

NBER WORKING PAPER SERIES

HOW POVERTY FELL

Vincent J. Armentano
Paul Niehaus
Tom Vogl

Working Paper 35163
<http://www.nber.org/papers/w35163>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2026

We are grateful to Cole Dreier, Radhika Goyal, Jie Liao, and Gongyu Zhou for research assistance. Thanks to Karthik Muralidharan as well as seminar participants at Central European University, NBER, Rice University, UC San Diego, and University of Warwick for their feedback, and especially to Federico Rossi and Tommaso Porzio for many early discussions that helped shape the direction of this project. We gratefully acknowledge funding from the National Science Foundation under award #2116944. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Vincent J. Armentano, Paul Niehaus, and Tom Vogl. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

How Poverty Fell
Vincent J. Armentano, Paul Niehaus, and Tom Vogl
NBER Working Paper No. 35163
April 2026
JEL No. I32, N30, O12

ABSTRACT

The share of the global population living in extreme poverty fell dramatically from an estimated 36% in 1990 to 9% in 2015. We describe how this decline happened: the extent to which changes within as opposed to between cohorts contributed to poverty declines, and the key changes in the lives of households as they transitioned out of (and into) poverty. We do so using cross-sectional and panel sources that are representative or near-representative of five countries that collectively accounted for 75% of global poverty decline between 1990 and 2015. The data show that overlapping birth cohorts experienced the decline of poverty together over time, such that poverty decline can be viewed as a primarily within-cohort phenomenon. Within cohorts, the data reveal substantial churn, casting the challenge of escaping poverty as a “slippery slope” more than a long-term trap. The data also illustrate a diversity of pathways out of poverty: sectoral transitions, migration, and changing occupational choices and female labor force participation can all account for some part of poverty reduction, but in all but a handful of cases, a majority of households exiting poverty did so without experiencing these changes.

Vincent J. Armentano
University of California, San Diego
varmenta@ucsd.edu

Paul Niehaus
University of California, San Diego
Department of Economics
and NBER
pniehaus@ucsd.edu

Tom Vogl
University of California, San Diego
Department of Economics
and NBER
tvogl@ucsd.edu

1 Introduction

One of humanity’s great recent achievements has been the reduction in the share of the global population living in extreme poverty, from an estimated 41% in 1981, to 36% in 1990, to 9% in 2015.¹ Development economists have made great strides toward measuring the decline of poverty (Chen and Ravallion, 2010) and characterizing the lives of people who remain poor in spite of it (Banerjee and Duflo, 2007). But systematic descriptions of *how* it happened—of the anatomy of poverty decline—are rare. This paper aims to provide a description that is as comprehensive as the available data will allow.

By “how,” we mean in part what happened to the livelihoods of individuals and households as they moved out of (and often, as we will see, back into) extreme poverty. Did they find work in a factory? Start their own business? Move to a city? And so on. We will also emphasize that changes like these *within* the life of any one person, or cohort or generation, may be more or less important than changes across cohorts or generations, as the less poor young replace their poorer parents and grandparents in the population. So our initial concern will be to understand the demographic structure of poverty decline.

We focus on five countries—China, India, Indonesia, Mexico, and South Africa—that collectively accounted for 75% of global poverty decline during 1990-2015. We include these countries because they have available both (a) 3 or more waves of cross-sectional household surveys spanning 15 or more years, and (b) 3 or more waves of panel household surveys spanning 9 or more years, all spanning large portions of the national population and most fully representative of it. They also provide a degree of global representativeness in the sense that they span the major regions of the developing world, and—in the case of China, India, and Indonesia—include some of the most important contributors to global poverty decline. Altogether, we harmonize 49 distinct waves of repeated cross-sectional surveys and 26 distinct waves of panel surveys. The data have many imperfections, certainly relative to data from wealthier countries, and these will occupy us for a number of pages below. But they are generally high quality by the standards pertinent to work on extreme poverty and thus, in our view, afford a valuable opportunity to better understand a crucial episode in human history.

We study poverty in the now-conventional sense of living on less than \$2.15 per day in 2017 PPP dollars (World Bank, 2023), using both consumption- and income-based measures as available.² Consumption has the advantage that it reflects whatever smoothing of intertemporal shocks to earnings is feasible. Consumption is also thought to be better-measured than income in many surveys, though to be precise, this advantage refers only to measures of non-durable consumption.³

¹These rates are based on the \$2.15 2017 PPP poverty line (Chen and Ravallion, 2010; World Bank, 2023). See <https://pip.worldbank.org/>. In 2025, the World Bank revised its estimate of global extreme poverty slightly upwards, to 10.5%, but this change largely reflects changes in underlying (real) national poverty lines (Foster et al., 2025).

²The \$2.15 threshold is appropriate given that our data run through 2017 at latest. See Foster et al. (2025) for the rationale for applying a revised threshold of \$3.00 in 2021 PPP dollars to more recent data.

³It is unclear whether careful measurement of non-durable consumption compensates for the usual omission of the value people derive from major durables such as housing. We reluctantly follow this precedent in our main exercises

Income has the complementary advantages that it includes money used to purchase all types of consumption and that—crucially for our purposes—we can identify and study its sources.

We use the repeated cross-sectional data to investigate how the pace of poverty decline within cohorts compares with that between cohorts. This exercise is in part motivated by recent results from [Porzio et al. \(2020\)](#) suggesting that cohort effects play an important role in structural change during the development process. We confirm that as poverty has fallen in aggregate, successive birth cohorts have entered adulthood at progressively lower poverty rates. However, we also find that the pace of this between-cohort process roughly matches the within-cohort pace of poverty decline. Poverty rates are similar across the age distribution at a point in time; aggregate poverty decline manifests in downward parallel shifts of a flat cross-sectional age profile. This fact holds true whether we use consumption or income, where available, as well as whether we focus on household heads, as is common practice in the economics literature, or take the whole population into account.

The downward march of flat cross-sectional age profiles has two noteworthy implications. First, changes in the age structure of the population—aging, for example—cannot account for poverty decline. With flat cross-sectional age profiles, the aggregate poverty rate is invariant to the age structure, at least in an accounting sense. Second, most poverty decline occurs over people’s lives, rather than through successive cohort replacement. Cohorts are interchangeable in their poverty rates at a point in time, so cohort replacement plays little role in the trajectory of poverty rates. To quantify this intuition, we develop a cohort-time decomposition of poverty change and take it to the repeated cross-sectional data, confirming that the lion’s share of poverty decline accrues within cohorts rather than between them. We do not take a stand on whether this dominant within-cohort variation reflects innovation over time or a stable age pattern, as for example [Deaton and Paxson \(1994\)](#) emphasize in their lifecycle interpretation of analogous patterns they document for *average* consumption and income in Taiwan circa 1980. But compared with their data on averages in a high-income setting, our data on extreme poverty are more likely to reflect liquidity constraints, which make innovations over time more relevant.⁴

Panel data let us examine these within-cohort changes in more detail. We emphasize income-based measures here, as our focus is on understanding what changes enabled households to increase their living standards, as opposed to simply measuring those living standards per se. We also focus primarily on households present at the start and end of the panel, though conclusions are broadly similar if we follow all households through the rounds for which they are present.

We first document that progress out of poverty, while substantial, was not irreversible. Many households that were initially poor exited poverty, but many that were *not* poor at some point subsequently entered it. The share exiting poverty conditional on starting in it ranged from 45% in Rural India to as high as 67% in China. But the share entering poverty conditional on starting non-poor was also substantial, ranging from 13% in Indonesia to as high as 34% in Mexico. Estimates using consumption, where they are available, tell the same basic story. These conclusions do

but also explore the role of housing to the extent possible.

⁴Furthermore, the lifecycle theory does not explain the pronounced tracking of consumption with income that emerges in both their data and ours.

not change meaningfully as we add noise to the data, suggesting that measurement error in the original figures is not the main driver. And, perhaps most importantly, round-to-round probabilities of transitioning out of poverty were generally stable or—in the case of China and India—shifted advantageously over time. In short, while poverty traps may have delayed some households’ progress out of poverty, what stands out in the data is how slippery the slope was for many that *did* at some point exit poverty. This finding mirrors, systematically and on a large scale, a point emphasized in several earlier studies of individual panel and pseudo-panel datasets (Baulch and Hoddinott, 2000; Narayan et al., 2009; Krishna, 2010; Merfeld and Morduch, 2024).⁵

We then examine changes in the livelihoods of households that exited or entered poverty—the activities in which they engaged, and the income they obtained from them. Overall, no one pathway out of poverty predominates. For almost all countries and dimensions of activity, only a minority of the households that exited poverty did so while changing their primary means of earning a living. And households that entered poverty often changed their activities in similar ways. For example, the share of poverty-*entering* households that switched from agriculture to non-agriculture is 60%–100% of the corresponding share of poverty-exiting households. One important exception, as we will shortly note, applies in the more advanced economies of Mexico and South Africa, where transitions to wage work played a predominant role.

With respect to sector, transitions out of agriculture played a limited role. They did not account for the largest share, let alone the majority, of transitions out of poverty in any country. In fact, an accounting decomposition credits transitions *into* agriculture with a meaningful share of poverty reduction in most countries.⁶ More generally, in every country, households that stayed in the same sector contributed more to poverty decline than households that changed sector. In relation to classic debates over the relative importance of raising productivity within agriculture (Schultz, 1964; Lipton, 1977; Mellor, 2017) versus moving people out of it (Lewis, 1954; Rostow, 1991), the results are consistent with both perspectives, but perhaps more with the within-agriculture view.⁷ In this sense, the results keep with the “pro-poor growth” literature of the early 2000s, which tended to find stronger national-level associations between poverty reduction and growth in agriculture specifically (see for example Dollar and Kraay, 2002; Kraay, 2006).⁸

Rural-to-urban migration, in the three countries (Mexico, Indonesia and South Africa) where the panel surveys tracked migrants, similarly accounted for a limited share of poverty decline.⁹ These

⁵Balboni et al. (2021) provide evidence of poverty traps at the household level, although Kraay and McKenzie (2014) argue that traps are most plausible at the regional level. See also Carter and Barrett (2006) and the special issue it introduces for related discussion.

⁶One reaches the opposite conclusion if one ignores the panel structure of the data and applies older, cross-sectional decomposition techniques such as that in Ravallion and Huppi (1991).

⁷These views are not mutually exclusive. Barrett et al. (2010) argue that development within agriculture is a prerequisite for movement out of agriculture, but both are important for poverty reduction. Dercon and Gollin (2014) similarly argue that both are important, but under different market conditions.

⁸See also related recent work on the interpretation of the agricultural productivity gap by Gollin et al. (2014).

⁹One might worry that the absence of migrant data from China and India limits what we can learn not only about migration but also about sectoral transitions. The evidence from the three panels that do track migration, however, suggests that the vast majority of sectoral transitions occurred in the absence of migration. Of households that switched between the agricultural and non-agricultural sectors, only 1.0% (Mexico), 3.7% (South Africa), and

migrants did exit poverty at relatively high rates in Mexico and South Africa, but in Indonesia they did worse than those that did not move, and in all countries they were rare.¹⁰ Consequently, only between 0% (Mexico) and 11% (Indonesia) of households that exited poverty did so while moving from a rural to an urban location. Notably, however, rural-to-*rural* migrants accounted for a third of all net poverty decline in Indonesia. This result is consistent with Bryan and Morten’s (2019) finding—using different data—that the majority of Indonesian migration took place within rather than across urban/rural category, and with Beegle et al.’s (2011) point that “any movement has the potential to be welfare improving.”¹¹

With respect to occupational choice, patterns are noticeably different in the less-developed economies (China, India and Indonesia) relative to the more-developed ones (Mexico and South Africa). In the former group, those who stayed or became self-employed contributed the most to poverty reduction: reminiscent of, for example, the occupational hierarchy induced by credit constraints in Banerjee and Newman (1993). In the latter group, on the other hand, transitions within and into wage work account for the bulk of poverty decline. This is as one might expect given the shift towards wage work that often coincides with development (Jensen, 2022), and is consistent with the idea that many of the self-employed in such countries are “entrepreneurs” only out of necessity (e.g. Schoar, 2010). The contrast between the two groups of countries underscores the importance of Dupas et al.’s (2026) point that poor households are increasingly located in non-poor countries, which may reduce the conceptual distance between the study of development and the study of poverty dynamics in rich countries.¹²

Poverty transitions had a nuanced relationship with women’s labor force participation. Except (narrowly) in India, households in which a woman entered the workforce exited poverty at the fastest rates. But female labor force entry was uncommon in all countries except South Africa. And in China, nearly half of households *withdrew* women from the labor force, while nevertheless accounting for nearly half of all net poverty decline. Chinese women may have withdrawn from the labor force once their households could afford it, as in Goldin’s (1995) argument that “women’s work is often implicitly bought by the family” as economies initially develop.¹³

Finally, the panel data shed light on the role of redistribution. Redistribution could in principle contribute to poverty decline in two distinct ways: indirectly, by stimulating growth in income from other sources (as for instance Best et al. (2026) find in Brazil), or directly if it becomes more generous over time. In the data, we see that transfers (both public and private) played a largely

16.1% (Indonesia) also migrated.

¹⁰See also Hamory et al. (2020) and Lagakos et al. (2020), who estimate and interpret changes in mean earnings associated with rural-to-urban migration in several countries, including Indonesia.

¹¹It also calls to mind Eckert and Peters’s (2025) finding that rural wages grew in the United States even during periods of rapid structural change from agriculture to manufacturing.

¹²See also Gindling and Newhouse (2014) for descriptive statistics on self- and wage-employment in a cross-section of low- and middle-income countries.

¹³Goldin attributes this pattern to a combination of income and price effects and to technology-driven changes in the nature of agricultural work. Boserup (1970) first drew widespread attention to the idea that economic development might limit opportunities for women. More recent scholarship has suggested that each tends to encourage the other (Duflo, 2012) at current levels of development, but also that cultural barriers loom large (Jayachandran, 2015).

“protective” role: they increased as a share of income for households that entered poverty, and decreased (or increased only slightly) for those that exited it, but did not become meaningfully more generous over time. India is a notable exception in this regard, however: in India, the gap between pre- and post-transfer poverty rates widened meaningfully over time, as public transfer generosity increased.¹⁴ One can thus see in the micro-data India’s exceptionalism as a country that chose generous redistribution—which [Kuznets \(1955\)](#) and others have associated with later stages of development—relatively early ([Muralidharan, 2024](#)).

We see our analysis of how poverty fell situated within the literature in three ways. First, our analysis provides a dynamic counterpart to cross-sectional descriptions of extreme poverty, such as that provided by [Banerjee and Duflo \(2007\)](#). Second, our analysis complements exercises that seek to account for differences in per capita GDP, either in the cross-section ([Caselli, 2005](#)) or over time ([Jones, 2016](#)). Those exercises pose different questions and characterize different households, focusing on an “average household,” while ours focuses on those near the poverty line. These households turn out to be quite different: in the average country-year in our data, the average household is found at the 74th percentile of the per capita consumption distribution, while the poverty line is at the 43rd percentile. Third, our panel microdata allows us to examine individual transitions between non-agriculture and agriculture, rural areas and cities, self-employment and wage work, and women’s non-participation and participation in the labor force, which are not visible in aggregate or repeated cross-sectional data.

Our inquiry into the process of poverty decline follows in a long tradition of descriptive economics research on inequality and poverty, which has shaped subsequent research agendas. In wealthier countries, this literature includes contributions by [Katz and Murphy \(1992\)](#) on how inequality relates to the supply and demand of skilled labor, by [DiNardo et al. \(1996\)](#) on how it relates to labor market institutions, by [Chetty et al. \(2014\)](#) on how it persists across generations, and by [Piketty et al. \(2018\)](#) on how it accrues to capital versus labor, before versus after taxes. In poorer countries, the literature includes [Chen and Ravallion’s \(2010\)](#) estimates of extreme poverty around the world, [Deaton’s \(2005\)](#) exploration of the relationship between economic growth and poverty decline, and [Banerjee and Duflo’s \(2007\)](#) static account of the lives of poor people in a wide range of contexts. We aim to complement these canonical studies by setting out comparable, systematic facts on what happens within and between households as poverty rates decline. We hope these facts help guide the design and interpretation of studies estimating the effects of development interventions. We conclude with some concrete examples along these lines in [Section 5](#).

2 Data

We compile data on consumption poverty and income poverty from five large countries. For each country, we draw on two types of household surveys: repeated cross-sectional surveys and panel surveys. In this section, we discuss the choice of countries, the survey data available for each

¹⁴In South Africa, this gap was also meaningful, but it did not widen over time.

country, and the methods we use for measuring poverty in a comparable way across countries over time. We then describe aggregate poverty trends in each country and relevant aspects of the economic historical context.

2.1 Household surveys

For the range of analyses we undertake, repeated cross-sectional surveys and panel surveys each have their strengths and limitations. The repeated cross-sectional surveys tend to have large samples and span many years. Because each wave has a separate sample with its own design, these surveys also stay up to date with changing population composition. However, these surveys do not allow us to track individuals or households as they exit or enter poverty. In contrast, the panel surveys tend to have more limited samples and study periods, but they do allow tracking. The two survey types thus complement each other. The repeated cross-sections are useful for charting the demographic structure of poverty decline at the population level, while the panels are useful for scrutinizing household-level dynamics.

Table 1 lists the five study countries and their associated surveys, along with key survey details. We selected countries with surveys of both types that are near or fully nationally representative and span many years.¹⁵ We require at least three waves of data collection, which will allow us to study volatility.

These criteria directly lead us to four of our five countries: India, Indonesia, Mexico, and South Africa. All four have high-quality, representative cross-section and panel household surveys on consumption, income, and their determinants. The cross-sectional surveys were designed to be representative of the national population when they were fielded. We apply the sampling weights provided by the survey organizations, so that we can interpret our estimates as representative of the national population at the time of the survey. The panel surveys were designed to be representative of the target population in the baseline wave. The target population is national in all cases but the Indonesia Family Life Survey (IFLS), which drew its baseline sample from 13 provinces containing 83% of the Indonesian population (Strauss and Witoelar, 2022). We apply the baseline household sampling weights, which are available for all four surveys.¹⁶

To add China, a leading contributor to global poverty decline during this period (Chen and Ravallion, 2010), we relax the representativeness requirement slightly. Like the other countries, China has both types of data spanning large geographies over many years. However, the Chinese surveys are less certain to be representative of a meaningful target population. The cross-sectional survey, the Chinese Household Income Project (CHIP), drew separate samples of urban and rural households in each wave. Samples were drawn by the State Statistical Bureau, but details on representativeness are scarce.¹⁷ We weight the urban and rural samples to match the urban and

¹⁵All repeated cross-sections span at least fifteen years, and all panels span at least nine.

¹⁶Some surveys include attrition-adjusted sampling weights, but these are highly correlated with the baseline weights and not available for all surveys.

¹⁷Later waves of CHIP included a separate, small sample of rural-to-urban migrants. We do not use this sample because of its lack of comparability with earlier waves.

rural population shares in the nearest census. The panel survey, the China Health and Nutrition Survey (CHNS), drew its baseline sample from 8 provinces, which contained 40% of the Chinese population in the contemporaneous 1990 census of China. Researchers did not have access to provincial sampling frames, so they drew a stratified random sample of counties and cities within each province, followed by a simple random sample of communities within each county or city, followed by a simple random sample of households within each community (Popkin et al., 2010). No sampling weights are provided.

Beyond China, sample specifics vary across the panel surveys. The Indonesian, Mexican, and South African surveys track migrants, while the Chinese and Indian surveys do not. In effect, we can investigate the role of household migration for three countries only. Additionally, part of the India Human Development Survey (IHDS) sample is a follow-up of rural households from the earlier Human Development Profile of India, here too limited to non-migrants. This linkage provides a longer-term view of rural households, from 1993-2012, as compared with urban households' exclusive appearance in the IHDS, from 2005-2012.

