady states and it is these estimates I use to form the average threshold in Figure B.8.

7.4 Additional Results: Correlates, Ability, and Persistence

Appendix Section C presents additional analyses on the correlates of poverty traps, dispersion in latent ability, and productivity persistence. Larger households are more likely to face multiple steady states, and among those with a threshold, households below it tend to be younger, have a higher share of women, and possess lower levels of education. I find little evidence that the size of households' thresholds correlates with observable characteristics, and at the study level, few characteristics predict the presence of poverty traps.

Turning to latent ability and productivity persistence, I estimate moderate dispersion in household ability: a 0.58 log-point difference between the 10th and 90th percentiles. This dispersion is smaller than previous estimates (1.25–1.9) in Gollin and Udry (2021) and Bils et al. (2021), possibly because accounting for production non-convexities prevents overstatement of ability differences. Estimated productivity persistence is comparable to prior work, with an average annualized persistence of 0.299—well within the 0.2–0.4 range reported by Maue et al. (2020) for agricultural settings in Africa.

7.5 Summary

Taken together, these results show that poverty-trap dynamics are common in the aggregate but do not determine the fate of all households. Production non-convexities and endogenous consumption responses generate multiple steady states in many settings, yet only a small minority of households reside in the productivity–asset region where such dynamics bind. This helps explain why empirical evidence for poverty traps has been elusive: the mechanisms exist, but affect relatively few households.

Appendix D demonstrates that these findings are robust to a range of alternative estimation choices across both the production and consumption steps. While study-specific conclusions vary slightly, the broader patterns remain: widespread production non-convexities, non-monotonic consumption functions, and multiple steady states in average asset dynamics, but few households actually trapped. Across all robustness checks, the median share of trapped households ranges from 10% to 26% (Appendix Tables (D.24)-(D.27)).

The next section builds on this heterogeneity to move from descriptive evidence to counterfactual analysis, quantifying how alternative anti-poverty programs affect household growth and what optimal policy would look like in these environments.

8 A Calibrated Growth Model and Optimal Policy

In this section, I solve and calibrate household growth models for each study. The models generate structural consumption policy rules and asset dynamics calibrated against the

estimated semi-structural production and consumption functions in Section 7. The advantage of this structural approach is that it allows me to simulate household responses to counterfactual transfer policies, such as changes in the future income stream or stochastic environment that semi-structural models cannot accommodate.

First, I use this model to validate the estimated thresholds and trap dynamics of the semi-structural model. Next, I solve for the social planner's optimal transfer policy. Finally, I discuss implications on the targeting and design of common cash transfer programs.

8.1 Household Growth Model

8.1.1 Model and Model Solution

The household growth model is a special case of the model described previously in Equations (1)-(3), where the general utility function and productivity process are specialized respectively to a CRRA utility function and the Ornstein-Uhlenbeck (OU) process, a continuous time analogue of the AR(1) process[26]:

$$V(k_{i0}, \eta_i, \omega_{i0}) = \max_{c_{it}} E_0 \int_0^\infty e^{-\rho t} \frac{c_{it}^{1-\gamma} - 1}{1-\gamma} dt \tag{39}$$

$$dk_{it} = f(k_{it}; \omega_{it}, \eta_i) - c_{it} - \delta k_{it} \quad dt \tag{40}$$

$$d\omega_{it} = -\theta \omega_{it} dt + \sigma dW_{it} \tag{41}$$

The model's Hamilton-Jacobi-Bellman (HJB) equation is given by:

$$\rho V(k, \eta, \omega) = \max_c \left\{ \frac{c^{1-\gamma} - 1}{1-\gamma} + V_k \cdot [f(k; \omega, \eta) - c - \delta k] + V_\omega \cdot (-\theta \omega) + \frac{1}{2} V_{\omega\omega} \cdot \sigma^2 \right\} \tag{42}$$

However, the non-convexities in the estimated production functions $f(\cdot)$ mean that the value function is no longer strictly concave and has a convex kink, therefore I use a finite-difference upwind scheme described in Achdou et al. (2022) to solve the model with an irregular grid chosen to ensure sufficient coverage near the estimated non-convexities.[27]

To calibrate the model, I use the estimated production functions from Section 7, along with estimates of the productivity process, which leaves only two free parameters to estimate per study: the discount rate, ρ, and the coefficient of relative risk aversion, γ. I estimate these two parameters using a minimum distance estimator. I choose $\{\rho, \gamma\}$ to minimize the distance between the estimated consumption policy function and asset

[26] To maintain consistency with standard notation for continuous time models, I use ρ for the discount rate, γ for the coefficient of relative risk aversion, and θ for the mean reversion parameter in the productivity process. Previously, ρ denoted the AR(1) persistence parameter in the statistical model.

[27] There could potentially be multiple convex kinks, however in practice I typically only find 'one' non-convexity in the production technology.

transition equation from the dynamic panel model to the structural model's predicted consumption policy function and asset transition equation. I discuss further implementation details, including model calibration and counterfactual simulation, in Appendix H.

8.2 Traps and Transfers

First, I use the calibrated household growth model to validate the estimated poverty trap thresholds and dynamics of the semi-structural model. Appendix Table A.11 shows similar evidence of poverty traps to the semi-structural model, with evidence of traps 'on aggregate' in many studies and only a small proportion of households 'trapped' below the threshold. The definition of 'trapped' is slightly different in the structural model, since it uses the estimated Markov transition matrix to calculate the probability a household with a given set of initial conditions (k_0, ω_0, η_0) will remain below the threshold in the future – I define a household as trapped if there's at least a 50% chance they will remain below the threshold without a big-push transfer, but there exists a transfer, less than $1,000, that will move them above the threshold such that they have at least a 50% chance of remaining above the threshold in two years time.

8.3 Optimal Transfer Policy

A key question policymakers face is how to optimally allocate scarce resources to alleviate poverty. In the presence of poverty traps, a tension arises between "targeting impact versus deprivation" (Haushofer et al. 2025) – that is, should transfers be targeted at the households most likely to show the largest impact, which in a poverty trap setting are those relatively *richer*, higher productivity households on the cusp of crossing the poverty trap threshold, or should transfers target the very poorest households, which have the highest marginal utility of an additional dollar but lowest productivity?

8.3.1 Transfer Policy Optimization

To answer this question, I use the calibrated value functions to solve for the planner's optimal transfer policy, which maximizes the discounted expected utility of all households in a study, subject to a budget constraint, B:

$$T^* = \arg\max_T \int V(k + T(k, \eta), \eta, \bar{\omega}) dF_{k,\eta} \tag{43}$$

$$\text{s.t. } B_\eta = \int T(k, \eta, \bar{\omega}) dF_{k,\eta} \tag{44}$$

To maintain computational tractability and interpretability, I focus on solving the optimal transfer policy by initializing households' idiosyncratic shock to the study-specific

median, $\bar{\omega}$, uniformly distributing households across the empirical distribution of initial assets and fixing the planner's budget, B_η, to a fixed amount per productivity type, η. By fixing the budget per type I can solve the planner's problem separately for each productivity type, significantly reducing the dimensionality of the optimization problem.

This leads to the Lagrangian and first-order condition:

$$\mathcal{L}(T, \lambda, \eta) = \sum_{i}^{N} \mu_i V(k_i + T_i; \eta, \bar{\omega}) + \lambda \left(B_\eta - \sum_{i}^{N} \mu_i T_i \right) \tag{45}$$

$$\frac{\partial \mathcal{L}}{\partial T_i} = \mu_i V'(k_i + T_i; \eta, \bar{\omega}) - \lambda \mu_i = 0 \tag{46}$$

Where μ_i is the weight vector over the initial asset distribution, T_i the transfer to households at position i on the discretized state grid. For studies with fixed costs in the production function, the non-convexities generate a kink in the value function, $V(\cdot)$, which complicates the planner's optimization problem – there are potentially multiple T_is that satisfy the first-order condition, Equation 46.

Intuitively, this corresponds to the planner deciding between a "big-push" transfer, T_+, that equates costs (shadow price λ) and benefits on the upper region of the value function, V'_+, or a smaller transfer, T_-, that equates costs and benefits on the lower region of the value function, V'_-. Rather than enumerate all possible combinations of $\{T_{i,+}, T_{i,-}\}$ for each household i, I restrict the class of transfer policies the planner can choose such that the planner can only choose an interval of assets, $[k_-, k_+]$, where the planner picks the higher transfer, $T_{i,+}$. I discuss further details of the optimization in Appendix I.

8.3.2 Transfer Policy Results

Figure 9 illustrates the optimal transfer policies in two studies, based on a $500-per-person budget across three productivity types: low, medium, and high. In Mali Beaman et al. (2023), transfers concentrate on the poorest households and decline as assets increase. For high-productivity households, transfers begin slightly higher at low asset levels but phase out earlier; otherwise, differences across types are minor. In contrast, the Bangladesh policy Balboni et al. (2022) is more complex. Low- and high-productivity households follow similar declining schedules, while medium-productivity households face a non-monotonic path: a small initial transfer, a sharp increase near $100 in assets, then a decline. This design reflects a planner's trade-off between targeting deprivation and maximizing impact—allocating small transfers to those with the highest marginal utility and larger transfers to those with higher marginal returns.

On average, the optimal transfer policy decreases with assets: the poorest households receive the largest transfers. Appendix Table A.20 reports the share of policies that are monotonically decreasing by ability decile. Across all budgets, over 95% of optimal

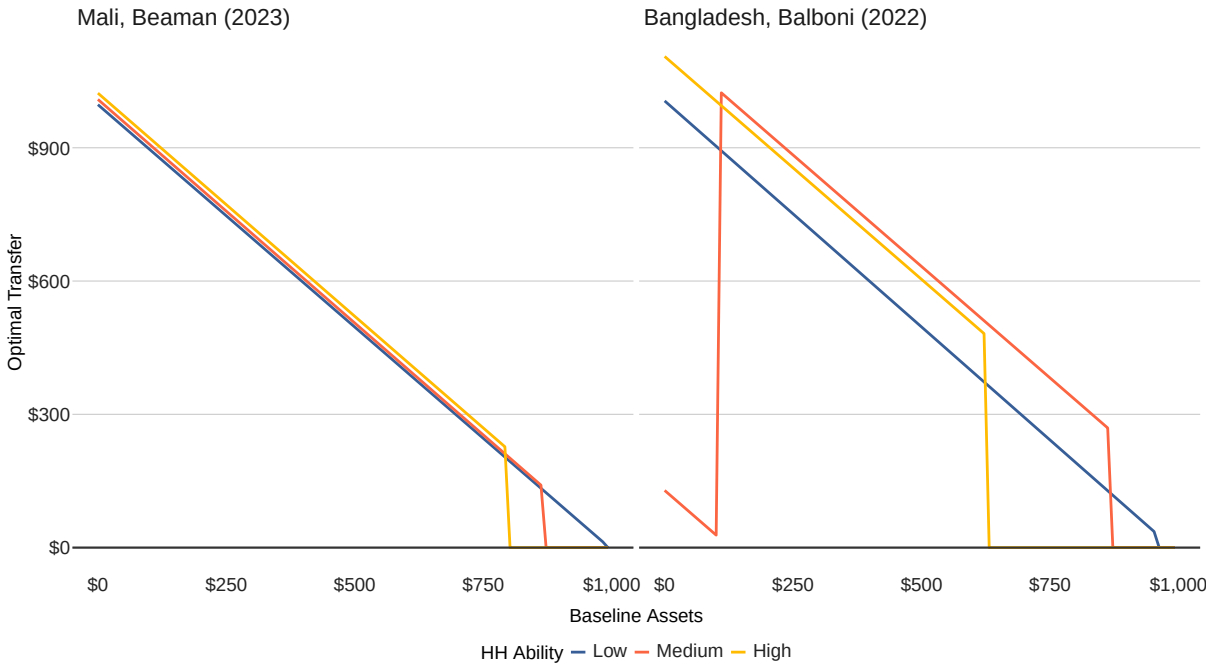


Figure 9: Optimal Transfer Policies in Mali and Bangladesh

Notes: This figure shows the optimal transfer policies in two studies. Panel (a) shows the optimal transfer policy in Mali, Beaman et al. (2023), while panel (b) shows the optimal transfer policy in Bangladesh, Balboni et al. (2022). The blue lines represent the optimal transfer policy using a total budget per head of $25, and the yellow lines represent the optimal transfer policy using a total budget per head of $500.

policies are monotonic. As predicted by the model, policies are most often monotonic for low- and high-productivity households, while those in the middle are slightly more likely to prefer non-monotonic transfers.

8.3.3 Optimal Policy Targeting and Welfare Effects

Next, I compare the welfare effects of the optimal transfer policy—which maximizes the present discounted value of household consumption, accounting for future asset dynamics—with two simpler targeting rules. The first is a flat, untargeted transfer in which all households receive the same amount. The second targets current consumption, maximizing immediate utility while ignoring future asset accumulation.

Appendix Table A.21 reports the percentage increase in welfare, measured as the gain in discounted consumption relative to no transfers. On average, with a budget of \$100 per person, the optimal policy yields an 23.2% welfare gain. By comparison, the flat transfer delivers a 14.7% gain—about 37% less than the optimized policy. Surprisingly, the current-consumption targeting rule achieves a nearly equivalent welfare improvement of 18%, despite ignoring long-run asset dynamics.

Though initially counterintuitive, this result arises from two key factors. In a world with fixed costs, high consumption can signal two distinct household types. First, it may reflect high-productivity households that accumulate enough assets to exceed the fixed cost threshold. These households will reach a high-asset steady state regardless and thus have little need for transfers. Second, it may signal low-productivity households that consume heavily because they see little chance of escaping a poverty trap and therefore have little incentive to save.

In both cases, high consumption implies a low marginal return on transfers—either because the household is already on a path to self-sufficiency or because it is unlikely to become self-sufficient. Conversely, low consumption suggests a higher marginal return: the household either has high marginal utility of consumption or is productively investing in its future, though not yet able to escape poverty. As a result, targeting current consumption—despite ignoring future asset dynamics—can serve as an effective heuristic.

Appendix Table A.22 supports this. In a sample of 100 households, the average rank difference between the optimal and consumption-targeting policies is just 6.6 places at a higher budget (\$1000 per head), widening modestly to 12.2 places at a smaller budget (\$100 per head).

8.3.4 Optimal Program Design: Big Push Best, Safety-Net Second

Finally, I evaluate the welfare effects of three common program designs: a one-time "big-push" transfer, a monthly stream of smaller transfers, and a safety-net that guarantees a minimum level of income. Intuitively, big-push transfers may be most effective in the

presence of poverty traps, while safety nets could promote investment by reducing the risk of future deprivation.

Appendix Table A.23 summarizes the optimal program for each study across four budget levels: $10, $100, $500, and $1000. In most cases and at most budget levels, the one-time big-push transfer yields the greatest welfare gains. The safety-net is optimal only in a few cases, and the monthly transfer is never optimal.

9 Conclusion

At the turn of the millennium, one in three people lived in extreme poverty. Today, that number is closer to one in ten. The fight against global poverty is relentless, and while much progress has been made, much remains to be done. This paper provides evidence that poverty traps remain a barrier to prosperity in many contexts around the world. I find that non-convexities in production technology, plausibly driven by fixed costs to investment such as large livestock, appear more widespread than previously thought. These non-convexities create kinks in the production function that generate multiple steady states in the asset accumulation equation and therefore the possibility of poverty traps in many contexts.

However, a key takeaway of the paper is that the existence of poverty traps is not a binary condition, but rather a matter of degree. The evidence suggests that some households are trapped in poverty, initial conditions do matter, and that large, one-time transfers can have sustained effects on poverty alleviation. Many households are not trapped, but face such low returns to investment that large transfers are no guarantee of escaping poverty.

While these findings may be new to economists, it appears that poor households are already aware of these dynamics. I find evidence that households are forward-looking, and that they respond to indivisibilities in investment by smoothing consumption and saving in anticipation of approaching investment thresholds. This implies that households are aware of the potential for poverty traps, that they respond to them in a rational manner, and naïve comparisons of consumption in randomized cash transfer experiments may underestimate the welfare effects of such programs.

Overall, my findings suggest that there is no anti-poverty silver bullet. While large cash transfers, such as those provided by GiveDirectly, can help many households escape poverty traps, others may require more sustained interventions to boost returns to investment and human capital. One example of such an intervention are "Graduation" programs that combine large, one-time transfers with complementary support to boost human capital and consumption/savings behavior. My findings suggest that such programs are likely to be effective for two reasons. First, they provide a "big-push" that can shift trapped households to a higher-asset equilibrium. Second, the complementary

interventions can help boost returns to investment for households with single, low-asset steady states, plausibly increasing the effectiveness of the one-time transfer and reducing the likelihood of falling back into poverty.

This policy implication highlights an avenue for future work. In the current framework, household ability is treated as time-invariant, which prevents me from studying how interventions that build skills or productivity might interact with initial asset support. Many graduation programs combine multiple components, and extending the model to allow for evolving ability would enable researchers to quantify both the individual and complementary effects of such interventions.

Another limitation of the paper is the reliance on short-panels with randomized transfers that are typically limited to a few points of support. While internal lagged instruments provided by the panel structure help with identification, longer panels and rich variation in transfer sizes would help sharpen estimates of the nonparametric production and consumption functions. An even better approach might be to directly measure the risks and returns of large, indivisible assets such as livestock, obviating the need to rely on nonparametric production functions to conclude fixed costs generate production function non-concavities.

This paper follows a long theoretical and empirical tradition in development economics that seeks to identify the presence of poverty traps–one of the field's oldest and most enduring ideas. While the evidence presented here is suggestive of poverty traps in many contexts, much work remains to be done to better understand the mechanisms that generate them, and the policies that can best alleviate them. I would like to finish by expressing my gratitude to the researchers and field teams who designed and implemented the original experiments that made this study possible, and to the households who participated in them. Their willingness to share their experiences continues to expand our understanding of poverty and the potential paths out of it.

References

Yves Achdou, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *The Review of Economic Studies*, 89(1):45–86, January 2022. ISSN 0034-6527, 1467-937X. doi: 10.1093/restud/rdab002.

S. Rao Aiyagari. Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics*, 109(3):659–684, 1994. ISSN 0033-5533. doi: 10.2307/2118417.

Christophe Andrieu, Arnaud Doucet, and Roman Holenstein. Particle Markov Chain Monte Carlo Methods. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 72(3):269–342, June 2010. ISSN 1369-7412. doi: 10.1111/j.1467-9868.2009.00736.x.

Manuel Arellano and Stephen Bond. Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations. *The Review of Economic Studies*, 58(2):277–297, April 1991. ISSN 0034-6527. doi: 10.2307/2297968.

Manuel Arellano and Olympia Bover. Another look at the instrumental variable estimation of error-components models. *Journal of Econometrics*, 68(1):29–51, July 1995. ISSN 0304-4076. doi: 10.1016/0304-4076(94)01642-D.

Manuel Arellano, Richard Blundell, and Stéphane Bonhomme. Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework. *Econometrica*, 85(3):693–734, 2017. ISSN 1468-0262. doi: 10.3982/ECTA13795.

Manuel Arellano, Richard Blundell, Stéphane Bonhomme, and Jack Light. Heterogeneity of consumption responses to income shocks in the presence of nonlinear persistence. *Journal of Econometrics*, 240(2):105449, March 2024. ISSN 0304-4076. doi: 10.1016/j.jeconom.2023.04.001.

Costas Azariadis and John Stachurski. Chapter 5 Poverty Traps. In Philippe Aghion and Steven N. Durlauf, editors, *Handbook of Economic Growth*, volume 1, pages 295–384. Elsevier, January 2005. doi: 10.1016/S1574-0684(05)01005-1.

Sarah Baird, Craig McIntosh, Berk Özler, and Utz Pape. Asset transfers and anti-poverty programs: Experimental evidence from Tanzania. *Journal of Development Economics*, 166:103182, January 2024. ISSN 0304-3878. doi: 10.1016/j.jdeveco.2023.103182.

Clare Balboni, Oriana Bandiera, Robin Burgess, Maitreesh Ghatak, and Anton Heil. Why Do People Stay Poor? *The Quarterly Journal of Economics*, 137(2):785–844, April 2022. ISSN 0033-5533, 1531-4650. doi: 10.1093/qje/qjab045.

Oriana Bandiera, Nidhi Parekh, Barbara Petrongolo, and Michelle Rao. Men are from Mars, and Women Too: A Bayesian Meta-analysis of Overconfidence Experiments. *Economica*, 89(S1):S38–S70, 2022. ISSN 1468-0335. doi: 10.1111/ecca.12407.

Abhijit Banerjee, Esther Duflo, Nathanael Goldberg, Dean Karlan, Robert Osei, William Parienté, Jeremy Shapiro, Bram Thuysbaert, and Christopher Udry. A multifaceted program causes lasting progress for the very poor: Evidence from six countries. *Science*, 348(6236):1260799, May 2015. doi: 10.1126/science.1260799.

Abhijit Banerjee, Emily Breza, Esther Duflo, and Cynthia Kinnan. Can Microfinance Unlock a Poverty Trap for Some Entrepreneurs?, October 2019.

Abhijit V. Banerjee and Andrew F. Newman. Occupational Choice and the Process of Development. *Journal of Political Economy*, 101(2):274–298, 1993. ISSN 0022-3808.

Faisal Bari, Kashif Malik, Muhammad Meki, and Simon Quinn. Asset-Based Microfinance for Microenterprises: Evidence from Pakistan. *American Economic Review*, 114(2): 534–574, February 2024. ISSN 0002-8282. doi: 10.1257/aer.20210169.

Christopher B. Barrett, Michael R. Carter, and Jean-Paul Chavas, eds. *The Economics of Poverty Traps*. National Bureau of Economic Research Conference Report series., 2018.

Vittorio Bassi, Raffaela Muoio, Tommaso Porzio, Ritwika Sen, and Esau Tugume. Achieving Scale Collectively. *Econometrica*, 90(6):2937–2978, 2022. ISSN 1468-0262. doi: 10.3982/ECTA18773.

Lori Beaman, Dean Karlan, Bram Thuysbaert, and Christopher Udry. Selection Into Credit Markets: Evidence From Agriculture in Mali. *Econometrica*, 91(5):1595–1627, 2023. ISSN 1468-0262. doi: 10.3982/ECTA18916.

Mark Bils, Peter J. Klenow, and Cian Ruane. Misallocation or Mismeasurement? *Journal of Monetary Economics*, 124:S39–S56, November 2021. ISSN 0304-3932. doi: 10.1016/j.jmoneco.2021.09.004.

Christopher Blattman and Stefan Dercon. The Impacts of Industrial and Entrepreneurial Work on Income and Health: Experimental Evidence from Ethiopia. *American Economic Journal: Applied Economics*, 10(3):1–38, July 2018. ISSN 1945-7782. doi: 10.1257/app.20170173.

Christopher Blattman, Nathan Fiala, and Sebastian Martinez. The Long Term Impacts of Grants on Poverty: 9-Year Evidence from Uganda's Youth Opportunities Program, April 2019.

Hoyt Bleakley and Joseph Ferrie. Shocking Behavior: Random Wealth in Antebellum Georgia and Human Capital Across Generations *. *The Quarterly Journal of Economics*, 131(3):1455–1495, August 2016. ISSN 0033-5533. doi: 10.1093/qje/qjw014.

Richard Blundell and Stephen Bond. Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics*, 87(1):115–143, November 1998. ISSN 0304-4076. doi: 10.1016/S0304-4076(98)00009-8.

Richard Blundell and Stephen Bond. GMM Estimation with persistent panel data: An application to production functions. *Econometric Reviews*, 19(3):321–340, January 2000. ISSN 0747-4938. doi: 10.1080/07474930008800475.

Richard Blundell, Xiaohong Chen, and Dennis Kristensen. Semi-Nonparametric IV Estimation of Shape-Invariant Engel Curves. *Econometrica*, 75(6):1613–1669, 2007. ISSN 1468-0262. doi: 10.1111/j.1468-0262.2007.00808.x.

Richard Blundell, Luigi Pistaferri, and Ian Preston. Consumption Inequality and Partial Insurance. *American Economic Review*, 98(5):1887–1921, December 2008. ISSN 0002-8282. doi: 10.1257/aer.98.5.1887.

Emily Breza, Arun G. Chandrasekhar, and Davide Viviano. Generalizability with ignorance in mind: Learning what we do (not) know for archetypes discovery, July 2025.

Francisco J. Buera. A dynamic model of entrepreneurship with borrowing constraints: Theory and evidence. *Annals of Finance*, 5(3):443–464, June 2009. ISSN 1614-2454. doi: 10.1007/s10436-009-0121-2.

Pedro Carneiro, Lucy Kraftman, Giacomo Mason, Lucie Moore, Imran Rasul, and Molly Scott. The Impacts of a Multifaceted Prenatal Intervention on Human Capital Accumulation in Early Life. *American Economic Review*, 111(8):2506–2549, August 2021. ISSN 0002-8282. doi: 10.1257/aer.20191726.

Bob Carpenter, Andrew Gelman, Matthew D. Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A Probabilistic Programming Language. *Journal of Statistical Software*, 76:1–32, January 2017. ISSN 1548-7660. doi: 10.18637/jss.v076.i01.

Michael R. Carter and Travis J. Lybbert. Consumption versus asset smoothing: Testing the implications of poverty trap theory in Burkina Faso. *Journal of Development Economics*, 99(2):255–264, November 2012. ISSN 0304-3878. doi: 10.1016/j.jdeveco.2012.02.003.

David Cass. Optimum Growth in an Aggregative Model of Capital Accumulation1. *The Review of Economic Studies*, 32(3):233–240, July 1965. ISSN 0034-6527. doi: 10.2307/2295827.

Matias D. Cattaneo, Richard K. Crump, Max H. Farrell, and Yingjie Feng. On Binscatter. *American Economic Review*, 114(5):1488–1514, May 2024. ISSN 0002-8282. doi: 10.1257/aer.20221576.

Victor Chernozhukov and Han Hong. An MCMC approach to classical estimation. *Journal of Econometrics*, 115(2):293–346, August 2003. ISSN 0304-4076. doi: 10.1016/S0304-4076(03)00100-3.

