

How Poverty Fell*

Vincent Armentano[†]

Paul Niehaus[‡]

Tom Vogl[§]

February 18, 2025

Abstract

The share of the global population living in extreme poverty fell dramatically from an estimated 44% in 1981 to 9% in 2019. We describe *how* this 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 countries that collectively accounted for 70% of global poverty decline since 1990. The repeated cross-sections show that all birth cohorts experienced the decline of poverty over time in parallel, such that poverty decline can be viewed as a primarily within-cohort phenomenon. The panels show substantial within-cohort churn: gross transitions out of poverty were much larger than net changes, as many households also lapsed back into poverty. The overall picture is of a “slippery slope” rather than a long-term trap. The role of sectoral transitions varied across countries, though progress within sectors generally played a larger role than transitions between sectors.

JEL codes: I32, N30, O12

Keywords: poverty, poverty dynamics, structural change

*We are indebted to Cole Dreier and Radhika Goyal for their insights and dedicated efforts, without which this project would not have been possible. We also thank Gongyu Zhou and Jie Liao for research assistance. We gratefully acknowledge funding from the National Science Foundation under award #2116944.

[†]University of California, San Diego

[‡]University of California, San Diego

[§]University of California, San Diego

1 Introduction

One of humanity’s great achievements of the modern era has been the reduction in the share of the global population living in extreme poverty—from an estimated 44% in 1981, to 38% in 1990, to 9% in 2019.¹ Lay people in high-income countries often misperceive that this rate has held flat or even increased (Rosling et al., 2018), and even economists are perhaps prone to emphasizing the (very real) challenges facing people living in extreme poverty more than their successes in getting out of it. But empirically, the basic picture has been very positive.

This paper aims to provide a description of *how* this decline happened, as systematically and comprehensively as the available data will allow. By “how,” we mean in part what happened in the lives of individuals and households as they moved out of (and often, as we will see, back into) extreme poverty. Did they plant a new cash crop on their farm? Find work in a factory? Start their own business? Move to a city? and so on. But 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 have collectively accounted for 70% of global poverty decline since 1990. 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. The data available for them 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. 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.² 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.

¹These rates are based on the \$2.15 2017 PPP poverty line (Chen and Ravallion, 2010; World Bank, 2023). See <https://pip.worldbank.org/>.

²It 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 but also explore the role of housing to the extent possible.

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, or take all household members 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 accrues over people’s lives, rather than between their lives and the next cohort’s. Between years, one cohort enters the population, one cohort exits, and a multitude of surviving cohorts experiences within-cohort change. 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.

Panel data then allow us to 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 households that were initially non-poor also entered it. Using income-based measures, the probability of exiting poverty conditional on starting in it ranged from 26% in South Africa to as high as 57% in Rural India. But the probability of entering poverty conditional on starting out non-poor was also substantial, ranging from 15% in South Africa to as high as 34% in Mexico. Estimates using consumption, where they are available, tell the same basic story. Transition probabilities generally appear flat or—in the case of China and India—shift advantageously over time. Overall, the global picture is one in which poverty falls at a moderate pace not primarily because individual households remain stuck in it, but because the rate at which they exit it is moderately above the rate at which others fall back into it. This takeaway confirms systematically and on a large scale a point that several earlier studies of individual panel and pseudo-panel datasets have noted.³

We then examine changes in the livelihoods of households that exited or entered poverty. We

³See for example the studies summarized in Baulch and Hoddinott (2000) and the results from retrospective surveys conducted in communities around the world by Narayan et al. (2009) and Krishna (2010), among others.

consider both the mix of activities in which working household members engaged—which we use to classify households as working primarily in agriculture or not, primarily in self- or wage-employment, and so on—as well as the realized shares of income that households obtained from different sources.

Considering these measures together, 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 ways similar to those that exited. For example, the share of poverty-entering households that switched from agriculture to non-agriculture is 53%-114% the corresponding share of poverty-exiting households. An exception applies in the more advanced economies of Mexico and South Africa, where progress out of poverty was closely associated with transitions from self- to wage-employment and entry into poverty with the opposite.