To unify our approach to the panel datasets while maximizing the length of follow-up, we follow only the households present at baseline, as defined by each survey. Some of the panel surveys add households in later waves, for example when original household members split off and form new households, but we do not include them. The sizes of the original households may vary from wave to wave, but their identification does not. South Africa is an exception, as the National Income Dynamics Study (NIDS) tracked individuals rather than households, so we mimic household tracking by following first-wave household heads and the people with whom they are living in a given round. The consequence is a consistent analytic approach for all five panel datasets.

2.2 Poverty measurement

We follow World Bank (2023) procedures for defining and measuring poverty. Four key questions arise. First, what measure of economic well-being do we compare with the poverty line? Second, whose poverty rate do we calculate, households or individuals? Third, what poverty line do we use? Fourth, how do we adjust prices to make the poverty standard comparable across space and time? As Appendix B.9 reports, our core conclusions are generally not sensitive to reasonable alternatives; we discuss exceptions there.

Consumption versus income We use both income-based and consumption-based definitions of poverty. Consumption is conventionally preferred for poverty measurement in low- and middle-income countries. One reason is ease of measurement. Poor households may self-produce or barter in-kind for much of what they consume, so income alone fails to capture their access to resources. In contrast, the consumption modules in all but one survey (CHNS) have detailed questions about expenditures and consumption of self-produced goods.

Another reason for preferring consumption to income for the measurement of poverty, less specific to low- and middle-income countries, is that consumption more directly measures households'

material conditions. If households smooth their consumption, then both consumption and wellbeing will be less volatile than income. In this case, low consumption may be the more appropriate marker of material deprivation. Even so, typical consumption aggregates omit important durables—most notably housing—limiting their advantage in this regard.¹⁸

At the same time, income has the advantage of being assignable to specific household members and specific activities, which is crucial to our panel analysis. Specifically, much of the panel analysis concerns the economic sources of within-cohort gains, which can only be studied using income. Rising consumption cannot be sustained in the long run without rising income, and we wish to understand the sources of that income. The questions of the panel analysis thus lend themselves to income-based measures.

We therefore emphasize consumption poverty in the repeated cross-sectional analysis and income poverty in the panel analysis. Beyond the conceptual rationale above, this approach best leverages the strengths and weaknesses of each data type. All of the cross-sectional surveys have detailed consumption expenditure modules; many also collect data on the incomes of household members, but coverage is incomplete.¹⁹ Conversely, all of the panel surveys collect income, while only some collect consumption.²⁰ Where possible, we check the sensitivity of our results to using the complementary poverty measure.

Households versus individuals Because household members share resources, typical poverty measures aggregate consumption or income at the household level and then adjust for household size. Following common practice (e.g., [Chen and Ravallion, 2010](#)), we do so by simply dividing consumption or income by household size. One can then compare to a poverty line by calculating either the share of *households* that are poor or the share of *individuals* living in poor households.²¹ We take the individual-level approach in the repeated cross-sectional analysis and the household-level approach in the panel analysis.

The individual-level approach is fitting for the repeated cross-sectional analysis because it uses a population concept that does not depend on household formation decisions that may change over time. It ensures that our description of poverty change is representative of a well-defined population, though it differs from the conventional approach in the economics literature on life-cycle processes. Researchers in this literature typically analyze households and assign them the demographic characteristics of their heads (e.g. [Aguiar and Hurst 2013](#); [Deaton 1997](#)). In our application, this approach would amount to studying how the probability of living in a poor household relates to age, period, and cohort among household heads only. But the probability of being a household head varies markedly by age and sex (Appendix Figure [A.3](#)), so the composition of

¹⁸We explore durable consumption in the panel datasets, several of which include questions about housing values.

¹⁹In India, some NSS rounds collect weekly wage and salary income from a separate sample of households, but a small share report this type of income, and researchers have doubted whether the responses provide reliable information on household living standards ([Shukla, 2010](#)); in Indonesia, early waves of SUSENAS lack income altogether.

²⁰In China, the CHNS lacks consumption altogether; in India, the HDPI (the rural-only survey preceding the IHDS) also lacks consumption.

²¹The [World Bank \(2023\)](#) defines the poverty headcount ratio in the second way, as the share of individuals living in poor households.

household heads born in a given year changes over time.²² We instead study how the probability of living in a poor household relates to age, period, and cohort among *all* individuals, which bypasses this issue while still accurately describing the extent of poverty decline during people’s lifetimes. We verify in Appendix A that we obtain similar results when we restrict the sample to household heads.

The household-level approach befits the panel analysis for both practical and conceptual reasons. Practically, some of the panel datasets track only households, not individuals, making an individual-level analysis for all five countries impossible. Conceptually, the panel analysis asks how a household’s location or mix of economic activities relates to its poverty transitions, which inherently requires us to view the household as an economic unit. As such, we treat the household as the unit of analysis in the panels.

Poverty line We adopt the “\$1-a-day” extreme poverty line (Ravallion et al., 1991) that served as the global standard throughout our study period, updated to \$2.15 per person per day in 2017 prices. Both the original \$1 line and the updated \$2.15 line were based on averaging the national poverty lines of poor countries during the study period. We choose this line because it is transparent, well-known, and the basis of prominent statements about poverty decline around the world (Chen and Ravallion, 2010; World Bank, 2023).²³

Price adjustment For temporal and international comparability, we adjust prices in a two-step procedure. First, we convert consumption and income to 2017 prices in local currency units using the national consumer price index. Second, we convert 2017 local currency units to 2017 international dollars using consumption purchasing power parities (PPP) from the 2017 International Comparison Program (ICP). The application of PPPs in a single year avoids issues of non-comparability across ICP rounds.²⁴

2.3 Aggregate poverty decline in sample countries

Although we chose the sample countries mainly based on the joint availability of repeated cross-section and panel data, they cover a range of world regions and account for much of the world’s poverty decline. Appendix Figure A.1 charts the number of people living in consumption-poor households as a share of the world population, within each sample country and across the world. As a simple validation of our data sources, we carry out this exercise first using World Bank data and then using our repeated cross-section data.

²²A good example is South Africa, where 74% of 50-year-old heads were male in 1995, but only 47% of 65-year-old heads were male in 2010.

²³More recently, the World Bank raised its international extreme poverty line to \$3 in 2021 prices, outpacing inflation, because poor countries have raised their national poverty lines in real terms (Foster et al., 2025).

²⁴In our main approach, we do not adjust for spatial variation in prices within countries, nor for variation in consumption baskets between the poor and non-poor (Deaton and Dupriez, 2011). In Appendix B.2, we include a sensitivity check that adjusts for price differences between urban and rural areas.

The World Bank series find that sample countries account for 75% of global poverty decline during 1990-2015, with most of it accruing to China, India, and Indonesia. The series based on our repeated cross-sectional datasets are similar, but they highlight the exact study period for each country, as well as the period of overlap. Full overlap starts with the first South African survey in 1995 and ends with the last Indian survey in 2011. However, the individual country series extend up to a decade in either direction.

Figure 1 zooms in on the poverty experience of each sample country. We plot up to four series for each country, for a range of sample definitions and poverty measures that will be relevant to the subsequent analyses. For all countries, we estimate (i) the share of all individuals living in consumption-poor households and (ii) the share of household heads living in consumption-poor households. For China, Mexico, and South Africa, we also estimate the analogous shares using income rather than consumption. Share (i) is the individual or population poverty rate, while share (ii) is the household poverty rate. As described in Section 2.2, our repeated cross-sectional analyses focus on the population rate, while our panel analyses focus on the household rate.

Every series in Figure 1 indicates that poverty fell. Population and household poverty rates show similar trends but moderately different levels. The level of the population rate tends to exceed that of the household rate, consistent with a positive correlation between poverty and household size. Consumption and income poverty rates also show similar trends wherever both are available. In levels, income poverty is less prevalent than consumption poverty in China but equally prevalent in Mexico and South Africa.

Overall, Figure 1 makes clear that poverty declined more dramatically in China, India, and Indonesia than in Mexico and South Africa. This takeaway is consistent with the stacked series in Appendix Figure A.1, where the first three countries swamped the last two. Due to lower initial poverty rates and smaller populations, Mexico and South Africa account for little of the decline of global poverty. Nevertheless, they add contextual variation to our sample and will be useful for understanding whether the anatomy of poverty decline in Asia extends to other parts of the world.

2.4 Historical context

To contextualize the rest of the analysis, we briefly describe significant economic and social developments in the countries and time periods spanned by the data. Appendix Table A.1 reports changes in major macroeconomic indicators during those intervals.

All countries generally saw sustained growth in real per capita incomes during the sample periods. Notable exceptions including a major recession in Indonesia in 1998 following the Asian Financial Crisis, and significant recessions in Mexico in 1986, 1995, and 2009.²⁵ Overall, average growth rates were fastest in China, India, and Indonesia—the poorer countries—and slower in Mexico and South Africa, particularly during the periods spanned by the panel sources. This

²⁵The 1998 recession in Indonesia is not evident in the poverty series in Figure 1 because we do not use the 1998 round of the SUSENAS, for which survey weights are not available. China was not substantially affected by the crisis and indeed contributed to bailing out some of its neighbors.

pattern lines up with the slower rates of poverty decline in the less-poor countries, noted above. Average inflation ranged from the mid- to high-single digits, with the exception of Mexico, where the higher average reflects the inflation spikes of the 1980s and 1990s.

As economies grew, labor generally shifted out of agriculture and into services. Patterns are more mixed with respect to industry, which employed an increased share of the workforce in China, India, and Indonesia but a reduced share in South Africa and Mexico during the periods for which we have panel data. South Africa’s labor market was notable for its high rate of unemployment, which hovered at approximately 20% through the late 2000s and by the end of our series in 2017 had edged even higher to 24%.²⁶ South Africa was also unusual in that it was hit particularly hard by the HIV/AIDS epidemic: the prevalence of HIV in the age 15–49 population grew from an estimated 7% in 1995 to 19% in 2017.²⁷ South Africa has typically ranked among the countries with the highest prevalence rates in the world.

In terms of the policy environment, the data span several episodes of economic liberalization. These include India’s relaxation of a wide range of economic regulations starting in 1991; ongoing market-oriented reforms in China during the 1990s such as the passage of the first Company Law in 1993 and a push towards privatizing state-owned enterprises in 1998–2000; the ratification of the North American Free Trade Agreement (directly affecting Mexico) in 1994; and China’s accession to the World Trade Organization in 2001. South Africa, meanwhile, undertook wide-ranging reforms focused on increasing equity following the end of apartheid in 1994. These included, for example, land redistribution under the Restitution of Land Rights Act of 1994, labor market reform via the Labor Relations Act of 1995, and policies to promote black business ownership. Overall, the data overlap with some of the most significant episodes of policy reform in recent history, at least as concerns poverty reduction.

3 Poverty decline within and between cohorts

We use the repeated cross-sectional data to explore the relative roles of *intra*-generational poverty decline, with the initially poor exiting poverty during their lifetimes, and *inter*-generational decline, with their non-poor children replacing them in the population. Because parent-child linkages are unavailable, generations are inexactly defined in our data, so our analysis focuses on birth cohorts, or groups of individuals born in the same year. We return later to the connection between cohort change and generational change. Two questions motivate our analysis of birth cohorts in the repeated cross-sectional data. First, what share of poverty decline occurred within cohorts, as opposed to between them? Second, was the intra-cohort rate of poverty decline faster or slower than the inter-cohort rate?

These questions are closely related, but the first is more delicate because it more obviously

²⁶World Bank Open Data, <https://data.worldbank.org/indicator/SL.UEM.TOTL.ZS?locations=ZA>, accessed 14 January 2025.

²⁷World Development Indicators, <https://databank.worldbank.org/reports.aspx?source=2&series=SH.DYN.AIDS.ZS&country=ZAF>, accessed 31 July 2024.

pushes us to take a stand on the role of the lifecycle. Within a cohort, poverty decline may reflect a stable age pattern or innovations over time, and these forces are not separately identified in the data.²⁸ For consumption poverty, a lifecycle consumption model could interpret within-cohort decline as the evolution of preferences with age, although to treat this interpretation as unique, one must assume away uncertainty or borrowing constraints, which are common in low- and middle-income economies. For income poverty, a lifecycle human capital model could interpret it as the return to experience or variation in on-the-job training. If these age forces play out similarly over the lives of all cohorts, then they can only contribute to aggregate poverty decline if the age structure of the population changes. A cohort-age approach thus largely rules out within-cohort progress as a source of aggregate poverty decline. Since we do not want to rule out a role for shared progress over time, we take a cohort-time approach. That is to say, we attribute aggregate poverty decline to either cohort replacement or the passing of time, acknowledging that the passing of time represents both age and period forces.

To this end, we begin in Section 3.1 by developing a cohort accounting framework and using it to derive our main cohort-time decomposition of poverty decline. We then use the repeated cross-sectional data to lay out patterns of poverty by age, period, and cohort in Section 3.2, followed by estimates of the decomposition in Section 3.3. The decomposition provides an intuitive way to attribute poverty decline to within- and between-cohort forces, and it also directly compares their relative pace. Our main takeaway, observable in several ways, will be that poverty declined between birth cohorts at roughly the same pace as it did within them. In our decomposition, this fact implies that the lion's share of poverty decline accrued over time within cohorts.

3.1 Cohort accounting framework

Consider a population made up of successive cohorts born in year b and observed at age a , or equivalently year $t = b + a$. At the end of each year t , the cohort aged A exits, and a new cohort aged 0 enters. For expositional simplicity, assume a rectangular population structure in which cohorts have common size and longevity. At the end of the section, we describe how to adjust our decomposition to allow for variable population shares. In Appendix A, we find that the adjustment leads to similar empirical results.

Let y be consumption or income, with distribution $F_b^t(y)$ in cohort b in year t . Let $\pi_b^t \equiv F_b^t(\bar{y})$ be the share below poverty line $\bar{y}$. From a lifecycle perspective, π_b^t is the poverty rate that cohort b experiences at age $a = t - b$. Because the number of living cohorts in any year is A , the overall poverty rate is:

$$\Pi^t = \frac{1}{A} \sum_{b=t-A}^{t-1} \pi_b^t \quad (1)$$

The equal size and longevity assumptions greatly simplify Equation (1), as each cohort is weighted by the constant $\frac{1}{A}$ instead of its unique population share.

²⁸One can separately identify them by restricting age or period effects, but theory suggests no general restrictions to apply to both consumption and income.

Equation (1) implies that one can express the change in aggregate poverty from $t - 1$ to t as the average of age-specific changes:

$$\Pi^t - \Pi^{t-1} = \frac{1}{A} \left(\sum_{b=t-A}^{t-1} \pi_b^t - \pi_b^{t-1} \right) \quad (2)$$

This expression is an age-time decomposition of poverty change, in the sense that it considers how year-to-year poverty decline accrues across ages. It asks how much poorer this year's 30-year-olds are than last year's 30-year-olds, and so on. These age-specific differences (and hence also aggregate poverty decline) may be small even if every cohort experiences substantial reductions in poverty during its lifecycle.

Rearranging and differencing out lagged terms yields our main cohort-time decomposition, which is useful for studying cohort dynamics:

$$\Pi^t - \Pi^{t-1} = \frac{1}{A} \left(\underbrace{\pi_{t-1}^t - \pi_{t-2}^{t-1}}_{\text{entry of by new cohorts}} + \sum_{b=t-A}^{t-2} \left\{ \underbrace{\pi_b^{t-1} - \pi_{b-1}^{t-1}}_{\text{replacement by surviving cohorts}} + \underbrace{\pi_b^t - \pi_b^{t-1}}_{\text{within-cohort change}} \right\} \right) \quad (3)$$

The first term captures the interannual difference in the youngest cohort's poverty rate. The second term captures the successive replacement of older cohorts by younger, after the oldest cohort disappears between $t - 1$ and t . The third term captures the trajectories of cohorts alive in both $t - 1$ and t , as they age and experience changing economic conditions. The second and third terms cleanly isolate between- and within-cohort change. The first term is less interpretable, mixing these two effects (because lagged poverty rates are unavailable for new cohorts), and is in a sense an artifactual result of observing outcomes only in discrete time; it disappears in the continuous-time limit. For these reasons, we will emphasize the second and third terms, which decompose poverty decline in surviving cohorts.

This cohort-time approach allows us to speak to some stylized economic possibilities. To build intuition for these possibilities, Appendix Figure A.2 visually compares how poverty rates play out over time and over the lifecycle in four stylized examples. In the first, each cohort enters the population at a lower poverty rate than its predecessor but then experiences a constant poverty rate over time.²⁹ Economically, this scenario could reflect rising schooling or other pre-labor market human capital investments, without any subsequent returns to experience or technological progress. In this case, the 'replacement by surviving cohorts' term in Equation (3) is negative, while the 'within-cohort' term is zero.

In the other three examples of Appendix Figure A.2, poverty falls within cohorts over time, so that in each case, the 'within-cohort' term in Equation (3) is negative. This circumstance could reflect technological progress or capital inflows, for instance, that affect all cohorts similarly. The

²⁹If poverty declines within but not between cohorts, then aggregate poverty does not change; poverty has a lifecycle pattern but no aggregate trend.

sign of the ‘replacement by surviving cohorts’ term varies, however. First, if each new cohort enters the population at a *lower* poverty rate than surviving cohorts, then the ‘replacement by surviving cohorts’ term is also negative. Here, the pace of poverty decline across cohorts exceeds the pace of poverty decline within cohorts. Economically, this pattern could reflect a combination of pre-labor market human capital investment and aggregate technological trends. Second, if each new cohort enters at the *same* poverty rate as surviving cohorts, then the ‘replacement by surviving cohorts’ term becomes zero. Here, the between-cohort pace of poverty decline matches the within-cohort pace. Third, if each new cohort enters at a *higher* poverty rate than the surviving cohorts, then the ‘replacement by surviving cohorts’ term turns positive. Here, the between-cohort pace lags the within-cohort pace. In sum, the sign of the ‘replacement by surviving cohorts’ term tells us whether the rate of improvement across cohorts exceeds, equals, or falls short of the rate of improvement as people get older.

We have assumed a rectangular population in which all cohorts have common size and longevity, but this assumption has little bearing on our results because, as we will show, changing population structure has played no role in poverty decline. If we allow population shares α_b^t to vary across cohorts and over time, then we can rewrite Equations (2)-(3) to have a poverty change component and a population change component. The poverty change component is the same as now, except that changes in poverty rates are weighted by lagged population shares α_{b-1}^{t-1} instead of the constant $\frac{1}{A}$. The population change component takes the form $\sum_b \pi_b^t (\alpha_b^t - \alpha_{b-1}^{t-1})$. The data will show that cohorts tend to have similar poverty rates π_b^t at a point in time. Since changes in population shares sum to zero across cohorts, it follows that the population change component will be close to zero. In Appendix A, we verify empirically that it is, and furthermore that the decomposition results do not change if we reweight the sum in Equation (3).

3.2 Tracking poverty by age, year, and cohort

Equation (2) asks how poverty decline was distributed across ages, while Equation (3) asks how it was distributed across birth cohorts. Figure 2 illuminates both perspectives by plotting how consumption poverty varied with age, both at a point in time and within a single cohort. The grey curves are cross-sectional age profiles, comparing differently aged individuals in the same year. These cross-sectional profiles tell us how young and old differ at a point in time, not about how the probability of living in a poor household changes over an individual’s lifecycle. The colored curves with markers are cohort age profiles, following individuals born in the same year—as they age through their respective lifecycles. To keep the figure uncluttered, we plot only every fifth birth cohort. The cohorts move diagonally across the cross-sectional age profiles as the latter shift vertically.

The most remarkable feature of the cross-sectional age profiles is their flatness. South Africa’s shows a moderate positive slope, but the others are broadly flat. At a point in time, the young, middle-aged, and old had similar propensities to be poor. In some countries, this propensity fell considerably over the sample period, but the decline manifested in parallel shifts of the cross-

sectional profile. Similar results obtain when we restrict the sample to household heads (Appendix Figure A.4), and when we use income poverty instead of consumption poverty in the three countries where data are available (Appendix Figures A.5 – A.6).