Drew Creal. A Survey of Sequential Monte Carlo Methods for Economics and Finance. *Econometric Reviews*, 31(3):245–296, May 2012. ISSN 0747-4938. doi: 10.1080/07474938.2011.607333.

Tommaso Crosta, Dean Karlan, Finley Ong, Julius Rüschenpöhler, and Christopher R. Udry. Unconditional Cash Transfers: A Bayesian Meta-Analysis of Randomized Evaluations in Low and Middle Income Countries, August 2024.

Jesse M Cunha, Giacomo De Giorgi, and Seema Jayachandran. The Price Effects of Cash Versus In-Kind Transfers. *The Review of Economic Studies*, 86(1):240–281, January 2019. ISSN 0034-6527. doi: 10.1093/restud/rdy018.

Partha Dasgupta. Nutritional status, the capacity for work, and poverty traps. *Journal of Econometrics*, 77(1):5–37, March 1997. ISSN 0304-4076. doi: 10.1016/S0304-4076(96)01804-0.

Suresh De Mel, David McKenzie, and Christopher Woodruff. One-Time Transfers of Cash or Capital Have Long-Lasting Effects on Microenterprises in Sri Lanka. *Science*, 335(6071):962–966, February 2012. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1212973.

Stefano DellaVigna and Elizabeth Linos. RCTs to Scale: Comprehensive Evidence From Two Nudge Units. *Econometrica*, 90(1):81–116, January 2022. ISSN 0012-9682. doi: 10.3982/ECTA18709.

Eric Edmonds and Caroline Theoharides. The Short Term Impact of a Productive Asset Transfer in Families with Child Labor: Experimental Evidence from the Philippines. *Journal of Development Economics*, 146, September 2020. ISSN 03043878. doi: 10.1016/j.jdeveco.2020.102486.

Dennis Egger, Johannes Haushofer, Edward Miguel, Paul Niehaus, and Michael Walker. General Equilibrium Effects of Cash Transfers: Experimental Evidence From Kenya. *Econometrica*, 90(6):2603–2643, 2022. ISSN 1468-0262. doi: 10.3982/ECTA17945.

Marcel Fafchamps, David McKenzie, Simon Quinn, and Christopher Woodruff. Microenterprise growth and the flypaper effect: Evidence from a randomized experiment in Ghana. *Journal of Development Economics*, 106:211–226, January 2014. ISSN 0304-3878. doi: 10.1016/j.jdeveco.2013.09.010.

Paul J. Gertler, Sebastian W. Martinez, and Marta Rubio-Codina. Investing Cash Transfers to Raise Long-Term Living Standards. *American Economic Journal: Applied Economics*, 4(1):164–192, January 2012. ISSN 1945-7782. doi: 10.1257/app.4.1.164.

Vilas J. Gobin, Paulo Santos, and Russell Toth. No Longer Trapped? Promoting Entrepreneurship through Cash Transfers to Ultra-poor Women in Northern Kenya. *American Journal of Agricultural Economics*, 99(5):1362–1383, October 2017. ISSN 00029092.

Douglas Gollin and Christopher Udry. Heterogeneity, Measurement Error, and Misallocation: Evidence from African Agriculture. *Journal of Political Economy*, 129(1):1–80, January 2021. ISSN 0022-3808. doi: 10.1086/711369.

Prankur Gupta, Daniel Stein, Kyla Longman, Heather Lanthorn, Rico Bergmann, and Emmanuel Nshakira-Rukundo. Cash Transfers amid Shocks: A Large, One-Time, Unconditional Cash Transfer to Refugees in Uganda Has Multidimensional Benefits after 19 Months. *World Development*, 173, January 2024. ISSN 0305750X. doi: 10.1016/j.worlddev.2023.106339.

Robert E. Hall and Frederic S. Mishkin. The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households. *Econometrica*, 50(2):461–481, 1982. ISSN 0012-9682. doi: 10.2307/1912638.

Johannes Haushofer and Jeremy Shapiro. The Short-term Impact of Unconditional Cash Transfers to the Poor: Experimental Evidence from Kenya*. *The Quarterly Journal of Economics*, 131(4):1973–2042, November 2016. ISSN 0033-5533. doi: 10.1093/qje/qjw025.

Johannes Haushofer, Paul Niehaus, Carlos Paramo, Edward Miguel, and Michael Walker. Targeting Impact versus Deprivation. *American Economic Review*, 115(6):1936–1974, June 2025. ISSN 0002-8282. doi: 10.1257/aer.20221650.

Chang-Tai Hsieh and Peter J. Klenow. Misallocation and Manufacturing TFP in China and India*. *The Quarterly Journal of Economics*, 124(4):1403–1448, November 2009. ISSN 0033-5533. doi: 10.1162/qjec.2009.124.4.1403.

Reshmaan Hussam, Natalia Rigol, and Benjamin N. Roth. Targeting High Ability Entrepreneurs Using Community Information: Mechanism Design in the Field. *American Economic Review*, 112(3):861–898, March 2022. ISSN 0002-8282. doi: 10.1257/aer.20200751.

Munenobu Ikegami, Michael Carter, Christopher B. Barrett, and Sarah Janzen. Poverty Traps and the Social Protection Paradox, October 2016.

Sid Kankanala. Generalized Bayes in Conditional Moment Restriction Models, October 2025.

Aart Kraay and David McKenzie. Do Poverty Traps Exist? Assessing the Evidence. *Journal of Economic Perspectives*, 28(3):127–148, September 2014. ISSN 0895-3309. doi: 10.1257/jep.28.3.127.

Michael Kremer, Stephen P. Luby, Ricardo Maertens, Brandon Tan, and Witold Więcek. Water Treatment And Child Mortality: A Meta-Analysis And Cost-effectiveness Analysis, January 2023.

Yoonseok Lee. Nonparametric Estimation of Dynamic Panel Models with Fixed Effects, August 2007.

Lizhen Lin and David B. Dunson. Bayesian Monotone Regression using Gaussian Process Projection, June 2013.

Robert E. Lucas. On the mechanics of economic development. *Journal of Monetary Economics*, 22(1):3–42, July 1988. ISSN 0304-3932. doi: 10.1016/0304-3932(88)90168-7.

Karen Macours, Patrick Premand, and Renos Vakis. Transfers, Diversification and Household Risk Strategies: Experimental evidence with lessons for climate change adaptation. 2012.

Casey C. Maue, Marshall Burke, and Kyle J. Emerick. Productivity Dispersion and Persistence Among the World's Most Numerous Firms. Technical Report w26924, National Bureau of Economic Research, April 2020.

Craig McIntosh and Andrew Zeitlin. Using Household Grants to Benchmark the Cost Effectiveness of a USAID Workforce Readiness Program, September 2020.

David J. McKenzie and Christopher Woodruff. Do Entry Costs Provide an Empirical Basis for Poverty Traps? Evidence from Mexican Microenterprises. *Economic Development and Cultural Change*, 55(1):3–42, October 2006. ISSN 0013-0079. doi: 10.1086/505725.

Rachael Meager. Understanding the Average Impact of Microcredit Expansions: A Bayesian Hierarchical Analysis of Seven Randomized Experiments. *American Economic Journal: Applied Economics*, 11(1):57–91, January 2019. ISSN 1945-7782. doi: 10.1257/app.20170299.

Rachael Meager. Aggregating Distributional Treatment Effects: A Bayesian Hierarchical Analysis of the Microcredit Literature. *American Economic Review*, 112(6):1818–1847, June 2022. ISSN 0002-8282. doi: 10.1257/aer.20181811.

Benjamin Moll. Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation? *American Economic Review*, 104(10):3186–3221, October 2014. ISSN 0002-8282. doi: 10.1257/aer.104.10.3186.

Kevin M. Murphy, Andrei Shleifer, and Robert W. Vishny. Industrialization and the Big Push. *Journal of Political Economy*, October 1989. ISSN 0022-3808. doi: 10.1086/261641.

Whitney K. Newey and James L. Powell. Instrumental Variable Estimation of Nonparametric Models. *Econometrica*, 71(5):1565–1578, 2003. ISSN 0012-9682.

Ugo Pica-Ciamarra, Derek Baker, John Chassama, Mohamadou Fadiga, and Longin Nsiima. Linking Smallholders to Livestock Markets in Tanzania: Combing Market and Household Survey Data. 2011.

Danny T. Quah. Twin Peaks: Growth and Convergence in Models of Distribution Dynamics. *The Economic Journal*, 106(437):1045–1055, 1996. ISSN 0013-0133. doi: 10.2307/2235377.

Gabriel Riutort-Mayol, Paul-Christian Bürkner, Michael R. Andersen, Arno Solin, and Aki Vehtari. Practical Hilbert space approximate Bayesian Gaussian processes for probabilistic programming. *Statistics and Computing*, 33(1):17, December 2022. ISSN 1573-1375. doi: 10.1007/s11222-022-10167-2.

P. N. Rosenstein-Rodan. Problems of Industrialisation of Eastern and South-Eastern Europe. *The Economic Journal*, 53(210/211):202–211, 1943. ISSN 0013-0133. doi: 10.2307/2226317.

Jeffrey D. Sachs, John W. McArthur, Guido Schmidt-Traub, Margaret Kruk, Chandrika Bahadur, Michael Faye, and Gordon McCord. Ending Africa's Poverty Trap. *Brookings Papers on Economic Activity*, 2004(1):117–216, 2004. ISSN 0007-2303.

Paulo Santos and Christopher B. Barrett. Persistent poverty and informal credit. *Journal of Development Economics*, 96(2):337–347, November 2011. ISSN 0304-3878. doi: 10.1016/j.jdeveco.2010.08.017.

Richard Sedlmayr, Anuj Shah, and Munshi Sulaiman. Cash-plus: Poverty impacts of alternative transfer-based approaches. *Journal of Development Economics*, 144:102418, May 2020. ISSN 0304-3878. doi: 10.1016/j.jdeveco.2019.102418.

A. K. Skiba. Optimal Growth with a Convex-Concave Production Function. *Econometrica*, 46(3):527–539, 1978. ISSN 0012-9682. doi: 10.2307/1914229.

Liangjun Su and Xun Lu. Nonparametric dynamic panel data models: Kernel estimation and specification testing. *Journal of Econometrics*, 176(2):112–133, October 2013. ISSN 0304-4076. doi: 10.1016/j.jeconom.2013.04.020.

Chad Syverson. What Determines Productivity? *Journal of Economic Literature*, 49(2): 326–365, June 2011. ISSN 0022-0515. doi: 10.1257/jel.49.2.326.

A Appendix Tables

Table A.1: Sample Counts by Region

Region	Number of Studies	Number of Countries
East Asia & Pacific	1	1
Latin America & Caribbean	5	4
South Asia	6	3
Sub-Saharan Africa	15	9

Notes: This table reports the number of studies per region in the final harmonized sample and number of unique countries represented in each region.

Table A.2: Sample Counts by Country Income Group

Statistic	Low income	Lower middle income	Upper middle income	High income
World Bank # of countries	26	50	54	87
Sample # of countries	5	10	2	0
Fraction of WB countries (%)	12	23	24.9	40.1
Fraction of sample countries (%)	29.4	58.8	11.8	0
Number of households	15,416	40,358	19,037	0
Number of observations	49,580	130,121	83,591	0

Notes: This table reports the statistics on the number of countries in each World Bank income group compared to the number of studies in the final harmonized sample in each income group.

Table A.3: Treatment Statistics

Program	Fraction Treated	Median Transfer (USD PPP)	Median Instalments
Mexico, Cunha (2019)	0.39	21.43	15
Nigeria, Carneiro (2021)	0.47	22.00	24
Nicaragua, Macours (2012)	0.72	25.83	12
Ethiopia, Banerjee (2015)	1.00	25.99	41
Ghana, Banerjee (2015)	0.26	30.00	13
India, Banerjee (2021)	0.53	30.40	41
India, Banerjee (2015)	0.52	35.04	10
Kenya, Haushofer (2016)	0.27	44.89	8
Uganda, Sedlmayr (2020)	0.08	77.56	2
India, Hussam (2022)	0.34	100.00	1
Ghana, Fafchamps (2014)	0.50	120.00	1
Mexico, Gertler (2012)	0.52	121.63	16
Mali, Beaman (2023)	0.27	140.00	1
Peru, Banerjee (2015)	1.00	143.33	10
Rwanda, Mcintosh (2020)	0.47	164.08	2
Kenya, Gobin (2017)	0.67	237.97	1
Somalia, Abdullahi (2023)	1.00	250.00	1
Sri Lanka, Demel (2009)	0.52	289.27	1
Tanzania, Baird (2024)	0.56	350.00	1
Honduras, Banerjee (2015)	0.33	376.96	2
Sri Lanka, Demel (2012)	0.48	395.57	1
Uganda, Blattman (2019)	0.49	400.00	1
Philippines, Edmonds (2020)	0.50	518.00	1
Bangladesh, Balboni (2022)	0.21	541.71	1
Kenya, Egger (2022)	0.49	860.00	3
Ethiopia, Blattman (2018)	0.30	1030.00	1
Uganda, Gupta (2024)	0.51	1286.29	3
Overall	0.44	441.11	2

Notes: This table reports statistics on the treatment arms in each study. 'Fraction Treated' is the average treatment status of all households in the study, averaging over survey waves. 'Median Transfer' is the median non-zero transfer administered to households - not the total transfer received by a household if they were treated. 'Median Instalments' is the median number of installments in which the transfer was disbursed to households. The 'Overall' row pools all studies to give statistics averaging over all studies. The fraction treated in Ethiopia and Peru (Banerjee et al., 2015) and Somalia (Abdullahi et al., 2023) is equal to 1 because all households in the control group received some level of consumption support, therefore to be conservative, I classify all households in these studies as treated.

Table A.4: Outcome Statistics

Program	Sample Size			Median (USD PPP)		
	Households	Observations	Survey Length (Months)	Assets	Income	Consumption
Mexico, Gertler (2012)	11,610	66,458	73	-	90	140
Bangladesh, Balboni (2022)	6,273	29,613	132	279	87	150
Nicaragua, Macours (2012)	6,539	15,695	36	-	316	281
Mali, Beaman (2023)	4,048	14,770	84	1,142	24	156
Kenya, Egger (2022)	6,841	13,682	32	546	35	198
Nigeria, Carneiro (2021)	4,533	13,123	48	-	15	1,434
Uganda, Sedlmayr (2020)	3,804	10,845	36	462	8	545
Uganda, Blattman (2019)	2,677	10,707	108	144	183	1,059
Mexico, Cunha (2019)	5,148	10,296	25	86	18	276
Ghana, Banerjee (2015)	2,598	7,794	40	335	16	219
Honduras, Banerjee (2015)	2,403	7,209	53	889	167	296
Peru, Banerjee (2015)	2,279	6,837	48	4,860	99	575
Sri Lanka, Demel (2009)	617	5,770	36	2,884	145	320
India, Hussam (2022)	1,345	5,354	10	431	355	152
Sri Lanka, Demel (2012)	408	5,304	68	2,820	172	-
Kenya, Gobin (2017)	1,729	5,062	16	221	98	189
India, Banerjee (2021)	991	4,955	120	-	1,831	51
Ghana, Fafchamps (2014)	791	4,746	16	279	84	163
Philippines, Edmonds (2020)	2,296	4,584	24	4,071	480	772
Ethiopia, Blattman (2018)	645	3,179	25	-	23	135
India, Banerjee (2015)	987	2,961	38	165	81	162
Kenya, Haushofer (2016)	1,419	2,831	18	343	19	153
Ethiopia, Banerjee (2015)	915	2,745	40	1,585	46	178
Uganda, Gupta (2024)	1,264	2,528	29	607	52	772
Rwanda, Mcintosh (2020)	1,268	2,457	14	563	36	93
Somalia, Abdullahi (2023)	795	2,349	45	500	50	105
Tanzania, Baird (2024)	588	1,438	32	-	87	78
Overall	74,811	263,292	132	595	81	196

Notes: This table reports outcome summary statistics and the number of total observations for each study in the final harmonized sample. 'Assets' refers to total household assets, 'Income' is household income per month, and 'Consumption' is household consumption per month. All monetary values are sample medians in 2020 USD.

Table A.5: Proxy Outcome Statistics

Program	Sample Size			Median (USD PPP)		
	Households	Observations	Survey Length (Months)	Assets	Income	Consumption
Mexico, Gertler (2012)	11,610	66,458	73	129	90	140
Bangladesh, Balboni (2022)	6,273	29,613	132	279	87	150
Nicaragua, Macours (2012)	6,539	15,695	36	95	316	281
Mali, Beaman (2023)	4,048	14,770	84	1,097	24	156
Kenya, Egger (2022)	6,841	13,682	32	546	35	198
Nigeria, Carneiro (2021)	4,533	13,123	48	0	15	1,434
Uganda, Sedlmayr (2020)	3,804	10,845	36	372	8	545
Uganda, Blattman (2019)	2,677	10,707	108	144	183	1,059
Mexico, Cunha (2019)	5,148	10,296	25	86	18	276
Ghana, Banerjee (2015)	2,598	7,794	40	335	16	219
Honduras, Banerjee (2015)	2,403	7,209	53	889	167	296
Peru, Banerjee (2015)	2,279	6,837	48	4,860	99	575
Sri Lanka, Demel (2009)	617	5,770	36	2,884	139	320
India, Hussam (2022)	1,345	5,354	10	431	355	152
Sri Lanka, Demel (2012)	408	5,304	68	2,820	175	-
Kenya, Gobin (2017)	1,729	5,062	16	221	98	189
India, Banerjee (2021)	991	4,955	120	460	1,811	51
Ghana, Fafchamps (2014)	791	4,746	16	279	84	163
Philippines, Edmonds (2020)	2,296	4,584	24	3,323	404	772
Ethiopia, Blattman (2018)	645	3,179	25	519	23	135
India, Banerjee (2015)	987	2,961	38	165	81	162
Kenya, Haushofer (2016)	1,419	2,831	18	343	19	153
Ethiopia, Banerjee (2015)	915	2,745	40	1,585	46	178
Uganda, Gupta (2024)	1,264	2,528	29	607	52	772
Rwanda, Mcintosh (2020)	1,268	2,457	14	563	36	93
Somalia, Abdullahi (2023)	795	2,349	45	500	50	105
Tanzania, Baird (2024)	588	1,438	32	280	87	78
Overall	74,811	263,292	132	335	81	196

Notes: This table reports outcome summary statistics and the number of total observations for each study in the final harmonized sample using 'proxy' outcomes pooled across income and asset measures. 'Assets' refers to the proxy asset measure, 'Income' is the best available proxy for household income per month, and 'Consumption' is household consumption per month. All monetary values are sample medians in 2020 USD.

Table A.6: Covariate Statistics

	Median		
Program	Household Size	Age (Years)	Years Education
Mexico, Gertler (2012)	6	40	3
Bangladesh, Balboni (2022)	4	42	4
Nicaragua, Macours (2012)	5	42	4
Mali, Beaman (2023)	7	33	-
Kenya, Egger (2022)	4	42	7
Nigeria, Carneiro (2021)	7	32	-
Uganda, Sedlmayr (2020)	6	17	5
Uganda, Blattman (2019)	6	25	7
Mexico, Cunha (2019)	4	43	6
Ghana, Banerjee (2015)	7	-	-
Honduras, Banerjee (2015)	6	-	-
Peru, Banerjee (2015)	5	-	-
Sri Lanka, Demel (2009)	5	42	11
India, Hussam (2022)	4	40	9
Sri Lanka, Demel (2012)	1	41	10
Kenya, Gobin (2017)	6	31	6
India, Banerjee (2021)	4	40	-
Ghana, Fafchamps (2014)	3	35	10
Philippines, Edmonds (2020)	6	43	10
Ethiopia, Blattman (2018)	4	21	10
India, Banerjee (2015)	4	-	-
Kenya, Haushofer (2016)	5	31	9
Ethiopia, Banerjee (2015)	5	-	-
Uganda, Gupta (2024)	8	38	7
Rwanda, Mcintosh (2020)	4	23	7
Somalia, Abdullahi (2023)	6	-	-
Tanzania, Baird (2024)	5	-	-
Overall	5	37	6

Notes: This table reports median household size, median household age, and median years of education of adult household members for each study in the final harmonized sample.

Table A.7: Sample Balance – Final Harmonized Sample

Variable	# Studies	# Observations	Wider Sample Mean	Subsample Difference
Proxy Income	48	314,391	266.3	-189.7 [0.348]
Total Transfer	59	458,893	99.5	77.8 [0.052]
Consumption	49	356,377	246.3	-136.3 [0.231]
Income	45	295,091	67.5	-5.1 [0.851]
Wage Income	15	55,782	226.9	-158.5 [0.243]
Revenue	19	150,759	923.3	1,482.1 [0.508]
Profit	26	113,987	197.3	-166.1 [0.282]
Savings	25	140,882	3,644.7	-3,271.6 [0.445]
Debt	32	159,724	169.0	1,726.8 [0.307]
Lending	4	5,418	40.0	-16.3 [0.226]
Total Hours Worked	31	136,876	77.1	300.7 [0.242]
Number of Cows	25	257,768	0.2	2.6 [0.209]
Number of Goats	22	228,337	0.3	4.3 [0.192]
Number of Pigs	16	168,516	0.2	0.4 [0.036]
Metal Roof	22	146,083	0.5	0.4 [0.397]
Age (Years)	45	355,040	36.0	2.8 [0.411]
Years of Education	35	293,119	5.4	-1.2 [0.391]
Gender = Female	54	390,611	0.6	-0.1 [0.723]
Household Size	53	431,760	5.2	0.4 [0.376]

Notes: This table reports the sample mean of outcomes and covariates for the wider cleaned sample and the estimated difference in means between the final harmonized sample and the wider cleaned sample. p-values are from two-sided t-tests of equality of means in square brackets.

Table A.8: Poverty Trap Simulation Metrics – Single Steady State

		Dynamic Panel Model			Nonparametric Asset Transition Regression		
Ability Quantile	True # Steady States	Estimated # Steady States	% Correct # Steady States	% Correct # Steady States > 1	Estimated # Steady States	% Correct # Steady States	% Correct # Steady States > 1
0.0	1.0	1.0	100	100	3.5	18	18
0.05	1.0	1.0	100	100	3.5	18	18
0.11	1.0	1.0	100	100	3.5	18	18
0.16	1.0	1.0	100	100	3.5	18	18
0.21	1.0	1.0	100	100	3.5	18	18
0.26	1.0	1.0	100	100	3.5	18	18
0.32	1.0	1.0	100	100	3.5	18	18
0.37	1.0	1.04	98	98	3.5	18	18
0.42	1.0	1.08	96	96	3.5	18	18
0.47	1.0	1.04	98	98	3.5	18	18
0.53	1.0	1.2	90	90	3.5	18	18
0.58	1.0	1.2	90	90	3.5	18	18
0.63	1.0	1.24	88	88	3.5	18	18
0.68	1.0	1.2	90	90	3.5	18	18
0.74	1.0	1.24	88	88	3.5	18	18
0.79	1.0	1.08	96	96	3.5	18	18
0.84	1.0	1.0	100	100	3.5	18	18
0.89	1.0	1.0	100	100	3.5	18	18
0.95	1.0	1.0	100	100	3.5	18	18
1.0	1.0	1.0	100	100	3.5	18	18
Overall	-	1.07	97	97	3.5	18	18

Notes: This table shows the performance metrics of the poverty trap estimator on simulated data generated using the structural model without production technology non-convexities and therefore only a single steady state. Each figure reports the average result over 100 simulated datasets and each row corresponds to a different ability percentile. The overall average is reported in the final row. 'True # Steady States' is the total number of steady states in the asset transition equation using the model, 'Estimated # Steady States' is the average number of steady states estimated by each trap estimator. '% Correct # Steady States' is the proportion of simulations where the estimator correctly estimates the number of steady states. '% Correct # Steady States > 1' is the proportion of simulations where the estimator correctly estimates more than one steady state if there are multiple steady states present. 'Dynamic Panel Model' is the estimator presented in this paper, 'Nonparametric Asset Transition Regression' uses the nonparametric regression of future assets on current assets plus the randomized transfer.

Table A.9: Poverty Trap Simulation Metrics – Multiple Steady States

		Dynamic Panel Model			Nonparametric Asset Transition Regression		
Ability Quantile	True # Steady States	Estimated # Steady States	% Correct # Steady States	% Correct # Steady States > 1	Estimated # Steady States	% Correct # Steady States	% Correct # Steady States > 1
0.0	1.0	1.0	100	100	3.33	27	27
0.05	1.0	1.06	94	94	3.33	27	27
0.11	1.0	1.0	100	100	3.33	27	27
0.16	1.0	1.0	100	100	3.33	27	27
0.21	1.0	1.0	100	100	3.33	27	27
0.26	3.0	1.02	1	1	3.33	38	73
0.32	3.0	1.52	26	26	3.33	38	73
0.37	3.0	2.6	80	80	3.33	38	73
0.42	3.0	2.96	98	98	3.33	38	73
0.47	3.0	2.78	89	89	3.33	38	73
0.53	3.0	1.86	43	43	3.33	38	73
0.58	1.0	1.08	96	96	3.33	27	27
0.63	1.0	1.04	98	98	3.33	27	27
0.68	1.0	1.0	100	100	3.33	27	27
0.74	1.0	1.0	100	100	3.33	27	27
0.79	1.0	1.0	100	100	3.33	27	27
0.84	1.0	1.0	100	100	3.33	27	27
0.89	1.0	1.0	100	100	3.33	27	27
0.95	1.0	1.0	100	100	3.33	27	27
1.0	1.0	1.0	100	100	3.33	27	27
Overall	-	1.35	86	86	3.33	30	41

Notes: This table shows the performance metrics of the poverty trap estimator on simulated data generated using the structural model with production technology non-convexities and therefore multiple steady states. Each figure reports the average result over 100 simulated datasets and each row corresponds to a different ability percentile. The overall average is reported in the final row. 'True # Steady States' is the total number of steady states in the asset transition equation using the model, 'Estimated # Steady States' is the average number of steady states estimated by each trap estimator. '% Correct # Steady States' is the proportion of simulations where the estimator correctly estimates the number of steady states. '% Correct # Steady States > 1' is the proportion of simulations where the estimator correctly estimates more than one steady state if there are multiple steady states present. 'Dynamic Panel Model' is the estimator presented in this paper, 'Nonparametric Asset Transition Regression' uses the nonparametric regression of future assets on current assets plus the randomized transfer.