We also see no cases in which changes in transfers (from public and private sources) played a dominant role. Among households that exited poverty, the share of income they obtained from transfers either rose slightly or fell substantially. Among those that entered poverty, the share generally rose substantially or fell slightly. Overall, the data are consistent with progressive redistribution, but not with transfer income accounting directly for a major share of the income gains that moved households above the poverty line. In this sense, the households that left poverty did so largely on their own.

Our final exercise quantifies the extent to which the observed changes in livelihood activities can account for poverty decline—how much, for example, of Indonesia’s net within-cohort poverty decline was attributable to households shifting from working primarily in agriculture to non-agriculture. We decompose the overall change in the poverty rate into a weighted average of the net changes experienced by households that changed their activities and those that did not. The decomposition reveals a few noteworthy patterns.

Transitions out of agriculture accounted for a limited role. They did not account for the largest share, let alone the majority, of transitions out of poverty in any country. And the decomposition credits transitions *into* agriculture with a poverty reducing role: the opposite of the conclusion we reach if we ignore the panel structure of the data and apply older, cross-sectional decomposition techniques to it. More broadly, in every country, households that stayed in the same sector contributed more to poverty decline than households that changed sector. Migration, particularly rural-to-urban migration, also accounts for a limited amount of poverty decline in the three countries (Mexico, Indonesia and South Africa) for which migrants were tracked, with the one notable exception that rural-to-*rural* migrants accounted for a third of all net poverty decline in Indonesia.⁴

With respect to occupational choice, patterns are quite different in the more developed economies (Mexico and South Africa) relative to the less developed ones (China, India and Indonesia). In the former group, transitions within and into wage work account for the bulk of poverty decline, while in the latter, those who stayed or became self-employed contributed the most. These patterns

⁴Related work by Bryan and Morten (2019) estimates that reducing barriers to internal migration in Indonesia would yield modest but meaningful aggregate productivity gains.

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). Meanwhile, in the more developed economies, 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.

Finally, poverty transitions have a nuanced relationship with women’s participation in the labor force. In most countries, households in which a woman entered the labor force contributed meaningfully to poverty decline, while those in which a woman exited the labor force experienced either an increase in poverty or at best a lower-than-average decrease. These patterns are consistent with the mechanical contributions to household living standards that one would expect from having an additional income earner. China is the exception. In China, households in which a woman began working exited poverty at the highest rate of any group, but they were greatly outnumbered by households in which a woman left the labor force, which accounted for nearly half of all net poverty decline. This result suggests a stronger selection effect in China than elsewhere, wherein women withdrew from the labor force when their households could afford it.⁵

We see our analysis of how poverty fell situated within the literature in three ways. First, it provides a dynamic counterpart to cross-sectional descriptions of extreme poverty, such as that provided by Banerjee and Duflo (2007). Second, it complements exercises that seek to account for differences in the cross-section (Caselli, 2005) or changes over time (Jones, 2016) in per capita GDP. These exercises characterize an “average household,” while ours focuses attention on those near the poverty line, which turn out to be quite different: on average in our data, the average household is found at the 66th percentile of the per capita consumption distribution, while the poverty line is at the 20th percentile. Third, our disaggregated approach lets us 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 the aggregate. Over all, our goal is to establish a set of comparable, systematic facts that can guide the design and interpretation of studies assessing the effects of specific development policies and interventions.

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 country, and the methods we use for measuring poverty in a comparable way across countries over time.

⁵Goldin (2005) emphasizes this income effect in accounting for the general tendency for women’s labor force participation to fall initially as an economy develops.

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 individual 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.⁶ We require at least three waves of data collection, which will (among other things) allows 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 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.⁷

To add China, a leading contributor to global poverty decline (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.⁸ We weight the urban and rural samples to match the urban and 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

⁶All repeated cross-sections span at least fifteen years, and all panels span at least nine.