From this point of view, poverty decline was a shared experience, benefiting the young and old. One can see this point more directly by following Equation (2) and computing changes in age-specific poverty rates. Appendix Figure A.7 plots short and long differences in age-specific rates. The overall trend toward lower poverty, as well as the occasional transitory increase, was similar across the age distribution. The same is true in the heads-only sample (Appendix Figure A.8), with the exception of the youngest heads in Indonesia, ages 20-24, who saw less poverty decline but were highly selected; 20-24 year-olds were 13-35% as likely to be household heads as individuals in any other 5-year age group.

The flat cross-sectional profiles may seem at odds with the common conception that people are most likely to be poor early in their lives. The cohort lifecycle profiles clarify that this conception is true if we follow a group of individuals as they age. In all countries—and especially China, India, and Indonesia, where aggregate poverty fell most dramatically—most cohorts experienced pronounced declines in poverty over their lifecycles. Each cohort also has a lower intercept than its predecessor, so that the negatively-sloped lifecycle profiles shift downward across cohorts.

These patterns are the flip side of shared progress. Cohorts follow parallel trajectories as they age, and each cohort’s lifecycle trajectory lies beneath its predecessor’s. The cross-sectional profiles are flat because within-cohort change exactly keeps pace with between-cohort change. This combination is characteristic of lifecycle consumption profiles in developing economies that are experiencing growth. For example, Deaton and Paxson (1994) find offsetting slopes and intercept shifts in Taiwanese lifecycle consumption profiles from late 20th-century data. There as here, the young and old experienced similar household economic outcomes at a given point in time.

Equation (3) projects the cohort trajectories on time rather than age, which transposes them such that their slopes remain unchanged, but their level shifts shrink considerably. Figure 3 demonstrates this point, with year on the horizontal axis and the share poor on the vertical axis. The cohort profiles trend downward at the same rate, but they are now clustered together vertically. In contrast to the cohort-age trajectories in Figure 2, where intercept shifts implied a large role for cohort effects, the cohort-time trajectories here display no major level differences across cohorts, indicating a much smaller role for cohort effects.

The small inter-cohort differences at a point in time have nuanced implications for the role of generational replacement. Cohorts born 25 years apart, to use a conventional definition of a generation, did start their lives at very different rates of poverty. But by the time a younger cohort entered the population, its older counterpart had already closed the gap. As long as the two cohorts coexisted in the population, they experienced poverty decline in tandem. Put another way, the younger cohort was less poor over its lifetime but equally poor at the same moment in time.

3.3 Cohort-time decomposition

We next implement the cohort-time decomposition from Equation (3). Because most of the surveys are not annual, we adjust the decomposition to sum surviving cohorts from $t - A$ to $t - s - 1$ and new cohorts from $t - s$ to $t - 1$, where s is the most recent survey year that preceded t . This adjustment implies that the ‘entry of new cohorts’ component, which we do not emphasize because it has an ambiguous interpretation, will be larger in countries with longer intervals between surveys. The comparison of the ‘replacement by surviving cohorts’ component and the ‘within-cohort’ component is informative irrespective of inter-survey intervals.

For each country, we first compute the decomposition separately for each pair of adjacent surveys, then normalize by the number of years between surveys, and finally average across pairs, weighting by the number of years. Figure 4 reports the results in four separate panels, first for consumption poverty in the full sample, then for consumption poverty in the household heads, and finally for income poverty in both samples where possible. Each panel plots total annualized poverty decline in purple, followed by Equation (3)’s three components in other colors.

In all countries and all sample/outcome combinations in Figure 4, the ‘replacement by surviving cohorts’ component is small relative to the ‘within-cohort’ component, consistent with flat cross-sectional profiles that decline in parallel for all cohorts. One apparent anomaly is China, where the ‘entry of new cohorts’ component is consistently negative and fairly large. This result is a byproduct of CHIP’s long intersurvey intervals, which result in a large number of cohorts being classified as new. Because the ‘entry of new cohorts’ component has an ambiguous interpretation, the decompositions for China remain consistent with a primary role for within-cohort decline.

The decomposition results suggest cohort-neutral poverty decline: the cross-cohort pace of decline matched the within-cohort pace of decline. The cohort-time decomposition attributes the lion’s share of poverty decline to within-cohort change. Across all samples and poverty measures, within-cohort change accounts for more than half of poverty decline in surviving cohorts. In our preferred combination of consumption poverty for all individuals, the share exceeds 87% in all countries (Appendix Table A.2).

The baseline decomposition does not weight cohorts by their size, raising the question of whether our finding on the importance of within-cohort change is robust to relaxing the rectangular population assumption. Appendix A presents supplementary results suggesting that it is. Appendix Figure A.9 replicates Figure 4 allowing cohort population shares to vary. As expected given the flat cross-sectional poverty-age profiles, changing population structure does not contribute to poverty decline. And when we weight cohorts by their lagged population shares, we still find a large role for within-cohort change. In the main ‘Consumption, all’ decomposition, India displays the only notable change, with the within-cohort share among surviving cohorts going from 111% to 68% of poverty decline in surviving cohorts. Weighting matters more in India because some younger, larger cohorts made less progress early in their lives, a pattern visible in Figure 2. Nevertheless, India’s within-cohort component remains large, and other countries’ remain larger.³⁰

³⁰In the ‘Consumption, heads’ decomposition, India exhibits a larger drop in the within-cohort share, as does South

To more clearly summarize the role of within-cohort progress, Appendix Table A.2 compares within-cohort shares in various decompositions. Keeping with our emphasis on the components that afford a clean interpretation, the table reports the within-cohort component as a share of poverty decline in surviving cohorts only. Except for the ‘Consumption, heads’ decompositions in India and South Africa, within-cohort progress always accounts for a majority of poverty decline in both the unweighted and weighted decompositions.

The decomposition results also raise the question of whether the importance of within-cohort change is specific to poverty decline or more broadly felt across the distribution of consumption or income. To this end, Appendix Table A.2 also calculates the cohort-year decomposition for the logarithm of mean consumption or income, instead of the poverty rate. The within-cohort share remains large. On average too, living standards improved over people’s lives.

Our cohort-time decomposition sidesteps the role of the lifecycle, but we can partially address it by asking whether changes in measured cohort characteristics can explain the decline in age-specific poverty rates, as documented in Appendix Figures A.7 and A.8. We conduct such an analysis in Appendix Table A.3, finding that rising cohort completed education does not explain a meaningful part of age-specific poverty decline in the countries that experienced large declines. Specifically, we regress cohort-year poverty rates on calendar year as well as age fixed effects, with and without controlling for the shares completing primary and secondary school. Annualized rates of age-specific poverty decline, as estimated by the coefficient on calendar year, are not sensitive to the inclusion of cohort completed education. As before, this finding holds for consumption and income poverty in the full and heads-only samples.

In sum, most poverty decline—whether measured using consumption or income, whether indexed by the demographic characteristics of household heads or all household members—accrued within cohorts over time. The similarity of the results for consumption and income poverty mirrors the tracking of consumption to income observed in range of contexts, which researchers have taken to reject the most basic formulation of the lifecycle model, without uncertainty or borrowing constraints (Attanasio and Weber, 2010). If one incorporates these elements, then theory predicts that the timing of income matters for the timing of consumption; preferences do not fully determine the trajectory of consumption poverty within a household over time. In the next section, we will turn to panel data to investigate what happened to households’ earnings as they exited poverty. We will see substantial changes in income over the households’ lives, which will motivate us to examine their sources.

4 Within-cohort poverty dynamics

We turn next to examining within-cohort changes using panel sources. As discussed in Section 2, we emphasize income-based measures here, as we seek to understand what changes enabled households to increase their living standards, as opposed to simply measuring those living standards per se.

Africa, but these results are subject to concerns about selection into headship, as exhibited in Appendix Figure A.3.

Where possible, we also report Appendix results using a consumption-based definition of poverty.

We focus on the share poor among households rather than individuals, as above. In the panel surveys, we must deal with the added complexity that household composition and headship status may change across survey rounds. As our default approach, we limit the analysis to the original households from the baseline survey wave whom we also observe in the final survey wave. As Appendix Table A.4 reports, four of the five panels successfully tracked 75% or more of initial households in their final wave, while that in China (52%) tracked fewer.³¹ As noted earlier, the results in this sample are generally similar to those we obtain if we use all households from the baseline survey and study changes through the last round in which we observe them.

As for changes in the composition of the households we do observe, one simple way to assess how consequential they may have been for changes in poverty is to examine how much of the latter can be explained by the former in a statistical sense. Appendix Table A.5 reports results from regressions of changes in poverty status on changes in household size and the shares of household members in several age and sex categories. With one exception (using an income-based definition in Mexico, which leaves little poverty change to explain) the demographic shares explain a negligible share of overall poverty decline, and in some cases they explain a negative share (i.e. they predict that poverty increased when in fact it decreased). We discuss some further sensitivity checks on a case-by-case basis below.

4.1 Poverty transition rates

The net change in the poverty rate over time equals the sum of flows into and out of poverty. Suppressing cohort subscripts b (as the households in a given panel sample can be construed as constituting a single cohort), we can write this change as:

$$\Delta\pi^t = \pi^t - \pi^{t-1} = (1 - \pi^{t-1})p_0^t - \pi^{t-1}(1 - p_1^t) \quad (4)$$

where p_s^t denotes the probability of being poor in period t conditional on having poverty status s in period $t - 1$. Thus, $(1 - p_1^t)$ is the probability that a poor household escapes poverty, and p_0^t the probability that a non-poor household falls back into it.

Both probabilities are estimated to be quite high in the data; poor households often escaped poverty, and non-poor households often fell back into it. Figure 5a plots the estimated probabilities p_1^t and p_0^t of being poor in the last round conditional on poverty status in the first round using the income definition. For initially poor households, these estimates are substantially less than one in every country, ranging from 55% in rural India to as low as 33% in China. This result indicates widespread poverty exit. But the rates of entry into poverty among households that were *not* initially poor are also non-trivial: between 13% in Indonesia and 34% in Mexico. Looking at consumption (Figure 5b) does not fundamentally change this picture. If we take a less binary look

³¹In India, to be precise, the third wave of the panel successfully interviewed 91% of the rural households from the HDPI that had been selected for follow-up, as well as 83% of the IHDS baseline sample.

at the data and examine flows among multiple income (or consumption) categories, we again see noticeable movement in both directions, as Appendix Figures A.10a and A.10b illustrate.

A corollary of this “churn” is that aggregate poverty rates would have declined more had poverty exits been permanent. Table 2 quantifies this point, comparing the share of initially poor households who were still poor at endline with the share who were poor in *every* survey wave, who would have remained poor even if we eliminated the possibility of poverty re-entry. When calculated using income, the share always poor is substantially lower than the share poor at endline. In Indonesia, for example, 27% of the initially poor remained poor at endline, but only 11% remained poor throughout the panel. Thus, an additional 16% of the initially poor would have been non-poor at endline had poverty escapes been permanent. Income data may over-state churn in standards of living to the extent that incomes are volatile, and households smooth their consumption over time. But the picture is broadly the same when we use non-durable consumption, for the data sources that allow it. In Indonesia, 20% of the initially consumption-poor remained consumption-poor at endline, while 9% were persistently poor throughout the panel.

Accounting for durables is a harder problem. One might expect the flow of services that households receive from consumer durables—especially housing—to be less volatile than their consumption of non-durables. The common practice of omitting durables from measures of living standards may lead us to overstate poverty churn. To examine this issue further, we construct a novel measure of housing poverty. We focus on housing specifically because all of our panel datasets provide some measure of the value of the houses or housing services that households consumed, and because that value is substantial: even at a discount rate of 10% per annum, which we view as conservative, the value of housing services represents 23% to 60% of households’ total consumption on average.³²

We calibrate housing poverty lines for each survey and each {renter, owner} category such that the share of households that was housing-poor in the final round equals the share that was consumption-poor, and then use those lines to calculate how many households were housing-poor in other rounds of the survey.³³ Appendix B.6 defines these calculations precisely; their virtue is that they give us measures of the *relative* volatility of non-durable and durable consumption poverty without forcing us to take a stand on the discount rate, the methodological hurdle which has often deterred poverty researchers from including durables in their consumption aggregates. By this definition, housing poverty did indeed fluctuate somewhat less than (non-durable) consumption poverty (Columns 5–6 of Table 2). Revisiting Indonesia, for example, 13% of initially housing-poor households were always housing poor, as opposed to 22% who were housing poor at endline. This stability does not negate the volatility of non-durable consumption poverty, nor are the results in tension: one would expect consumption of housing to be less volatile, given the costs of adjusting it.

³²Appendix Table B.1 reports figures for this and other, higher interest rates, at which the estimated housing share is also higher. As a benchmark, the median (mean) rate in the MIX Market Intelligence database of microcredit lending rates during 2000–2019 was 17% (21%).

³³For China, we use income rather than consumption poverty for calibration, since the CHNS lacks data on consumption.

Another potential nuance is that the basic story—that households in poverty had a high probability of exiting it—may have been less true for some than for others. Some may have had more opportunity to make progress, while others were truly stuck. Because the panels extend over several rounds, we can examine this issue: if it were true, we would see the conditional probability of escaping poverty tending to fall over time, as more and more of those capable of exiting would have already done so.

The data say otherwise. Appendix Figure A.11a plots estimates of round-to-round transition probabilities over time in the sample of households observed in all rounds (see Appendix Table A.4 for sample composition information). Wherever the estimated probabilities change (in China, India, and to a lesser extent Indonesia), they move advantageously: the probability of exiting poverty conditional on starting in it rose, and the probability of entering poverty conditional on starting out of it fell. The other series are fairly flat.³⁴

Appendix Figure A.12 presents a more continuous view of the data, plotting local linear regressions of poverty in a given round conditional on income in the previous round, separately for each round. This approach removes any artifacts due to discretization—if most of the poor in one round were just slightly below the poverty line, for example, then the exit rate after that round might look artificially high. In practice, however, the general pattern is the same as in Appendix Figure A.11a: conditional probabilities of being poor given *any* initial level of income fell in China and India, but were roughly constant elsewhere. The local linear regression estimates are monotonically declining in initial income and display no particular change in slope near the poverty line.

One might reasonably wonder about the sensitivity of this exercise to measurement error. While the surveys we use are generally known for their high-quality fieldwork, even the most carefully implemented surveys are susceptible to measurement error. We cannot quantitatively correct for measurement error without making strong assumptions about its form. As a simple and transparent alternative, we examine the sensitivity of our conclusions to *adding* noise. In Appendix Figure A.13, we report how adding mean-zero noise to income and consumption affects the estimated shares poor at endline and in every wave, corresponding to the first two columns of Table 2.³⁵ We find that these statistics are not sensitive to added noise. This insensitivity has two intuitive explanations. First, the relationship between living standards in one round and poverty in the next round is not particularly steep in the neighborhood of the poverty threshold (see Appendix Figure A.12). Second, conditional on being non-poor in at least one round, most households were non-poor in several.

Another natural question is how much changes in household composition predict poverty transitions. The birth of a child will typically increase household size without increasing household income, for example, and so could cause a household to fall below the poverty line on a per capita basis. The results in Appendix Table A.5, described above, suggest that this phenomenon did not play a major role in long-run poverty change. To gauge how much it matters for round-by-round

³⁴Consumption is available in too few rounds to be of much use for this exercise, though we report the corresponding results in Appendix Figure A.11b for completeness.

³⁵Appendix B.7 describes the protocol for adding mean-zero noise.

transitions, we fix the actual trajectory of household income and consumption, but construct two counterfactual household sizes: the size of the household in the preceding period, or that quantity plus the *average* change in household size between the two periods. Appendix Figures [A.14a](#) and [A.14b](#) show that the levels of and trends in poverty transition rates are not substantially different under either of these counterfactual assumptions.

4.2 Changing livelihoods

We next examine in what ways the livelihoods of the households that exited (or entered) poverty changed. We find it helpful to think of this question in two steps: the mix of activities the household undertook in order to earn income, and the income they actually obtained from those activities. The two will tend to move together but need not always, as some efforts to earn money will have been more successful than others. A household member might seek employment, for example, but find none. And so on.

We code activities in a simple way, classifying each household based on the primary way it reported seeking to earn a living. This approach will facilitate exposition and also provide a logical basis for decomposing changes in poverty rates in Section [4.3](#) below. We classify households with respect to sector (agricultural or non-agricultural), occupation type (wage work or self-employment), migration (between urban and rural areas), and female labor force participation. The sector and occupation classifications are—with one exception that we note below—based solely on respondents’ own descriptions of the kinds of income-generating activities they primarily pursued, as opposed to how successful they were in doing so. For example, if a majority of household members reported working (or looking for work) in agriculture, then we classify the household as primarily agricultural, even if the majority of its realized income came from non-agricultural activities. Appendix [B.8](#) provides a full description of each of these classification procedures; in brief,

- The sectoral classification is based on the sectors within which individuals were reported to be economically active whenever that information is available, grouping agriculture together with other primary sector activities (e.g. forestry, mining). On average, these sector codes are missing for 6% of households, in which case we use the sector from which they obtained the majority of their realized income (see Table [B.2](#)).
- The occupational classification labels a household as primarily wage-seeking if (a) at least one adult is in the labor force, and (b) a majority of adults in the work force are either employed for a wage or seeking wage employment, and as primarily self-employed otherwise. The latter category includes a minority of households in which *no* adult is in the labor force; see Appendix [B.8.3](#) for further discussion. Both wage- and self-employment are common in the data: pooling across rounds and averaging across surveys, 50% of households have at least one member who identifies as self-employed, and 61% have at least one who identifies as a wage worker.
- The migrant classification (in the panels from Indonesia, Mexico, and South Africa that track

migration) first notes whether the household was located in different administrative regions in the first and final rounds of the panel, and then classifies them into four categories based on whether their places of origin and destination were rural or urban. Along with households that did not relocate from a rural or an urban place, this coding yields six categories in total.

- The female labor force classification distinguishes between households with at least one working-age woman in the labor force and those without. The latter category includes households with no working-age women; see Appendix B.8.4 for further discussion.

To characterize realized income, we calculate the shares of income that households obtained from different sources, doing so along dimensions that largely parallel how we classify activities. The sectoral split distinguishes between income from the agricultural and non-agricultural sectors. The factor split distinguishes between labor income, capital income, owned-enterprise income, and transfer income. One should think of owned-enterprise income as a mix of labor and capital income, which we cannot separate given the structure of the surveys and the bookkeeping standards of micro-enterprises. The gender split distinguishes between income earned by male and female household members, as well as income from sources such as a family-run farm that cannot be unambiguously attributed. Again, these decompositions involve some judgment calls, the details of which are in Appendix B.5.

Table 3 reports (changes in) the values of these quantities for the subset of households that exited poverty (odd-numbered columns) and those that entered it (even-numbered columns).³⁶ A few patterns are worth highlighting.

First, no one pathway out of poverty predominates. Only a minority of households that exited poverty did so while changing their status on any one of the margins we examine—sector, occupation, location, or female labor force participation. The one exception to this rule is China, where 55% of households that exited poverty experienced changes in female labor participation. As we discuss further below, in this case women mostly *withdrew* from the labor market, so the change is more likely to be a consequence of poverty decline than a cause.

Second, households that entered poverty often did so while changing their livelihoods activities in ways similar to households that exited it. For example, the share of households entering poverty while switching from agricultural to non-agricultural employment is about 2/3 or more of the share of households exiting poverty that did the same. But there are some exceptions. In Mexico and South Africa, for example, households that exited poverty were much more likely to switch from self- to wage-employment, while those that entered poverty were much more likely to do the opposite. This pattern with respect to occupational choice was directionally similar in other countries, but less pronounced.