Table A.10: Population Identification Simulations

Instruments	Pct. Significant (1%)	Pct. Significant (5%)	Pct. Significant (10%)
Difference instruments, no transfer	1.6	6.9	12.5
Difference instruments, with transfer	1.5	7.3	13.0
Level instruments, no transfer	0.0	0.2	1.1
Level instruments, with transfer	0.0	0.2	1.0
Current assets (Endogenous)	59.8	76.0	82.5
Lagged income (Endogenous)	90.0	99.3	99.8

Notes: This table shows results from simulations using the structural model calibrated against an example study. Each column shows the proportion of p-values that reject the null hypothesis of mean independence of the instruments and error terms at the corresponding significance level. Under the null hypothesis, valid instruments should reject at the nominal rate, whilst endogenous instruments should reject at a much higher rate. Rows 1-4 show results using the correct, model-based instruments, whilst rows 5-6 show results using endogenous lagged variables as instruments. Rates are calculated over 500 simulated draws of 100 households over 5 time periods.

Table A.11: Estimated Poverty Trap Summary Statistics

Program	Threshold (USD)	Below Threshold (%)	Trapped (%)
Ghana, Banerjee (2015)	138	64.2	25.3
Tanzania, Baird (2024)	163	89.9	30.7
Somalia, Abdullahi (2023)	247	48.6	24.6
Honduras, Banerjee (2015)	261	71.1	5.6
Kenya, Egger (2022)	317	0	0
Bangladesh, Balboni (2022)	471	98.4	24.3
Uganda, Sedlmayr (2020)	505	90.9	37.5
Peru, Banerjee (2015)	550	95.5	55.4
Mexico, Gertler (2012)	747	100	0
Kenya, Haushofer (2016)	890	98.7	0
Ethiopia, Banerjee (2015)	907	100	0
Nigeria, Carneiro (2021)	981	0	0
Mali, Beaman (2023)	-	-	-
Ethiopia, Blattman (2018)	-	-	-
Uganda, Blattman (2019)	-	-	-
Mexico, Cunha (2019)	-	-	-
Ghana, Fafchamps (2014)	-	-	-
Uganda, Gupta (2024)	-	-	-
Rwanda, Mcintosh (2020)	-	-	-

Notes: This table reports poverty trap statistics, estimated using the calibrated household growth model, for each study site. The first column reports the program name, whilst the 'Threshold' column reports the average asset threshold a household must cross to move from a low- to high-steady state in 2020 USD. The 'Below Threshold (%)' column shows the proportion of households at baseline with assets below this threshold. The 'Trapped (%)' column shows the proportion of households below the threshold that are 'trapped' – that is, the proportion of households that are productive enough such that if they were given a transfer to cross the threshold, they would continue to accumulate assets, but are not so productive that they are predicted to cross this threshold without a transfer.

Table A.12: Production Indivisibility and Consumption Nonmontonicity Confusion Matrix

	Consumption monotonic	Consumption non-monotonic
No production kinks	50.0%	50.0%
Production kinks	30.8%	69.2%

Notes: This table shows the confusion matrix between production indivisibility and consumption non-monotonicity classifications. I exclude studies where the production function is entirely flat (Blattman and Dercon 2018) as predictions about poverty traps are unclear in this case.

Table A.13: Productivity Dispersion

Study	SD of η	$\log\Delta$ 75-25	$\log\Delta$ 90-10	Ratio 75-25	Ratio 90-10
Tanzania, Baird (2024)	0.42	0.55	1.10	1.74	3.00
Ghana, Fafchamps (2014)	0.38	0.50	0.96	1.65	2.62
Ghana, Banerjee (2015)	0.35	0.47	0.91	1.60	2.48
India, Banerjee (2021)	0.35	0.45	0.90	1.56	2.47
India, Hussam (2022)	0.32	0.40	0.81	1.50	2.24
Nicaragua, Macours (2012)	0.32	0.42	0.79	1.52	2.20
Mexico, Cunha (2019)	0.27	0.33	0.70	1.39	2.02
Philippines, Edmonds (2020)	0.23	0.33	0.65	1.39	1.92
Nigeria, Carneiro (2021)	0.26	0.37	0.65	1.45	1.91
Ethiopia, Banerjee (2015)	0.26	0.34	0.63	1.41	1.88
Bangladesh, Balboni (2022)	0.24	0.32	0.61	1.38	1.85
India, Banerjee (2015)	0.25	0.30	0.61	1.35	1.84
Uganda, Gupta (2024)	0.22	0.30	0.57	1.35	1.77
Somalia, Abdullahi (2023)	0.27	0.23	0.56	1.26	1.75
Honduras, Banerjee (2015)	0.22	0.28	0.56	1.32	1.74
Uganda, Sedlmayr (2020)	0.23	0.25	0.50	1.29	1.65
Ethiopia, Blattman (2018)	0.21	0.28	0.46	1.33	1.58
Peru, Banerjee (2015)	0.17	0.22	0.45	1.24	1.56
Kenya, Haushofer (2016)	0.17	0.24	0.43	1.27	1.54
Mali, Beaman (2023)	0.16	0.22	0.41	1.25	1.51
Mexico, Gertler (2012)	0.16	0.21	0.40	1.23	1.50
Kenya, Gobin (2017)	0.15	0.20	0.38	1.23	1.47
Rwanda, Mcintosh (2020)	0.15	0.19	0.36	1.21	1.43
Uganda, Blattman (2019)	0.14	0.18	0.36	1.20	1.43
Kenya, Egger (2022)	0.13	0.17	0.32	1.19	1.38
Unconditional	0.26	0.28	0.58	1.32	1.78

Notes: This table shows summary statistics of the estimated latent ability (fixed productivity) distribution for each study using the stochastic EM algorithm. 'SD of η' is the standard deviation of the estimated ability distribution, '$\log\Delta$ 75-25' is the difference between the 75th and 25th percentiles of the estimated ability distribution on the log scale, and '$\log\Delta$ 90-10' is same difference between the 90th and 10th percentiles. 'Ratio 75-25' and 'Ratio 90-10' are the corresponding ratios on the raw (i.e. exponential) scale. 'Unconditional' shows the same statistics pooling all studies together.

Table A.14: Production Parameters

Study	ρ	σ_ξ
Somalia, Abdullahi (2023)	0.446 [0.199, 0.604]	0.378 [0.187, 0.564]
Tanzania, Baird (2024)	0.351 [0.318, 0.383]	0.448 [0.158, 0.759]
Bangladesh, Balboni (2022)	0.478 [0.47, 0.486]	0.998 [0.957, 1.034]
Ethiopia, Banerjee (2015)	0.349 [0.302, 0.392]	0.488 [0.192, 0.678]
Ghana, Banerjee (2015)	0.357 [0.326, 0.384]	0.904 [0.342, 1.153]
Honduras, Banerjee (2015)	0.107 [0.095, 0.118]	0.842 [0.802, 0.88]
India, Banerjee (2015)	0.265 [0.212, 0.316]	0.869 [0.458, 1.109]
Peru, Banerjee (2015)	0.398 [0.361, 0.428]	1.173 [1.118, 1.217]
India, Banerjee (2021)	0.03 [0.003, 0.083]	1.415 [1.341, 1.49]
Mali, Beaman (2023)	0.275 [0.266, 0.285]	0.932 [0.896, 0.968]
Ethiopia, Blattman (2018)	0.023 [0.002, 0.076]	0.604 [0.263, 0.894]
Uganda, Blattman (2019)	0.566 [0.545, 0.586]	0.928 [0.356, 1.18]
Nigeria, Carneiro (2021)	0.349 [0.124, 0.547]	0.766 [0.346, 1.101]
Mexico, Cunha (2019)	0.595 [0.523, 0.666]	0.927 [0.367, 1.201]
Sri Lanka, Demel (2009)	0.143 [0.125, 0.162]	0.45 [0.277, 0.584]
Sri Lanka, Demel (2012)	0.105 [0.089, 0.123]	0.391 [0.212, 0.558]
Philippines, Edmonds (2020)	0.356 [0.107, 0.531]	0.556 [0.211, 0.901]
Kenya, Egger (2022)	0.183 [0.164, 0.201]	0.85 [0.336, 1.112]
Ghana, Fafchamps (2014)	0.108 [0.093, 0.126]	0.438 [0.231, 0.591]
Mexico, Gertler (2012)	0.161 [0.155, 0.169]	0.989 [0.969, 1.011]
Kenya, Gobin (2017)	0.02 [0.014, 0.026]	0.932 [0.867, 0.988]
Uganda, Gupta (2024)	0.103 [0.073, 0.148]	1.469 [0.46, 1.743]
Kenya, Haushofer (2016)	0.266 [0.245, 0.285]	0.88 [0.286, 1.151]
India, Hussam (2022)	0.007 [0.003, 0.012]	0.427 [0.196, 0.579]
Nicaragua, Macours (2012)	0.655 [0.604, 0.698]	0.323 [0.158, 0.466]
Rwanda, Mcintosh (2020)	0.245 [0.208, 0.28]	0.893 [0.357, 1.14]
Uganda, Sedlmayr (2020)	0.453 [0.442, 0.465]	1.155 [1.102, 1.205]
Overall	0.299 [0.007, 0.654]	0.794 [0.248, 1.489]

Notes: This table shows the estimated productivity persistence parameter, ρ_j, approximately annualized to account for differences in panel length across studies ($\rho_j^{12/\text{Survey gap (months)}}$). The right column shows the estimated productivity innovation standard deviation, σ_{ξ_j}. The point estimate is the posterior mean with 95% credible intervals in square brackets. Overall shows the hierarchical mean estimates across all studies.

Table A.15: Correlates of Household-level Poverty Traps: Thresholds

Dependent Variable:	Threshold Exists				Threshold Size (USD PPP)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Logit	Logit	Logit	Logit	OLS	OLS	OLS	OLS
Below Median HH Size	-0.402***	-0.401***	-0.397***	-0.397***	2.24	2.27	2.03	1.99
	(0.102)	(0.101)	(0.100)	(0.100)	(6.08)	(6.09)	(5.98)	(5.99)
Female		-0.320	-0.320	-0.326		-3.57	-3.55	-2.22
		(0.197)	(0.198)	(0.198)		(8.30)	(8.32)	(8.47)
Below Median Age			-0.030	-0.024			1.96	0.795
			(0.058)	(0.061)			(6.86)	(6.98)
Below Median Education				0.042				-6.52**
				(0.066)				(2.71)
Number of Studies	11	11	11	11	10	10	10	10
Observations	39,381	39,381	39,381	39,381	1,491	1,491	1,491	1,491

Notes: This table shows two household-level regressions: a logistic regression of whether a household faces multiple steady states in their asset transition equation, Columns (1)-(3), on household-level characteristics; and an OLS regression of the threshold size on household-level characteristics, conditional on facing a threshold, shown in Columns (4)-(6). 'Female' is the fraction of household members that are female, 'Below Median Age' is whether the average age of adult household members is below the study median, and 'Below Median Education' is whether the average years of education of adult household members is below the study median. All regressions include study fixed effects. Standard errors clustered at the study level in parentheses. Stars denote significance at the 10%, 5%, and 1% levels respectively.

Table A.16: Correlates of Household-level Poverty Traps: HH Trapped

Dependent Variable:	HH Trapped Below Threshold			
	(1)	(2)	(3)	(4)
Below Median HH Size	-1.21×10^{-5}	-0.0007	-0.004	-0.004
	(0.036)	(0.037)	(0.036)	(0.035)
Female		0.096**	0.096**	0.083**
		(0.034)	(0.033)	(0.029)
Below Median Age			0.028	0.040*
			(0.021)	(0.020)
Below Median Education				0.065**
				(0.023)
Number of Studies	10	10	10	10
Observations	1,491	1,491	1,491	1,491

Notes: This table shows a logistic regression of a household being classified as 'trapped' (i.e. faces an asset transition equation with multiple steady states and is below the asset threshold) on household-level characteristics. 'Female' is the fraction of household members that are female, 'Below Median Age' is whether the average age of adult household members is below the study median, and 'Below Median Education' is whether the average years of education of adult household members is below the study median. All regressions include study fixed effects. Standard errors clustered at the study level in parentheses. Stars denote significance at the 10%, 5%, and 1% levels respectively.

Table A.17: Study-level Production Kink Correlates

Dependent Variable:	Production Kinks				
	(1)	(2)	(3)	(4)	(5)
Number of Households (100s)	0.003**	0.003*	0.003	0.002	0.002
	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
Panel Rounds		0.049	0.044	-0.004	-0.001
		(0.073)	(0.078)	(0.078)	(0.074)
Panel Length (months)			0.0006	0.001	0.0004
			(0.004)	(0.004)	(0.004)
Mean Transfers (100s)				-0.034**	-0.036***
				(0.013)	(0.012)
Any Livestock					0.226
					(0.214)
Intercept	0.403***	0.250	0.246	0.598*	0.484
	(0.126)	(0.247)	(0.254)	(0.292)	(0.312)
Observations	25	25	25	25	25

Notes: This table shows a study-level regression of whether a study has a kink in the estimated production function on study-level characteristics. 'Number of HH (100s)' is the number of households in the study sample in hundreds, 'Panel Rounds' is the number of survey rounds in the study, 'Panel Length (months)' is the length of time between baseline and the final endline survey in months, 'Mean Transfers (100s)' is the average transfer size in hundreds of 2020 USD, and 'Any Livestock' is an indicator for whether the study collects data on livestock ownership. Stars denote significance at the 10%, 5%, and 1% levels respectively.

Table A.18: Study-level Consumption Nonmonotonicity Correlates

Dependent Variable:	Non-Monotonic Consumption				
	(1)	(2)	(3)	(4)	(5)
Number of Households (100s)	0.002*	0.001	0.0006	0.0006	0.0006
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
Panel Rounds		0.132**	0.112*	0.132**	0.133*
		(0.056)	(0.054)	(0.060)	(0.065)
Panel Length (months)			0.003	0.002	0.002
			(0.002)	(0.002)	(0.003)
Mean Transfers (100s)				0.014	0.013
				(0.015)	(0.014)
Any Livestock					0.136
					(0.229)
Intercept	0.558***	0.147	0.130	-0.011	-0.080
	(0.124)	(0.249)	(0.253)	(0.313)	(0.325)
Observations	25	25	25	25	25

Notes: This table shows a study-level regression of whether a study has a nonmonotonic average consumption rule on study-level characteristics. 'Number of HH (100s)' is the number of households in the study sample in hundreds, 'Panel Rounds' is the number of survey rounds in the study, 'Panel Length (months)' is the length of time between baseline and the final endline survey in months, 'Mean Transfers (100s)' is the average transfer size in hundreds of 2020 USD, and 'Any Livestock' is an indicator for whether the study collects data on livestock ownership. Stars denote significance at the 10%, 5%, and 1% levels respectively.

Table A.19: Study-level Multiple Asset Transition Correlates

Dependent Variable:	Multiple Steady States				
	(1)	(2)	(3)	(4)	(5)
Number of Households (100s)	0.003***	0.003**	0.0009	0.0010	0.001
	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
Panel Rounds		0.086	0.035	0.083	0.082
		(0.074)	(0.071)	(0.057)	(0.061)
Panel Length (months)			0.007***	0.006***	0.006***
			(0.002)	(0.002)	(0.002)
Mean Transfers (100s)				0.034***	0.035***
				(0.010)	(0.009)
Any Livestock					-0.077
					(0.180)
Intercept	0.302**	0.034	-0.012	-0.361	-0.322
	(0.125)	(0.248)	(0.244)	(0.216)	(0.251)
Observations	25	25	25	25	25

Notes: This table shows a study-level regression of whether a study has an average asset transition equation with multiple steady states on study-level characteristics. 'Number of HH (100s)' is the number of households in the study sample in hundreds, 'Panel Rounds' is the number of survey rounds in the study, 'Panel Length (months)' is the length of time between baseline and the final endline survey in months, 'Mean Transfers (100s)' is the average transfer size in hundreds of 2020 USD, and 'Any Livestock' is an indicator for whether the study collects data on livestock ownership. Stars denote significance at the 10%, 5%, and 1% levels respectively.

Table A.20: Proportion of Monotonic Optimal Transfer Policies

Ability Decile	100 USD (1)	500 USD (2)	1000 USD (3)
1	95.8	100.0	100.0
2	100.0	100.0	100.0
3	100.0	100.0	100.0
4	95.8	91.7	100.0
5	100.0	91.7	100.0
6	95.8	91.7	100.0
7	91.7	95.8	91.7
8	95.8	95.8	95.8
9	95.8	95.8	91.7
10	100.0	95.8	91.7
All	97.1	95.8	97.1

Notes: This table shows the proportion of optimal transfer policies that are monotonically decreasing in assets across productivity deciles for three budget levels, \$100, \$500, and \$1000 per head, results average across all studies.

Table A.21: Targeting and Welfare Effects of Transfer Programs

	100 USD			1000 USD		
Program	Untargeted (1)	Consumption Targeted (2)	Optimal Transfer (3)	Untargeted (4)	Consumption Targeted (5)	Optimal Transfer (6)
Bangladesh, Balboni (2022)	18.2	19.4	19.4	157.7	155.8	160.9
Ethiopia, Blattman (2018)	17.3	17.4	17.5	160.5	160.8	162.0
Uganda, Gupta (2024)	18.2	20.2	20.2	154.7	149.7	154.7
Philippines, Edmonds (2020)	17.9	21.0	21.0	140.4	140.0	140.7
Kenya, Gobin (2017)	17.5	20.7	20.7	120.5	121.9	122.2
Peru, Banerjee (2015)	17.3	20.6	20.6	117.3	119.5	119.8
Uganda, Blattman (2019)	15.5	17.4	18.1	117.3	117.4	117.5
India, Banerjee (2021)	16.8	20.9	21.3	109.4	109.9	111.1
Mexico, Gertler (2012)	15.4	17.5	17.8	112.5	113.6	113.9
Uganda, Sedlmayr (2020)	15.7	16.7	20.1	100.6	88.1	101.7
Somalia, Abdullahi (2023)	14.8	17.5	18.7	99.7	99.8	100.1
Mexico, Cunha (2019)	15.1	19.0	19.6	94.8	96.4	96.7
Honduras, Banerjee (2015)	15.9	21.3	22.9	91.0	92.1	92.5
Nicaragua, Macours (2012)	12.8	15.0	15.4	87.2	88.4	88.7
Ghana, Banerjee (2015)	15.2	20.3	22.5	77.8	78.5	79.1
Kenya, Egger (2022)	13.3	16.4	17.0	80.5	81.0	81.4
Mali, Beaman (2023)	14.1	19.5	19.6	75.9	77.2	77.3
Tanzania, Baird (2024)	14.1	19.2	19.6	73.3	74.3	74.7
India, Banerjee (2015)	12.7	17.9	18.7	63.6	64.5	64.8
Ethiopia, Banerjee (2015)	12.1	15.9	17.2	54.0	54.5	55.1
Ghana, Fafchamps (2014)	11.3	15.6	15.7	54.4	55.2	55.3
India, Hussam (2022)	11.3	15.7	16.4	51.6	52.3	52.5
Rwanda, Mcintosh (2020)	10.7	14.7	14.8	49.5	50.3	50.4
Kenya, Haushofer (2016)	9.3	12.2	12.7	33.1	33.3	33.5
All Programs	14.7	18.0	23.2	94.9	94.8	164.5

Notes: This table presents the welfare effects of three transfer targeting rules: the optimal transfer policy, a flat transfer policy, and a current-consumption targeting policy. Welfare effects are measured relative to the value function under a baseline of no transfers, equivalent to the percentage change in the present discounted value of future consumption streams. '100 USD' indicates a budget of 100 USD per head, '1000 USD' indicates a budget of 1000 USD per head.

Table A.22: Difference in Policy Targeting Ranks

	Average Absolute Rank Difference		
Budget (USD/per capita)	Consumption Targeted	Rank(Assets)	Rank(Consumption)
100	12.2	14.1	15.9
500	8.0	3.6	7.9
1000	6.6	2.4	6.6

Notes: This table shows the average difference in targeting ranks between the optimal transfer policy and alternative targeting policies, across different budget levels.

Table A.23: Best Transfer Policy by Program

	10 USD		100 USD		500 USD		1000 USD	
Program	Best Policy (1)	Welfare Increase (%) (2)	Best Policy (3)	Welfare Increase (%) (4)	Best Policy (5)	Welfare Increase (%) (6)	Best Policy (7)	Welfare Increase (%) (8)
Somalia, Abdullahi (2023)	Safety net	0.4	One-off	4.3	One-off	15.9	One-off	24.6
Tanzania, Baird (2024)	-	-	One-off	87.2	-	-	-	-
Bangladesh, Balboni (2022)	One-off	22.4	One-off	75.0	One-off	94.6	One-off	97.6
Ethiopia, Banerjee (2015)	One-off	7.7	One-off	47.1	-	-	-	-
Ghana, Banerjee (2015)	One-off	9.8	One-off	48.7	One-off	82.4	One-off	90.9
Honduras, Banerjee (2015)	One-off	1.3	-	-	One-off	36.7	One-off	52.3
India, Banerjee (2015)	One-off	50.3	One-off	88.6	-	-	-	-
Peru, Banerjee (2015)	Safety net	2.2	Safety net	11.4	Safety net	24.4	One-off	36.7
India, Banerjee (2021)	One-off	3.5	One-off	18.4	One-off	43.1	One-off	55.5
Mali, Beaman (2023)	One-off	26.6	One-off	70.1	One-off	89.4	One-off	94.3
Ethiopia, Blattman (2018)	One-off	1.3	One-off	6.6	-	-	-	-
Uganda, Blattman (2019)	Safety net	6.6	Safety net	7.5	Safety net	10.9	Safety net	14.4
Nigeria, Carneiro (2021)	-	-	One-off	11.6	One-off	16.8	-	-
Mexico, Cunha (2019)	Safety net	1.3	Safety net	12.2	-	-	-	-
Kenya, Egger (2022)	Safety net	6.8	Safety net	9.1	Safety net	17.9	Safety net	24.1
Mexico, Gertler (2012)	One-off	2.8	One-off	17.0	One-off	40.4	One-off	52.6
Uganda, Gupta (2024)	One-off	2.3	One-off	13.8	One-off	34.7	One-off	45.8
Kenya, Haushofer (2016)	One-off	5.1	One-off	34.9	One-off	68.9	-	-
Nicaragua, Macours (2012)	One-off	3.2	One-off	22.3	One-off	54.0	One-off	69.3
Rwanda, Mcintosh (2020)	One-off	9.9	One-off	43.0	One-off	73.9	One-off	84.0
Uganda, Sedlmayr (2020)	Safety net	1.4	One-off	7.7	One-off	24.9	One-off	36.0

Notes: This table shows the optimal transfer policy between three types: 'One-off' is a one-time transfer at the start of the program, 'Safety Net' is a guaranteed income floor, chosen to match the budget constraint in expectation, and 'Stream' is a constant transfer every period. The welfare increase shows the percentage increase in consumption-equivalent welfare compared to no transfers.

B Appendix Figures

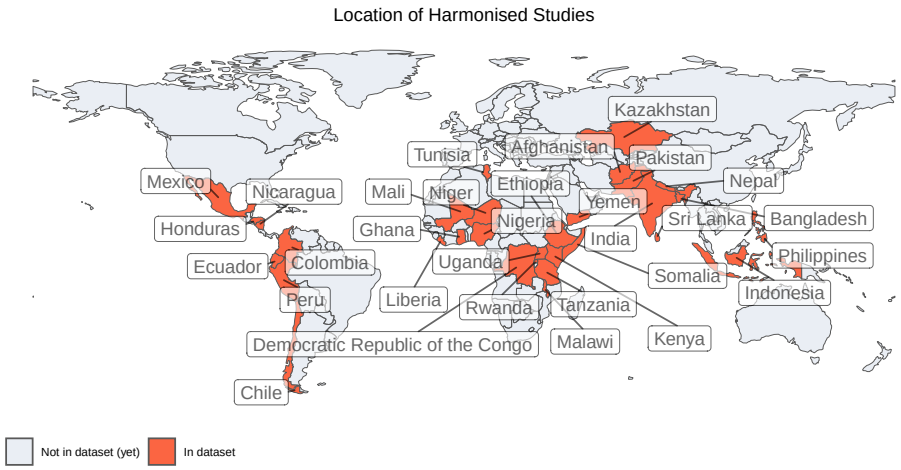


Figure B.1: Map of Included Studies – Wider Clean Sample

Notes: This figure shows the map of the wider harmonized sample of 58 RCTs.

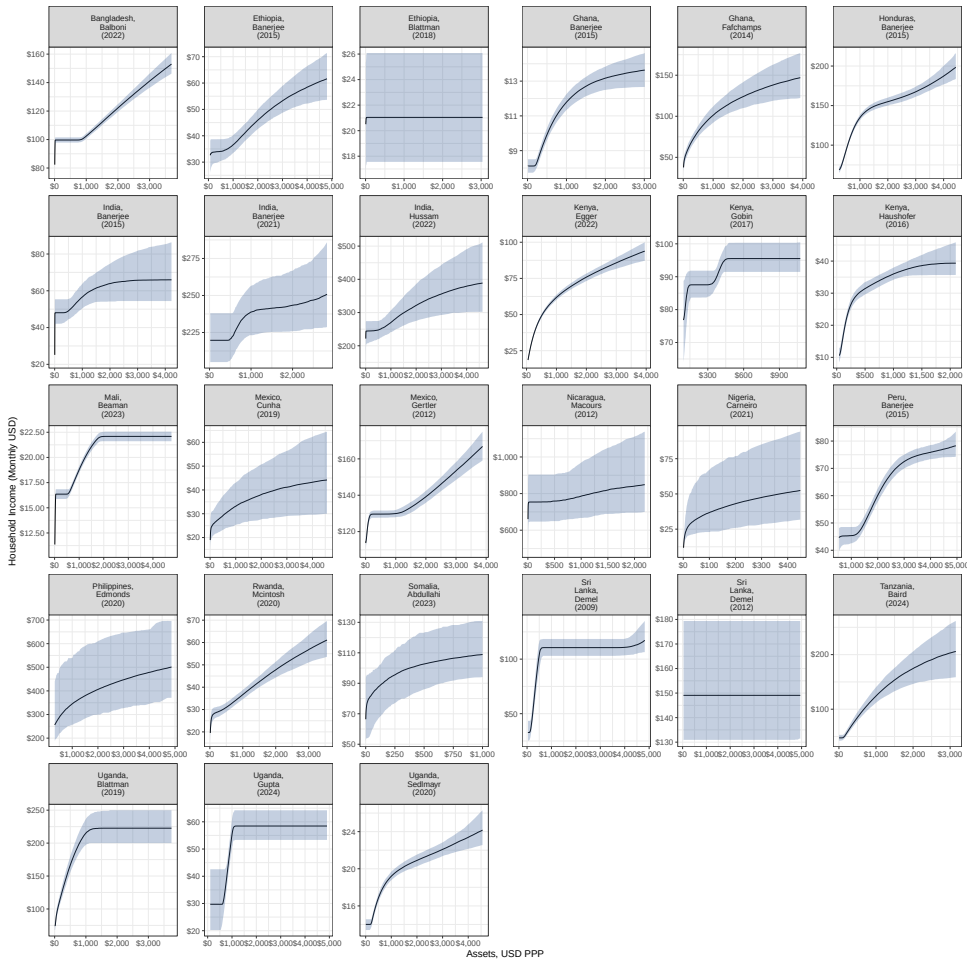


Figure B.2: Production Function Estimates

Notes: This figure shows the estimated production function estimates using the quasi-Bayesian estimator described in Section 6. The blue line shows the posterior mean production function, while the blue shaded area shows the 95% credibility interval. Each panel shows results for a different study.