⁷Some surveys include attrition-adjusted sampling weights, but these are highly correlated with the baseline weights and not available for all surveys.

⁸Later 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.

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 original households. We use each survey's own definition of a household, which varies little across surveys. In South Africa, the National Income Dynamics Study (NIDS) explicitly followed individuals rather than households, so we emulate household tracking by following first-wave household heads and their associated household outcomes, wherever they are living. 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. Three key questions arise. First, what measure of economic well-being do we compare with the poverty line? Second, what poverty line do we use? Third, how do we adjust prices to make the poverty standard comparable across space and time?

Consumption versus income We use both income-based and consumption-based definitions of poverty. Consumption is conventionally preferred for research on poverty 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, 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, consumption may be the more appropriate measure of material deprivation. Consumption may also be less prone to fluctuating above or below the poverty line, an issue especially important in longitudinal analyses of poverty change.

For both practical and conceptual reasons, we emphasize consumption poverty in the repeated cross-sectional analysis and income poverty in the panel analysis. In practical terms, 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.⁹ Conversely, all of the panel surveys collect income, while only some collect consumption.¹⁰ In conceptual terms, the questions of the repeated cross-sectional analysis lend themselves to consumption-based measures, while those of the panel analysis lend themselves to income-based measures. The former concern how living standards rise between and within cohorts; the latter concern the economic sources of within-cohort gains. Consumption arguably captures living standards more completely; income can be assigned to specific sources. Where possible, we check the sensitivity of our results to using the complementary poverty measure.

Poverty line We adopt the current international poverty line of \$2.15 per person per day in 2017 prices, an update of the “\$1-a-day” line originally proposed by Ravallion et al. (1991). Both the original \$1 line and the updated \$2.15 line are based on averaging the locally-determined poverty lines of poor countries. 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).

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.¹¹

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 and then using our repeated cross-section data.

The World Bank series find that sample countries account for 70% of global poverty decline since 1990, 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

⁹In India, some NSS rounds collect weekly wage and salary income from a separate sample of households, but only 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.

¹⁰In China, the CHNS lacks consumption altogether; in India, the HDPI (the rural-only survey preceding the IHDS) also lacks consumption.

¹¹For simplicity, we do not adjust for spatial variation in prices within countries, for example between urban and rural areas, nor for variation in consumption baskets between the poor and non-poor (Deaton and Dupriez, 2011).

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. Most of our analyses focus on the household rate, but some robustness checks also consider the population 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. 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 2009.¹² Overall, average growth rates were fastest in China, India, and Indonesia—the poorer countries—and slower in Mexico and South Africa, particular during the periods spanned by the panel sources. This pattern lines up with their relatively slow poverty rate declines, 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

¹²The 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.

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 in 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%.¹³ 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 6.9% in 1995 to 18.8% in 2017,¹⁴ and South Africa has typically been 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 assess the how the pace of *intra*-generational poverty decline, with the initially poor exiting poverty during their lifetimes, compares with the pace of *inter*-generational decline, with their non-poor children replacing them in the population. To improve the mapping to the data, we conceptualize generations as cohorts, or groups of individuals born in the same year. Our question then involves the relative pace of within- versus between-cohort progress against poverty.

We begin in Section 3.1 by developing a cohort accounting framework and using it to derive a cohort decomposition of poverty decline. We then use the repeated cross-section 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. For many readers, the decomposition may evoke well-known issues about the difficulty of separately identifying age, period, and cohort effects, but the question we pose does not raise these issues. Instead, we aim to clarify what one can learn about the relative pace of within- versus between-cohort change given the linear dependence of age, period, and cohort. Our main takeaway, observable in several ways, will be that poverty has declined between birth

¹³World Bank Open Data, <https://data.worldbank.org/indicator/SL.UEM.TOTL.ZS?locations=ZA>, accessed 14 January 2025.

¹⁴World Development Indicators, <https://databank.worldbank.org/reports.aspx?source=2&series=SH.DYN.AIDS.ZS&country=ZAF>, accessed 31 July 2024.

cohorts at roughly the same pace as it has within them. We will see that if one separates poverty decline into a cohort component and a time component, then this fact implies that the lion's share of poverty decline has accrued over time within cohorts.