Third, while the activities and income shares data generally tell similar stories, they move in opposite directions in a few interesting cases. In China, for example, households that exited poverty

³⁶If we define poverty transitions using consumption (not reported), the picture looks broadly similar. One cannot classify livelihoods using consumption data; the way households spent money is not informative about how they earned it. But changes in livelihoods activities and realized income shares look similar when we identify those who entered and exited poverty based on the amounts they consumed.

were roughly two and a half times more likely to do so while transitioning from primarily wage employment to primarily self-employment than in the opposite direction. Yet the average share of income that households obtained from wage income rose, while that from self-employment fell. In Indonesia, slightly more households exited poverty while transitioning into than out of agriculture, and yet the income share from agriculture fell substantially. These juxtapositions suggest that prices (e.g. changing wages) likely played an important role.

Fourth, outside of India, transfers largely played a “protective” role; rising transfer generosity was not a meaningful driver of poverty reduction. This conclusion is visible in the shares of income from public and private transfers, as documented in the bottom panel of Table 3. Among households that exited poverty, the share of income from both public and private transfers generally fell or rose slightly. Among those that entered poverty, on the other hand, the share of income from both public and private transfers always rose, in some cases substantially. In most countries, the data are thus consistent with redistribution playing a “safety net” role, mitigating the effects of churn on poverty. In India, however, the income share of transfers rose substantially for households that *exited* poverty, even more so than for those that entered. One might imagine that this pattern reflects the introduction of the landmark National Rural Employment Guarantee Scheme in 2005, but in fact the NREGS accounted for only 14% of the increase in generosity, with the rest coming from a range of pensions, scholarships, and other public assistance programs.

Appendix Table A.8 examines the role of public transfers more directly. The top panel reports the poverty rates we actually observe, alongside those we obtain if we deduct public transfers from household income. The bottom panel reports measures of the overall generosity and targeting accuracy of public transfers. Two generalizations emerge. First, transfers to the poor were typically more generous in the richer countries. In both Mexico and South Africa, poverty rates were substantially higher pre-transfer than post-transfer, indicating that redistribution lowered poverty in an accounting sense. In China and Indonesia, on the other hand, pre- and post-transfer rates were essentially identical. Second, redistribution did not become sufficiently *more* generous or better-targeted over time to account for much of the incremental change in poverty.³⁷ These generalizations have two notable exceptions. In South Africa, increased public transfer generosity accounted for a meaningful 3.8 out of a total 18.6 percentage point decline in the poverty rate. And—consistent with the results above—India defies both patterns. Despite being one of the poorer countries, it issued transfers sufficient (and sufficiently well-targeted) to meaningfully lower the poverty rate. And by the final round of the IHDS, actual poverty rates were 6.7 percentage points lower than they would have been without public transfers. India thus provides an important example in which a policy of increased redistribution contributed meaningfully to poverty reduction.

³⁷Public transfers did increase in all countries but Mexico, but targeting (as measured by the share of these transfers accruing to households below the \$2.15 threshold) worsened, perhaps because poverty targeting became harder as poverty rates fell.

4.3 Accounting for poverty decline

The results above give a general sense of what changes in livelihoods were associated with exit from or entry into poverty. We move next to quantifying how much net poverty reduction the observed changes in livelihoods activities can, in a purely accounting sense, explain. For example, we would like to know how much net poverty decline in Indonesia was attributable to households that shifted from working primarily in the agricultural to the non-agricultural sector, to those that stayed in agriculture, and so on.

We define our accounting framework as follows. Let $\alpha_{ss'}$ be the share of households that were of type s in one period and s' in a subsequent period, where for the sake of concreteness s might indicate working primarily in the agricultural ($s = a$) or non-agricultural ($s = n$) sector. A household that did not transition out of agriculture would be marked with the subscript aa , a household that transitioned from agriculture to non-agriculture would be marked with an , and so on. Then the expression for the overall rate of poverty transition in Equation (4) can be further decomposed as:

$$\Delta\pi = \underbrace{\alpha_{aa} \times \Delta\pi_{aa}}_{\text{intra: ag} \rightarrow \text{ag}} + \underbrace{\alpha_{nn} \times \Delta\pi_{nn}}_{\text{intra: non} \rightarrow \text{non}} + \underbrace{\alpha_{an} \times \Delta\pi_{an}}_{\text{inter: ag} \rightarrow \text{non}} + \underbrace{\alpha_{na} \times \Delta\pi_{na}}_{\text{inter: non} \rightarrow \text{ag}} \quad (5)$$

where the key $\Delta\pi_{ss'}$ terms are the net rates of poverty escape for households in the ss' category. This expression says that households can be grouped into four classes based on their sectoral transitions; the overall rate of poverty change is the weighted average of the rates within each of these groups. It is related to the decomposition introduced by Ravallion and Huppi (1991, Equation 4) (henceforth RH), the key difference being that Equation (5) uses panel data to directly measure quantities that they sought to infer indirectly from repeated cross-sections.³⁸ We report results from the RH decomposition in Appendix A for the sake of comparison.

Figure 6 reports the results obtained by applying Equation (5) to the data along the sectoral dimension. It also introduces a format that we will use repeatedly as we decompose poverty changes along various dimensions. For each country we plot four bars, corresponding to the four terms in Equation (5). The width of each bar represents the population share $\alpha_{ss'}$ in the corresponding

³⁸In fact, RH requires only (when applied to structural change, for example) that we know the share poor in the agricultural and non-agricultural sectors at two points in time. This simplicity may seem too good to be true, and indeed one can with a bit of algebra derive interpretations that are at odds with common sense. Using our α -notation for population shares, and denoting initial and final poverty rates in sector s by π_s^0 and π_s , respectively, the RH decomposition can be written as

$$\Delta\pi = \underbrace{(\alpha_{aa} + \alpha_{an})(\pi_a - \pi_a^0)}_{\text{intrasectoral: ag}} + \underbrace{(\alpha_{nn} + \alpha_{na})(\pi_n - \pi_n^0)}_{\text{intrasectoral: non}} + \underbrace{\pi_n^0(\alpha_{an} - \alpha_{na}) + \pi_a^0(\alpha_{na} - \alpha_{an})}_{\text{intersectoral shifts}} + \underbrace{\sum_s (\alpha_{ss'} - \alpha_{s's})(\pi_s - \pi_s^0)}_{\text{interaction}} \quad (6)$$

which can yield misleading inferences. For example, even if no households changed poverty status while changing sectors ($\Delta\pi_{an} = \Delta\pi_{na} = 0$), the intersectoral term will be non-zero provided only that initial poverty rates in the two sectors differed and there was net population movement between them.

group, and the height represents the net rate $\Delta\pi_{ss'}$ at which members of that group changed poverty status, with negative values indicating poverty exit. As a result, the product of the two, i.e. the area of the bar, indicates the total contribution of group ss' to poverty change. This approach lets us visually represent situations like, for example, when a group exited poverty at a high rate but was too small to make a meaningful contribution to aggregate decline, or vice versa.

With respect to structural change specifically, perhaps the most striking pattern concerns the role of transitions out of agriculture. As noted above, all of the countries we study saw substantial shares of their labor forces leave agriculture: a hallmark of the development process more generally. But households leaving agriculture do not account for the largest share of poverty decline—let alone the majority—in any of the countries we study. China comes closest, but even in China, more households exited poverty while remaining in agriculture, and the net rate of poverty exit was actually slightly higher among households that remained in agriculture than among those that left it. In Indonesia, the same comparisons hold true. In India, the net exit rate was slightly higher among households that left agriculture, but few such households existed, so that households that stayed in or out of agriculture accounted for the bulk of the decline. And in Mexico and South Africa, households working in agriculture no longer accounted for a substantial share of the population even at baseline.

Related to this point, households that stayed in the same sector contributed more to poverty decline than households that switched in every country. In China, India, and Indonesia, these households primarily stayed in agriculture. Continuing within agriculture was thus a viable route out of poverty for many millions of people. In Mexico and South Africa, on the other hand, households that stayed out of agriculture played the largest role, reflecting these economies' less agricultural starting points.

These patterns do not necessarily imply that structural change played little *causal* role in poverty decline. It is possible, for example, that households made progress out of poverty *within* the agricultural sector *because* of the growth of the non-agricultural sector, through structural transformation or demand channels. But the descriptive facts imply some bounds on the underlying economic possibilities. In order for these scenarios to fit the data, the supply or demand for labor in agriculture would need to be relatively inelastic. If growth in the non-agricultural sector pulled workers out of agriculture, then agricultural labor demand would need to be inelastic for agricultural wages to rise substantially. If growth in the non-agricultural sector raised the demand for agricultural goods and therefore agricultural labor, then agricultural labor supply would likewise need to be inelastic for agricultural wages to rise substantially.

Another interesting feature of Figure 6 is that households that shifted *into* agriculture also contributed to poverty reduction, Mexico being the one exception. Here, our ability to track specific households over time is crucial. If we ignored the panel structure and applied the RH decomposition instead, we would interpret these transitions as having *increased* poverty in every country except Indonesia (see Appendix Table A.6). The latter approach forces us to infer that households transitioning from a sector with lower to a sector with higher poverty were more likely

to themselves become poor, while of course economic logic suggests that they may have transitioned precisely because they saw some potential advantage to doing so.

Turning to occupational choice in Figure 7, we again see markedly different patterns in China, India, and Indonesia, as opposed to Mexico and South Africa. In the latter group, households that transitioned into wage work or stayed within wage work accounted for essentially all of the poverty reduction, while those that left wage work or continued not doing wage work contributed to poverty increases. In South Africa, the latter were not numerous but in Mexico they were, and they fared very poorly. Some of those that exited wage work become entirely economically inactive, but even those that become self-employed contributed to poverty increase (Appendix Figure A.15). In the less-advanced economies, meanwhile, the majority of progress out of poverty was among households that were *not* primarily employed for a wage, or that transitioned out of being so. This is especially so when we distinguish those who switched into self-employment, as opposed to market inactivity (Appendix Figure A.15). Overall, these patterns suggest that at relatively low levels of development, transitioning into self-employment can be a marker of progress, as for example in Banerjee and Newman (1993); in the somewhat more advanced economies, on the other hand, the self-employed are more likely to be entrepreneurs out of necessity (Schoar, 2010), and progress is more closely associated with getting and holding a good job.

Interregional migration generally played a small role in the countries for which we observe it, with one interesting exception, as shown in Figure 8. Typically, one thinks of rural-to-urban migration as a consequential but costly step in the process of development. The South African data look closest to this view: rural-to-urban migrants experienced a larger rate of net poverty reduction than households migrating between within urban or rural regions, but they constituted too small a share (6%) of the population to make a substantial contribution to aggregate poverty decline. In Indonesia, on the other hand, rural-to-rural migrants had the highest net rate of poverty exit of any group and accounted for nearly a third of all poverty decline in the sample. An important lacuna here is the absence of migrant tracking in the panels from the two countries that contributed the most to global poverty decline, China and India. The role of migration in those countries remains an open question.

Finally, we consider female labor force participation in Figure 9. Participation rates vary a great deal, both across countries (from 16.3% in India to 66.1% in South Africa; see Appendix Table B.7) and over time, with between 18% (India) and 54% (China) of households changing status between baseline and endline. In all countries, households in which a woman entered the workforce had among the highest rates of net poverty exit. In South Africa they were numerous enough to contribute meaningfully to poverty change, accounting for 43% of net exits. Households in which a woman *exited* the labor force contributed *negatively* to poverty decline in Mexico, and in all countries, they had a lower net exit rate than households in which a woman began to work.³⁹ Over all, for countries other than China, female labor force entry was uncommon but associated

³⁹India had relatively low rates of female labor force participation generally (Appendix Table B.7), as noted by Pande et al. (2017).

with poverty reduction, as one would expect given the mechanical increase in cash earnings it would yield.⁴⁰

China, however, is unusual. In China, households in which a woman began working still exited poverty at the highest rate of any group, but they were a small minority. They were greatly outnumbered by households in which a woman left the labor force, which contributed substantially to poverty decline. In fact, nearly half of the total reduction in poverty was attributable to households in which a woman stopped working.⁴¹ This result suggests selection effects strong enough to offset any loss of cash income, reminiscent of the pattern of the declining female labor force participation that [Goldin \(1995\)](#) documented during the early stages of development, and in contrast to other countries where poverty either declined little or increased among such households.⁴²

5 Conclusion

Advances in data availability and measurement have allowed researchers to document the momentous decline of poverty across the world since the early 1990s ([Chen and Ravallion, 2010](#); [World Bank, 2023](#)). We build on these advances to investigate the anatomy of this decline in five large countries that have especially extensive data. Our repeated cross-section results suggest that within-cohort forces played an important role in poverty decline. Our panel results suggest substantial within-cohort churn and in most cases a modest role for changes in sector, occupation, location, and female labor force participation in within-cohort poverty decline. These results contribute a dynamic portrait of the lives of the poor, complementing [Banerjee and Duflo’s \(2007\)](#) classic cross-sectional portrait.

The parameters we estimate here are not causal; they do not separate the effects of choices from selection into them. Nor could they: it would hardly be realistic to try to accomplish in one paper what has occupied much of the field of development economics in recent decades. But it may be fruitful to consider the relationship between the descriptive facts here and causal work, and in particular which causal work to prioritize in the future.

We think of this issue as follows. If households were always able to do what is in their best interest, then neither descriptive nor causal work would be useful for improving households’ lives. But households living in or near poverty are often unable to make the decisions that would be best for them because they lack crucial information, cannot access well-functioning capital markets, or face various other hurdles and constraints. Viewed through this economic lens, descriptive results

⁴⁰Here, as elsewhere, measured poverty accounts for the home production of marketable commodities but not for the value of home-produced services—such as cooking, childcare, cleaning, and elder care—which are often produced disproportionately by women.

⁴¹Female labor force participation has been declining in China more broadly, partly due to declining employment in the state-owned sector ([Xu and Zuo, 2025](#)).

⁴²A small share (16% in South Africa and less than 10% in the other countries) of households had no working-age woman. As noted above, [Figure 9](#) groups these households together with those that had working-age women who did not work. If instead we restrict to the sample of households that always had working-age female members, results are similar for all countries except South Africa. In South Africa, the restriction eliminates the poverty-reducing role of households that never had woman in the labor force, as most of these never had a working-age woman.

yield implications of (at least) three sorts.

First, observing people making meaningful progress while adopting a given approach rejects the extreme view that it is simply not possible. While such views are perhaps less common today, the history of thought about development contains some well-known examples. Analysts worried that certain economic pathways—subsistence agriculture, for example (Lewis, 1954; Rostow, 1991), or specialization in primary products (Prebisch, 1950)—were effectively dead ends that could not lead to further meaningful development. If nothing else, results documenting widespread progress out of poverty within sectors suggest that experimental work seeking to accelerate that progress need not be a fool’s errand.

Second, observing rapid progress by people who adopted a given strategy even while few others did likewise suggests that (financial) market failures bind. The very high poverty exit rate of the few South African households that moved from rural to urban locations, for example, could be read in this way. Other explanations may also apply: heterogeneity (in returns, preferences, or unobserved costs) could yield this pattern in the absence of frictions. But observing it shifts our posteriors towards the existence of meaningful frictions—frictions that experimental work could then seek to relax in the specific contexts where observational returns are high.

Third, observing people making progress *without* making any visible change—as we often do here—suggests a few possibilities. We may simply be looking at the wrong decisions. We have emphasized decisions that are widely measured and that development economics has, in our reading, tended to emphasize, but others may be relevant. In a related vein, people may have escaped poverty thanks to a multitude of small changes, not one big one. For example, the key may be not whether one leaves agriculture for manufacturing, but rather the composite benefits of changes to crop mix, input use, planting technique, and so on. This reading calls to mind Banerjee and Duflo’s (2011) concluding argument that “small changes can have big effects,” proposing a research agenda devoted to finding many small but reliable improvements. And general equilibrium effects may have been important. This interpretation argues for an increased emphasis on quantifying the general equilibrium consequences of intervention, whether via reduced-form analysis of large scale experiments (Muralidharan and Niehaus, 2017) or structural work that traces out the aggregate consequences of rigorously-identified microeconomic parameters (Buera et al., 2023).

References

- Aguiar, Mark and Erik Hurst**, “Deconstructing life cycle expenditure,” *Journal of Political Economy*, 2013, 121 (3), 437–492.
- Amendola, Nicola and Giovanni Vecchi**, “Durable goods and welfare measurement,” *Journal of Economic Surveys*, 2022, 36 (4), 1179–1211.
- Atamanov, Aziz, Christoph Lakner, Daniel Gerszon Mahler, Samuel Kofi Tetteh-Baah, and Judy Yang**, “The Effect of New PPP Estimates on Global Poverty: A First Look,” Global Poverty Monitoring Technical Note 12, The World Bank May 2020.
- Attanasio, Orazio P and Guglielmo Weber**, “Consumption and saving: models of intertemporal allocation and their implications for public policy,” *Journal of Economic literature*, 2010, 48 (3), 693–751.
- Badan Pusat Statistik (BPS)**, “SUSENAS Study Description,” <https://microdata.fao.org/index.php/catalog/1528/study-description>. Accessed 1 February 2026.
- Balboni, Clare, Oriana Bandiera, Robin Burgess, Maitreesh Ghatak, and Anton Heil**, “Why Do People Stay Poor?,” *The Quarterly Journal of Economics*, 12 2021, 137 (2), 785–844.
- Banerjee, Abhijit and Andrew Newman**, “Occupational choice and the process of development,” *Journal of Political Economy*, 1993, 101 (2), 274–298.
- **and Esther Duflo**, “The economic lives of the poor,” *Journal of Economic Perspectives*, 2007, 21 (1), 141–168.
- Banerjee, Abhijit V. and Esther Duflo**, *Poor Economics: A Radical Rethinking of the Way to Fight Global Poverty*, New York: PublicAffairs, 2011.
- Barrett, Christopher B, Michael R Carter, and C Peter Timmer**, “A century-long perspective on agricultural development,” *American Journal of Agricultural Economics*, 2010, 92 (2), 447–468.
- Baulch, Bob and John Hoddinott**, “Economic mobility and poverty dynamics in developing countries,” *Journal of Development Studies*, 2000, 36 (6), 1–24.
- Beegle, Kathleen, Joachim De Weerdt, and Stefan Dercon**, “Migration and economic mobility in Tanzania: Evidence from a tracking survey,” *Review of Economics and Statistics*, 2011, 93 (3), 1010–1033.
- Best, Michael C, Felipe Lobel, and Valdemar Pinho Neto**, “Cash Transfers and Productive Inclusion: Evidence from Bolsa Familia,” Working Paper 35006, National Bureau of Economic Research March 2026.
- Boserup, Ester**, *Woman’s Role in Economic Development*, London: Allen & Unwin, 1970.
- Bryan, Gharad and Melanie Morten**, “The aggregate productivity effects of internal migration: Evidence from Indonesia,” *Journal of Political Economy*, 2019, 127 (5), 2229–2268.
- Buera, Francisco J., Joseph P. Kaboski, and Robert M. Townsend**, “From Micro to Macro Development,” *Journal of Economic Literature*, June 2023, 61 (2), 471–503.