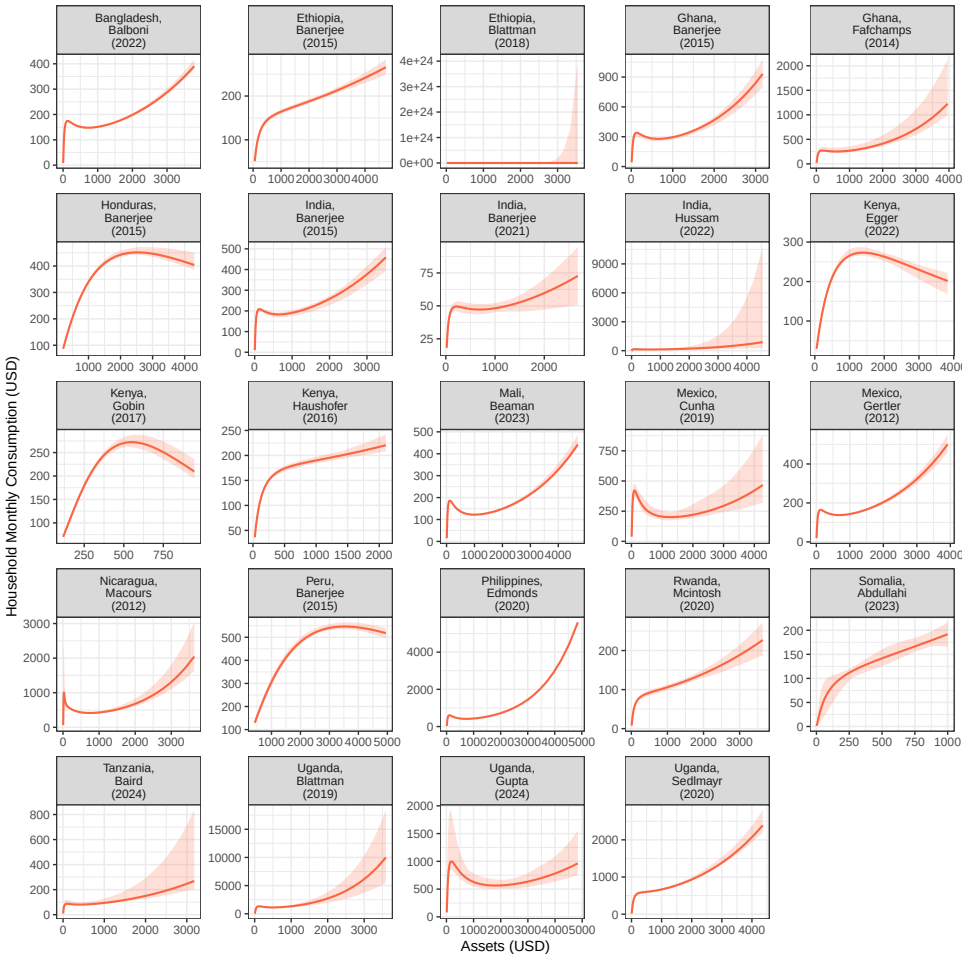


Figure B.3: Average Consumption Function Estimates

Notes: This figure shows average consumption function estimates using the stochastic EM algorithm described in Section 6.2. The blue line shows the posterior mean consumption function, while the blue shaded area shows the 95% confidence interval. Each panel shows results for a different study. Confidence intervals are calculated using a non-parametric bootstrap of households with 50 bootstrap samples.

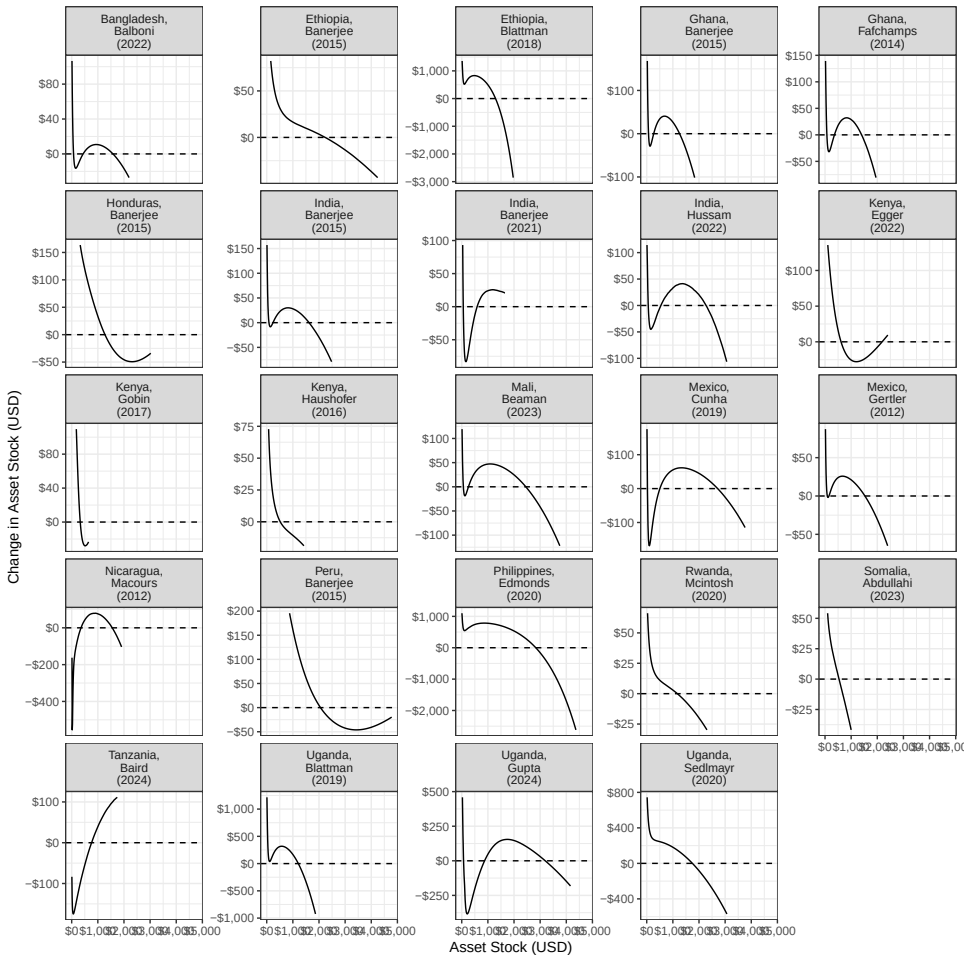


Figure B.4: Average Asset Transition Equation Estimates

Notes: This figure shows average asset transition equation estimates using the quasi-Bayesian estimator for the production function and the stochastic EM algorithm described in Section 6. The black line shows the expected asset transition equation. Each panel shows results for a different study.

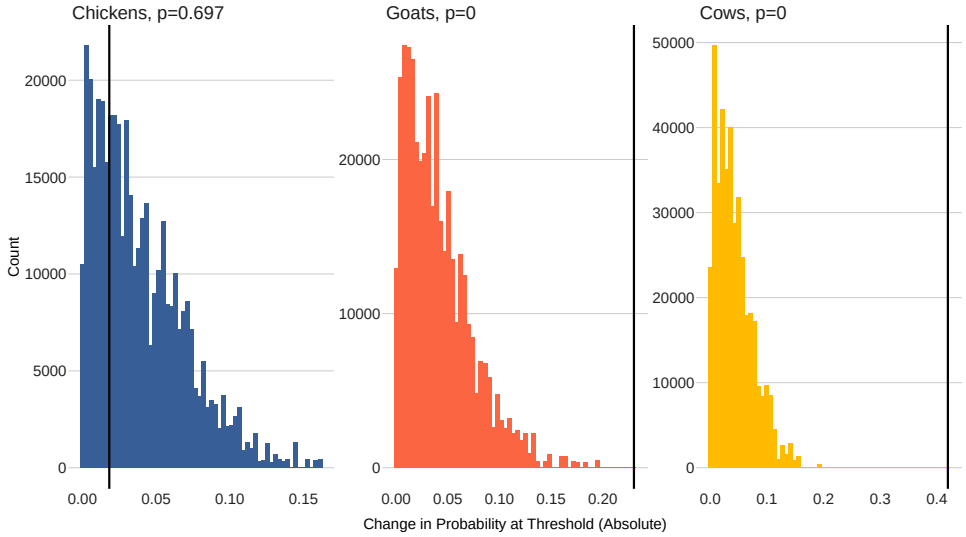


Figure B.5: Livestock Ownership Placebo Estimates

Notes: This figure shows a histogram of the estimated differences in livestock ownership at the household poverty trap threshold for placebo thresholds randomly drawn from the distribution of household assets. The vertical black line shows the estimated difference at the true household poverty trap threshold. I plot the absolute value of the point estimates and compare how extreme the true estimate is relative to the placebo distribution to calculate a p-value. I draw distance to the placebo threshold from a uniform distribution centered at zero with bounds of −5000 and 5000 in USD across 1000 simulations.

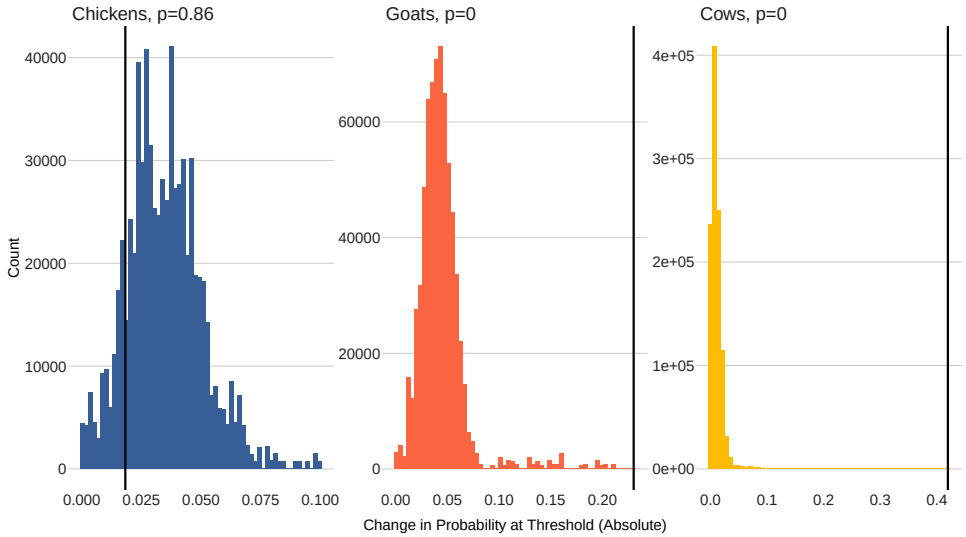


Figure B.6: Livestock Ownership Randomization Inference

Notes: This figure shows a histogram of the estimated differences in livestock ownership at the household poverty trap threshold for random reassignments of the distance to the threshold, within a program, I randomly reassign the distance to threshold for each household. The vertical black line shows the estimated difference at the true household poverty trap threshold. I plot the absolute value of the point estimates and compare how extreme the true estimate is relative to the randomization distribution to calculate a p-value. I draw 1000 random reassignments of the distance to threshold.

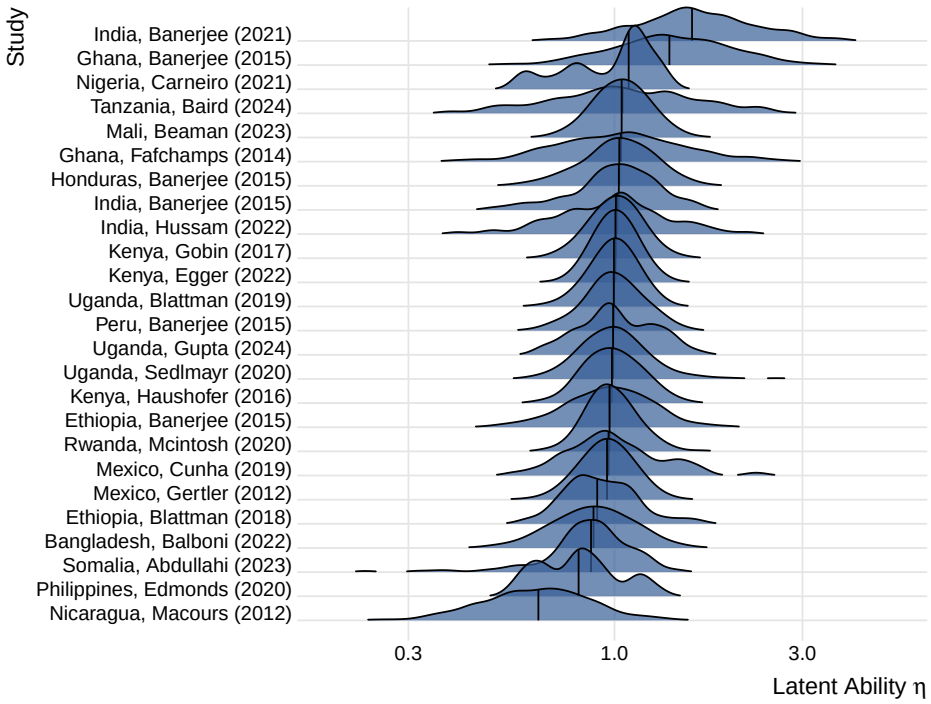


Figure B.7: Density of Household Ability (Fixed Productivity) by Study

Notes: This figure shows the estimated density of household ability, η, for each study. I plot the e exponential of latent ability, so that the x-axis shows multiplicative productivity differences. The vertical black line shows the median ability level in each study. Results are from the posterior mean of the ability fixed effect estimated using the stochastic EM algorithm described in Section 6.

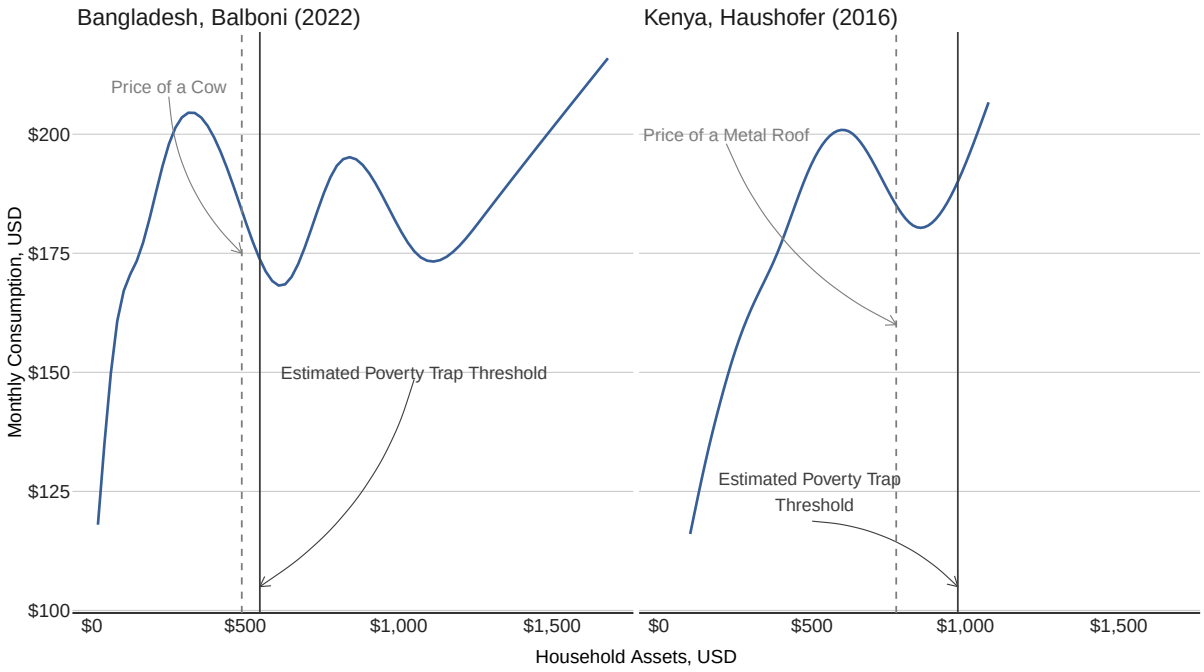


Figure B.8: Consumption Decreases Near Estimated Poverty Trap Thresholds in Bangladesh, Balboni et al. (2022) and Kenya, Haushofer & Shapiro (2016)

Notes: This figure shows a nonparametric regression of household monthly consumption on assets in Balboni et al. (2022) and Haushofer and Shapiro (2016) respectively. The price of the fixed-cost asset is taken directly from the original papers: \$488 for a cow in Bangladesh and \$774 for a metal roof in Kenya (after adjusting for inflation from 2012 PPP reported in Haushofer and Shapiro (2016) to 2020 USD). The vertical dashed line shows the estimated poverty trap threshold from the nonparametric dynamic panel model in Bangladesh (\$547) and the structural model in Kenya (\$974). I resort to the structural model estimates in Kenya since I do not find any households with multiple steady states using the dynamic panel model in Kenya, where households marginally avoid crossing the y-axis. The structural model results for Bangladesh (\$468) are similar to the nonparametric results and are omitted for visual clarity.

C Additional Results

C.1 Correlates of Traps

Next, I explore whether the prevalence of poverty traps and their mechanisms, both at the household level and on aggregate at the study level, are correlated with observable characteristics.

First, I find that larger households are more likely to face multiple steady states but little evidence that gender, age, or education predict the existence of a threshold. Since multiple steady states arise when households are neither too productive nor too unproductive, this pattern is consistent with larger households being able to pool resources and afford larger fixed costs, whereas smaller households may be too unproductive to ever face multiple steady states. I find little evidence that the threshold size is related to household size, age, or gender composition of the household, and significant but economically small effects of education (-6.52, $p = 0.0395$), Appendix Table A.15.

Next, amongst households facing a threshold, I find that trapped households - i.e. those that face multiple steady states and are below the threshold $\mathbf{1}_{\text{Trapped},j}(k, \eta) = 1$, have a greater share of women in the households, are younger, and less educated, Appendix Table A.16. Having below median education is associated with an 6.5 percentage point increase in the probability of being trapped ($p = 0.020$), while a greater share of women in the household is associated with an 8.3 percentage point increase in the probability of being trapped ($p = 0.02$). Interpreting these results requires caution, given the small sample size – only 1,491 households across 10 studies face a threshold and report covariates.

Unfortunately, the small sample size of 25 studies limits my ability to assess the impact of study-level characteristics, shown in Appendix Tables (A.17)-(A.19). I find the presence of livestock has a large but imprecise effect on the probability of observing a production kink (+22.6ppt, $p = 0.306$). Non-monotonicities in consumption are more common in studies with more survey waves, (+13.3ppt, $p = 0.055$) perhaps reflecting the greater ability to detect nonlinearities as the panel dimension increases. I find studies with longer follow-up periods and larger transfers are more likely to exhibit multiple steady states but these effects are small – a one month increase in the final endline is associated with a 0.6% increase in the probability of finding multiple steady states ($p = 0.003$), similarly a $100 increase in the average transfer is associated with a 3.5% increase ($p = 0.002$).

C.2 Dispersion in Ability and Productivity Persistence

Finally, I explore dispersion in latent ability, η, and the persistence of productivity shocks, ρ.

In Appendix Figure B.7 I plot the estimated distribution of time-invariant latent ability, η, for each study. While I do find substantial differences in the productivity of house-

holds, my estimates of productivity dispersion, shown in Appendix Table A.13 are slightly smaller than those found previously in the literature, including more recent papers which account for measurement error, such as (Gollin and Udry 2021; Bils et al. 2021; Maue et al. 2020). For instance, pooling across all studies, I estimate a difference in log productivity of 0.58 between the 10th and 90th percentiles, compared to estimates of 1.25-1.9 in the existing literature. My estimates of productivity dispersion are closer to, but still below, those of the United States in Hsieh and Klenow (2009), who estimate a value of 1.19 for US manufacturing firms in 2005. It is possible these lower estimates reflect the fact that households in my sample face kinked production functions, and accounting for these nonlinearities leads to lower dispersion in productivity residuals.[28] However, many other differences could explain these results, such as differences in data quality or the unit considered. For instance, my sample consists of household surveys of assets, rather than detailed firm-level production data, and typically only a single microentrepreneur or farmer, rather than a firm employing multiple workers.

Turning to productivity persistence, Appendix Table A.14 shows the annualized estimated persistence of the idiosyncratic productivity shock, ρ, for each study. I find an average annualized persistence of 0.299 with a range from 0.007 to 0.655 – while below values typically found in high-income countries (Syverson 2011), these are in line with estimates focusing on agricultural productivity in Africa (Maue et al. 2020).

[28] Maue et al. (2020) use a random forest estimator to measure productivity on agricultural plots, however due to shorter panel lengths, are unable to use dynamic panel methods and must regress inputs on outputs directly.

D Robustness Checks

D.1 Alternative Production Function Estimators

In this section, I present robustness checks using alternative production function estimators. First, I re-estimate the production function using a version of the quasi-Bayesian estimator that fixes the Gaussian process hyperparameters after a burn-in period rather than estimating them simultaneously with the production function parameters. Second, I use remove the hierarchical prior over the length scale parameters. Third, I present results using only studies that have longer panels of at least 3 time periods and incorporate measurement error in the production function equation. Finally, I re-estimate the model without the level moments, which require assumptions about initial conditions and the relationship between household ability and changes in income/assets.

I find very little evidence of differences in the estimated production functions across these alternative specifications. In most cases, the posterior mean is virtually identical to the preferred specification in the main text, with occasional differences in the width of the credible intervals. Figures (D.9)-(D.12) present the production function estimates for each alternative specification in turn.

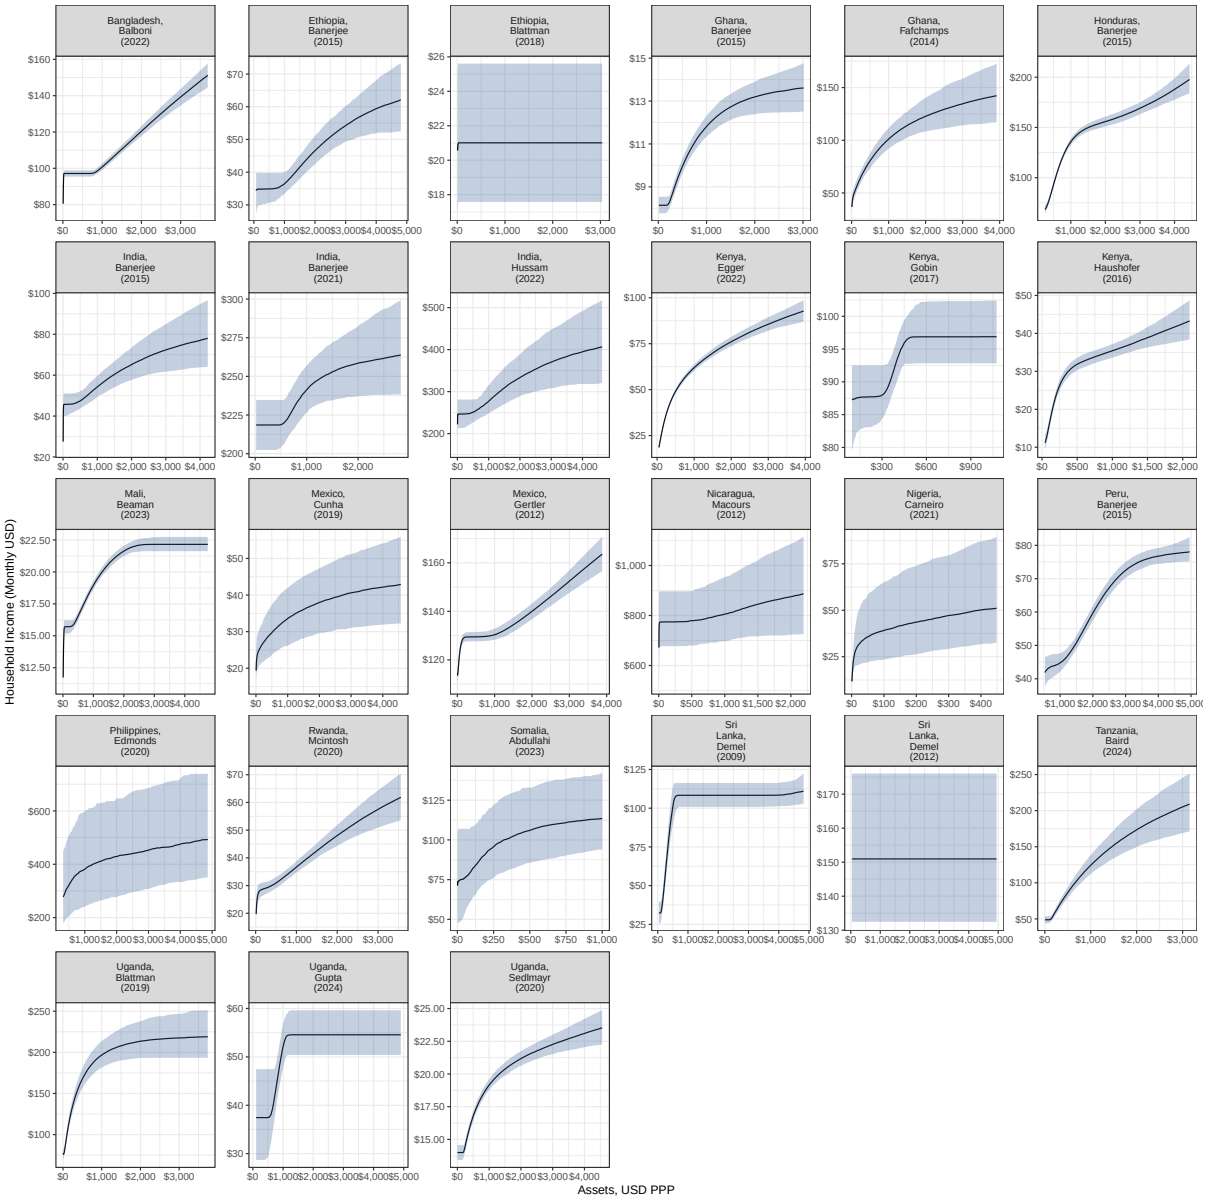


Figure D.9: Production Function Estimates: Fixed Gaussian Process Hyperparameters

Notes: This figure shows production function estimates using the quasi-Bayesian estimator that fixes the Gaussian process hyperparameters after a burn-in period, rather than estimating them simultaneously with the production function parameters. The blue line shows the posterior mean production function, while the blue shaded area shows the 95% credible interval. Each panel shows results for a different study.