In our analysis of the repeated cross-sectional data, we focus on consumption poverty while also reporting results for income poverty where possible. In part, this prioritization is practical: all of our cross-sectional surveys have consumption, while only some have income. But it also has conceptual appeal because consumption poverty more closely reflects the standard of living. Nevertheless, it raises some theoretical ambiguities. In a lifecycle model with no uncertainty or borrowing constraints, the trajectory of consumption poverty over time reflects shifting preferences, while its starting level reflects lifetime wealth. From this point of view, discussions about poverty change should focus only on lifetime wealth; variation within a lifetime reflects choice rather than resource availability. However, wherever both consumption and income are available, we will find that consumption poverty follows a similar path to income poverty, and the two measures yield similar decomposition results. We will argue that this similarity suggests that rising incomes play an important role in the decline of poverty over people's lives, motivating our investigation of income dynamics and their sources in the panel data.

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, we 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 the Appendix, 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.

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

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

This expression is an *age* 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 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* decomposition, which is useful for studying generational dynamics:

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

The first term captures the difference between years 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 mixes them because lagged poverty rates are unavailable for new cohorts.

Within this cohort decomposition, one can distinguish four types of poverty decline. Appendix Figure A.2 visually compares how they play out over time and over the lifecycle. 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.¹⁵ While unlikely, this scenario highlights a sharp case in which our cohort decomposition entirely attributes poverty decline to cohort replacement. The ‘replacement’ terms in Equation (3) are negative, while the ‘within’ term is zero.

In the other three cases, poverty falls within cohorts over time, so that the ‘within’ term in Equation (3) is negative. The cases depend on the signs of the ‘replacement’ terms. First, if each new cohort enters the population at a *lower* poverty rate than surviving cohorts, then the ‘replacement’ terms are also negative. Here, the pace of poverty decline across cohorts at the start of the lifecycle exceeds the pace of poverty decline within cohorts over the lifecycle. 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’ terms turn positive. Here, the between-cohort pace lags the within-cohort pace. In sum, the signs of the ‘replacement’ terms tell us whether the rate of improvement across cohorts exceeds, equals, or falls short of the rate of improvement as people get older.

The time and lifecycle representations of the same underlying variation highlight the linear dependence of age, period, and cohort. After conditioning on cohort, one cannot distinguish age-specific patterns from shared progress over time, and indeed the cohort decomposition does not attempt to disentangle them. For example, suppose that age-specific poverty rates decline in parallel over time. One can interpret this pattern as a common age profile with an intercept that

¹⁵If poverty declines within but not between cohorts, then aggregate poverty does not change; poverty has a lifecycle pattern but no aggregate trend.

falls across cohorts or as a shared decline in poverty over time with no inter-cohort change. A cohort-age decomposition would render the between- versus within-cohort question trivial: at least with a stable age distribution, it would attribute all aggregate poverty decline to cohort rather than age. In contrast, the cohort-time decomposition in Equation (3) allows one to directly ask whether the between-cohort rate of change is slower or faster than the within-cohort rate of change.

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 the Appendix, we verify empirically that it is, and furthermore that the decomposition results do not change if we reweight the sums in Equations (2)-(3).

Cohort and age are individual-level characteristics, but we measure poverty at the household level, raising the question of how to map individual-level demographic characteristics to household-level economic outcomes. As our main approach, we follow the convention of assigning households the demographic characteristics of their heads (e.g. Aguiar and Hurst 2013). This approach amounts to studying how the probability of living in a poor household relates to age, period, and cohort among household heads. Because headship may be endogenous to economic circumstance, we verify that we obtain similar results when relate this probability to age, period, and cohort among *all* individuals.