- Carolina Population Center**, “China Health and Nutrition Survey: Survey Design,” <https://www.cpc.unc.edu/projects/china/about/design/sample>. Accessed 1 February 2026.
- Carter, Michael R. and Christopher B. Barrett**, “The economics of poverty traps and persistent poverty: An asset-based approach,” *The Journal of Development Studies*, 2006, 42 (2), 178–199.
- Caselli, Francesco**, “Accounting for cross-country income differences,” in Philippe Aghion and Steven N. Durlauf, eds., *Handbook of Economic Growth*, Vol. 1, Elsevier, 2005, pp. 679 – 741.
- Chen, Shaohua and Martin Ravallion**, “The developing world is poorer than we thought, but no less successful in the fight against poverty,” *Quarterly Journal of Economics*, 2010, 125 (4), 1577–1625.
- Chetty, Raj, Nathaniel Hendren, Patrick Kline, Emmanuel Saez, and Nicholas Turner**, “Is the United States still a land of opportunity? Recent trends in intergenerational mobility,” *American Economic Review*, 2014, 104 (5), 141–147.
- Deaton, Angus**, *The Analysis of Household Surveys*, Baltimore: Johns Hopkins University Press, 1997.
- , “Measuring poverty in a growing world (or measuring growth in a poor world),” *Review of Economics and statistics*, 2005, 87 (1), 1–19.
- **and Christina Paxson**, “Saving, growth, and aging in Taiwan,” in David Wise, ed., *Studies in the Economics of Aging*, Chicago: University of Chicago Press, 1994, pp. 331–362.
- **and Olivier Dupriez**, “Purchasing power parity exchange rates for the global poor,” *American Economic Journal: Applied Economics*, 2011, 3 (2), 137–166.
- **and Valerie Kozel**, “Data and dogma: the great Indian poverty debate,” *The World Bank Research Observer*, 2005, 20 (2), 177–199.
- Dercon, Stefan and Douglas Gollin**, “Agriculture in African Development: Theories and Strategies,” *Annual Review of Resource Economics*, 2014, 6 (1), 471–492.
- Desai, Sonalde and Reeve Vanneman**, “India Human Development Survey-II (IHDS-II): Individual Schedule Questionnaire,” <https://ihds.umd.edu/sites/default/files/2022-08/ihds2isq.pdf> 2012. Accessed 1 February 2026.
- DiNardo, John, Nicole M Fortin, and Thomas Lemieux**, “Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach,” *Econometrica*, 1996, 64 (5), 1001–1044.
- Dollar, David and Aart Kraay**, “Growth is good for the poor,” *Journal of Economic Growth*, 2002, 7 (3), 195–225.
- Duflo, Esther**, “Women Empowerment and Economic Development,” *Journal of Economic Literature*, December 2012, 50 (4), 1051–79.
- Dupas, Pascaline, Pinelopi Koujianou Goldberg, and Rohini Pande**, “The changing landscape of international development: An introduction,” in Pascaline Dupas, Pinelopi Koujianou Goldberg, and Rohini Pande, eds., *Handbook of Development Economics*, Vol. 6, Elsevier, 2026.

- Eckert, Fabian and Michael Peters**, “Spatial Structural Change,” Technical Report, UC San Diego Jun 2025.
- Ferreira, Francisco H.G., Shaohua Chen, Andrew Dabalen, Yuri Dikhanov, Nada Hamadeh, Dean Jolliffe, and Ambar Narayan**, “A Global Count of the Extreme Poor in 2012: Data Issues, Methodology and Initial Results,” *Journal of Economic Inequality*, 2016, 14 (2), 141–172.
- Foster, Elizabeth Mary, Dean Mitchell Jolliffe, Gabriel Lara Ibarra, Christoph Lakner, and Samuel Kofi Tetteh Baah**, “Global Poverty Revisited Using 2021 PPPs and New Data on Consumption,” Policy Research Working Paper Series 11137, The World Bank Jun 2025.
- Gindling, T.H. and David Newhouse**, “Self-Employment in the Developing World,” *World Development*, 2014, 56, 313–331.
- Goldin, Claudia**, “The U-Shaped Female Labor Force Function in Economic Development and Economic History,” in T. Paul Schultz, ed., *Investment in Women’s Human Capital and Economic Development*, University of Chicago Press, 1995.
- Gollin, Douglas, David Lagakos, and Michael E. Waugh**, “The agricultural productivity gap,” *Quarterly Journal of Economics*, 2014, 129 (2), 939–993.
- Hamory, Joan, Marieke Kleemans, Nicholas Y Li, and Edward Miguel**, “Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata,” *Journal of the European Economic Association*, 11 2020, 19 (3), 1522–1555.
- Jayachandran, Seema**, “The Roots of Gender Inequality in Developing Countries,” *Annual Review of Economics*, 2015, 7 (Volume 7, 2015), 63–88.
- Jensen, Anders**, “Employment Structure and the Rise of the Modern Tax System,” *American Economic Review*, January 2022, 112 (1), 213–34.
- Jolliffe, Dean Mitchell, Daniel Gerszon Mahler, Christoph Lakner, Aziz Atamanov, and Samuel Kofi Tetteh-Baah**, “Assessing the Impact of the 2017 PPPs on the International Poverty Line and Global Poverty,” Policy Research Working Paper Series 9941, The World Bank Feb 2022.
- Jones, Chad I.**, “The facts of economic growth,” in John B. Taylor and Harald Uhlig, eds., *Handbook of Macroeconomics*, Vol. 2, Elsevier, 2016, pp. 3 – 69.
- Katz, Lawrence F and Kevin M Murphy**, “Changes in relative wages, 1963–1987: supply and demand factors,” *Quarterly Journal of Economics*, 1992, 107 (1), 35–78.
- Kraay, Aart**, “When is growth pro-poor? Evidence from a panel of countries,” *Journal of Development Economics*, 2006, 80 (1), 198–227.
- **and David McKenzie**, “Do Poverty Traps Exist? Assessing the Evidence,” *Journal of Economic Perspectives*, September 2014, 28 (3), 127–48.
- Krishna, Anirudh**, *One Illness Away: Why People Become Poor and How They Escape Poverty*, Oxford: Oxford University Press, 2010.
- Kuznets, Simon**, “Economic Growth and Income Inequality,” *The American Economic Review*, 1955, 45 (1), 1–28.

- Lagakos, David, Samuel Marshall, Ahmed Mushfiq Mobarak, Corey Vernot, and Michael E. Waugh**, “Migration costs and observational returns to migration in the developing world,” *Journal of Monetary Economics*, 2020, 113, 138–154.
- Lewis, W. Arthur**, “Economic Development with Unlimited Supplies of Labour,” *The Manchester School*, 1954, 22 (2), 139–191.
- Lipton, Michael**, *Why Poor People Stay Poor: Urban Bias in World Development*, Australian National University Press, 1977.
- Luxembourg Income Study**, “LIS Documentation: China – CHIP,” <https://www.lisdatacenter.org/wp-content/uploads/our-lis-documentation-by-cn02-survey.pdf> 2002. Accessed 1 February 2026.
- , “LIS Documentation: Mexico – ENIGH,” <https://www.lisdatacenter.org/wp-content/uploads/our-lis-documentation-by-mx10-survey.pdf> 2010. Accessed 1 February 2026.
- Mellor, John**, *Agricultural Development and Economic Transformation* Palgrave Studies in Agricultural Economics and Food Policy, Palgrave MacMillan, 2017.
- Merfeld, Joshua and Jonathan Morduch**, “Poverty at Higher Frequency,” Technical Report, NYU August 2024.
- Ministry of Statistics and Programme Implementation, Government of India**, “National Sample Survey: Variable Documentation,” <https://microdata.gov.in/NADA/index.php/catalog/85/variable/V18>. Accessed 1 February 2026.
- Muralidharan, Karthik**, *Accelerating India’s Development: A State-Led Roadmap for Effective Governance*, Gurugram: Viking, 2024.
- **and Paul Niehaus**, “Experimentation at scale,” *Journal of Economic Perspectives*, 2017, 31 (4), 103–124.
- Narayan, Deepa, Lant Pritchett, and Soumya Kapoor**, *Moving Out of Poverty: Volume 2. Success from the Bottom Up*, Washington, DC: World Bank, 2009.
- Pande, Rohini, Charity Troyer Moore, and Erin K Fletcher**, “Women and work in India: Descriptive evidence and a review of potential policies,” CID Working Paper 339 2017.
- Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman**, “Distributional national accounts: methods and estimates for the United States,” *Quarterly Journal of Economics*, 2018, 133 (2), 553–609.
- Popkin, Barry M, Shufa Du, Fengying Zhai, and Bing Zhang**, “Cohort Profile: The China Health and Nutrition Survey—monitoring and understanding socio-economic and health change in China, 1989–2011,” *International Journal of Epidemiology*, 2010, 39 (6), 1435–1440.
- Porzio, Tommaso, Federico Rossi, and Gabriella Santangelo**, “The human side of structural transformation,” CEPR Discussion Paper DP15110 2020.
- Prebisch, Raúl**, “The Economic Development of Latin America and Its Principal Problems,” ECLA Document, United Nations, Economic Commission for Latin America, New York 1950.


- RAND Corporation**, “Indonesian Family Life Survey (IFLS): Documentation and Questionnaires,” <https://www.rand.org/health/surveys/FLS/IFLS/download.html>. Accessed 1 February 2026.
- Ravallion, Martin and Monika Huppi**, “Measuring changes in poverty: A methodological case study of Indonesia during an adjustment period,” *World Bank Economic Review*, 1991, 5 (1), 57–82.
- **and Shaohua Chen**, “China’s (uneven) progress against poverty,” *Journal of Development Economics*, 2007, 82 (1), 1–42.
- **, Gaurav Datt, and Dominique Van de Walle**, “Quantifying absolute poverty in the developing world,” *Review of Income and wealth*, 1991, 37 (4), 345–361.
- Rostow, W. W.**, *The Stages of Economic Growth* Cambridge Books, none ed., Cambridge University Press, September 1991.
- Rubalcava, Luis and Graciela Teruel**, “Mexican Family Life Survey (MxFLS),” <https://catalog.ihnsn.org/catalog/7077/related-materials>. Accessed 1 February 2026.
- Schoar, Antoinette**, “The Divide between Subsistence and Transformational Entrepreneurship,” *Innovation Policy and the Economy*, 2010, 10, 57–81.
- Schultz, T. W.**, *Transforming traditional agriculture.*, New Haven: Yale University Press, 1964.
- Shukla, Rajesh**, *How India Earns, Spends and Saves: Unmasking the Real India*, New Delhi: Sage, 2010.
- Southern Africa Labour and Development Research Unit (SALDRU)**, “NIDS Wave 5 Panel User Manual,” Technical Report, University of Cape Town 2018. Accessed 1 February 2026.
- **, “National Income Dynamics Study (NIDS),”** <https://microdata.worldbank.org/index.php/catalog/2595/related-materials>. Accessed 1 February 2026.
- Statistics South Africa**, “Income and Expenditure Survey 2010/2011: Household Questionnaire,” <https://isibaloweb.statssa.gov.za/metadata/IES/2011/IES%20HHQuestionnaire.pdf> 2011. Accessed 1 February 2026.
- Strauss, John and Firman Witoelar**, “Indonesia Family Life Survey,” in Danan Gu and Matthew E. Dupre, eds., *Encyclopedia of Gerontology and Population Aging*, Springer, 2022, pp. 2600–2605.
- World Bank**, “Consumer Price Index (2010 = 100),” <https://data.worldbank.org/indicator/FP.CPI.TOTL.ZG>. Series FP.CPI.TOTL. Accessed [date].
- **, “PPP Conversion Factor, Private Consumption,”** <https://data.worldbank.org/indicator/PA.NUS.PRVT.PP>. Series PA.NUS.PRVT.PP. Accessed [date].
- World Bank**, “Poverty and Inequality Platform Methodology Handbook,” Available at <https://datanalytics.worldbank.org/PIP-Methodology/> 2023.
- Xu, Yuli and Sharon Xuejing Zuo**, “Gendered Impacts of Privatization: A Life Cycle Perspective from China,” SSRN Working Paper 5129928 2025.

Figure 1: Poverty rates over time, by sample and measure




Note: Repeated cross-sectional data. A household is defined as poor if consumption or income is below \$2.15 per person per day. The “all” series represent the share of all individuals living in poor households. The “heads” series represent the share of household heads living in poor households.

Figure 2: Consumption poverty over age, by year and cohort




Note: Repeated cross-sectional data. The grey curves are cross-sectional age profiles. The connected scatterplots illustrate poverty over the lifecycle as experienced by cohorts born in years ending in 0 and 5.

Figure 3: Consumption poverty over year, by cohort



Note: Repeated cross-sectional data. The connected scatterplots illustrate poverty over time as experienced by as experienced by cohorts born in years ending in 0 and 5. The figure is different from Figure 2 because cohort poverty rates are projected on time rather than age.


Figure 4: Cohort decomposition of poverty decline




Note: Repeated cross-sectional data. The figure implements the cohort-year decomposition in Equation (3) separately for each country.

Figure 5: Poverty transition probabilities

(a) Income definition




(b) Consumption definition




This figure reports estimated transition probabilities between the first and last round of each panel survey: specifically, the probability that a household was poor in the final round conditional on being poor or not poor, respectively, in the initial round. For comparability with Figures A.11a and A.11b the sample includes only households observed in all survey rounds. Poverty is defined in the top (bottom) panel as having income (consumption) per capita per day below \$2.15 in 2017 USD.

Figure 6: Structural change and poverty status




This figure reports changes in poverty headcounts between the first and last round of the panel surveys, broken down by groups defined based on whether the household's primary source of income was in the agricultural or non-agricultural sector in each round. For example, the left-most bar represents Chinese households that began and ended the panel in the agricultural sector. Its width is the share of the total sample these households represent, while its height is the net rate at which they exited poverty. Where space permits, we also annotate each bar with the percent of the total net change in headcount poverty within the given country for which the group accounted. In cases where poverty increased in some groups, the sum of these terms may exceed 100%. The sample consists of households observed in both the first and last waves of each survey. See Appendix B.8.1 for details on the classification of households into sectors.

Figure 7: Occupational choice and poverty status




This figure reports changes in the poverty rate broken down as in Figure 6, but with groups defined based on whether or not the household earned its income primarily from wage employment. See Appendix B.8.3 for details on this classification. Note that households in which no member reported being in the labor force are grouped here with others that were not primarily wage-employed; Figure A.15 reports a more detailed breakdown that further distinguishes between these groups.

Figure 8: Interregional migration and poverty status



This figure reports changes in the poverty rate broken down as in Figure 6, but with groups defined based on whether or not the household changed location via migration and if so whether to or from a rural or urban area. See Appendix B.8.2 for details on this classification. Note that migration data are not available for China or India.

Figure 9: Female outside employment and poverty status



This figure reports changes in the poverty rate broken down as in Figure 6, but with groups defined based on whether or not the household had at least one adult female member working outside of the home. See Appendix B.8.4 for details on this classification.

Table 1: Data sources

Country	Panel surveys		Repeated cross-section surveys	
	Name	Description	Name	Description
China	China Health & Nutrition Survey (CHNS)	10 waves, 1989-2015, 8/22 provinces, 1,922 households	Chinese HH Income Project (CHIP)	4 waves, 1988-2013, national, 15,417 households
India	India Human Development Survey (IHDS)	3 waves, 1993-2012, national*, 34,639 households	National Sample Survey (NSS), thick rounds	5 waves, 1988-2012, national**, 115,023 households
Indonesia	Indonesia Family Life Survey (IFLS)	5 waves, 1993-2015, 13/27 provinces, 6,019 households	National Socio-economic Survey (SUSENAS)	21 waves, 1994-2015, national, [†] 254,528 households
Mexico	Mexico Family Life Survey (MxFLS)	3 waves, 2002-2012, national, 7,182 households	National Survey of HH Income & Spending (ENIGH)	14 waves, 1989-2014, national, 16,275 households
South Africa	National Income Dynamics Study (NIDS)	5 waves, 2008-2017, national, 5,391 households	Income/Expenditure & Living Conditions Surveys (IES & LCS)	5 waves, 1995-2015, national, [‡] 25,129 households

This table describes the coverage of the datasets in our core sample. Notes are as follows:

* The 1993 round is the 1993-4 Human Development Profile of India, which was rural-only but covered the same households as the rural portion of the subsequent 2005 and 2012 rounds of the IHDS.

** We omit the 1999 NSS due to controversies over its design; see [Deaton and Kozel \(2005\)](#).

[†] We omit the 1998 SUSENAS because it lacks sampling weights.

[‡] We omit the 2005 IES because it categorizes age in 5-year intervals.

Table 2: Subsequent poverty among initially poor households

Country	Waves	Income definition		Consumption definition		Housing definition	
		At Endline	Always	At Endline	Always	At Endline	Always
China	10	29%	4%			32%	20%
India	2	50%	39%				
Indonesia	5	27%	11%	20%	9%	22%	13%
Mexico	3	50%	37%	69%	57%	69%	72%
South Africa	5	36%	19%	46%	30%	42%	41%

This table reports the share of initially poor households who were still poor in later panel waves. It is intended to illustrate the consequences of churn for aggregate poverty reduction. The sample includes initially poor households observed in all panel waves. Poverty is defined on the basis of income, consumption of non-durables, or the value of housing services consumed. Details on the construction of the latter measure are in Appendix B.6. By construction, it yields actual poverty rates identical to those based on non-durable consumption (or, in China, income) except in cases where mass points in the underlying distributions prohibit an exact match. Consumption-based estimates are omitted for China because the CHNS did not collect non-durables consumption data and for India because we observe only two rounds of consumption data.


Table 3: Livelihood changes among households whose poverty status changed

	China			India		Indonesia		Mexico		South Africa	
	Exits	Entries		Exits	Entries	Exits	Entries	Exits	Entries	Exits	Entries
Share of sample	54	5		27	9	48	5	22	20	29	11
Sector	Any change	34		30	29	29	31	15	12	13	10
	Ag to not	37		21	15	15	12	10	6	8	8
	Not to ag	3		8	14	14	19	5	6	5	2
Occupation	Any change	44		23	25	31	40	38	41	39	30
	Self to wage	13		10	7	12	10	32	6	37	19
	Wage to self	31		12	18	18	31	6	35	2	11
Migration	Any change										
	Rural to urban					13	11	1	1	5	3
	Urban to rural					11	11	0	0	2	1
Female LFP	Any change	55		22	16	29	22	29	28	42	39
	LFP to not	47		9	10	12	13	4	21	3	10
	Not to LFP	8		12	5	18	9	25	7	38	30
Income shares	Ag	-26	-9	-10	7	-36	-24	2	0	-2	-2
	Wage	42	1	-7	-9	-10	-23	22	-21	31	-19
	Self Employment	-23	-16	-20	-8	-4	-1	-7	1	-4	2
	Capital	3	1	0	0	1	0	-1	-1	1	-1
	Public Transfers	-2	10	21	17	2	4	-9	11	-23	14
	Private Transfers	3	25	6	0	11	19	-36	10	-4	5
	Male	6	-12	-25	-14	-8	-34	34	-25	19	-14
Female	Female	4	-20	0	1	3	-1	14	4	11	-2
	Unattributed	-10	32	25	13	5	35	-47	22	-30	15

This table reports changes in the livelihoods of households whose poverty status changed between the baseline and endline rounds of the panel surveys. Columns labelled “Exits” refer to households that exited poverty between those rounds, and those labelled “Entries” refer to households that entered poverty between those rounds. The top panel indicates the share of all households interviewed in the baseline round that fall into these two categories. The middle panel reports the share of households in the given column category that experienced the indicated change (e.g. among households in China that exited poverty 40% also changed sector). The bottom panel reports changes in the share of income that households in the given column category received from various sources (e.g. among households in China that exited poverty the average share of income obtained from agriculture fell by -26%). Blank entries imply the absence of data. For example, only the IHDS and NIDS have questions about non-rental capital income. A small share of households with poverty classification changes report any capital income, 1% in the IHDS and 4% in the NIDS data. Details on the classification of income and of primary income-generating activities are in Appendices B.5 and B.8, respectively.


A Additional exhibits

Figure A.1: Poverty decline in sample countries and the world




Note: Series are stacked in increasing order of their contribution to the world poverty headcount. Poverty is defined as living in a household with consumption per capita below \$2.15 per day. The left panel plots World Bank estimates; the right panel replaces the World Bank's country-specific estimates with our estimates from repeated cross-sectional data.

Figure A.2: Typology of cohort poverty decline




Note: Four types of cohort poverty decline, all consistent with the same aggregate poverty decline. The first column plots cohort poverty rates over the lifecycle, with lighter colors indicating later cohorts. Grey curves correspond to cross-sectional poverty-age profiles. The second column plots cohort poverty rates over time, with the same color scheme. The third column plots aggregate poverty rates over time. By construction, the third column is identical across all scenarios.