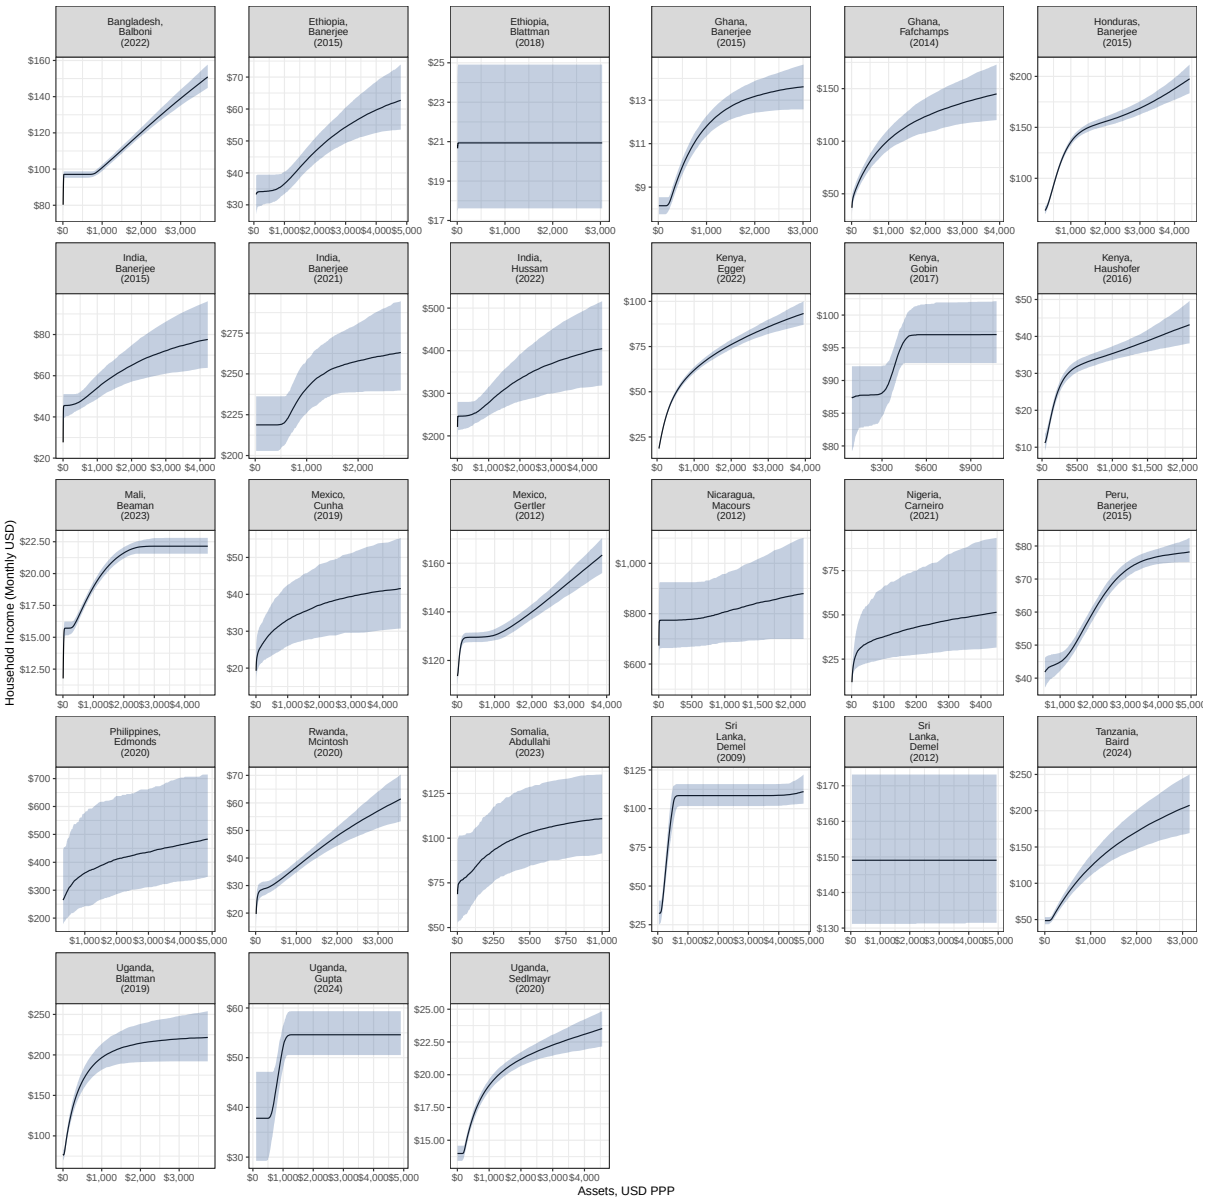


Figure D.10: Production Function Estimates: No Hierarchical Model

Notes: This figure shows production function estimates using the same quasi-Bayesian estimator as in the main text, but removes the hierarchical prior over the length scale parameter and estimates each length scale independently. The blue line shows the posterior mean production function, while the blue shaded area shows the 95% credible interval. Each panel shows results for a different study.

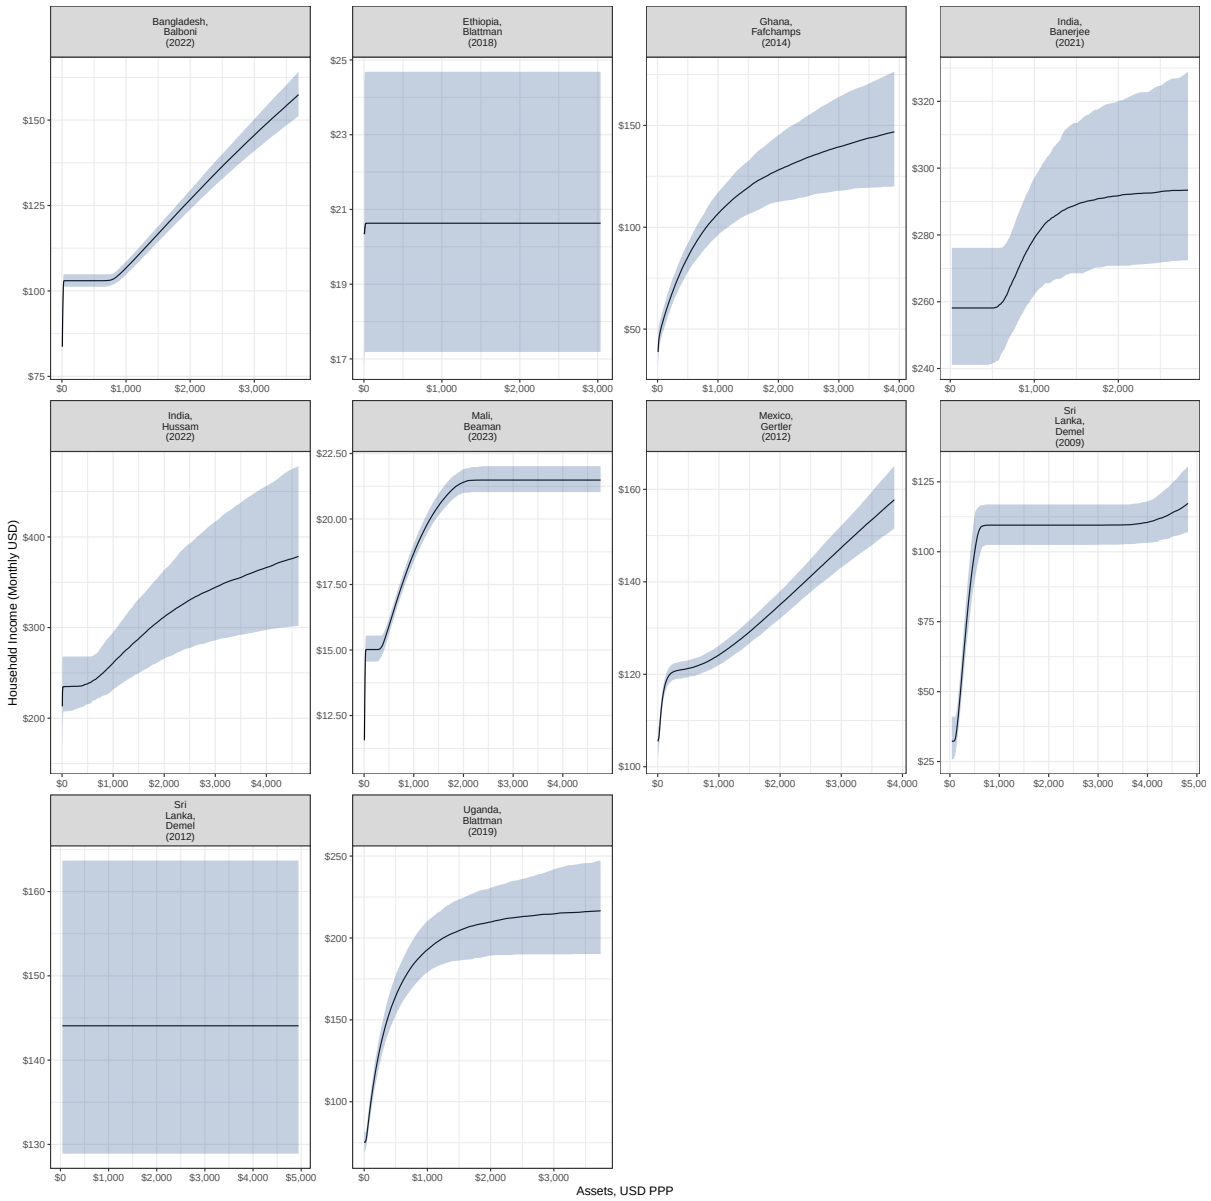


Figure D.11: Production Function Estimates: Dynamic Panel with Measurement Error

Notes: This figure shows production function estimates using the quasi-Bayesian estimator that incorporates measurement error in the production function equation. The blue line shows the posterior mean production function, while the blue shaded area shows the 95% credible interval. Each panel shows results for a different study.

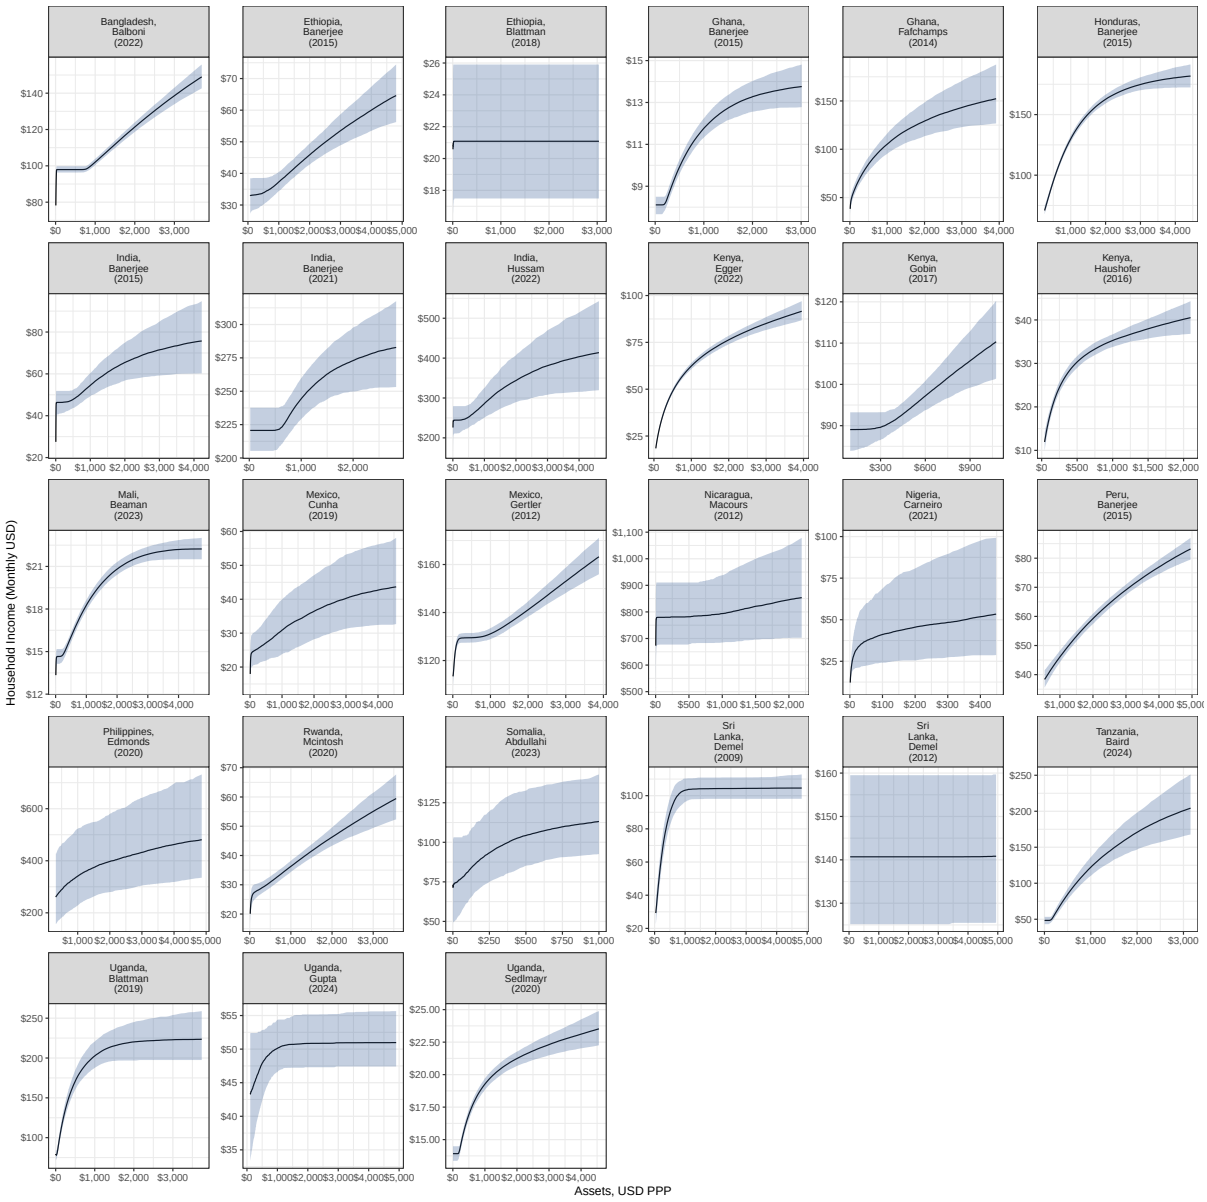


Figure D.12: Production Function Estimates: No Level Moments

Notes: This figure shows production function estimates using the quasi-Bayesian estimator that removes the level moments from the model. The blue line shows the posterior mean production function, while the blue shaded area shows the 95% credible interval. Each panel shows results for a different study.

D.1.1 Alternative Consumption Function Estimators

In this section, I present results that re-estimate the consumption function by directly using estimates of households' ability and productivity shocks from the production function step, rather than integrating over their distribution using the stochastic EM algorithm.

Figure D.13 presents the average consumption function results. I find very similar patterns to the main text specification, with some minor discrepancies in the number of studies with non-monotonic consumption functions. In both the main text and alternative specification, I find 9 studies with monotonic consumption functions. Overall, the results suggest that the main findings are robust to this alternative estimation approach.

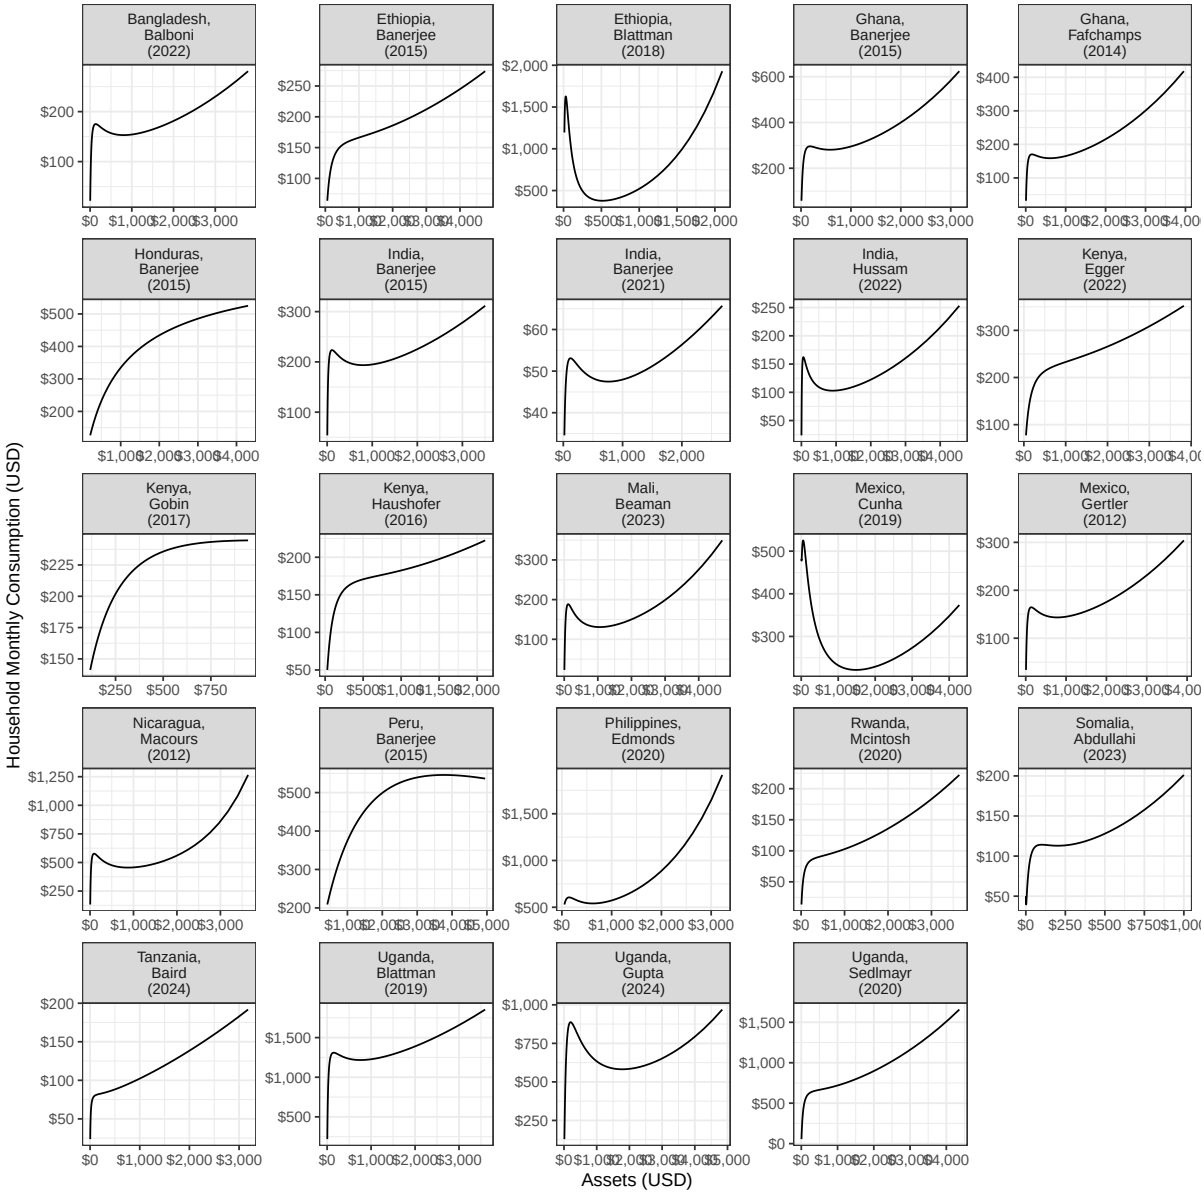


Figure D.13: Average Consumption Function Estimates: Fixed Effects Approach

Notes: This figure shows average consumption function estimates where household ability and productivity shocks are first estimated using a fixed effects approach and the production function residuals, and then consumption is regressed on these estimates and assets. The black line shows the average consumption function. Each panel shows results for a different study.

D.1.2 Poverty Trap Estimates

Finally, I present poverty trap estimates using the alternative production and consumption estimation approaches outlined above.

Table D.24: Poverty Trap Thresholds: Fixed GP Hyperparameters

Program	Threshold (USD PPP) (1)	Distance Below Threshold (USD PPP) (2)	$\frac{\text{Threshold}}{\text{Consumption}}$ (3)	Transfer (USD PPP) (4)	HHs Below Threshold (%) (5)	HHs Trapped Below Threshold (%) (6)
Ethiopia, Blattman (2018)	159	154	3	1,030	19.23	13.46
Ghana, Banerjee (2015)	229	141	7	406	38.25	3.75
Mali, Beaman (2023)	248	201	11	140	24.50	24.50
Ghana, Fafchamps (2014)	302	193	6	120	54.25	21.00
Mexico, Gertler (2012)	307	251	12	379	59.25	32.50
India, Banerjee (2015)	399	328	9	788	61.75	48.50
Kenya, Gobin (2017)	420	296	12	238	81.00	2.50
Uganda, Sedlmayr (2020)	429	268	4	155	39.75	0.25
Nicaragua, Macours (2012)	491	436	9	485	82.00	71.25
Rwanda, Mcintosh (2020)	522	316	24	328	51.75	0.00
Bangladesh, Balboni (2022)	555	431	12	542	54.00	14.50
Tanzania, Baird (2024)	601	493	40	350	75.00	30.75
Uganda, Blattman (2019)	669	548	3	400	80.00	79.50
Mexico, Cunha (2019)	711	614	11	21	78.38	78.38
India, Banerjee (2021)	753	657	55	547	77.00	23.75
India, Hussam (2022)	780	562	15	100	79.19	34.23
Uganda, Gupta (2024)	842	736	9	2,967	55.13	55.13
Kenya, Egger (2022)	2,436	1,774	49	1,871	96.25	25.00
Peru, Banerjee (2015)	3,795	1,950	33	1,433	74.50	73.50
Honduras, Banerjee (2015)	4,321	3,268	81	754	99.00	0.25
Somalia, Abdullahi (2023)	-	-	-	500	-	-
Ethiopia, Banerjee (2015)	-	-	-	338	-	-
Nigeria, Carneiro (2021)	-	-	-	528	-	-
Philippines, Edmonds (2020)	-	-	-	518	-	-
Kenya, Haushofer (2016)	-	-	-	404	-	-
Overall	539	433	10	338	55.13	21.00

Notes: This table presents poverty trap thresholds estimated using the quasi-Bayesian production function estimator that fixes the Gaussian process hyperparameters after a burn-in period, rather than estimating them simultaneously with the production function parameters.

Table D.25: Poverty Trap Thresholds: No Hierarchical Model

Program	Threshold (USD PPP) (1)	Distance Below Threshold (USD PPP) (2)	$\frac{\text{Threshold}}{\text{Consumption}}$ (3)	Transfer (USD PPP) (4)	HHs Below Threshold (%) (5)	HHs Trapped Below Threshold (%) (6)
Somalia, Abdullahi (2023)	186	179	10	500	14.29	0.71
Ghana, Banerjee (2015)	226	139	7	406	37.75	14.50
Mali, Beaman (2023)	239	193	11	140	24.25	24.25
Mexico, Gertler (2012)	304	248	12	379	59.00	32.50
Philippines, Edmonds (2020)	336	319	3	518	39.58	29.17
Ethiopia, Blattman (2018)	384	371	8	1,030	28.85	26.92
Ghana, Fafchamps (2014)	389	253	7	120	63.25	48.50
India, Banerjee (2015)	390	321	9	788	61.00	47.75
Kenya, Gobin (2017)	427	303	12	238	81.00	1.75
Tanzania, Baird (2024)	504	415	33	350	73.00	31.00
Nicaragua, Macours (2012)	507	451	9	485	82.25	73.25
Uganda, Blattman (2019)	524	422	3	400	75.50	75.50
Mexico, Cunha (2019)	534	451	8	21	73.87	73.87
Rwanda, Mcintosh (2020)	571	346	27	328	55.50	1.50
Bangladesh, Balboni (2022)	606	461	13	542	56.00	16.50
Uganda, Gupta (2024)	646	577	7	2,967	48.72	48.72
India, Banerjee (2021)	768	668	56	547	77.25	25.75
India, Hussam (2022)	791	571	16	100	79.19	31.21
Peru, Banerjee (2015)	3,826	1,962	33	1,433	75.50	74.25
Ethiopia, Banerjee (2015)	-	-	-	338	-	-
Honduras, Banerjee (2015)	-	-	-	754	-	-
Nigeria, Carneiro (2021)	-	-	-	528	-	-
Kenya, Egger (2022)	-	-	-	1,871	-	-
Kenya, Haushofer (2016)	-	-	-	404	-	-
Uganda, Sedlmayr (2020)	-	-	-	155	-	-
Overall	504	371	9	338	55.50	25.75

Notes: This table presents poverty trap thresholds estimated using the quasi-Bayesian production function estimator without the hierarchical prior over the Gaussian process parameters.