3.2 Tracking poverty by age, year, and cohort

Equation (2) asks how poverty decline is distributed across ages, while Equation (3) asks how it is distributed across birth cohorts. Figure 2 illuminates both perspectives by plotting how poverty varies with age, both at a point in time and within a single cohort. The grey curves are cross-sectional age profiles, comparing households with different aged heads in the same year. These cross-sectional profiles tell us how households with young and old heads differ at a point in time, not about how the probability of being poor changes over a given household’s lifecycle. The colored curves with markers are cohort age profiles, following households with heads 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.¹⁶ At a point in time, households

¹⁶India’s and Indonesia’s cross-sectional age profiles have positive slopes in the 20s, when headship rates are low

with young, middle-aged, and old heads have similar propensities to be poor. In some countries, this propensity fallen considerably over the sample period, but the decline manifested in parallel shifts of the cross-sectional profile. Similar results obtain when we plot the share of all individuals (rather than household heads) living in consumption-poor households (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 has been a shared experience, benefiting young and old households. One can see this 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, is similar across the age distribution.¹⁷

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 households 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 predecessors’. The cross-sectional profiles are flat because within-cohort change exactly keeps pace with between-cohort change. This combination of features is characteristic of lifecycle consumption profiles in developing economies that are experiencing growth. For example, Deaton (1997) finds offsetting slopes and intercept shifts in Taiwanese lifecycle consumption profiles from late 20th-century data. There as here, the young and old experience 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, implying a much smaller role for cohort effects and therefore generational replacement. Cohorts born 25 years apart, to use a conventional definition of a generation, do tend to start their lives at very different rates of poverty. But by the time a younger cohort enters the population, its older counterpart will have already closed the gap.

(Appendix Figure A.3), so that selection into headship likely plays a role. When we plot the share poor among all individuals rather than just household heads, the pattern disappears (Appendix Figure A.4).

¹⁷An exception is Indonesia, where age-specific declines appear smaller for households with very young heads. As before, the pattern disappears when we redraw the figure using all individuals rather than just household heads (Appendix Figure A.8), consistent with selection into headship at very young ages.

3.3 Cohort decomposition

We next quantify this intuition by implementing the cohort 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 ‘new cohorts’ component will be larger in countries with longer intervals between surveys. Between the dependence on the intersurvey interval and the mixing of between- and within-cohort variation, the ‘new cohorts’ component lacks a clear interpretation. In contrast, the ‘surviving cohorts’ and ‘within-cohort’ components cleanly isolate the two types of variation.

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 then average across pairs, weighting by the number of years. We report the results in four separate panels, first for consumption poverty in the sample of household heads, then for consumption poverty in the full sample, and finally for income poverty in both samples where possible. Each plots total annualized poverty decline in blue, followed by Equation (3)’s three components in other colors.

Figure 4 reports the results of the decomposition. In all countries and all sample/outcome combinations, the ‘replacement’ components are small relative to the ‘within-cohort’ component, consistent with poverty declining at roughly equal paces between cohorts and within them. Matching the cohort time series in Figure 3, the ‘replacement by surviving’ components are negative for Indonesia and South Africa, but the magnitudes are modest. Furthermore, this result is specific to the heads-only sample; if we include all household members in the decomposition, then the magnitudes shrink further and even flip sign in South Africa. The one other anomaly is China, where the ‘replacement by 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 ‘replacement by new’ component mixes within- and between-cohort variation, the decompositions for China remain consistent with a primary role for within-cohort decline.

Looking across samples and poverty measures, the decomposition results suggest cohort-neutral poverty decline: the cross-cohort pace of decline matches the within-cohort pace of decline. Because surviving cohorts dominate the population, the decomposition attributes the lion’s share of poverty decline to within-cohort variation. This result is not specific to the rectangular population assumption. 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 reweighting cohorts by their lagged population shares leaves the remainder of the decomposition unchanged.

In sum, most poverty decline—whether measured using consumption or income, whether indexed by the demographic characteristics of household heads or all household members—accrues 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 features, then 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. We will emphasize income-based measures here, as our focus will be on understanding what changes enabled households to increase their living standards, as opposed to simply measuring those living standards per se. Rising consumption cannot be sustained in the long run without rising income, and we wish to understand the sources of that income. Where possible, we will also report Appendix results using a consumption-based definition of poverty. Recall that the panel surveys provide no consumption data for China, and only two rounds' worth for India.