Figure A.3: Household headship by age and sex




Note: Pooled repeated cross-sectional data. Share of individuals who are household heads by age and sex.

Figure A.4: Consumption poverty over age, heads only, by year and cohort




Note: Counterpart to Figure 2, restricting to household heads only. Repeated cross-sectional data. The grey curves are cross-sectional age profiles. The connected scatterplots illustrate poverty over the lifecycle as experienced by cohorts born in years ending in 0 and 5.

Figure A.5: Income poverty over age, by year and cohort




Note: Counterpart to Figure 2 using data on income poverty. Repeated cross-sectional data. The grey curves are cross-sectional age profiles. The connected scatterplots illustrate poverty over the lifecycle as experienced by cohorts born in years ending in 0 and 5.

Figure A.6: Income poverty over age, heads only, by year and cohort




Note: Counterpart to Figure 2 using data on income poverty, restricting to household heads only. Repeated cross-sectional data. The grey curves are cross-sectional age profiles. The connected scatterplots illustrate poverty over the lifecycle as experienced by cohorts born in years ending in 0 and 5.

Figure A.7: Changes in consumption poverty by age




Note: Repeated cross-sectional data. Curves represent changes between adjacent surveys. Scatterplot represents the change from the first to the last survey wave.

Figure A.8: Changes in consumption poverty by age, heads only



Note: Counterpart to Figure A.7, restricting to household heads only. Repeated cross-sectional data. Curves represent changes between adjacent surveys. Scatterplot represents the change from the first to the last survey wave.


Figure A.9: Cohort decomposition with variable population shares



Note: Counterpart to Figure 4 allowing cohorts to have variable population shares. Repeated cross-sectional data.

Figure A.10: Transitions in living standards


(a) Income definition



(b) Consumption definition


China

No expenditure data collected by the CHNS.




This figure illustrates transitions between different living standards categories between the first and last rounds of each panel survey (conducted in the years denoted at the bottom of each pane). The sample includes households observed in both of those rounds. Living standards are measured using income (consumption) per capita per day in the top (bottom) panel. Living standards categories are: at or below the \$2.15 international poverty line (IPL), between 1x and 2x the IPL, and above 2x the IPL.

Figure A.11: Evolution of transition probabilities over time




This figure reports the round-by-round probabilities of being poor in the subsequent round of each panel survey, conditional on poverty status in a given round. (Note that this implies that transition probabilities out of the final round of the survey are not observed.) The sample includes only households observed in all survey rounds. Poverty is defined in the top (bottom) panel as having income (consumption) per capita per day below \$2.15 in 2017 USD. Note that no consumption data are available in the 1993–4 Human Development Profile of India or in any round of the CHNS, so that corresponding transition probabilities are not reported.

Figure A.12: Transition probabilities as non-parametric functions of initial income



This figure reports the round-by-round probabilities of being poor in the subsequent round of each panel survey, as a smooth function of income in a given round. Results for each survey wave are plotted as a distinct line. Nonparametric fits were estimated using local linear regression applied to the full data from each pair of rounds, but for the sake of legibility are plotted here over a limited range from the 0th to 95th percentiles of the initial income distribution. Vertical lines indicate the international poverty line, i.e. \$2.15/day in 2017 USD.


Figure A.13: Sensitivity of poverty rates to added noise




This figure reports the estimated shares poor at endline, and in every wave of the survey, after adding mean-zero noise to household per-capita income following the protocol defined in Appendix B.7. Figures labelled “0%” correspond to those in the first two columns of Table 2; those labeled “10%” and “20%” are the average values obtained from 1,000 simulations which introduced a corresponding amount of error in absolute value as a share of the overall mean.

Figure A.14: Transition probabilities with counterfactual household sizes

(a) Income definition




(b) Consumption definition



This figure reports both the actual round-by-round transition probabilities observed in the data, and counterfactual probabilities based on alternative definitions of household size. Specifically, the “Lagged HH Size” is constructed by setting the size of each household in each round to its size in the previous round, and the “Lagged size + mean growth” series is constructed by adding to that quantity the mean change in household size between rounds.

Figure A.15: Occupational choice and poverty status, splitting non-workers from non-wage workers



Note: Statistics were calculated on the subset of households observed in both the first and last survey waves. Households were classified into mutually exclusive categories of being primarily self employed, wage employed or not having any classifiable members and this standing was allowed to change over time which enabled the observation of category-switchers.

Table A.1: Macroeconomic context

Country	Data type	Date range	%Δ GDPPC	%Δ CPI	%Δ labor shares		
					Agr.	Ind.	Srv.
China	Cross-sectional	1988–2013	8.4%	4.8%	−29% [†]	+9% [†]	+19% [†]
	Longitudinal	1989–2015	8.5%	4.1%	−32% [†]	+8% [†]	+23% [†]
India	Cross-sectional	1983–2009	3.8%	7.3%	−11% [†]	+6% [†]	+5% [†]
	Longitudinal	1993–2012	4.6%	7.1%	−16%	+10%	+6%
Indonesia	Cross-sectional	1994–2018	3.0%	8.8%	−15%	+4%	+11%
	Longitudinal	1993–2015	3.0%	9.5%	−13%	+4%	+8%
Mexico	Cross-sectional	1984–2014	0.6%	18.2%	−12% [†]	+1% [†]	+11% [†]
	Longitudinal	2002–2012	0.6%	4.2%	−3%	−3%	+6%
South Africa	Cross-sectional	1995–2010	2.1%	5.6%	−6%	−5%	+12%
	Longitudinal	2008–2017	0.3%	5.4%	−1%	−2%	+4%

This table summarizes macroeconomic changes in the countries and during the periods we study. “%Δ GDPPC” is the average annualized percentage change in GDP per capita, based on data from the World Development Indicators. “%Δ CPI” is the average annualized percentage change in a consumer price index based on data from the World Bank Open Data platform (data.worldbank.org, series `FP.CPI.TOTL`) accessed 30 July 2024. The columns headed “%Δ labor shares” report the total changes in the share of the labor force employed in agricultural, industry, and services, respectively, based on data from the World Bank Open Data platform (series `SL.AGR.EMPL.ZS`, `SL.IND.EMPL.ZS` and `SL.SRV.EMPL.ZS`), which are in turn based on modeled estimates provided by the International Labor Organization (<https://ilostat.ilo.org/data/>). Entries marked with a [†] are changes from 1991, the first year in which labor share data are available, until the last year of the survey.

Table A.2: Within-cohort share in alternative decompositions

	Unweighted poverty	Weighted poverty	Unweighted growth
A. Consumption, all			
China	104%	114%	102%
India	111%	68%	117%
Indonesia	95%	96%	94%
Mexico	88%	83%	110%
South Africa	130%	180%	131%
B. Consumption, heads			
China	121%	116%	115%
India	100%	49%	97%
Indonesia	85%	86%	77%
Mexico	71%	71%	63%
South Africa	54%	29%	57%
C. Income, all			
China	102%	112%	101%
India			
Indonesia			
Mexico	89%	73%	114%
South Africa	120%	144%	114%
D. Income, heads			
China	121%	117%	112%
India			
Indonesia			
Mexico	87%	73%	108%
South Africa	91%	80%	92%

This table reports the within-cohort share of economic progress across alternative cohort-time decompositions. The first column is based on the baseline poverty decomposition in Equation 3, as reported in Figure 4. The second column is based on the extension of Equation 3 that allows for variable cohort population shares, as described methodologically at the end of Section 3.1 and reported in Figure A.9. The third column repeats the unweighted version of Equation 3 but with log mean consumption or income as the dependent variable, rather than a poverty indicator.

Table A.3: Cohort educational attainment and the annual rate of decline in age-specific poverty

	Coefficient on year		Cohort-year cells
	Age FE	Age FE + education shares	
A. Consumption, all			
China	-3.16 (0.06)	-3.44 (0.16)	183
India	-1.22 (0.03)	-1.30 (0.12)	305
Indonesia	-3.27 (0.02)	-3.18 (0.05)	1,281
Mexico	-0.27 (0.01)	-0.11 (0.02)	854
South Africa	0.26 (0.06)	0.60 (0.10)	183
B. Consumption, heads			
China	-2.98 (0.06)	-3.49 (0.15)	183
India	-0.80 (0.04)	-0.76 (0.14)	305
Indonesia	-3.27 (0.02)	-3.10 (0.05)	1,281
Mexico	-0.28 (0.01)	-0.12 (0.03)	854
South Africa	0.26 (0.07)	0.91 (0.12)	183
C. Income, all			
China	-2.93 (0.04)	-2.44 (0.07)	183
India			
Indonesia			
Mexico	-0.33 (0.01)	-0.16 (0.03)	854
South Africa	-1.24 (0.06)	-1.19 (0.12)	183
D. Income, heads			
China	-2.73 (0.04)	-2.53 (0.08)	183
India			
Indonesia			
Mexico	-0.32 (0.01)	-0.14 (0.03)	854
South Africa	-1.14 (0.08)	-0.97 (0.17)	183

This table reports results from regressions of the share living in a poor household (measured in percentage points) on calendar year with age fixed effects. The coefficient on calendar year corresponds to the annualized rate of poverty decline. The unit of observation is the cohort-year cell; the regressions weight cells by their share of the population in the relevant year. The sample excludes surveys lacking questions about completed education: China 1988 and South Africa 1995 and 2000. Panels A and B rely on consumption poverty, while Panels C and D rely on income poverty where available. Panels A and C rely on all household members, while Panels B and D rely on heads only. The rightmost column additionally controls for the share of individuals or heads who left school after completing primary school and the share who completed secondary school or higher.

Table A.4: Missingness in panel data sources

Country	Waves	Household-year observations	Unique households present		
			At baseline	At baseline and endline	In all rounds
China	10	27,903	3,697	1,922	1,403
India (National)	2	69,278	34,639	34,639	34,639
India (Rural Only)	3	31,432	10,792	9,848	9,848
Indonesia	5	32,910	6,975	6,019	5,704
Mexico	3	22,562	7,869	7,182	6,824
S. Africa	5	29,952	7,195	5,391	4,487

This table describes tracking in the panel surveys. In most of our analysis we use the sample of households observed at both baseline and endline; in some cases we use the more restricted sample observed in all rounds. One case requires some further explanation: for the India (rural) sample, households present “at baseline” refers to those surveyed in 2005 as part of wave 1 of the IHDS, and not to the larger rural sample of 33,230 households surveyed by the HDPI in 1993–4. The difference is largely accounted for by the fact that the IHDS intentionally set out to survey a subset of the HDPI households, but we do not have access to details on this sub-sampling which would let us calculate how many HDPI households not surveyed by the IHDS in 2005 were not surveyed due to attrition. We therefore treat the 2005 IHDS sample as the effective baseline sample.

Table A.5: Changes in household demographics and changes in poverty status

	Income definition					Consumption definition				
	China	India	Indonesia	Mexico	S. Africa	China	India	Indonesia	Mexico	S. Africa
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Males under 18	-0.20** (0.09)	-0.02 (0.03)	0.02 (0.04)	0.08 (0.07)	0.002 (0.06)		-0.05** (0.03)	-0.03 (0.04)	-0.04 (0.06)	-0.04 (0.05)
Males 18-70	-0.24*** (0.08)	-0.41*** (0.02)	-0.23*** (0.04)	-0.31*** (0.06)	-0.20*** (0.05)		-0.28*** (0.02)	-0.13*** (0.04)	-0.09* (0.05)	-0.12*** (0.04)
Males over 70	-0.20* (0.11)	-0.14*** (0.04)	-0.13** (0.06)	-0.10 (0.08)	-0.30*** (0.08)		-0.15*** (0.04)	0.07 (0.06)	-0.11 (0.07)	-0.10 (0.07)
Females 18-70	-0.29*** (0.09)	-0.31*** (0.02)	-0.09** (0.04)	-0.10* (0.05)	-0.20*** (0.05)		-0.24*** (0.02)	-0.01 (0.04)	-0.11** (0.05)	-0.09** (0.04)
Females over 70	-0.14 (0.11)	-0.16*** (0.04)	-0.01 (0.06)	-0.02 (0.07)	-0.24*** (0.06)		-0.18*** (0.04)	0.03 (0.05)	-0.05 (0.07)	-0.07 (0.06)
Females under 18										
Household Size	0.01 (0.01)	0.01*** (0.001)	0.01*** (0.003)	-0.03*** (0.005)	0.02*** (0.004)		0.04*** (0.001)	0.03*** (0.003)	0.01* (0.004)	0.06*** (0.004)
Constant	-0.45*** (0.02)	-0.15*** (0.003)	-0.41*** (0.01)	0.002 (0.01)	-0.18*** (0.01)		-0.16*** (0.003)	-0.40*** (0.01)	0.06*** (0.01)	-0.09*** (0.01)
Mean poverty change	-0.49	-0.18	-0.44	-0.01	-0.19		-0.19	-0.43	0.06	-0.08
$\vec{\beta} \cdot \vec{X}$	-0.03	-0.02	-0.02	-0.01	-0.01		-0.01	0	-0.01	0
Observations	1,888	34,638	5,961	6,761	5,388		34,638	5,961	6,761	5,388
Adjusted R ²	0.01	0.03	0.02	0.01	0.02		0.05	0.03	0.001	0.08

Note: This table reports the results of regressions that estimate associations between changes in poverty status (the dependent variable) and changes in demographic composition (the independent variables). Each column reports the results of separate regression, and each row reports the coefficient on a regressor defined as the change in the number of household members in the category stated between the first and last rounds of the survey. The dependent variable in all regressions is the change in an indicator for being poor based on income (Columns 1-5) or consumption (Columns 6-10). Robust standard errors are in parenthesis, and statistical significance is denoted as: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$. The row $\vec{\beta} \cdot \vec{X}$ reports the inner product of the estimated coefficients (sans the intercept) and the mean of the independent variables, so that by definition this quantity plus the estimated intercept must equal the mean of the dependent variable, which is the overall change in the poverty rate in the given sample.

Table A.6: Structural change: Ravallion & Huppi (1991) comparison

Country	Category	HPF Share	RH Share
China	Ag $\rightarrow$ Ag	39%	73%
	Ag $\rightarrow$ Not	38%	53%
	Not $\rightarrow$ Ag	3%	-38%
	Not $\rightarrow$ Not	21%	21%
	RH interaction term		-10%
India	Ag $\rightarrow$ Ag	43%	66%
	Ag $\rightarrow$ Not	25%	32%
	Not $\rightarrow$ Ag	5%	-19%
	Not $\rightarrow$ Not	27%	26%
	RH interaction term		-6%
Indonesia	Ag $\rightarrow$ Ag	40%	55%
	Ag $\rightarrow$ Not	14%	-1%
	Not $\rightarrow$ Ag	14%	1%
	Not $\rightarrow$ Not	32%	45%
	RH interaction term		0%
Mexico	Ag $\rightarrow$ Ag	13%	24%
	Ag $\rightarrow$ Not	26%	50%
	Not $\rightarrow$ Ag	-6%	-29%
	Not $\rightarrow$ Not	67%	57%
	RH interaction term		-2%
S. Africa	Ag $\rightarrow$ Ag	2%	16%
	Ag $\rightarrow$ Not	9%	10%
	Not $\rightarrow$ Ag	8%	-10%
	Not $\rightarrow$ Not	82%	87%
	RH interaction term		-3%

Note: The sample includes all households observed in both baseline and endline survey waves. See Appendix for details of the protocol used to classify households as primarily agricultural versus non-agricultural.

Table A.7: Livelihood changes across all households and exiters

		China		India		Indonesia		Mexico		South Africa	
		All	Exits	All	Exits	All	Exits	All	Exits	All	Exits
Share of sample		100	54	100	27	100	48	100	22	100	29
Sector	Any change	39	40	27	30	29	29	13	15	11	13
	Ag to not	35	37	18	21	14	15	8	10	8	8
	Not to ag	4	3	10	8	15	14	5	5	4	5
Occupation	Any change	44	44	22	23	33	31	33	38	33	39
	Self to wage	13	13	9	10	11	12	16	32	29	37
	Wage to self	31	31	13	12	21	18	17	6	3	2
Migration	Any change										
	Rural to urban					13	13	1	1	5	5
	Urban to rural					12	11	0	0	2	2
Female LFP	Any change	54	55	19	22	30	29	27	29	41	42
	LFP to not	47	47	9	9	13	12	12	4	5	3
	Not to LFP	7	8	9	12	16	18	16	25	37	38
Income shares	Ag	-24	-26	-15	-10	-30	-36	1	2	-2	-2
	Wage	27	42	-7	-7	-13	-10	-1	22	8	31
	Self Employment	-22	-23	-29	-20	-3	-4	-4	-7	-1	-4
	Capital	3	3	-1	0	1	1	0	-1	2	1
	Public Transfers	-2	-2	32	21	3	2	0	-9	-6	-23
	Private Transfers	8	3	4	6	13	11	-9	-36	-3	-4
	Male	2	6	-12	-25	-16	-8	2	34	3	19
	Female	2	4	-4	0	2	3	7	14	5	11
	Unattributed	-4	-10	16	25	14	5	-9	-47	-8	-30

Note: This table reports changes in the livelihoods of all households, in contrast with Table 3 which just examines those whose poverty status changed between the baseline and endline rounds of the panel surveys. The top panel reports the share of households in the given column category that experienced the indicated change (e.g. among households in China 39% changed sector). The bottom panel reports changes in the share of income that households in the given column category received from various sources (e.g. among households in China the average share of income obtained from agriculture fell by 24%). Blank entries imply the absence of data. For example, only the IHDS and NIDS have questions about capital income. Details on the classification of income and of primary income-generating activities are in Appendices B.5 and B.8, respectively.

Table A.8: Poverty rates and transfer policy

	China			India			Indonesia			Mexico			S. Africa		
	First	Last	Diff	First	Last	Diff	First	Last	Diff	First	Last	Diff	First	Last	Diff
<i>Poverty rate</i>															
Actual	76.2	27.2	-49.0	54.0	37.2	-16.8	65.9	21.2	-44.8	43.1	41.5	-1.6	44.5	25.9	-18.6
Without public transfers	76.3	27.8	-48.5	56.9	44.0	-12.9	65.9	22.1	-43.9	52.7	47.6	-5.1	60.3	45.5	-14.8
Difference	-0.05	-0.6	-0.5	-2.9	-6.7	-3.9	0.0	-0.9	-0.9	-9.5	-6.0	3.5	-15.8	-19.5	-3.8
<i>Public transfer policy</i>															
Generosity (transfer per capita per day)	0.07	5.7	5.6	4.2	30.7	26.5	0.0	20.1	20.1	101.7	70.5	-31.2	92.9	131.4	38.5
Targeting (% accruing to the otherwise-poor)	36.9	12.2	-24.7	48.4	37.7	-10.7	75.9	23.6	-52.3	59.8	55.1	-4.6	84.3	65.9	-18.4

Note: This table reports poverty rates including and excluding income from public transfers, and metrics that characterize the public transfer policies in place, in the first and last waves of each panel survey. “Generosity” refers to the mean total value of public transfers per person per day in the entire sample; a policy of giving every individual a basic income set at the extreme poverty line, for example, would correspond to a generosity of \$215. “Targeting” refers to the value of transfers received by households whose pre-transfer income places them below the poverty line, divided by the total value of transfers received by all households. A policy that perfectly targeted transfers to reduce the poverty gap, for example, would yield a targeting value of 100.

B Methodology

This appendix provides a more detailed description of the procedures we follow to clean, classify, and interpret the data, with intermediate results to illustrate the consequences of these choices.