Table D.26: Poverty Trap Thresholds: Measurement Error

Program	Threshold (USD PPP) (1)	Distance Below Threshold (USD PPP) (2)	$\frac{\text{Threshold}}{\text{Consumption}}$ (3)	Transfer (USD PPP) (4)	HHs Below Threshold (%) (5)	HHs Trapped Below Threshold (%) (6)
Ethiopia, Blattman (2018)	93	90	2	1,030	13.46	13.46
Uganda, Blattman (2019)	183	144	1	400	51.00	2.50
Mali, Beaman (2023)	228	184	10	140	30.75	17.00
Mexico, Gertler (2012)	283	230	11	379	61.00	27.50
Bangladesh, Balboni (2022)	354	267	7	542	53.75	3.75
India, Banerjee (2021)	611	540	45	547	66.75	16.75
India, Hussam (2022)	642	453	13	100	73.49	33.89
Ghana, Fafchamps (2014)	-	-	-	120	-	-
Overall	283	230	5	338	52.37	15.11

Notes: This table presents poverty trap thresholds estimated using the quasi-Bayesian production function estimator restricting the sample to studies with at least 3 time periods and incorporating measurement error in the production function equation.

Table D.27: Poverty Trap Thresholds: Consumption Fixed Effects

Program	Threshold (USD PPP) (1)	Distance Below Threshold (USD PPP) (2)	$\frac{\text{Threshold}}{\text{Consumption}}$ (3)	Transfer (USD PPP) (4)	HHs Below Threshold (%) (5)	HHs Trapped Below Threshold (%) (6)
Philippines, Edmonds (2020)	142	140	1	518	29.17	8.33
Uganda, Sedlmayr (2020)	178	114	2	155	15.50	0.25
Ethiopia, Blattman (2018)	304	293	6	1,030	26.92	5.77
Ghana, Banerjee (2015)	404	256	12	406	54.50	9.50
Mali, Beaman (2023)	429	312	19	140	34.00	31.50
India, Banerjee (2021)	429	387	31	547	53.50	10.50
Kenya, Gobin (2017)	436	311	13	238	81.75	6.50
India, Banerjee (2015)	455	374	10	788	66.50	34.00
Mexico, Gertler (2012)	472	397	19	379	67.50	2.25
Uganda, Blattman (2019)	472	379	2	400	73.75	26.25
Ghana, Fafchamps (2014)	480	322	9	120	67.00	26.00
Bangladesh, Balboni (2022)	655	492	14	542	57.50	5.75
Kenya, Egger (2022)	659	352	13	1,871	54.50	3.50
Nicaragua, Macours (2012)	724	653	13	485	85.25	49.00
Mexico, Cunha (2019)	744	643	11	21	78.38	62.16
Tanzania, Baird (2024)	768	624	51	350	82.75	29.75
India, Hussam (2022)	818	594	16	100	80.20	26.51
Uganda, Gupta (2024)	942	820	10	2,967	58.97	52.56
Honduras, Banerjee (2015)	4,074	3,033	76	754	98.75	5.00
Peru, Banerjee (2015)	4,266	2,149	37	1,433	86.00	9.25
Somalia, Abdullahi (2023)	-	-	-	500	-	-
Ethiopia, Banerjee (2015)	-	-	-	338	-	-
Nigeria, Carneiro (2021)	-	-	-	528	-	-
Kenya, Haushofer (2016)	-	-	-	404	-	-
Rwanda, Mcintosh (2020)	-	-	-	328	-	-
Overall	476	383	8	338	66.75	10.00

Notes: This table presents poverty trap thresholds estimated using the consumption function estimated via a fixed effects approach that first estimates household ability and productivity shocks from the production function residuals, and then regresses consumption on these estimates and assets.

Table D.28: Poverty Trap Thresholds: No Level Moments

Program	Threshold (USD PPP) (1)	Distance Below Threshold (USD PPP) (2)	$\frac{\text{Threshold}}{\text{Consumption}}$ (3)	Transfer (USD PPP) (4)	HHs Below Threshold (%) (5)	HHs Trapped Below Threshold (%) (6)
Uganda, Blattman (2019)	191	149	1	400	52.25	12.25
Ghana, Banerjee (2015)	254	157	8	406	40.25	0.00
Mali, Beaman (2023)	282	231	13	140	24.75	24.75
Ghana, Fafchamps (2014)	284	179	5	120	53.25	19.75
Mexico, Gertler (2012)	294	240	12	379	58.00	33.50
Philippines, Edmonds (2020)	308	294	2	518	35.42	4.17
India, Banerjee (2015)	412	340	9	788	62.50	49.00
Nicaragua, Macours (2012)	502	446	9	485	82.25	65.00
Rwanda, Mcintosh (2020)	522	316	24	328	51.75	0.00
Uganda, Sedlmayr (2020)	528	338	5	155	49.50	0.25
Tanzania, Baird (2024)	573	471	38	350	75.00	28.50
Bangladesh, Balboni (2022)	611	464	13	542	56.00	15.25
Ethiopia, Blattman (2018)	613	575	12	1,030	50.00	48.08
Mexico, Cunha (2019)	620	531	9	21	75.68	73.87
Uganda, Gupta (2024)	665	594	7	2,967	49.36	48.72
India, Hussam (2022)	765	551	15	100	78.86	36.24
Kenya, Gobin (2017)	834	667	24	238	98.75	66.75
India, Banerjee (2021)	892	781	65	547	81.00	30.00
Peru, Banerjee (2015)	3,836	1,965	33	1,433	75.75	74.00
Somalia, Abdullahi (2023)	-	-	-	500	-	-
Ethiopia, Banerjee (2015)	-	-	-	338	-	-
Honduras, Banerjee (2015)	-	-	-	754	-	-
Nigeria, Carneiro (2021)	-	-	-	528	-	-
Kenya, Egger (2022)	-	-	-	1,871	-	-
Kenya, Haushofer (2016)	-	-	-	404	-	-
Overall	528	446	9	338	51.75	19.75

Notes: This table presents poverty trap thresholds estimated using the quasi-Bayesian production function estimator removing the level moments from the model. Consumption is estimated using the stochastic EM algorithm as in the main text.

E Identification

E.1 Production Function

Below I lay out the identification assumptions for the production function, noting that this problem is a simpler version of the models studied in Su and Lu (2013); Lee (2007).

I drop the study level subscript j and focus on identification in a single study. Consider the dynamic panel model:

$$y_{it} = \rho\, y_{i,t-1} + f(k_{it}) - \rho f(k_{i,t-1}) + (1-\rho)\eta_i + \xi_{it}, \tag{47}$$

where η_i denotes a time invariant individual effect and ξ_{it} is an idiosyncratic innovation. The goal is to identify $\{\rho, f(\cdot), \sigma_\xi^2, \sigma_\eta^2\}$ from panel data $\{y_i^T, k_i^T, F_i^T\}$, where $x_i^t = (x_{i1}, \ldots, x_{it})$ denotes the history of variable x for individual i up to time t. y, k, and F denote log income, log assets, and randomized transfers, respectively.

Assumption 1 (Sampling and stability) *(i) Individuals $i = 1, \ldots, N$ are independent and identically distributed. (ii) The time dimension $T \geq 3$ is fixed while $N \to \infty$. (iii) The process is stationary with finite second moments: $|\rho| < 1$, $E[\xi_{it}] = 0$, and $E[\xi_{it}^2] < \infty$.*

Intuition. These conditions ensure well defined population moments and a stationary AR(1) component, as is standard in dynamic panel models.

Assumption 2 (Timing and shocks) *The innovations $\{\xi_{it}\}$ are serially uncorrelated and mean independent of past outcomes, past/current assets, and unobserved heterogeneity:*

$$E[\xi_{it} \mid y_i^{t-1}, k_i^t, \eta_i] = 0.$$

Assets are predetermined, i.e. chosen at $t-1$.

Intuition. Sequential exogeneity ensures that current shocks do not affect past choices, so lagged outcomes and other predetermined variables can serve as valid instruments.

Assumption 3 (Initial conditions) *The individual effects are mean independent of the initial changes in income and changes in assets across t:*

$$E[(1-\rho)\eta_i \mid \Delta y_{i2}] = 0, \tag{48}$$

$$E[(1-\rho)\eta_i \mid \Delta k_{it}] = 0. \tag{49}$$

Intuition. The initial change in income and subsequent changes in assets do not provide information about the expected value of the unobserved individual effect η_i.

Assumption 4 (Normalization) *The individual effects have mean zero:*

$$E[\eta_i] = 0.$$

Intuition. Because $(1-\rho)\eta_i$ and the level of $f(\cdot)$ enter additively the "intercept" of $f(\cdot)$ cannot be separately identified. Normalizing $E[\eta_i] = 0$ fixes the location of the model and pins down the level of $f(\cdot)$.

Assumption 5 (Randomized transfers/instruments) *There exist excluded variables F_{it} (randomized transfers F_{it}) that are independent of ability and shocks and affect outcomes only through assets:*

$$E[\xi_{it} \mid F_i^t] = 0, \qquad F_i^t \perp (\eta_i, \xi_{it}).$$

Intuition. Randomized transfers provide valid instruments for k_{it} and $k_{i,t-1}$.

Assumption 6 (Support and variation in assets) *The joint distribution of $(k_{it}, k_{i,t-1})$ has a density and exhibits sufficient variation such that $\Pr(k_{it} \neq k_{i,t-1}) > 0$.*

Intuition. This condition ensures that there is enough independent variation in current and lagged assets to recover the function $f(\cdot)$ from the differenced equation via the operator $T_\rho[f](k_{it}, k_{i,t-1}) = f(k_{it}) - \rho f(k_{i,t-1})$ once the fixed effect is normalized, Assumption 4.

Assumption 7 (Completeness of instruments) *Let W_{it} denote the set of valid instruments (lagged differenced variables and randomized transfers). The conditional distribution of W_{it} given $(k_{it}, k_{i,t-1})$ is boundedly complete, so that*

$$E\left[y_{it} - \rho y_{i,t-1} - (1-\rho)\eta_i - T_\rho[f](k_{it}, k_{i,t-1}) \mid W_{it}\right] = 0 \quad \Rightarrow \quad f = f^\star.$$

Intuition. Bounded completeness ensures that the conditional moment restriction uniquely determines $f(\cdot)$ for a given value of ρ (Newey and Powell 2003; Blundell et al. 2007).

Moments for identification. Under Assumptions 1–7, the model generates the following identifying moments:

$$E\left[y_{it} - \rho y_{i,t-1} - T_\rho[f](k_{it}, k_{i,t-1}) \mid W_{it}\right] = 0. \tag{50}$$

The remaining moments all depend on the structural function $f(\cdot)$, however, I suppress this dependence for notational simplicity:

$$E[\xi_{it} - \xi_{i,t-1} \mid y_{i,t-s}] = 0, \quad s \geq 2, \tag{51}$$

$$E[\xi_{it} - \xi_{i,t-1} \mid k_{i,t-s}] = 0, \quad s \geq 1, \tag{52}$$

which use lagged levels as instruments for the differenced equations.

Variance moments for $(\sigma_\xi^2, \sigma_\eta^2)$*:*

$$E[((1-\rho)\eta_i + \xi_{it})^2] = (1-\rho)^2\sigma_\eta^2 + \sigma_\xi^2, \tag{53}$$

$$E[((1-\rho)\eta_i + \xi_{it})((1-\rho)\eta_i + \xi_{i,t-1})] = (1-\rho)^2\sigma_\eta^2. \tag{54}$$

Together with the normalization $E[\eta_i] = 0$, these moments identify $\{\rho,\ f(\cdot),\ \sigma_\xi^2,\ \sigma_\eta^2\}$.

Joint identification of ρ and $f(\cdot)$. The conditional moment for $f(\cdot)$ and the orthogonality conditions for ρ form a joint system because both depend on the productivity innovation, $\xi_{it} = y_{it} - \rho y_{i,t-1} - T_\rho[f](k_{it}, k_{i,t-1}) - (1-\rho)\eta_i$. Identification occurs simultaneously: for any candidate value of ρ, Assumption 7 ensures that the function $f(\cdot)$ satisfying the conditional moment (50) is unique up to normalization, and the true value of ρ is the one for which all population orthogonality conditions are satisfied. The variance components $(\sigma_\xi^2, \sigma_\eta^2)$ are then identified from standard GMM moments exploiting lagged instruments and the covariance structure of the residual.

E.2 Consumption Function - Large T Identification Argument

To build intuition, this section presents a simple identification argument for the consumption function for long panels, $T \to \infty$ panels - I discuss identification with T fixed in Section E.3.

With large T, the key is that I can identify the latent productivity components directly from the income process. Once I have identified the latent components, I can treat them as observed regressors in the consumption equation, which becomes a standard nonparametric regression problem.

I treat the production function parameters as known, since they can be identified under the assumptions above. The joint process is:

$$y_{it} - f(k_{it}) = \eta_i + \omega_{it} \tag{55}$$

$$e_{it} = \eta_i + \omega_{it} \tag{56}$$

$$\omega_{it} = \rho\omega_{i,t-1} + \xi_{it} \tag{57}$$

$$c_{it} = g(k_{it}, \eta_i, \omega_{it}) + u_{it} \tag{58}$$

$$k_{i,t+1} = y_{it} - c_{it} + (1-\delta)k_{it} + v_{it} \tag{59}$$

I term e_{it} residual income, recognizing that $y_{it} - f(k_{it})$ is income after accounting for returns to assets. The function $g(\cdot)$ is the unknown consumption function to be identified. The other variables and parameters are as defined in the main text, u_{it} and v_{it} are i.i.d. measurement error terms independent of latent variables and observed assets.

E.2.1 Identification of Latent Income Components

The key to my argument is that the income process in Equations (63) and (64) is "invertible" because ρ is known.

Identification of Latent Components. Given Equations (63) and (64), and a panel with $T \to \infty$, the latent components η_i and ω_{it} (for all t) are identified.

Focusing on the observable term $e_{it} - \rho e_{i,t-1}$:

$$\begin{aligned} e_{it} - \rho e_{i,t-1} &= (\eta_i + \omega_{it}) - \rho(\eta_i + \omega_{i,t-1}) \\ &= \eta_i - \rho\eta_i + \omega_{it} - \rho\omega_{i,t-1} \\ &= \eta_i(1-\rho) + (\omega_{it} - \rho\omega_{i,t-1}) \end{aligned}$$

Substituting the AR(1) process from Equation (64):

$$e_{it} - \rho e_{i,t-1} = \eta_i(1-\rho) + \xi_{it} \tag{60}$$

Let $\alpha_i = \eta_i(1-\rho)$, which is a time-invariant, household-specific term. Equation (60) becomes:

$$e_{it} - \rho e_{i,t-1} = \alpha_i + \xi_{it}$$

Since ξ_{it} is i.i.d. with $E[\xi_{it}] = 0$, I can identify α_i by averaging over time. As $T \to \infty$, the sample mean converges to the true value by the law of large numbers:

$$\hat{\alpha}_i = \frac{1}{T-1}\sum_{t=2}^{T}(e_{it} - \rho e_{i,t-1}) \xrightarrow{p} \alpha_i$$

Given that α_i is identified and $\rho \neq 1$ is known, I can recover the permanent component η_i:

$$\eta_i = \frac{\alpha_i}{1-\rho}$$

Therefore, η_i is identified for each household i.

Once η_i is identified, recovering ω_{it} is simple. From the income equation (63):

$$\omega_{it} = e_{it} - \eta_i$$

Since e_{it} is observed and η_i is now identified, ω_{it} is also identified for all i and t.

E.2.2 Identification of the Consumption Function $g(\cdot)$

With the latent income components η_i and ω_{it} identified, they can be treated as observable regressors in the consumption equation.

The model in (65) is:

$$c_{it} = g(k_{it}, \eta_i, \omega_{it}) + u_{it}$$

Where:

- c_{it} is observed.
- k_{it} is observed.
- η_i is identified.
- ω_{it} is identified.
- u_{it} is an i.i.d. measurement error with $E[u_{it}] = 0$, assumed independent of the regressors.

This setup is a standard nonparametric regression problem. The function g is identified as the conditional expectation of c_{it} given the full set of (now identified) regressors:

$$g(k, \eta, \omega) = E[c_{it} | k_{it} = k, \eta_i = \eta, \omega_{it} = \omega] \tag{61}$$

This conditional expectation is nonparametrically identified from the joint distribution of $(c_{it}, k_{it}, \eta_i, \omega_{it})$.

E.2.3 Summary

The simplicity of this argument hinges entirely on the residual income specification $e_{it} = \eta_i + \omega_{it}$ and the ability to recover η_i by averaging over time.

In short panels with fixed T, this averaging strategy fails due to the incidental parameters problem and identification must use knowledge of the full distribution of residual income over time, as I discuss next.

E.3 Consumption - Fixed-Panel Identification Argument

I restate the model for household i in a panel of finite length T. The observed variables are income (y_{it}), residual income (e_{it}), consumption (c_{it}), and assets (k_{it}).
The residual income process is:

$$e_{it} = y_{it} - f(k_{it}) \tag{62}$$

$$e_{it} = \eta_i + \omega_{it} \tag{63}$$

where η_i is a time-invariant fixed effect and ω_{it} follows a known AR(1) process:

$$\omega_{it} = \rho \omega_{it-1} + \xi_{it} \tag{64}$$

The consumption function $g(\cdot)$ is nonparametric:

$$c_{it} = g(k_{it}, \eta_i, \omega_{it}) + u_{it} \tag{65}$$

The asset accumulation process is:

$$k_{it+1} = y_{it} - c_{it} + (1-\delta)k_{it} + v_{it} \tag{66}$$

With T fixed and finite, I cannot identify η_i by averaging over time. This would introduce an estimation error ($\frac{1}{T-1}\sum \xi_{it}$) that does not vanish, leading to inconsistent estimates (the incidental parameters problem).

Instead, I treat η_i and ω_{it} as unobserved latent variables and rely on a nonparametric instrumental variables (NPIV) strategy, similar to Arellano, Blundell, and Bonhomme (ABB). The identification argument is sequential. I first identify the distributions of the latent income components and then use those results to identify the consumption function g.

E.3.1 Assumptions

Some of these assumptions overlap with those in the production function identification section, but I restate them here for completeness.

Assumption 8 (Latent Components) *Latent ability η_i and the innovations ξ_{it}^T are mutually independent. The innovations ξ_{it} are i.i.d. over i and t with $E[\xi_{it}] = 0$. The panel length $T \geq 3$. ω_{i0} is drawn from the stationary distribution of the AR(1) process in Equation (64).*

Assumption 9 (Error Exogeneity) *The error terms u_{it} and v_{it} are i.i.d. over i and t, with $E[u_{it}] = 0$ and $E[v_{it}] = 0$. They are independent of the initial conditions (k_{i1}, ω_{i0}) and all latent variables (η_i, ξ_{is}^T).*

Assumption 10 (Initial Conditions) *The initial asset stock k_{i1} is pre-determined. It is independent of all future innovations $\{\xi_{it}, u_{it}, v_{it}\}_{t=1}^T$. The joint distribution of initial assets and latent factors $(k_{i1}, \eta_i, \omega_{i0})$ is unrestricted.*

Assumption 11 (Gaussian Residual Income Process)

1. *The latent components are mutually independent and normally distributed: $\eta_i \sim N(0, \sigma_\eta^2)$ and $\xi_{it} \sim N(0, \sigma_\xi^2)$.*
2. *The parameters $(\rho, \sigma_\eta^2, \sigma_\xi^2)$ are known or identified from the production process.*

Assumption 12 (Completeness) *Let $W_i = (k_{i1}, E_i)$ be the vector of pre-determined and instrumental observables. For any function $g(\cdot)$ such that $E[|g(\eta_i, \omega_{i1})|] < \infty$, the condition*

$$E[g(\eta_i, \omega_{i1})|W_i] = 0 \quad \textit{almost surely}$$

implies that $g(\eta_i, \omega_{i1}) = 0$ almost surely.

E.4 Proof of Identification (for $t = 1$)

The identification strategy is sequential and follows arguments in ABB (Section 4.2).

Step 1: Identification of the Residual Income Process.

By Assumption 11, the distribution of latent residual income is multivariate normal and its conditional distribution has an analytical form.

Step 2: Identification of the Consumption Function $g(\cdot)$.

The consumption equation at $t = 1$:

$$c_{i1} = g(k_{i1}, \eta_i, \omega_{i1}) + u_{i1}$$

In this equation, c_{i1} and k_{i1} are observed. The terms η_i and ω_{i1} are latent. I can treat this as an NPIV problem where:

- Outcome: c_{i1}
- Observed Regressor: k_{i1}
- Latent/Endogenous Regressors: (η_i, ω_{i1})
- Instruments: $E_i = (e_{i2}, \ldots, e_{iT})$

The instruments E_i are valid because they are correlated with the latent productivities (η_i, ω_{i1}) (by construction of the production process) and are independent of the error term u_{i1} (by Assumption 9). Let $W_i = (k_{i1}, E_i)$ be the full vector of instruments. Take the conditional expectation of c_{i1} given W_i:

$$E[c_{i1}|W_i] = E[g(k_{i1}, \eta_i, \omega_{i1}) + u_{i1}|W_i]$$

By Assumption 9, $E[u_{i1}|W_i] = 0$. Using LIE:

$$E[c_{i1}|W_i] = E\left[E[g(k_{i1}, \eta_i, \omega_{i1})|W_i, \eta_i, \omega_{i1}]\middle|W_i\right]$$

$$E[c_{i1}|k_{i1}, E_i] = \iint g(k_{i1}, \eta, \omega) \cdot f(\eta, \omega|k_{i1}, E_i)\, d\eta\, d\omega$$

This is a classic NPIV problem.

- $E[c_{i1}|k_{i1}, E_i]$, is identified from the joint distribution of observables.
- $f(\eta, \omega|k_{i1}, E_i)$ is identified. By Bayes' rule, this density is proportional to $f(E_i|k_{i1}, \eta, \omega) f(\eta, \omega|k_{i1})$. The first term, $f(E_i|k_{i1}, \eta, \omega)$, simplifies to $f(E_i|\eta, \omega)$ due to the strict exogeneity of initial assets with respect to future income shocks (Assumption 10) and is identified from the residual income process in Step 1.
  The second term, $f(\eta, \omega|k_{i1})$, is identified following the argument in Arellano et al. (2017). I observe the conditional density of assets given the income history,

$f(k_{i1}|E_i)$. This observed density relates to the latent variables via the integral:

$$f(k_{i1}|E_i) = \iint f(k_{i1}|\eta,\omega) f(\eta,\omega|E_i)\, d\eta\, d\omega$$

Since $f(\eta,\omega|E_i)$ is identified (Step 1) and assumed to be boundedly complete, I can invert this integral operator to uniquely recover the unknown density $f(k_{i1}|\eta,\omega)$. Therefore, the joint density $f(\eta,\omega,k_{i1})$ and the required conditional density $f(\eta,\omega|k_{i1})$ are identified.

- The unknown function is $g(k_{i1},\cdot,\cdot)$ for a fixed value of k_{i1}.

By Assumption 12, the linear operator defined by this integral is invertible. Therefore, the function $g(k,\eta,\omega)$ is nonparametrically identified for all (k,η,ω) in the support of the data.

Step 3: Sequential Identification for $t > 1$.

Once g is identified, the function h and the parameter δ in the asset accumulation equation (66) can be identified at $t = 1$:

$$k_{i2} - (1-\delta)k_{i1} = y_{i1} - c_{i1} + v_{i1}$$

$$E[k_{i2} - (1-\delta)k_{i1} - y_{i1} + c_{i1}|Z_i] = 0$$

where Z_i includes instruments for c_{i1}. Once g is identified, the argument can proceed sequentially to $t = 2$, where k_{i2} is now an endogenous regressor. The logic mirrors the full sequential argument in ABB (Section 4.2).

F Estimation Details

F.1 Production Function

The production equation in the model is given by:

$$y_{ijt} = f(k_{ijt}) + \eta_{ij} + \omega_{ijt} \tag{67}$$

$$\omega_{ijt} = \rho\omega_{ij,t-1} + \xi_{ijt} \tag{68}$$

which implies:

$$w_{ijt} = y_{ijt} - \rho y_{ijt-1} - \left[f(k_{ijt}) - \rho f(k_{ij,t-1})\right] = (1-\rho)\eta_{ij} + \xi_{ijt} \tag{69}$$

$$\Delta w_{ijt} = \xi_{ijt} - \xi_{ij,t-1} \tag{70}$$

To form the conditional mean of the residual function, as described in Kankanala (2025), I use a Hilbert Space Gaussian Process (HSGP) approximation to a Gaussian process prior over the function $f(\cdot)$. The HSGP approximation uses a basis function expansion to approximate the Gaussian process reducing the computational burden of standard Gaussian process regression.

F.1.1 Basis Function Approximation

I use 50 basis functions, which I generate in the probabilistic programming language Stan (Carpenter et al. 2017) using code adapted from Riutort-Mayol et al. (2022)[29] The covariance kernel, $k(x, x')$, is approximated as:

$$k(x, x') \approx \sum_{m=1}^{50} s_\theta\left(\sqrt{\lambda_m}\right)\phi_m(x)\phi_m(x') \tag{71}$$

$$\lambda_m = \left(\frac{m\pi}{2L}\right)^2 \tag{72}$$

$$\phi_m(x) = \sqrt{\frac{1}{L}}\sin\left(\sqrt{\lambda_m}(x+L)\right) \tag{73}$$

$$L = 2 \times \underbrace{1.2\left(\max(x) - \min(x)\right)}_{\text{rule of thumb}} \tag{74}$$

λ_j and ϕ_j are all computed in the transformed data block of Stan as they do not vary across iterations. A key tuning parameter is L which is the "boundary condition", I set this to twice the rule of thumb value recommended by Riutort-Mayol et al. (2023) and find this works well. At this point, I also standardize the input data by its study-specific mean and standard deviation.