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.2 reports, three of the five panels successfully tracked 90% or more of initial households in their final wave, while those in South Africa (78%) and especially China (52%) tracked fewer.¹⁸ 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.3 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 (Mexico, using an income-based definition) 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 thought of as

¹⁸In India, to be precise, the three-wave panel of rural households only tracked 91% of initial households to endline, while the two-wave nationally representative panel tracked 83%.

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 quite high in the data; poor households often escaped poverty, and non-poor households often fell back into it. Figure 5a plots the probabilities p_1^t and p_0^t of being poor in the last round of each panel survey using the income definition. For initially poor heads, these probabilities are substantially less than one in every country, ranging from 57% in rural India to as low as 26% in South Africa. This result indicates widespread poverty exit. But the rates of entry into poverty among households that were *not* initially poor are also substantial: between 15% in South Africa and 34% in Mexico. Looking at consumption (Figure 5b) does not fundamentally change this picture. If we take a less binary look at the data and examine flows among multiple income (or consumption) categories, we again see substantial movement in both directions, as Figures A.10a and A.10b illustrate.

A corollary of this “churn” is that aggregate poverty rates would have declined by 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, and who would thus 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, 37% of the initially poor remained poor at endline, but only 16% remained poor throughout the panel. Thus, an additional 21% 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 22% to 43% of households’ total consumption on average.¹⁹

¹⁹Appendix 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%).

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.²⁰ 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 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, 16% of initially housing-poor households were always housing poor, as opposed to 10% 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.

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.

This is largely not what the data say. Appendix Figure A.11a plots the evolution of round-to-round transition probabilities over time in the sample of households observed in all rounds (see Appendix Table A.2 for sample composition information). Wherever the probabilities change (in China, India, and to a lesser extent Indonesia), they move advantageously: the probability of poverty conditional on non-poverty in the prior round fell over time. The other series are fairly flat.²¹

Appendix Figure A.12 presents a more continuous view of the data, plotting local linear regressions of the probability of being poor in a given round conditional on income in the previous round, separately for each round. This 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 Figure A.11a: conditional probabilities of being poor given *any* initial level of income fell in China and India, but were roughly constant elsewhere. These local linear regression estimates are monotonically declining in initial income and display no particular change in slope near the poverty line.

Another natural question is how much changes over time in household composition influence transition probabilities. The birth of a child will typically increase household size without mechan-

²⁰For China, we use income rather than consumption poverty for calibration, since the CHNS lacks data on consumption.

²¹Consumption 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.

ically 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.3 described above suggest that this phenomenon was not a major factor in the long run. To get some sense how much it matters for the round-by-round transitions we are examining here, 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.13a and A.13b 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. The household might reduce the labor hours it allocates to farming, but it might still earn more from farming if prices improve. 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) and occupation type (wage work or self-employment). These 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. We also note whether the household changed locations from a rural to an urban area or vice versa, and whether any of its female members were in the labor force (i.e. working or actively seeking work). Appendix B.7 provides a full description of each of these classification procedures; in brief,

- The sectoral classification is based on the sector codes households self-report whenever these are available, grouping agriculture together with other primary sector activities (e.g. forestry, mining). On average, these self-reports are missing for 20% 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.7.3 for further discussion. Both wage- and self-employment are common in the data: pooling across rounds and averaging across surveys, 62% of households have at least

one member who identifies as self-employed, and 78% 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 a different administrative region in the final round of the panel than in the first round, 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 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.7.4 for further discussion.