B.1 Household definitions and size

We adopt each survey’s definition of households, which are as follows:

CHIP “Members live together and share income and expenditures” ([Luxembourg Income Study, 2002](#))

NSS “A group of persons normally living together and taking food from a common kitchen constitutes a household” ([Ministry of Statistics and Programme Implementation, Government of India, n.d.](#))

SUSENAS “A group of people whose members reside in the same dwelling and share food from the same cooking pot (the standard BPS definition)” ([Badan Pusat Statistik , BPS](#)).

ENIGH “A group of persons (not necessarily related by blood) who reside usually in the same private dwelling and share the meals. Domestic servants and guests who usually reside and share meals in the dwelling are considered as usual residents of the dwelling, but not as household members (they are part of the sample, but little information is collected for them)” ([Luxembourg Income Study, 2010](#))

IES & LCS “All persons who live together and provide themselves jointly with food or other essentials for living, or a single person living alone. Do not include persons in special dwellings, like boarding schools.” ([Statistics South Africa, 2011](#))

For the panel sources:

CHNS No explicit definition of a household. The questions establishing the household roster ask about residence, and in the roster section of each successive survey wave households are asked to explain any changes between rosters across waves. The CHNS records both the overall household size and the number of members surveyed; we use the former. ([Carolina Population Center, n.d.](#))

IHDS “All those who live under the same roof and share the same kitchen for 6+ months.” ([Desai and Vanneman, 2012](#))

IFLS “A person or group of persons who occupy a part of or an entire building and who usually live together and eat from the same kitchen. What is meant by eating from one kitchen is that the arrangement to fulfill daily necessities is jointly managed.” ([RAND Corporation, n.d.](#)) See Book K, Control Book of the IFLS questionnaire for more information.

MXFLS A household is defined as “a person or group of people, related or unrelated by biological bonds, who usually live together in a part of or in an entire building/dwelling and usually consume meals prepared with a common budget on the same stove/oven and even use the

same tools for preparing the meals.” (Rubalcava and Teruel, n.d.) See Book C, Control Book of the MxFLS questionnaire for more information.

NIDS A household is indirectly defined by the following household membership criteria: “You are a household member if: (i) You have lived under this ‘roof’ or within the same compound/homestead/stand at least 15 days during the last 12 months OR you arrived here in the last 15 days and this is now your usual residence and (ii) when you are together you share food from a common source with other household members and (iii) you contribute to or share in a common resource pool.” (Southern Africa Labour and Development Research Unit , SALDRU)

Overall the definitions that determine household size are fairly similar across sources, though there are some differences which may in turn induce difference in the poverty rates we calculate. For example, an individual is included in a household in NIDS if they spent only 15 out of the last 365 days sharing food with the others, while in the IFLS they must not have lived away from the household for more than six months. This means that long-term circular migrants would be considered a part of the household in NIDS but not in the IFLS. All definitions imply that household composition and size may change across survey rounds; see Table A.5 and the corresponding discussion at the opening of Section 4.

We also adopt each survey’s definition of household heads. We allow headship status to move across individuals throughout successive rounds of the panel data. In both the panel and repeated cross section data, When multiple household heads are indicated within a household, we take the oldest male among them to be the household head and the oldest female if there are no male heads. When no household head is indicated in the data, we assign one using the following decision rules:

1. Oldest non-dependent (age $\notin [18,70]$) male, if there is one, or if none,
2. Oldest non-dependent female, or if none,
3. Oldest dependent male, or if none,
4. Oldest dependent female

B.2 Accounting for price levels

We convert local currency units (LCU) to 2017 US dollars for comparability with the World Bank’s \$2.15/day Poverty Line. We do this in two steps, first converting nominal values to real 2017 LCU using CPI values provided by the World Bank (World Bank, n.d.a, Series FP.CPI.TOTL), and then converting these converted to 2017 USD using the PPP conversion factor for private consumption provided by the World Bank (World Bank, n.d.b, Series PA.NUS.PRVT.PP), again as in Chen and Ravallion (2010). We use a single CPI scaling value for each survey wave; some surveys were conducted in a period spanning more than one calendar year, in which case we use the year in which the last interviews were conducted.

As a sensitivity check, we also adjust for sub-national differences between rural and urban prices. For China, India and Indonesia, we use data from Jolliffe et al. (2022), drawn from Atamanov et al.

(2020). For Mexico, we use data from the National Council for the Evaluation of Social Development Policy (CONEVAL).⁴³ For South Africa we use data graciously provided by Statistics South Africa. In all cases, we adjust poverty lines following the methodology of Ferreira et al. (2016). The results of this sensitivity check are reported in Figures B.1 and B.2.

B.3 Measuring consumption

We calculate real consumption per capita as total annualized household consumption divided by household size (as defined above) and normalized by the relevant price index (as defined above). Total household consumption is the sum of consumption expenditure, i.e. spending on consumer goods and services, and the market value of self-produced goods which the household consumed. All consumption items are annualized, e.g. if households were asked about food expenditure during the previous week we scale this quantity by 365/7.

Durables. Consumption includes expenditure on consumer durables such as clothing, furniture, electronics, appliances, and vehicles for which data are available. It does not include expenditure on housing, given the well-known issues in accounting for the flow value of housing services consistently across both renters and owners; see the discussion below in Appendix B.6.

Rent. Consumption does not include expenditure on renting a dwelling, for the well-known reason that this makes renters incomparable with households that own their own home. We consider the value of housing services consumed by both renters and homeowners as part of a separate exercise described in Section B.6 below. Excluding rent from the consumption aggregate is generally straightforward, with the following exceptions:

- In the SUSENAS, the earliest eight waves (from 1994–2001) asked about housing expenditures as part of a bundle with other dwelling related expenditures (such as electricity, telephone, gas, kerosene, water and wood), while the following six waves (2002–2007) ask specifically about rent or estimated rent of a respondent’s dwelling as well as utilities and home maintenance separately. For the latter waves we exclude rent but leave in utilities and home maintenance. For the earlier waves we exclude the bundle of housing expenditures entirely. We prefer this approach because, in the latter waves in which both parts of the bundle are separately observable, rent accounted for (slightly) less than half (49.5%) of the total.

In most later years (2008–2011 and 2014–2017) questionnaires asked respondents about their rent, about lease payments, and about the value of any housing subsidies received, e.g. from an employer. We exclude all of these items from the calculation of the consumption aggregate. For three years (2012, 2013 and 2018) we do not have the questionnaires used but assume that variable names in the dataset match those from other rounds between 2008 and 2017, and so implement the same rule.

- In the ENIGH the first six waves (1984–1998) lump rent together with utilities and housing conservation expenditures. As in the SUSENAS case we drop this entire bundle, as we expect

⁴³<https://www.coneval.org.mx/Medicion/MP/Paginas/Lineas-de-Pobreza-por-Ingresos.aspx>

that rent is again the majority of this expenditure.

- In the NIDS we calculate aggregate consumption by summing pre-aggregated subcategories of consumption released with the data. The “non-food” subcategory includes an entry for “imputed rental income,” as this is equivalent to imputed consumption of owned housing; for consistency we subtract this item from the non-food aggregate (see Table 6.4 Components of aggregate household income in the ([Southern Africa Labour and Development Research Unit](#) , [SALDRU](#), NIDS Panel User Manual)). The household level income categories included are: labor market income, government grant income, other government income, investment income, remittance income, subsistence agricultural income and imputed rental income. To ensure that rent is not included we therefore subtract the rent item from the non-food subcategory before adding it to the total.

Self-produced goods. We include in the consumption aggregate the market value of goods that households produced themselves and consumed. In most cases households were asked directly to report the market value of such goods we use their responses. In the IHDS households were asked to report the quantities they consumed, the market prices of those goods, and (where relevant) their prices at subsidized rates made available through the Public Distribution System. We value goods using the market prices reported by the respondent when available, and using the median price in the primary sampling unit $\times$ wave when not.

Remittances. We include transfers out of the household in consumption, viewing them as one way in which households chose to use their disposable income. These are most often cash transfers sent to non-resident family members.

B.4 Measuring income

We measure total income as the sum of four subtotals described below. Where the data source provided pre-made aggregates for subcategories of income we worked from these, making adjustments as needed to conform with the principles laid out below. When the underlying questions about earnings were asked with a recall period less than a year (typically a month) we annualized these estimates without adjusting for seasonality.

- **Wage income** is income earned in return for providing labor to someone else. This includes both cash earnings (including any stipends for food or housing) and in-kind earnings (e.g. meals provided) wherever these are reported. Household members sometimes worked more than one job during the recall period covered; the underlying surveys were designed to be exhaustive of earnings from all jobs with the exception of the IFLS and MxFLS, which ask about the respondents primary and secondary jobs.
- **Capital income** is the sum of rent, dividends, interest and realized capital gains. In practices, all the panel sources we use report income from renting various forms of property. Only the IHDS and NIDS, however, report dividends, interest, and capital gains. We omit any reported revenue from sales of land and other property, as they represent gross flows and not net

capital gains and are probably best thought of as balance sheet transactions rather than income. This also avoids potential double-counting to the extent that capital gains from the sale of improved assets (e.g., livestock that have been reared) are already being reported in the available measures of enterprise profit.⁴⁴ Note that the information the MxFLS collects about asset rental income changes somewhat from wave to wave: in 2002 it includes rental income from all household assets, in 2005 it includes only non-agricultural household assets, and in 2009 it includes all non-agricultural household assets as well as tractors. It also reports the sum of payments from inheritance, lotteries, and life insurance schemes, which we treat as capital income.

- **Own-enterprise income** is the profit obtained from all enterprises owned and operated by the household (including both non-agricultural enterprises and farms), as well as the value of any goods or services produced by these enterprises that were not marketed but consumed by the household itself. Where the survey collected data on both revenues and costs, we difference these to obtain our measure of profit; this is the case for all enterprises in the NIDS, IFLS, and CHNS, and for farms in the MxFLS. Where the survey did not obtain separate measures of revenues and costs and only obtained a measure of profits, we use this measure; this applies to the IHDS and to non-farm enterprises in the MxFLS.

One should generally think of these measures as capturing a mixture of capital income and any economic rents. They are elicited net of wage payments to family workers, and family workers also separately report wages earned (including from any household enterprises) in a wage module. All income should thus be counted once, and the own-enterprise income measure should in principle exclude labor earnings. In practice, however, households may not think in terms of or report wage payments for family labor. One should thus think of own-enterprise income as likely also containing some labor earnings. We do not attempt to further disentangle these components.

We value self-produced goods as described in Section B.3 above, so that the number entering into the income aggregate is for the most part identical to the one that enters the consumption aggregate. One exception is food in the IHDS, which asks for the sum consumed in the last 30 days and then asks a categorical question about whether the food was purchased, produced, or both. In the cases where respondents said both (< 10% of the total) we cannot determine how much of the food consumed was self-produced, and so classify it entirely as purchased and count it as consumption but not income. In two surveys (the IFLS after round 1, and the MxFLS), we cannot distinguish between consumed food that was self-produced and food that was received as a transfer or gift; we classify this food as self-produced because we expect transfers to be a small share of the total.

- **Transfer income** is transfers received both in cash and in kind, from both private and public sources. Where the value of transfers in kind was reported we use that value; where it was not, we value the transfer using median reported price for that item in the enumeration area.

⁴⁴See <https://ihds.umd.edu/income-variables>.

As is standard in the literature (e.g. [Ravallion and Chen, 2007](#)) we do not attempt to price publicly provided services such as education for which appropriate reference prices are not available.⁴⁵

The public transfers captured are as follows:

- For China, the CHNS covers the value of generic public subsidies for gas, coal, electricity, childcare and the one-child policy, as well as income from poverty, disability and welfare transfers. It elicits participation in government-provided healthcare schemes but does not provide an estimated monetary value, so we exclude this. It does not contain a dedicated question about the Dibao program, a large means tested cash assistance program started in 1993, but income from this program is likely captured in general questions about income from public sources for welfare purposes. In contrast, it may not entirely capture transfers received as part of the Wu Bao (“Five Guarantees”) program that provides necessities to elderly individuals living in rural areas, as these transfers were in-kind and respondents would have needed to mentally estimate and account for their value in their responses to questions about income from public sources for purposes of welfare. It does not elicit the value of publicly provided health insurance under the New Cooperative Medicine Scheme.
- For India, the IHDS covers government pension income, old age and widow pension income, maternity benefits, disability income, scholarship support, and Annapurna food subsidies, as well as an ‘other public income’ category, at the individual level. It also covers, at the household level, the monetary value of assistance received from the Integrated Rural Development Program at the household level, and the quantity of rice, wheat, sugar and kerosene received from the Public Distribution System (PDS). We value the latter at the difference between their market and PDS unit prices. In the 2012 round the IHDS also covers individual wage income from the National Rural Employment Guarantee Scheme (NREGS, launched in 2006). It does not ask dedicated question about the Janani Suraksha Yojana program, which provides conditional cash transfers to individuals who gave birth in a hospital, but income from this program is likely captured in general questions about public maternity benefits. It does not elicit monetary values for government-provided health insurance (Rashtriya Swasthya Bima Yojana), for assistance such as supplementary nutrition provided through the Integrated Child Development Services, or for housing obtained through the Indira Awas Yojana program.⁴⁶
- For Indonesia, the first wave of the IFLS (1993) simply elicits the total value of all public transfers received by the household. The second wave (1997) asks about income from scholarships, retirement supplements and other grant income in its Household Non-Labor Income module but does not discern if the sources are public or private. So in 1997 we conservatively assign none of this income to the public transfer category. The third

⁴⁵If instead we include net transfers (i.e. net of outgoing transfers) rather than gross transfers in our income measure, poverty rates rise slightly, by an average of 1.0 percentage points across all rounds of all the panel surveys.

⁴⁶The 1993 HDPI does not contain any information on government transfers, but this is of no consequence for the analysis in Table 3 and thereafter since we there use only the two rounds of the IHDS.

and subsequent waves (2000, 2007, 2014) include a social safety net module covering the Bantuan Tunai Bersyarat (conditional cash transfers), Bantuan Langsung Sementara Masyarakat (unconditional cash transfers) and food subsidies for rice, sugar, cooking oil and an ‘other’ category (which could capture transfers under Raskin, among other schemes). It elicits participation in government-provided healthcare schemes but does not provide an estimated monetary value, so we exclude this. It does not capture interest-rate subsidies obtained through participation in the subsidized Kredit Usaha Rakyat micro-credit scheme for small businesses.

- For Mexico, the MxFLS usually records the value of transfers received via Progresa (conditional cash transfers), Procampo (direct agricultural support), Vivah (housing support), Word credit (microcredit program), Social coinvestment, PET (temporary employment), Alianza para el Campo (agricultural support), the National Fund for Social Enterprises (micro-enterprise support), Non-Progresa Educational support, and general pension income, as well as an ‘Other’ category. An exception is that in the second wave (2005) the list is modified to exclude Progresa. In the third wave (2009) the same list is employed but with slightly modified program names to reflect the way they would be known by the public. The values of subsidized milk, tortillas and government provided daycare could not be calculated and hence are excluded, along with other goods which may have been purchased from the state-subsidized Diconsa stores and which were not covered.
- For South Africa, the NIDS surveys covers the state sponsored old age pension, disability grants, child support grants, foster care grants, general aid, veterans’ pensions, care dependency grants as well as unemployment insurance and worker’s compensation. It also asks whether respondents benefited from government housing assistance but does not estimate the value of this assistance, so we exclude this. It does not cover short-term assistance in cash and in kind via the Social Relief of Distress program, though this may be captured by a catch-all “other public assistance;; question.

We treat all transfer payments that do not clearly come from government bodies as private transfers. This includes sources such as remittances from family members and assistance from NGOs, as follows:

- For China, the CHNS covers remittances and retirement support from non-resident family members as well as friends.
- For India, the IHDS asks about the value of assistance from NGOs, and remittances from non-resident family members.⁴⁷
- For Indonesia, the IFLS has detailed information on remittances: the values of cash remittances from parents, siblings, children and other non-resident family members, as well as assistance paying for tuition and health care costs and in-kind transfers of food stuff.

⁴⁷The 1993 HDPI provides no information about remittances or private transfers; as with public transfers, this is of no consequence for the analysis in Table 3 and thereafter since we there use only the two rounds of the IHDS.

Table B.1: Implied housing share of total consumption, by interest rate

	Alternative interest rates				
	$r = 0.10$	$r = 0.15$	$r = 0.20$	$r = 0.25$	$r = 0.30$
Indonesia	0.23 (0.18)	0.27 (0.2)	0.31 (0.21)	0.35 (0.22)	0.38 (0.23)
Mexico	0.43 (0.24)	0.49 (0.24)	0.54 (0.24)	0.57 (0.24)	0.6 (0.24)
South Africa	0.23 (0.21)	0.28 (0.23)	0.32 (0.24)	0.35 (0.25)	0.37 (0.25)

This table reports estimates of the share of total consumption represented by consumption of housing services in the panel sources. The flow value of housing services was calculated using the user cost approach as described in Section B.6, and using the range of interest rates noted in the columns and a depreciation rate of 1/30 throughout. It omits India because the IHDS does not report information about housing, and the China because the CHNS does not report information about non-housing non-durables consumption. Standard deviations are reported in parentheses.

- For Mexico, the MxFLS covers the value of money, aid, or donations to the household from non-resident family members and friends. We also include income from dowries in the value for private transfer income.
- For South Africa, the NIDS covers the value of any monetary or in-kind transfers received as gifts from non-household members.

B.5 Classifying income

Income of the four types listed above varies in the extent to which the survey data allow us to unambiguously attribute it to a particular sector, occupation, or gender.

- With respect to **sector**, all wage and enterprise income is attributable. The IFLS and MxFLS let us distinguish between capital income from the agricultural and non-agricultural sectors; otherwise we attribute capital income to the non-agricultural sector. We attribute all transfer income to the non-agricultural sector.
- With respect to **occupation**, all wage and enterprise income is (by definition) attributable. We treat capital income and transfer income as separate categories.
- With respect to **gender**, all wage income is attributable. Some enterprise income is attributable in cases in which the survey clearly identifies an enterprise as belonging to or being run by a single person in the household; outside these cases, we treat the income as gender-unattributable. We never attribute capital income or transfer income to a gender.

B.6 Estimating churn in the flow value of housing services

We observe three different kinds of information about housing value in our panel sources:

- Estimated values of houses owned by the household, in all surveys except the IHDS. We also observe estimates of the rent the household would need to pay to rent owned homes in a subset of the surveys (the CHNS, IFLS and NIDS) but use estimated values in our analysis because we observe these in all surveys.
- Estimated values of houses rented by the household, in the CHNS, MxFLS, and NIDS
- The value of rent paid by the household, in the CHNS and IFLS

Estimating the flow value of the housing services the household consumes in cases like the first two, where we observe a stock value, is a classic problem in living standards measurement; see [Amendola and Vecchi \(2022\)](#) for a review and discussion. To illustrate the consequences, Table B.1 reports estimates for home-owners of the share of total consumption that is housing services using the user cost approach, under (a) the assumption that houses depreciate over the course of 30 years, so that the annualized depreciation rate (using the straight-line method) is 1/30, and (b) a range of assumptions about the interest rate households face. As a point of reference for the latter, the median (mean) rate in the MIX Market Intelligence database of microcredit lending rates from 2,295 microfinance institutions during 2000–2019 was 17% (21%).⁴⁸ The housing share of consumption is substantial in all of these cases, but also quite varied across them.