The Gaussian process is then generated using a specific prior over the basis function

[29] I find little difference between 50 and 100 basis functions, below 30 the functions appear over-smoothed.

coefficients:

$$f_j(x) \approx \sum_{m=1}^{50} \left(s_\theta \left(\sqrt{\lambda_m} \right) \right)^{\frac{1}{2}} \phi_m(x) \tilde{\beta}_m \tag{75}$$

$$\tilde{\beta}_m \sim N(0,1) \tag{76}$$

$$\tag{77}$$

In practice, I generate $\tilde{\beta}$ and then scale it by $s_\theta(\cdot)$ which is the spectral density of eigenvalues for the 3/2 Matérn covariance function, again using code adapted from Riutort-Mayol et al. (2023).

$$\beta_m = \tilde{\beta}_m \times \left(s_\theta \left(\sqrt{\lambda_m} \right) \right)^{\frac{1}{2}} \tag{78}$$

$$f_j(x) \approx \sum_{m=1}^{50} \phi_m(x) \beta_m \tag{79}$$

Using this I generate three sets of GPs, using the precomputed basis functions:

$$f_j(k_{ijt}) = \sum_{m=1}^{50} \phi_m(k_{ijt}) \beta_{j,m} \tag{80}$$

$$f_j(k_{ijt-1}) = \sum_{m=1}^{50} \phi_m(k_{ijt-1}) \beta_{j,m} \tag{81}$$

$$f_j(k_{ijt-2}) = \sum_{m=1}^{50} \phi_m(k_{ijt-2}) \beta_{j,m} \tag{82}$$

To improve sampling efficiency, I find it helps to centre the prior mean of the GP on the line $y = x$ – this is effectively using a prior against poverty traps, as this reflects a constant returns to scale production function:

$$\tilde{f}_j(k_{ijt}) = \frac{f_j(k_{ijt})}{\sqrt{D}} + k_{ijt}$$

where $D = 10$ is the dimension of the first stage basis matrix, generated next.

F.1.2 First Stage Estimation

To flexibly model the first stage, I generate a matrix of Bernstein polynomials of degree 10 for each instrument + their respective interaction (so $k_{ijt-1} \times y_{ijt-2}$, $\Delta y_{ijt-2} \times \Delta k_{ijt-2}$, etc.). I use R's 'splines2' package which automatically handles scaling to the unit interval. The transfer instrument enters directly that is, I don't take any expansions of it. An observation i at time t therefore has the following basis vector, for $d = 10$ and $m = 3$

instruments, z_{it}:

$$\phi_{ijt} = \begin{bmatrix} b\left(z_{ijt}^{(1)}\right) \\ b\left(z_{ijt}^{(2)}\right) \\ \vdots \\ b\left(z_{ijt}^{(m)}\right) \\ F_{ijt} \end{bmatrix} \in \mathbb{R}^{(d+1)\cdot m+1}$$

Stacking each observation's polynomial basis into a single matrix of polynomials we have B^{diff} and B^{level} for the difference and level instruments + the transfer instrument respectively:

$$B = \begin{bmatrix} \phi_1^\top \\ \phi_2^\top \\ \vdots \\ \phi_n^\top \end{bmatrix} \in \mathbb{R}^{n\times((d+1)\cdot m+1)}$$

F.1.3 Conditional Mean and Objective Function

With the GPs in hand, I can now estimate the residual functions, r_{ijt}, and corresponding conditional mean functions, $\hat{m}$.

$$\Delta w_{ijt} = \left(y_{ijt} - \rho_j y_{ijt-1} - \left[\tilde{f}_j(k_{ijt}) - \rho_j \tilde{f}_j(k_{ijt-1})\right]\right) - \left(y_{ij,t-1} - \rho_j y_{ij,t-2} - \left[\tilde{f}_j(k_{ij,t-1}) - \rho_j \tilde{f}_j(k_{ij,t-2})\right]\right)$$
$$w_{ijt} = y_{ijt} - \rho_j y_{ijt-1} - \left[\tilde{f}_j(k_{ijt}) - \rho_j \tilde{f}_j(k_{ijt-1})\right]$$

If the panel is too short to estimate the differenced residual for a given household, I eliminate the moment condition for that household.

Next, I project these residual vectors onto the first stage basis matrices, B^{diff} and B^{level}, to get the conditional mean functions:

$$\beta^{\text{diff}} = \left(B^{\text{level}\top} B^{\text{level}}\right)^{-1} B^{\text{level}\top} \Delta w_{ijt} \tag{83}$$
$$\beta^{\text{level}} = \left(B^{\text{diff}\top} B^{\text{diff}}\right)^{-1} B^{\text{diff}\top} w_{ijt} \tag{84}$$
$$\hat{m}^{\text{diff}} = B^{\text{level}} \beta^{\text{diff}} \tag{85}$$
$$\hat{m}^{\text{level}} = B^{\text{diff}} \beta^{\text{level}} \tag{86}$$

This gives two N length vectors of conditional mean functions in addition to the moment vectors for the variances:

$$\hat{M} = \begin{bmatrix} \hat{m}^{\text{diff}} & \hat{m}^{\text{level}} & w_{ijt} w_{ijt} & w_{ijt} w_{ijt-1} \end{bmatrix}$$

To generate the objective function, I stack everything into a single vector, effectively using the identity weighting matrix, and compute:

$$\hat{g} = vec(\hat{M}) \tag{87}$$

$$Q = \frac{1}{N}\hat{g}'\hat{g} \tag{88}$$

$$\log \text{target} = -\frac{N}{2} \times Q + \log \text{prior density} \tag{89}$$

where priors are discussed below.

F.1.4 Priors

The length scale parameters have hierarchical priors:

$$\begin{aligned}
\ell_j \sim N(\mu_\ell, \sigma_\ell^2), \sigma_{\eta_j} \sim N^+\left(\mu_{\sigma_{\eta_j}}, \sigma^2_{\sigma_{\eta_j}}\right) \\
\sigma_{\xi_j} \sim N^+(\mu_{\sigma_{\xi_j}}, \sigma^2_{\sigma_{\xi_j}}) \\
\rho_j \sim N(\mu_\rho, \sigma_\rho^2)
\end{aligned}$$

With standard normal hyperpriors:

$$\begin{aligned}
\mu_\ell, \mu_{\sigma_\eta}, \mu_{\sigma_\xi} &\sim N^+(0,1) \\
\mu_\rho &\sim N(0,1) \\
\sigma_\ell, \mu_{\sigma_\eta}\sigma_{\sigma_\xi}, \sigma_\rho &\sim N^+(0,1)
\end{aligned}$$

F.1.5 Sampling and Output

I sample from the posterior using Hamiltonian Monte Carlo in Stan. I run 4 chains for 800 iterations each – the most I can run before hitting the server's time limit. I discard the first 400 iterations as warmup. I check convergence using the $\hat{R}$ statistic and effective sample sizes – all of which indicate good mixing and convergence across estimated results.

To generate posterior estimates of the production function, I report the predicted GP function evaluated at a grid of 100 points across the support of assets in each study, from the 1st to the 99th percentile.

F.2 Consumption and Asset Accumulation Equations

I use a tensor product of fourth order polynomial basis functions, $\varphi(\cdot)$, to flexibly estimate the consumption function, $g_j(\cdot)$:

$$g_j(k_{ijt}, \eta_{ij}, \omega_{ijt}) = \sum_{l=0}^{4}\sum_{m=0}^{4}\sum_{n=0}^{4} \beta_{j,lmn}\varphi_l\left(k_{ijt}\right)\varphi_m\left(\eta_{ij}\right)\varphi_n\left(\omega_{ijt}\right) \tag{90}$$

where I restrict the tensor-product basis to unique combinations of polynomial terms by imposing $l \geq m \geq n$, to remove duplicate interactions.

The initial asset level and asset transition is given by:

$$k_{ij0} = b_{0j} + b_{1j}\eta_{ij} + b_{2j}\omega_{ijt} + \nu_{ij0} \tag{91}$$

$$\exp(k_{ijt+1}) = \exp\left(f(k_{ijt}) + \eta_{ij} + \omega_{ijt} - g_j(k_{ijt}, \eta_{ij}, \omega_{ijt}) - \delta_{0j} + \nu_{ijt}\right) \tag{92}$$

This gives the parameter vector: $\theta_j = \left\{\beta_j, \sigma_u, \sigma_\nu, \sigma_e, \delta_{0j}, b_{0j}, b_{1j}, b_{2j}\right\}$ and $\mu = \left\{\sigma_{\ell j}, \sigma_{\eta j}, \sigma_{\xi j}, \rho_j\right\}$ are the production parameters estimated in the first stage. σ_e is a small amount of noise for the income equation to improve particle filter performance, alternatively, in the specification with measurement error, this is the measurement error standard deviation.

F.2.1 Likelihood

The likelihood is given by:

$$\begin{aligned}
&f(y_{ij}^{T_i}, c_{ij}^{T_i}, k_{ij}^{T_i}, \omega_{ij}^{T_i}, \eta_{ij} | F_{ij}^{T_i}; \mu, \theta) \\
&= \prod_{t\in T_i} f(y_{ijt}|k_{ijt}, \eta_{ij}, \omega_{ijt}, F_{ijt}; \mu) \\
&\times \prod_{t\in T_i} f(c_{ijt}|k_{ijt}, \eta_{ij}, \omega_{ijt}, F_{ijt}; \mu, \theta) \\
&\times \prod_{t\in T_i, t>t_i} f(\omega_{ijt}|\omega_{ijt-1}; \mu) \\
&\times f(\omega_{ijt_i}|\mu) \times f(\eta_{ij}|\mu) \\
&\times \prod_{t\in T_i, t>t_i} f(k_{ijt}|k_{ijt-1}, y_{ijt-1}, c_{ijt-1}, F_{ijt-1})
\end{aligned}$$

In practice, I fix the production parameters, μ, to their posterior mean estimates from the first stage and the production function, $f_j(\cdot)$, to its posterior mean function – estimating the consumption likelihood conditional on these estimates.[30]

[30] In future work, I plan to jointly estimate the production and consumption parameters however the production model currently uses Hamiltonian Monte Carlo in Stan and cannot be adapted to jointly fit the consumption model.

I initialize the consumption parameters, θ_j, using naive fixed effects OLS estimates of the consumption equation which directly estimate the latent states.

Next, I draw an initial η_{ij}^* for each household from its prior distribution.

F.2.2 Particle Marginal Metropolis-Hastings (PMMH) Sampler

Then, for a given household i in study j, I use a particle filter with $S = 50$ particles to target the posterior distribution: $p(\eta_{ij}, \omega_{ijt}, ..., \omega_{ijT}|x_{ij}^T)$ where $x_{ij}^T = \{y_{ij}^T, c_{ij}^T, k_{ij}^T\}$ are the observed data for household i in study j up to time T.

I use an auxiliary particle filter. The particle filter implementation is a guided Sequential Importance Resampling (SIR) filter designed for a state-space model with AR(1) dynamics.

The filter is initialized at $t = 0$ by drawing S particles $\{\omega_{i0}^s\}$ from the prior distribution $p(\omega_{0j})$ and performing an initial systematic resampling based on the prior log-densities. For each subsequent time step $t = 1...T$, the algorithm uses a guided proposal distribution, $\pi(\omega_{ijt}|\omega_{ijt-1}^s, e_{ijt}, \eta_{ij}^*)$, which incorporates the current observation e_{ijt} to steer particles toward regions of high likelihood.

This proposal is calculated analytically under a linear Gaussian assumption. Specifically, it computes the proposal mean μ_{ijt} and variance σ_{ijt}^2 using a one-step Kalman filter update, combining the AR(1) prior $E[\omega_{ijt}|\omega_{ijt-1}^s] = \rho\omega_{ijt-1}^s$ with the observation e_{ijt}, the resulting proposal variance is also inflated by a constant factor c. New particles $\{\omega_{ijt}^s\}$ are then sampled from this guided proposal, $N(\mu_{ijt}, \sigma_{ijt}^2)$. The incremental log importance weights are subsequently calculated using the standard correction:

$$\log w_{ijt}^s = \log p(x_{ijt}|\omega_{ijt}^s, \eta_{ij}^*) + \log p(\omega_{ijt}^s|\omega_{ijt-1}^s, \eta_{ij}^*) - \log \pi(\omega_{ijt}^s|\omega_{ijt-1}^s, e_{ijt}^s, \eta_{ij}^*)$$

Finally, the marginal log-likelihood is updated, and the particle set is resampled using systematic resampling before proceeding to the next time step.

The filter returns unbiased estimates of the marginal log-likelihood, $\log p(\eta_{ij}^*)$.

F.2.3 Metropolis-Hastings Update

With the particle filter in place, I use a Metropolis-Hastings (MH) step to update the latent fixed productivity, η_{ij}. Given the current draw η_{ij}^*, I propose a new draw η_{ij}' from a normal random walk proposal centered at η_{ij}^* with variance σ_{prop}^2.

I then run the particle filter to obtain an unbiased estimate of the marginal log-likelihood, $\log p(\eta_{ij}')$ and proposal particles $\{\omega_{ijt}'\}$. The MH acceptance probability is computed as:

$$\alpha = \min\left(1, \exp\left(\log p(\eta_{ij}') - \log p(\eta_{ij}^*) + \log p(\eta_{ij}') - \log p(\eta_{ij}^*)\right)\right)$$

If the proposal is accepted, I set $\eta_{ij}^* = \eta_{ij}', \omega_{ijt}^* = \omega_{ijt}'$; otherwise, I retain the current draw.

F.2.4 Stochastic EM Algorithm

I repeat the MH update for all households in study j for a total of 50 iterations, collecting the final draws of η_{ij}^* and ω_{ijt}^*, this constitutes the stochastic E-step of the algorithm.

Then, in the M-step I regress consumption and assets on the respective latent states using the tensor-product basis functions to update the consumption parameters, θ_j. For example, the consumption update is given by:

$$\begin{aligned} c_{ijt} &= g_j(k_{ijt}, \eta_{ij}^*, \omega_{ijt}^*) + \nu_{ijt} \\ \Rightarrow \theta_j &= \left(\Phi_j^\top \Phi_j\right)^{-1} \Phi_j^\top c_{ijt} \end{aligned}$$

where Φ_j is the matrix of tensor-product basis functions evaluated at $(k_{ijt}, \eta_{ij}^*, \omega_{ijt}^*)$ for all households and time periods in study j.

I repeat the E and M steps for 5000 total iterations, discarding the first 2500 as burn-in. After a warm-up period of 10 iterations, I use only 3 iterations of the MH update per E-step to speed up computation where the initial states are initialized at their previous draws.

G Monte-Carlo Simulations

This section provides three Monte Carlo exercises evaluating the proposed estimation strategy. I first examine estimator performance for the production function, then for the consumption function, and finally for the asset transition function.

G.1 Production Monte-Carlos

To evaluate the performance of the proposed estimation strategy, I conduct a series of Monte-Carlo simulations. I begin by calibrating a data-generating process (DGP) using the estimated parameters from the household growth model in Section 8.1.

G.1.1 Production Data Generating Process

Specifically, I use parameters from the household growth model I calibrated against the estimated production and consumption functions using data from Balboni et al. (2022). The calibrated growth model generates a Markov transition matrix, $\hat{P}$, for a (300×10×10) grid of assets, ability, and idiosyncratic productivity shocks. I use this transition matrix to simulate forwards the joint distribution of variables for a year, assuming the initial distribution is uniform over the grid. This joint distribution, $\tilde{\mu}$, is used to sample initial conditions for the Monte-Carlo simulations.

Therefore, for each simulation draw, s:

1. Sample initial assets, ability, and idiosyncratic productivity shocks from $\tilde{\mu}$ to form N observations $x^s_{i,0} = (k^s_{i0}, \eta^s_i, \omega^s_{i0})$.[31]
2. Simulate forward the initial conditions using the policy rules, budget constraint, and stochastic process for a warmup period of 10 months, generating $x^s_{i,10}$.
3. Simulate forward T periods of data, generating $\{x^s_{i,t}, y^s_{i,t}, c^s_{i,t}\}_{t=1}^T$ for $i = 1, \ldots, N$ households.

I vary the number of households, $N \in \{100, 500, 1000, 2000\}$ and the number of time periods, $T \in \{3, 5\}$, to evaluate the performance of the estimation procedure. I repeat this process for $S = 50$ simulation draws.

G.1.2 Estimation and Evaluation

I estimate the production function in each simulated dataset using the quasi-Bayesian procedure described in Section 6. For computational tractability, I run a single Markov

[31] I sample abilities and productivity shocks uniformly from between grid cells, to avoid bunching/unrealistic discretization in the simulated dataset. This weakens the initial condition correlation, which is why I simulate a warmup period in the next step.

chain with 2000 iterations and discard the first 1000 as burn-in. Unlike the empirical application, I treat the simulated data as coming from a single study and therefore omit the hierarchical prior on the persistence parameter.

To evaluate performance, I construct a grid of 100 evenly spaced points between the 5th and 95th percentiles of the simulated asset distribution. For each simulation s, I then interpolate the posterior mean production function, $\hat{f}^s(\cdot)$, and compute the metrics using deviations from the true log-production function, $f(\cdot)$, that generated the data.

G.1.3 Simulation Results

Figure G.14 shows true versus predicted production for 5000 random draws from the monte-carlo simulations using the calibrated DGP with $N = 2000$ households and $T = 5$ time periods. The predicted income closely matches the true income, indicating that the estimator accurately recovers the production function. Similarly, Table G.29 shows the performance metrics of the production function estimator on the simulated dataset. The RMSE decreases as the panel length and sample size increase, with a value of 0.742 in the smallest sample size and 0.28 in the very largest. The posterior credibility intervals are overly conservative, with coverage rates above 95% in all settings. The mean absolute error and mean absolute percentage error similarly decrease with larger sample sizes in most settings.

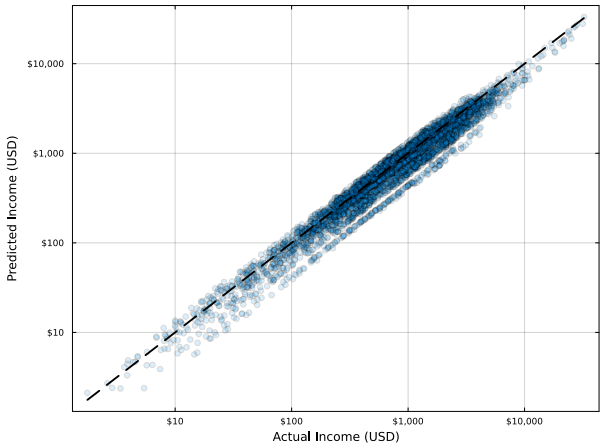


Figure G.14: True vs Predicted Production: Calibrated DGP

Notes: This figure shows 5000 random samples of true versus predicted income from the calibrated household production function simulations. Results use the DGP $N = 2000$ households and $T = 5$ time periods. The blue points show the predicted versus true log income for each household-period observation in the simulated dataset. The black dashed line shows the 45-degree reference line.

Table G.29: Log Production Function Simulation Metrics: Calibrated

T	N	RMSE	MAE	MAPE	Coverage
3	100	0.742	0.66	10.2	0.952
3	500	0.356	0.317	4.96	0.999
3	1000	0.412	0.393	6.12	0.993
3	2000	0.388	0.378	5.88	0.99
5	100	0.551	0.501	7.76	0.993
5	500	0.38	0.356	5.54	0.998
5	1000	0.36	0.344	5.34	0.988
5	2000	0.28	0.265	4.11	1

Notes: This table shows the performance metrics of the production function estimator on simulated data generated using the structural model calibrated against an example study. Each figure reports the average result over 50 simulated datasets where the integrated error is calculated over a grid of 100 points for each simulation draw. 'RMSE' is the root mean squared error, 'MAE' is the mean absolute error, 'MAPE' is the mean absolute percentage error, and 'Coverage' is the 95% confidence interval coverage.

G.2 Consumption Monte Carlos

To illustrate the importance of accommodating non-linearities in the consumption function, I construct three data generating processes (DGPs). The first two are designed to closely mimic the empirical environment, one with a standard, monotonic production technology and one with a non-monotonic "butterfly" technology that induces poverty trap dynamics. The third DGP provides a computationally simpler benchmark.

G.2.1 Data Generating Processes

The first DGP ("RCK") simulates the full structural model from Section 8 using a Cobb-Douglas production function, with variances and technology parameters chosen to approximately match the empirical distribution of assets, income, and consumption. The second DGP ("Skiba") uses the same structure but replaces the production function with a butterfly specification that generates a non-monotonic consumption response. Finally, the third DGP ("Random") generates assets, income, and consumption from randomly drawn production and consumption functions formed using a first-order tensor basis with standard normal coefficients.

I set:

$$\theta_{\text{RCK}} = \{A = 12.5,\ \alpha = 0.3,\ \rho_{AR} = 0.9,\ \rho = 0.05,\ \sigma_{\ln\eta} = 0.3,\ \sigma_{\ln\xi} = 0.2, \tag{93}$$

$$\sigma = 5.0,\ \sigma_u = 1.5,\ \gamma = 3.0\}, \tag{94}$$

$$\theta_{\text{Skiba}} = \{A_L = 12.5 \times \tfrac{3}{4},\ A_H = 12.5,\ F = 100.0,\ \alpha = 0.3,\ \rho_{AR} = 0.9,\ \rho = 0.05, \tag{95}$$

$$\sigma_{\ln\eta} = 0.3,\ \sigma_{\ln\xi} = 0.2,\ \sigma = 5.0,\ \sigma_u = 1.5,\ \gamma = 3.0\}, \tag{96}$$

$$\theta_{\text{Random}} = \{\sigma_\beta = 1.0,\ \sigma_\eta = 1.0,\ \sigma_\xi = 1.0,\ \sigma_u = 1.0,\ \sigma = 0.5,\ \rho_{AR} = 0.9\}. \tag{97}$$

Here, F denotes the fixed cost in the Skiba model, σ_u the standard deviation of measurement error in consumption, and σ the standard deviation of measurement error in

income.

For each DGP and each sample size $N \in \{50, 100, 500\}$,[32] I simulate $S = 50$ datasets. As in the production-function simulations, I draw initial conditions from a uniform distribution and simulate the model forward using the appropriate policy functions with an initial warmup period.

G.2.2 Estimation and Evaluation

To focus on the estimators' ability to recover the consumption function, I construct the income residual under the assumption that the production function is known. I then estimate the consumption function using three approaches:

- **Oracle:** Has access to the true latent variables (η_i, ω_{it}) and estimates the consumption function by regressing consumption on a basis expansion of assets and the latent states.
- **Fixed Effects (FE):** Estimates household ability using the production residual, imputes productivity shocks, and then regresses consumption on a basis expansion of assets, the estimated fixed effects, and imputed shocks.
- **Stochastic EM (S-EM):** Uses the stochastic EM algorithm described in Section 6, integrating over the distribution of latent states via PMCMC in the E-step.

The RCK and Random DGPs use fourth-order ordinary polynomial bases, while the Skiba DGP employs a Random Fourier Feature (RFF) expansion to better capture the non-monotonicity of the consumption response. For the Skiba DGP, I use 50 RFFs and concentrate frequency mass near the kink in the production function, which is approximately known from the production function's properties.

For each simulated dataset, I apply the stochastic EM algorithm from Section 6. Owing to its computational cost, I run a single initialization for 1000 iterations, discarding the first 500 as burn-in. The E-step uses 100 particles and 100 Metropolis-Hastings steps per iteration for the first two DGPs. For the third DGP, I adaptively tune the number of particles to maintain a standard deviation of the log-evidence near 0.5.

Since the consumption estimators occasionally return negative estimates for predicted consumption, in contrast to the production metrics I present results in levels rather than logs. Specifically, to evaluate estimator performance, I compute the root mean squared error (RMSE) of the estimated consumption function over simulated draws of assets and latent states. For the RMSE of latent variables, I use posterior means from the PMCMC procedure in the S-EM algorithm and the imputed fixed effects from the FE estimator. For the S-EM estimator, I also report the average coverage of the 80% posterior intervals for the latent variables.

[32] I use smaller sample sizes and a fixed T due to computational constraints. For the "Random" DGP, I additionally include a sample size of 1000.

G.2.3 Simulation Results

Figure G.15 plots the estimated consumption functions in the poverty trap DGP at three quantiles of the latent state distribution as a function of assets, with the true function shown as a black dashed line. Figure G.16 similarly compares true versus predicted consumption for each estimator. The S-EM estimator closely recovers the true consumption function and accurately captures the sharp non-linearities induced by the latent states. By contrast, the FE estimator produces substantially oversmoothed estimates that miss the true interaction between assets and latent variables. Even the Oracle estimator, despite having access to the true latent states, struggles to match the sharp local curvature of the consumption function due to the limited flexibility of the basis representation, indicating that more than 50 RFFs may be needed to fully capture the underlying structure.