To characterize realized income, we simply calculate changes in the share 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). 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, and—as we will discuss further below—in this case women mostly *withdrew* from the labor market, and 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 livelihood 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 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. The 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 were roughly three 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, in no cases did an increase in transfers (from either public or private sources) play a substantial role. Among households that exited poverty, the share of income from transfers rose slightly in China (4%) and everywhere else fell. Among households that entered poverty, the share of income from transfers rose substantially in China (36%), Mexico (24%) and slightly in South Africa (1%), while falling in India and Indonesia. By and large, the data look like what one would expect in an environment where redistribution is progressive and has not changed in overall generosity over time.

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. This is a classic question in the literature on structural change and poverty reduction.

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 agricultural activity to non-agricultural activity 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.

This decomposition is related to that introduced by Ravallion and Huppi (1991, Equation 4) (henceforth RH) and used in several applications for the same broad purpose, to attribute poverty

changes to changes within sectors and movements between them. An essential difference is that implementing Equation (5) requires panel data, while implementing the RH decomposition does not. In fact, RH requires only 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.²² We therefore focus on implementing Equation (5) while also reporting results from the RH decomposition in the Appendix 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 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 there were substantially fewer of them, so that households that stayed in agriculture or stayed 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, this primarily meant staying in agriculture. Continuing within agriculture was thus a viable route out

²²Using 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

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

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 sum of the intersectoral terms, $(\alpha_{an} - \alpha_{na})(\pi_n^0 - \pi_a^0)$, will be non-zero provided only that initial poverty rates in the two sectors differed and there was net population movement between them.

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' relatively non-agricultural starting points.

We emphasize that 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, which could put upward pressure on agricultural wages without creating many factory jobs for farm workers. But the descriptive facts imply some bounds on the underlying economic possibilities. In order for the scenario just described to fit the data, the supply and demand for labor in agriculture would need to be relatively inelastic.

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.4). 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.14). 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.14). 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 representative of this view: rural-to-urban migrants experienced by far the largest rate of net poverty reduction, 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% in India to 64% in South Africa) and over time, with between 18% (India) and 54% (China) of households changing status between baseline and endline (see Appendix Table B.7). In all countries, households in which a woman entered the workforce had among the highest rates of net poverty exit, and in some, they were numerous enough to contribute meaningfully to poverty change. Households in which a woman *exited* the labor force contributed *negatively* to poverty decline in Mexico and South Africa, and in all countries had a lower net exit rate than households in which a woman began to work.²³

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. This result suggests strong selection effects, in contrast to other countries where poverty either declined little or increased among such households.²⁴

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 an important role for intra-sectoral change in 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. This dynamic portrait can be useful for situating causal analyses of anti-poverty interventions. A randomized trial piloting an intervention that targets a particular pathway out of poverty can provide evidence on effectiveness, but broader facts on the prevalence of that pathway in national populations are crucial for determining next steps. If the pathway is already common, is there room to promote it further? If it is uncommon, which barriers to the pathway are most binding?

²³India had relatively low rates of female labor force participation generally (Appendix Table B.7), as noted and discussed by Pande et al. (2017).

²⁴A small share (18% in South Africa and less than 10% in the other countries) of households had no working-age woman. As noted above, the presentation in Figure 9 groups these households together with those that had working-age women who did not work. If instead we restrict to a sub-sample of households that always had working-age female members, results are similar for all countries except South Africa, where it eliminates the contribution to poverty reduction from households that never had woman in the labor force, as most of these never had a working-age woman.