One need not take a stand on this thorny issue, however, to answer the narrower question of how volatile living standards were as measured by housing services. We do so as follows. First, we calibrate a “housing poverty line” for households (both owners and renters) who report the value of the home they inhabit. We set this line so that the share of those households that are poor according to it is the same as the share poor according to the standard measure based on non-durables consumption in the final round of the survey. Formally, let $\bar{v}$ be defined by

$$\sum_{h : h \text{ owns a home}} \left[\mathbf{1} \left(c_h^{\bar{t}} < \$2.15 \right) - \mathbf{1} \left(v_h^{\bar{t}} < \bar{v} \right) \right] = 0 \quad (7)$$

where $v_h^{\bar{t}}$ is the value of the home inhabited by household h in final round $\bar{t}$ of a given survey. For households that report the rent they pay we similarly calibrate a “rental poverty line” such that the poverty rate among renters is the same using this definition as using non-durables consumption in the final round, i.e. defining a threshold $\bar{r}$ by

$$\sum_{h : h \text{ rents a home}} \left[\mathbf{1} \left(c_h^{\bar{t}} < \$2.15 \right) - \mathbf{1} \left(r_h^{\bar{t}} < \bar{r} \right) \right] = 0 \quad (8)$$

where $r_h^{\bar{t}}$ is the rent paid by household h in final round $\bar{t}$ of a given survey. We then use the calibrated values $\bar{v}$, $\bar{r}$ to classify households in all other rounds as either poor or non-poor. This includes households that switch between owning and renting status; we apply whichever threshold is relevant to their housing status in any given round. The result is a measure of headcount poverty

⁴⁸<https://databank.worldbank.org/source/mix-market>, accessed 30 December 2023. We define the interest rate as the ratio of interest income to the average gross loan portfolio.

that is guaranteed to be the same (up to integer constraints) as the measure based on non-durable consumption in the final round, but may be more or less volatile in previous rounds.

B.7 Assessing sensitivity to measurement error

Let $x_{h,t} > 0$ be the value of some measure of household h 's standard of living that we observe at time t (e.g. per capita consumption). To assess the sensitivity of our conclusions to noise in this measure, we introduce additional noise by defining

$$\hat{X}_{h,t} := x_{h,t} Z_{h,t} \quad (9)$$

$$Z_{h,t} \sim \text{lognormal}(\mu, \sigma^2) \quad (10)$$

for given parameters μ, σ . Henceforth to reduce clutter we will suppress subscripts. In order to make measurement error mean zero, we impose the parameter restriction

$$\mu = -\frac{\sigma^2}{2} \quad (11)$$

which implies that

$$\mathbb{E}[\hat{X}|x] = xe^0 = x \quad (12)$$

and that

$$\mathbb{V}[\hat{X}|x] = x^2(e^{\sigma^2} - 1) \quad (13)$$

which is increasing in σ^2 , so freely varying σ varies the degree of noise added. To calculate what this means for expected percent error,

$$\mathbb{E} \left[\left| \frac{\hat{X} - x}{x} \right| \middle| x \right] = \mathbb{E} \left[\left| \frac{x|Z - 1|}{x} \right| \middle| x \right] \text{ since } x > 0 \quad (14)$$

$$= \mathbb{E}[|Z - 1|] \text{ since } Z - 1 \text{ does not depend on } x \quad (15)$$

$$= \mathbb{P}(Z > 1)\mathbb{E}[Z|Z > 1] + \mathbb{P}(Z \leq 1)\mathbb{E}[Z|Z \leq 1] \quad (16)$$

which is straightforward to calculate using known closed-form expressions for the CDF and conditional expectation of a lognormal random variable.⁴⁹ We can thus proceed as follows:

1. Calculate expected percent error for a range of values of σ^2 , and use the results to select a set of values of σ^2 that correspond to a grid of expected percent error values of interest (e.g. 5% to 20% in 5% increments).
2. For each of these chosen values of σ^2 ,
 - (a) Generate some large number K of replicate datasets where each value $x_{h,t}$ has been replaced by a draw $\hat{X}_{h,t}$ as defined above. For each of these datasets, calculate the value

⁴⁹See https://en.wikipedia.org/wiki/Log-normal_distribution#Conditional_expectation.

Table B.2: Classification of households into the (non)agricultural sector

Country	Households			Classified household-years					
	Total	Unclassified	Classified	Total	By sector code	By income share	Ag. Share		
China	1,922	116	1,806	3,612	2,731 (76%)	881 (24%)			55%
India	34,639	109	34,530	69,060	57,616 (83%)	11,444 (17%)			48%
Indonesia	6,019	104	5,915	11,830	9,888 (84%)	1,942 (16%)			41%
Mexico	7,182	1,252	5,930	11,860	8,449 (71%)	3,411 (29%)			22%
S. Africa	5,391	165	5,226	10,452	10,023 (96%)	429 (4%)			10%

This table summarizes the classification of household $\times$ survey round observations in the panel sample as working primarily in agricultural or non-agricultural sector, as described in Section B.8.1. The first column indicates the total number of households observed in the first and last round of each survey, and which we attempted to classify; the next two columns indicate for how many of these we were unable and able, respectively, to classify them in both of those rounds. The remaining columns describe, for classified households, how many of the household-years we classified based on the sector codes in which household reported using (our more-preferred method) and based on actual income earned (our less-preferred method), as well as the overall final estimated share of observations for which we classified the household as working primarily in agricultural (“Ag. share”).

of some statistic of interest, e.g. the share of the population that is always poor or never poor.

- (b) Calculate the cross-replicate-dataset mean of these statistics, for comparison to its value when using the original data

B.8 Classifying households

Analyses based on Equations (5) and (6) and extensions thereof require that we classify households into categories based on the sector in which they primarily work, the type of occupation in which they primarily work, and so on. We construct these as described below. Note that for this analysis we focus on the sample of households observed in both the first and last round of each panel.

B.8.1 Classification by sector

We first classify as many households as possible using sectoral classifier codes. The IHDS and NIDS provide a single industry code; we classify the households as agriculture if this indicates that it is engaged in agricultural or another primary sector (e.g. forestry, mining) and as non-agricultural otherwise. The IFLS, CHNS and MxFLS allow households to indicate up to *two* sectors in which they are engaged; we classify the household as agricultural if either of these indicates agriculture or another primary sector, and as non-agricultural otherwise. This classifies between 10% and 55% of households in our balanced panels, depending on the survey (see Table B.2). We then classify the remaining households based on the source of their realized income (see below), labeling those that obtained more than half of their total income from agriculture (or other primary sector activities) as agricultural.

B.8.2 Classification by migration status

Migration analysis is possible only for Indonesia, Mexico and South Africa, as the CHNS and IHDS do not track households that move from their initial location between survey rounds. For those three countries, we classify a household as having migrated between the first and last rounds of the panel if it is located in a different administrative region in the last round than in the first round. We select administrative units for this calculation aiming to keep the geographic size of a unit roughly comparable across datasets:

- In Indonesia, we use Kecamatan of which 864 appear in the data across 227 Kabupatans and 20 provinces. In total, Indonesia has 7,252 Kecamatans with an average size of 250 square kilometers and 38,538 people.
- In Mexico, we use community of which 261 appear in the data across 204 municipalities and 24 states. In total, Mexico has 2,454 municipalities with an average size of 799 square kilometers and population of 52,924 people.
- In South Africa, we use municipality, of which 57 appear in the data across 9 states. In total, South Africa has 213 municipalities with an average size of 5,702 square kilometers and population of 272,527 people. South African municipalities are larger than Indonesian Kecamatans and Mexican Municipalities, which may explain the lower base rates of migration observed in South Africa in Figure 8.

We then further classify households that migrated based on whether they migrated from areas that the survey classified as {rural, urban} in the initial, and whether they migrated to areas that the survey classified as {rural, urban} in the final round. We observe some of all four possible cases.

B.8.3 Classification by occupational choice

We classify households in two steps, first classifying the occupational choices of their members and then defining an aggregate classification for the household as a function of these.

Table B.3 summarizes the classification of individuals. Each survey provides a field with occupational information for each person on the household roster, with some variation across surveys.⁵⁰ For the sake of exposition we have here collapsed these underlying occupational categories into the somewhat coarser categories listed in the column labeled “Partially aggregated occupational categories.” We then map each of these categories into one of three final classifications: self-employed, wage worker, or out of the labor force. We classify members listed as “unemployed” as wage workers on the grounds that being unemployed indicates that they wished to be employed by someone else. We classify members for whom no information is provided (“Unable to classify”) as out of the labor force; typically these are children and the elderly, though in some cases we do observe working-age

⁵⁰In the IHDS individuals are allowed to report multiple categories. We resolve that ambiguity as follows: we first classify anyone who reported doing farm work, animal work, or “business” as self-employed, then classify anyone who reported working as an agricultural laborer, a non-agricultural laborer, or working as part of the NREGS or food-for-work schemes as a wage worker; and then finally classify anyone who reported working in none of these categories but reported earning a positive salary as a wage worker.

adults in this category. Table B.4 reports the resulting classification of person $\times$ survey round observations. Taking a simple average across surveys, we classify an average of 30%, 35%, and 35% of working-age adults (18–70 years of age) as self-employed, wage workers, and out of the labor force, respectively. The corresponding figures for individuals outside of working age (i.e. less than 18 or more than 70 years of age) are 6%, 10%, and 84%.

We then classify households as follows. For those with at least one member who is not out of the labor force, we classify the household as self-employed or working for wages depending on which of those two categories more household members fall into. In the minority of cases (ranging from 3% in South Africa up to 12% in Indonesia) where there are ties we break the tie using the classification of the household head, provided the head is not economically inactive. In the handful of remaining cases (0% in China, 0.2% in India, 0.1% in Indonesia, 0.6% in Mexico and 0.1% in South Africa), we default to classifying the household as working for a wage.

This leaves a cross-survey average of 10.7% of households which report *no* members as being part of the labor force. (Note that 11% is likely a generous estimate, as some households have no members who report being engaged in the labor force but nevertheless report positive earned income.) We label these households “not in the labor force.” They tend to be smaller with household heads who are older and more likely to be a woman than their counterparts who are in the labor force (see Table B.6). For the sake of digestibility we default to grouping these households with the self-employed, so that we have two categories: households that were primarily wage-seekers, and all other households. This classification is the basis of Figure 7, for example, and is arguably the most informative view for examining the role of labor markets in the process of poverty exit and entry. Figure A.15 reports the more detailed three-category analogue in which we can further distinguish, for example, between transitions into active self-employment as opposed to no market activity.⁵¹

Table B.5 tabulates the resulting final classification of households. Overall, averaging across survey rounds, we classify 42.6%, 46.7%, and 10.7% of households as primarily self-employed, primarily in wage employment, and out of the labor force, respectively.

⁵¹Note that doing so is feasible within our approach, requiring a straightforward generalization of Equation (5), but infeasible using Equation (6) and repeated cross-sectional data (since that approach rests on the truism that people who entered one category must have left a uniquely identifiable complementary category).

Table B.3: Classification of individuals' occupations

Classification	Partially aggregated occupational categories	Percent reporting in				
		China	India	Indonesia	Mexico	S. Africa
Not in labor force	Inactive, labor force age (18-70)	16.2%	20.3%	21.6%	30.3%	15.8%
	Inactive, not labor force age (not 18-70)	20.6%	33.3%	37.1%	38.3%	39.2%
	Unable to classify	0.2%	0.0%	0.0%	0.7%	0.2%
Self-employment	Animal work		8.6%			
	Boss/employer				1.7%	
	Business		4.0%			2.3%
	Family worker			5.9%		
	Family worker in hh business- no remuneration				1.7%	
	Farm work		19.9%			1.8%
	Paid family worker	2.2%				
	Peasant on plot				2.1%	
	Self-employed, independent operator with no employees	26.0%		9.0%		5.1%
	Self-employed, owner-manager with employees	2.3%		0.6%		
Wage employment	Self-employed, owner-manager with temporary paid help			8.8%		
	Unpaid family worker	0.2%				
	Worker without remuneration on non-hh business				0.3%	
	Agriculture laborer		3.7%			
	Contractor with other people or enterprise	8.2%				
	Government worker			3.1%		
	Non-agricultural laborer		4.4%		16.9%	2.4%
	Nrega work		2.1%			
	Private worker			13.7%		
	Rural laborer				2.9%	
Unemployed	Temporary worker	13.1%				
	Unemployed					19.0%
	Wage/salary job		3.7%			19.2%
Works for another person or enterprise as permanent employee		11.0%				

Note: This table describes the bases on which we classified individuals appearing in the panel data sources as not in the labor force, self-employed, or having or seeking wage employment. "Classification" is the classification we ultimately assigned. "Partially aggregated occupational categories" describes the occupational classification in the underlying surveys after some basic initial aggregation, which we perform for expositional purposes. Note that some categories appear in more than one data source, while others are unique to a particular source. The remaining columns report the country-specific share of individual $\times$ survey round observations, for all individuals listed on household rosters, associated with each label. For example, the top-left cell indicates that 16.3% of person-year observations in the CHNS were labelled as "Inactive, Labor Force Age." Blank entries indicate that the given value was not a possible response in the given survey; 0.0% entries indicate that it was an option but that no one (up to rounding) selected it. The sample includes households observed in both the baseline and endline rounds of each panel.

Table B.4: Classification of individuals' occupations, by age group

Age group	Occupation	China	India	Indonesia	Mexico	S. Africa	Average
Labor force age (18–70)	Self-employment	48%	46%	39%	16%	6%	31%
	Wage employment	30%	21%	28%	31%	51%	32%
	Not in labor force	21%	33%	34%	52%	43%	37%
	Not in labor force	0%					0%
Not labor force age (not 18–70)	Self-employment	7%	11%	6%	2%	1%	5%
	Wage employment	11%	2%	3%	3%	4%	5%
	Not in labor force	81%	87%	91%	95%	95%	90%
	Not in labor force	0%					0%
Age unknown	Self-employment			14%	7%	2%	8%
	Wage employment			12%	15%	5%	11%
	Not in labor force		100%	74%	78%	93%	86%

Note: This table summarizes the classification of individuals into occupations that results from the mapping defined in Table B.3, by age group. The sample includes all person $\times$ survey round observations in the first and last rounds of each panel source. The numeric figures in each cell represent the share of observations, within a given country and age group, that were classified as having the occupation listed in the “Occupation” column.

Table B.5: Classification of households' primary occupation

Category	China	India	Indonesia	Mexico	S. Africa	Average
Primarily self-employed	53.3%	58.1%	57.9%	30.1%	13.7%	42.6%
Primary wage employment	44.6%	37.8%	37.4%	50.3%	63.3%	46.7%
Not in the labor force	2.1%	4.1%	4.8%	19.6%	23.1%	10.7%

Note: This table reports the final classification of households into occupation choice categories as of the first wave of each panel survey.

Table B.6: Characteristics of households classified as in versus not in the labor force

	China		India		Indonesia		Mexico		S. Africa	
	In	Not	In	Not	In	Not	In	Not	In	Not
Poor (income)	0.55	0.40	0.46	0.47	0.42	0.55	0.34	0.78	0.33	0.56
Poor (Consumption)			0.35	0.22	0.39	0.29	0.50	0.48	0.27	0.41
Female head	0.12	0.29	0.11	0.38	0.15	0.50	0.21	0.37	0.49	0.61
Head has primary education	0.50	0.62	0.22	0.31	0.27	0.25	0.43	0.34	0.50	0.35
Age of head	47.86	65.52	48.97	58.12	49.17	57.71	48.53	56.00	49.09	50.03
Household size	4.40	3.14	5.16	3.10	6.15	6.19	5.06	3.87	4.11	3.61

Note: This table compares characteristics of household observations classified as being in the labor force (columns labelled “In”) to those classified as being not in the labor force (columns labelled “Not”). The sample includes all household $\times$ survey round observations in the first and last rounds of each panel survey.

Table B.7: Classification of households by presence of female outside employment

Household type		Share by country				
Any working age woman	Any female labor force	China	India	Indonesia	Mexico	S. Africa
Yes	Yes	38.9%	16.2%	22.6%	25.6%	63.5%
Yes	No	51.8%	81.7%	70.2%	65.2%	20.8%
No	Yes	0.2%	0.1%	0.2%	0.1%	2.6%
No	No	9.1%	2.0%	7.0%	9.0%	13.0%
Final classification:						
Female outside employment		39.1%	16.3%	22.8%	25.8%	66.1%

Note: this table describes the classification of households in the panel surveys into those that did or did not have at least one female member in the labor force. The sample includes all household $\times$ survey round observations in the first and last rounds of each panel survey. Working age is defined as 18–70; note that a handful of households without working-age women still have a female member in the labor force, as some girls and elderly women worked.

B.8.4 Classification by female labor force participation

We start with the same classification of individuals into self-employed, wage worker, or out of the labor force as defined above and summarized in Table B.3. Combining this with data on the gender and ages of members this effectively defines four types of households, tabulated in Table B.7:

1. Those with working-age (18–70) women present, at least one of whom is not out of the labor force. 16% and 64% of households across countries and survey rounds. We classify these households as having a working woman (top row of Table B.7).
2. Those with working-age women present, all of whom are out of the labor force. These account for between 21% and 82% of households across countries and survey rounds. These households we classify as not having a working woman (second row).
3. Those without any working-age women present. These account for fewer than 10% of households in all countries and rounds except for South Africa, where they account for 16% of households. In almost all of these cases there is also no female member of *any* age working (fourth row), but in a handful there is a female member younger than 18 or older than 70 working (third row).

In our main results (Figure 9) we categorize households as having female outside employment if any female is in the labor force, regardless of whether or not a working-age woman is present.

B.9 Sensitivity checks

This section reports the sensitivity of core conclusions to the following methodological variations. In the cross-sectional results (Figure B.1), we examine the effects of


1. Using a poverty line of \$3.65 per person per day in 2017 USD, rather than \$2.15;

2. Using separate poverty lines for rural and urban areas that reflect the different prices in those locales, rather than a single nationwide deflator;
3. Using adult-equivalence scales to calculate household per capita consumption; and
4. Omitting purchases of durable consumer goods when calculating household consumption aggregates.

Results are generally not sensitive to these variations, with the sole exception that using a \$3.65 poverty line noticeably increases the share of poverty decline that our decomposition attributes to within-cohort changes in Indonesia, Mexico, and South Africa—reinforcing the conclusion we reach when we focus on the \$2.15 poverty line.


In the panel results (Figure [B.2](#)) we examine the effects of changes #1–#3 above (#4 does not apply since we focus in the panels on income-based measures). We observe a few noteworthy patterns. First, using a higher \$3.65 line to define poverty further reduces the gap between the share of initially poor who are also poor at endline, and the corresponding share of the initially non-poor. In other words, it suggests even more “churn” that when we use the \$2.15 line. Second, using regional as opposed to national price deflators has a few noticeable effects in Mexico. It implies more churn, again consistent with our earlier conclusion, and a smaller but still positive increase in the wage income share among poverty-exiting households. Beyond these points, however, the results are generally insensitive to the variations we examine.

Figure B.1: Sensitivity to alternative data preparation choices: cross-sectional results



Note: This figure assesses the sensitivity of conclusions to alternative data preparation choices for cross-sectional sources. Each row represents a different metric: the within-cohort share of poverty decline from Figure 4, the range of the poverty rates across all ages and cohorts from Figure 2, and the estimated slope of the age-poverty profiles from Figure 2 after conditioning on year-of-survey fixed effects. Each color represents a different sensitivity check: altering the poverty line to \$3.65 per person per day in 2017 USD, using separate poverty lines for rural and urban areas that reflect the different prices in those locales, altering the calculation of consumption per capita to account for differences in the consumption needs of adults and children and for economies of scale, and omitting purchases of durable consumer goods when calculating consumption. Each arrow represents the effect, in percentage points, of making the given change on the given metric.

Figure B.2: Sensitivity to alternative data preparation choices: panel results



Note: This figure assesses the sensitivity of conclusions regarding the cross-sectional data to alternative data preparation choices. Each row represents a different metric, as follows: the net poverty exit (entry) rates among households that were initially poor (not poor) from Figure 5a; the net share of all poverty exits attributable to households that switched occupational choice categories in Table 3; and the changes in income shares from different activities from Table 3. Each color represents a different sensitivity check, as follows: altering the poverty line to \$3.65 per person per day in 2017 USD, using separate poverty lines for rural and urban areas that reflect the different prices in those locales, and altering the calculation of income per capita to account for differences in the consumption needs of adults and children and for economies of scale. Each arrow represents the effect, in percentage points, of making the given change on the given metric.