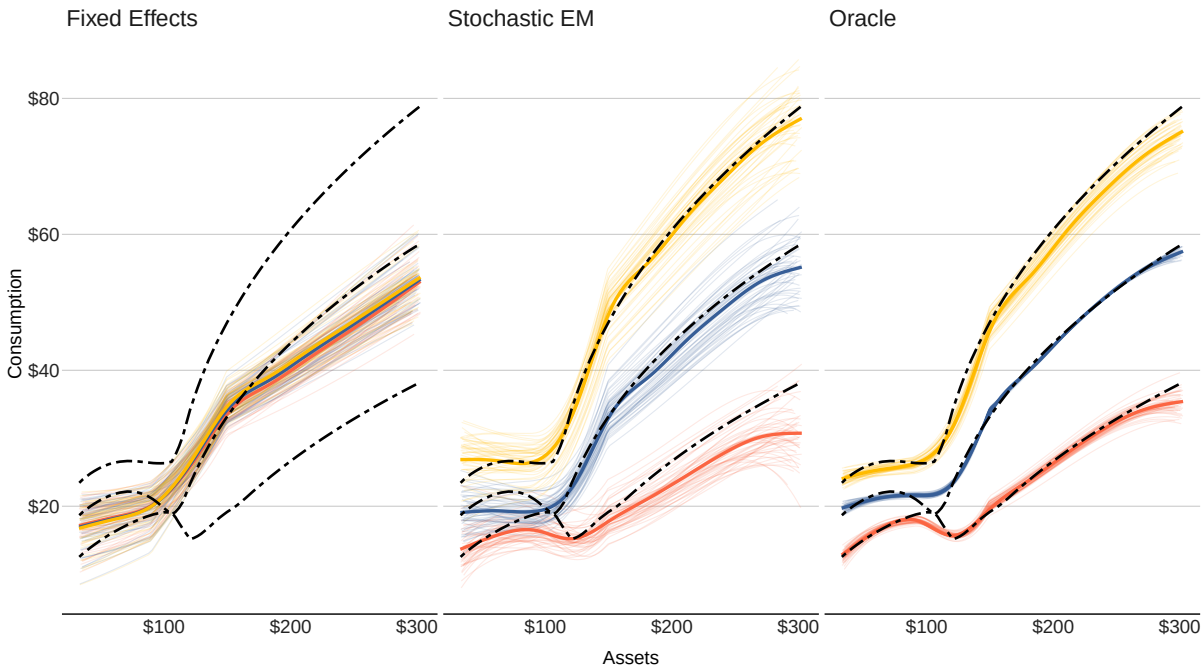


Figure G.15: Simulated Consumption Function Estimates: Poverty Trap DGP

Notes: This figure shows the estimated consumption functions from the poverty trap DGP with N=500 households and T=5 time periods. The light lines show 50 simulated estimates of the consumption function while the thicker line shows the average across simulations. The black dashed line shows the true consumption function used to generate the data. Each panel shows results for a different estimator: FE, S-EM, and the Oracle if ability and productivity were observed.

Tables G.30 and G.31 report performance metrics for the calibrated and non-calibrated DGPs. Across all settings, the S-EM estimator outperforms the FE estimator on both the consumption function and latent state recovery. In the calibrated DGPs, the S-EM estimator is roughly three times better than the FE in recovering the consumption function. As expected, the Oracle estimator achieves the best performance overall.

As sample size increases, S-EM performance improves for the consumption function but shows limited gains in recovering latent variables. Performance is similar in the calibrated

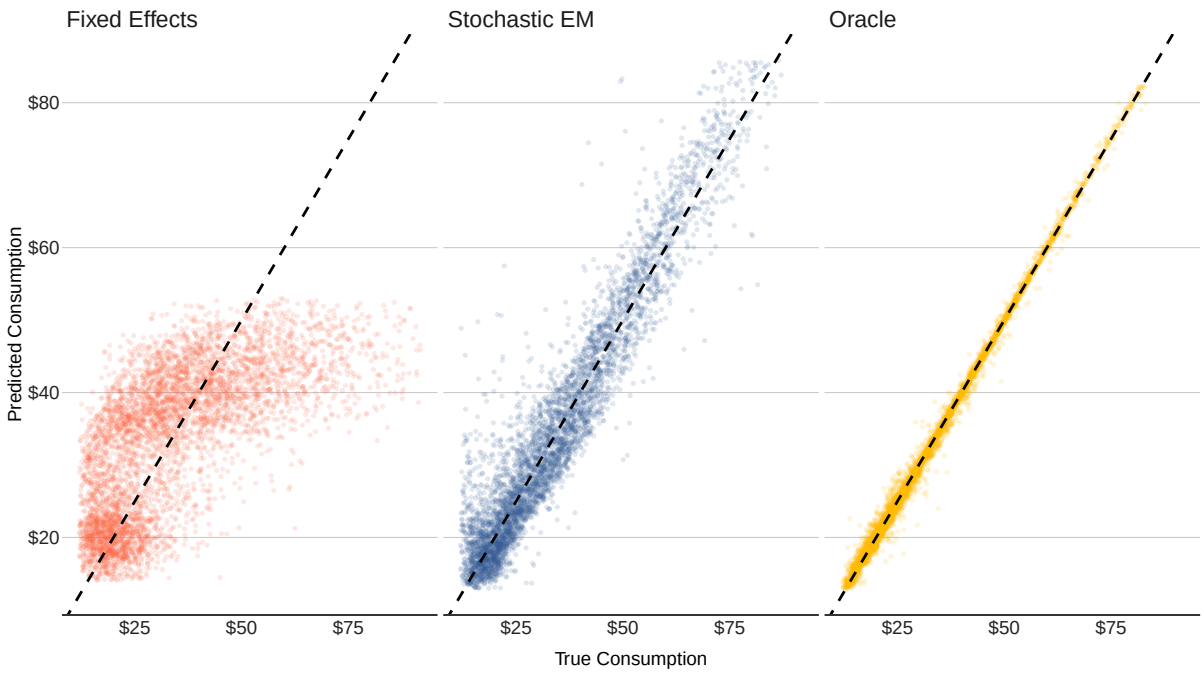


Figure G.16: True vs Predicted Consumption: Poverty Trap DGP

Notes: This figure shows the true versus predicted consumption from the poverty trap DGP with N=500 households and T=5 time periods. The light lines show 50 simulated estimates of the consumption function while the thicker line shows the average across simulations. The black dashed line shows the true consumption function used to generate the data. Each panel shows results for a different estimator: FE, S-EM, and the Oracle if ability and productivity were observed.

poverty trap and no poverty trap DGPs, suggesting that the S-EM estimator is robust to the presence of strong non-linearities in the consumption function.

Coverage of the 80% posterior intervals is poor in the calibrated DGPs, indicating that the number of particles, Metropolis-Hastings steps, or basis functions may be insufficient for accurate posterior approximation in these more complex environments. In the simpler, non-calibrated Random DGP, where the number of particles is tuned to maintain a stable log-evidence variance, coverage improves substantially.

Table G.30: Consumption Function Simulation Metrics: Calibrated

	Consumption Function, $g(\cdot)$			Ability, η			Productivity Shock, ω		
N	FE	S-EM	Oracle	FE	S-EM	S-EM 80% Coverage	FE	S-EM	S-EM 80% Coverage
No Poverty Trap (RCK Model)									
50	9.196	3.646	0.600	2.255	0.220	0.241	9.011	0.451	0.471
100	9.379	2.758	0.497	2.236	0.218	0.239	8.953	0.446	0.493
500	9.483	2.854	0.450	2.271	0.268	0.149	8.981	0.541	0.365
Poverty Trap (Skiba Model)									
50	14.315	16.654	1.610	2.255	0.246	0.249	8.984	0.515	0.442
100	13.992	6.743	1.188	2.221	0.273	0.330	8.944	0.562	0.526
500	13.451	6.371	1.037	2.272	0.321	0.354	8.989	0.655	0.527

Notes: This table shows performance metrics of the consumption function estimator on simulated data using the structural model calibrated against an example study. Results show the root mean squared error from 50 simulated datasets, with a fixed panel length of $T = 5$ periods, and varying sample sizes. The 'Oracle' estimator knows the true latent variables, whilst the 'S-EM' estimator marginalizes over the latent variables using the stochastic EM algorithm, 'FE' is the two-step fixed effects estimator that first estimates household ability as a fixed effect, and then estimates the consumption function. Metrics for the latent variables are calculated using the posterior means from the PMCMC algorithm used in the E-step of the S-EM algorithm. '80% Coverage' shows the average coverage of the 80% posterior intervals over the latent variables for the S-EM estimator. 'No Poverty Trap (RCK Model)' and 'Poverty Trap (Skiba Model)' refer to a stochastic Ramsey-Cass-Koopmans model with no production indivisibilities and endogenous household consumption, and a model with production indivisibilities and endogenous household consumption respectively.

Table G.31: Consumption Function Simulation Metrics: Non-Calibrated

	Consumption Function, $g(\cdot)$			Ability, η			Productivity Shock, ω		
N	FE	S-EM	Oracle	FE	S-EM	S-EM 80% Coverage	FE	S-EM	S-EM 80% Coverage
50	0.251	0.121	0.012	1.170	0.265	0.887	2.067	0.463	0.847
100	0.286	0.107	0.002	1.301	0.284	0.863	1.781	0.434	0.861
500	0.294	0.099	0.001	1.334	0.263	0.894	2.013	0.416	0.876
1000	0.291	0.089	0.001	1.371	0.265	0.882	2.104	0.410	0.872

Notes: This table shows performance metrics of the consumption function estimator on a simulated data generation process, with latent variables drawn from a standard normal distribution, rather than the structural model, the consumption rule is a 1st order tensor basis function of assets and the two latent variables, with standard normal coefficients. Results show the root mean squared error from 100 simulated datasets, with a fixed panel length of $T = 5$ periods, and varying sample sizes. The 'Oracle' estimator knows the true latent variables, whilst the 'S-EM' estimator marginalizes over the latent variables using the stochastic EM algorithm, 'FE' is the two-step fixed effects estimator that first estimates household ability as a fixed effect, and then estimates the consumption function. Metrics for the latent variables are calculated using the posterior means from the PMCMC algorithm used in the E-step of the S-EM algorithm. '80% Coverage' shows the average coverage of the 80% posterior intervals over the latent variables for the S-EM estimator.

G.3 Asset Transition Monte Carlos

To evaluate the performance of the production and consumption function estimators jointly, I conduct a series of Monte-Carlo simulations where I estimate the full model and form estimates of the asset transition function.

G.3.1 Data Generating Process

I use the same poverty trap DGP as in the consumption function simulations, with parameters given by θ_{Skiba}. For each simulation, I generate $S = 40$ datasets with varying samples sizes and a fixed panel length of $T = 5$ periods.

G.3.2 Estimation and Evaluation

I estimate the full model using the production and consumption function estimators described in Section 6. I run the production function quasi-Bayesian estimator with 1000 iterations, discarding the first 500 as burn-in. I then use the estimated production function to form income residuals and estimate the consumption function using the stochastic EM algorithm with only 200 iterations, discarding the first 100 as burn-in.

To evaluate performance, I construct a use the empirical distribution of assets, ability, and productivity shocks from the simulated datasets. I then compare the estimated asset transition function to the true transition function.

G.3.3 Simulation Results

Figure G.17 plots the true and predicted asset transitions for the calibrated DGP. With only $N = 100$ households, the predictions are noisy, but they sharpen as the sample grows and increasingly concentrate around the 45-degree line. Consistently, Table G.32 shows that estimator performance improves with sample size: the RMSE falls from 61.2 to 38.8 and the MAE from 44.4 to 26.4 as the sample increases from 100 to 10,000 households.

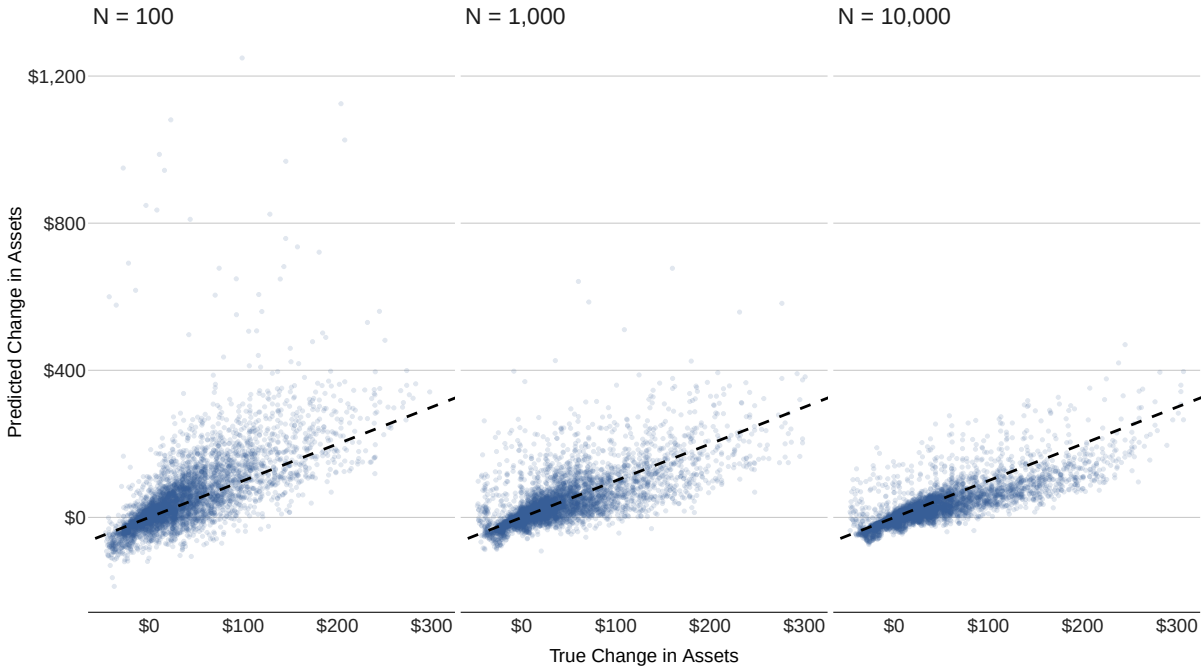


Figure G.17: True vs Predicted Asset Transitions: Calibrated DGP

Notes: This figure shows 5000 random samples of true versus predicted asset transitions from the calibrated household production function simulations. Each panel shows a different sample size, the blue points show the predicted versus true changes in assets for each household-period observation in the simulated dataset. The black dashed line shows the 45-degree reference line.

Table G.32: Asset Transition Function Simulation Metrics

	Root Mean Squared Error			Mean Absolute Error		
N	Production	Consumption	Change in Assets	Production	Consumption	Change in Assets
100	41.7	50.5	61.2	31.2	36.8	44.4
500	45.6	26.5	55.8	29.3	19.7	38.3
1000	46.7	17.0	49.4	29.6	13.3	32.7
2000	42.7	14.7	42.3	27.4	11.5	28.6
4000	44.9	12.0	42.6	28.0	9.3	27.9
5000	42.1	11.2	39.7	26.7	8.7	26.6
10000	40.2	11.5	38.8	25.4	8.9	26.4

Notes: This table shows the performance metrics of the asset transition function estimator on simulated data generated using a calibrated structural model. Each figure reports the average result over 40 simulated datasets where the integrated error is calculated over a grid of 100 asset points for each simulation draw and quantiles of latent productivity. 'RMSE' is the root mean squared error, 'MAE' is the mean absolute error.

H Household Growth Model Calibration and Solution Details

H.0.1 Model Calibration

To calibrate the model, I must estimate two parameters for each site j: the discount rate, ρ_j, and the coefficient of relative risk aversion, γ_j. First, I directly take the estimated production function, $\hat{f}_j(\cdot)$, to form the growth model's production function. Next, I match the mean and variance of the estimated AR(1) productivity process to its continuous time analogue, $\{\hat{\theta}_j, \hat{\sigma}_j\}$.[33] Finally, I use the quantiles of the estimated posterior means of the latent productivity variables to form a statespace grid covering the support of the estimated productivity process, $\{\eta_{ij}, \omega_{ijt}\}$, and plausible asset levels.[34]

To find the optimal parameters, I minimize a loss function by grid-searching over a set values for $\{\rho_j, \gamma_j\}$. To form the loss function, I solve the model for a given set of parameters, $\{\rho_j, \gamma_j\}$, and form the optimal consumption policy rule, $c_j(k, \eta, \omega)$ and predictions of next period's assets, k'_{ij}. The loss function is the weighted squared difference between the estimated consumption policy rule, $\hat{g}_j(k, \eta, \omega)$, and the model's optimal consumption policy, $c_j(k, \eta, \omega)$, in addition to the squared difference between model predicted assets,

[33]Specifically, I assume the AR(1) process is an Euler-Maruyama discretization of the continuous time OU process, which gives the following relationships: $\theta = -\log(\rho_{AR})/\Delta t, \sigma = \sqrt{2\theta\sigma^2_{AR}/(1-\exp(-2\theta\Delta t))}$

[34]Specifically, I use the 10th and 90th percentiles of the estimated productivity process for these state variables and a root-finding algorithm to find the highest investment level a household could plausibly make at $k^{q_{0.9}}, \eta^{q_{0.9}}, \omega^{q_{0.9}}$, where $q_{0.9}$ reflects the empirical 90th percentile of each variable. I set the upper bound of the asset grid to twice the maximum plausible investment level, which is usually many multiples of the observed empirical asset support.

$k_{beo,t+1}$ and observed assets, $\hat{k}_{beo,t+1}$:

$$L_j(\rho_j, \gamma_j) = \sum_{b=1}^{B} \sum_{e=1}^{E} \sum_{o=1}^{O} \hat{\mu}_{beo} \left[\left(c_j(k_b, \eta_e, \omega_o) - \hat{g}_j(k_b, \eta_e, \omega_o) \right)^2 + \left(k_{beo,t+1} - \hat{k}_{beo,t+1} \right)^2 \right] \tag{98}$$

where $\hat{\mu}_{beo}$ is the joint empirical density of assets and latent productivity projected onto the statespace grid.

H.1 Asset Density

In this section, I use the calibrated household growth model to simulate the distribution of assets across households over time. To demonstrate the importance of initial conditions, I simulate the model forwards using a uniform distribution of assets and latent productivity states a month at a time, up to one year, followed by estimates of the stationary distribution of assets. Figure H.18 shows the density of assets in Bangladesh, Balboni et al. (2022). The red lines show the estimated distribution of assets in the

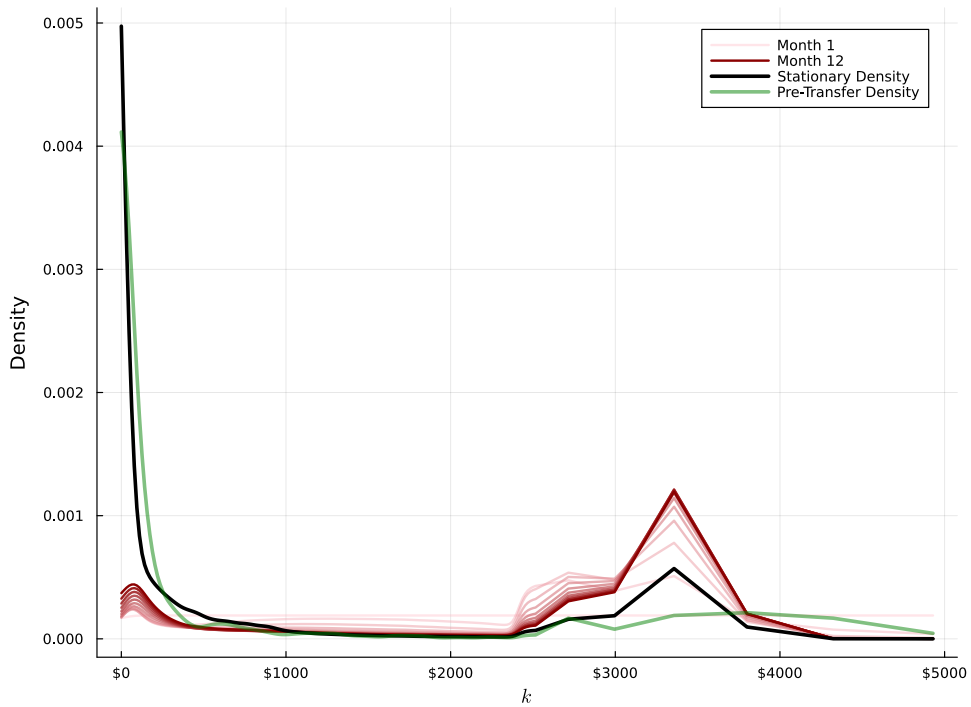


Figure H.18: Simulated Asset Density in Bangladesh, Balboni et al. (2022)

Notes: This figure shows the simulated asset density in Bangladesh using a model calibrated with data from Balboni et al. (2022). The red lines show the estimated density up to 12 months after initializing households with a uniform distribution of assets and latent productivity states. The black line shows the estimated stationary distribution of assets whilst the green line shows the empirical density at baseline. The blue line shows the empirical density of assets post-transfer.

study a month at a time, with darker lines corresponding to later months. The black line

shows the estimated stationary distribution of assets, whilst the green and blue lines show the empirical density of assets pre- and post-transfer, respectively. The figure demonstrates, in the case of Balboni et al. (2022) at least, that the model is able to reproduce the empirical density of assets – the stationary distribution captures the large density of households with close to zero assets, whilst also reproducing a second, smaller peak of assets around the second model steady state at approximately USD$ 3,500, similar to the empirical density pre/post-transfer.

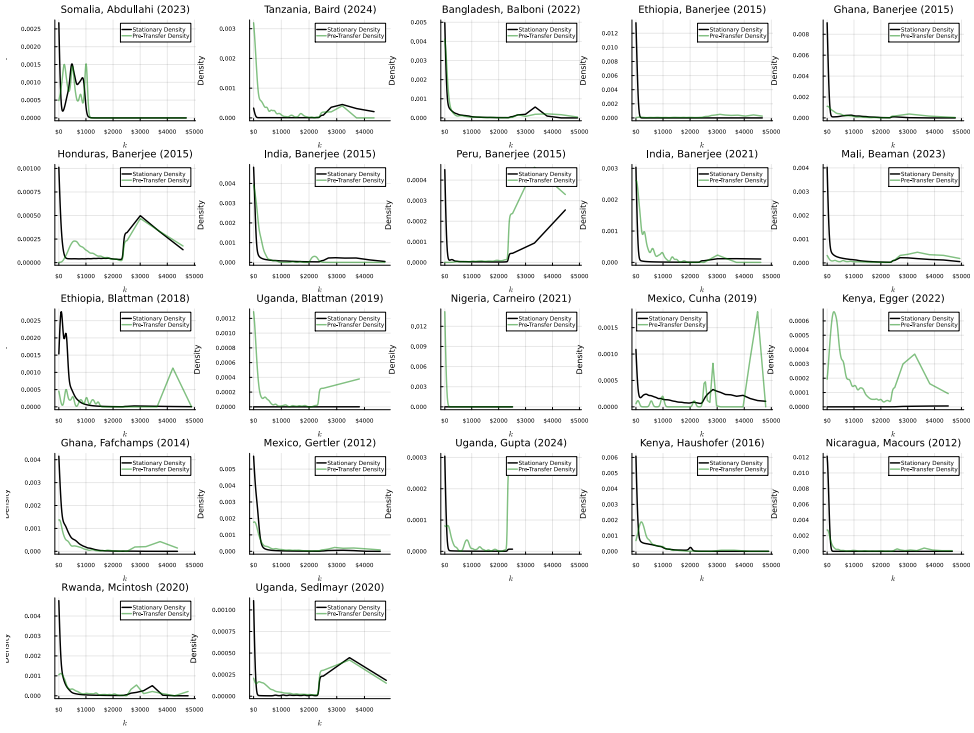


Figure H.19: Simulated Stationary and Baseline Asset Densities

Notes: This figure shows the simulated stationary distribution of assets, in black, compared to the baseline empirical asset density, in green, across all studies. The stationary asset distribution is calculated using the fixed point of the Fokker-Planck equation implied by the calibrated household growth model.

Figure H.19 shows similar results using the simulated stationary asset density for all studies, compared to the baseline pre-transfer asset density observed in each study. The figure again suggests in general, the model's stationary distribution aligns well with the pre-transfer empirical asset density, with the exception of a few studies. The fact that the model captures the empirical asset density well suggests that the optimal transfer policy I turn to in Section I is likely to be a good approximation of the true optimal policy.

I Transfer Policy Optimization

The planner seeks to solve:

$$T^* = \arg\max_T \int V(k + T(k,\eta), \eta, \bar{\omega}) dF_{k,\eta} \tag{99}$$

$$\text{s.t. } B_\eta = \int T(k,\eta,\bar{\omega}) dF_{k,\eta} \tag{100}$$

where $T(k,\eta)$ is the transfer policy as a function of initial assets and productivity type, $V(\cdot)$ is the value function from the household growth model, B_η is the planner's budget for productivity type η, and $\bar{\omega}$ is the median of the idiosyncratic productivity shock distribution.

This leads to the Lagrangian and first-order condition:

$$\mathcal{L}(T,\lambda,\eta) = \sum_i^N \mu_i V(k_i + T_i;\eta,\bar{\omega}) + \lambda\left(B_\eta - \sum_i^N \mu_i T_i\right) \tag{101}$$

$$\frac{\partial \mathcal{L}}{\partial T_i} = \mu_i V'(k_i + T_i;\eta,\bar{\omega}) - \lambda\mu_i = 0 \tag{102}$$

Where μ_i is the weight vector over the initial asset distribution, T_i the transfer to households at position i on the discretized state grid. For studies with fixed costs in the production function, the non-convexities generate a kink in the value function, $V(\cdot)$, which complicates the planner's optimization problem – there are potentially multiple T_is that satisfy the first-order condition.

Intuitively, this corresponds to the planner deciding between a "big-push" transfer, T_+, that equates costs (shadow price λ) and benefits on the upper envelope of the value function, V'_+, or a smaller transfer, T_-, that equates costs and benefits on the lower envelope of the value function, V'_-. One solution strategy could be to evaluate all possible combinations of $\{T_{i,+}, T_{i,-}\}$ for each household i and choose the combination that maximizes the planner's objective subject to the budget constraint. However, this is computationally infeasible since the number of combinations grows exponentially with the number of households, N.

Instead, my solution restricts the class of transfer policies the planner can choose such that when the planner faces two options, $\{T_{i,+}, T_{i,-}\}$, they can only choose an interval of assets, $[k_-, k_+]$, where the planner picks the higher transfer, $T_{i,+}$. Intuitively, fixing η the planner can order households by their initial assets, k_i, and choose a cutoff such that all households with assets above this cutoff receive the second transfer, $T_{i,+}$, and all households below this cutoff receive the smaller transfer, $T_{i,-}$. Many households will only have one solution, which will closely resemble $T_{i,-}$ or $T_{i,+}$ depending on their initial assets.

This reduces the planner's problem to two, nested one-dimensional problems. First, for

a given shadow price, λ, find the optimal interval $[k_-, k_+]$ that maximizes the planner's objective. Second, update λ until the budget constraint is satisfied.[35]

[35]It seems plausible that the planner's optimal transfer policy is not restricted by this interval assumption, however I provide no proof that this is the case. In simulations I find that naïve searches for combinations of $\{T_{i,+}, T_{i,-}\}$ do not yield higher welfare than the interval solution, although this is not surprising given the dimension of the problem.