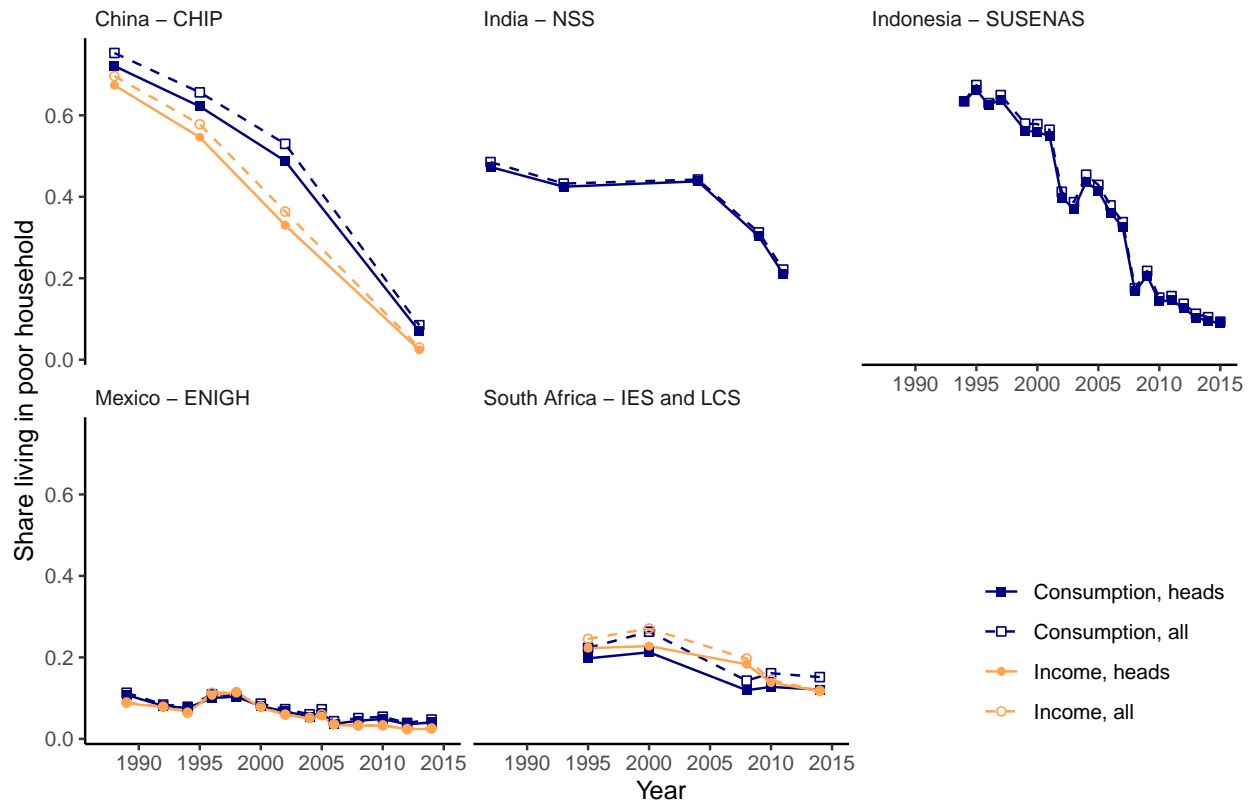
The work of harmonizing survey data collected in different places at different times is hard, but it results in a clearer picture of commonalities and differences. Chen and Ravallion (2010) made a great deal of progress on harmonizing cross-sectional poverty data, and the World Bank (2023) continues this important work. But even here, open issues remain—for example in the choice of income or consumption, and in the treatment of durables in the latter—and in most cases, the researchers default to the decisions of country statistics offices. And efforts to harmonize panel data on poverty remain nascent. We contribute on both fronts for five important low- and middle-income countries.

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.
- 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.
- 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.
- Baulch, Bob and John Hoddinott**, “Economic mobility and poverty dynamics in developing countries,” *Journal of Development Studies*, 2000, 36 (6), 1–24.
- Bryan, Gharad and Melanie Morten**, “The aggregate productivity effects of internal migration: Evidence from Indonesia,” *Journal of Political Economy*, 2019, 127 (5), 2229–2268.
- 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.
- Deaton, Angus**, *The Analysis of Household Surveys*, Baltimore: Johns Hopkins University Press, 1997.
- **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.
- 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, 2005.
- 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.
- Krishna, Anirudh**, *One Illness Away: Why People Become Poor and How They Escape Poverty*, Oxford: Oxford University Press, 2010.
- Narayan, Deepa, Lant Pritchett, and Soumya Kapoor**, *Moving Out of Poverty: Volume 2. Success from the Bottom Up*, Washington, DC: World Bank, 2009.

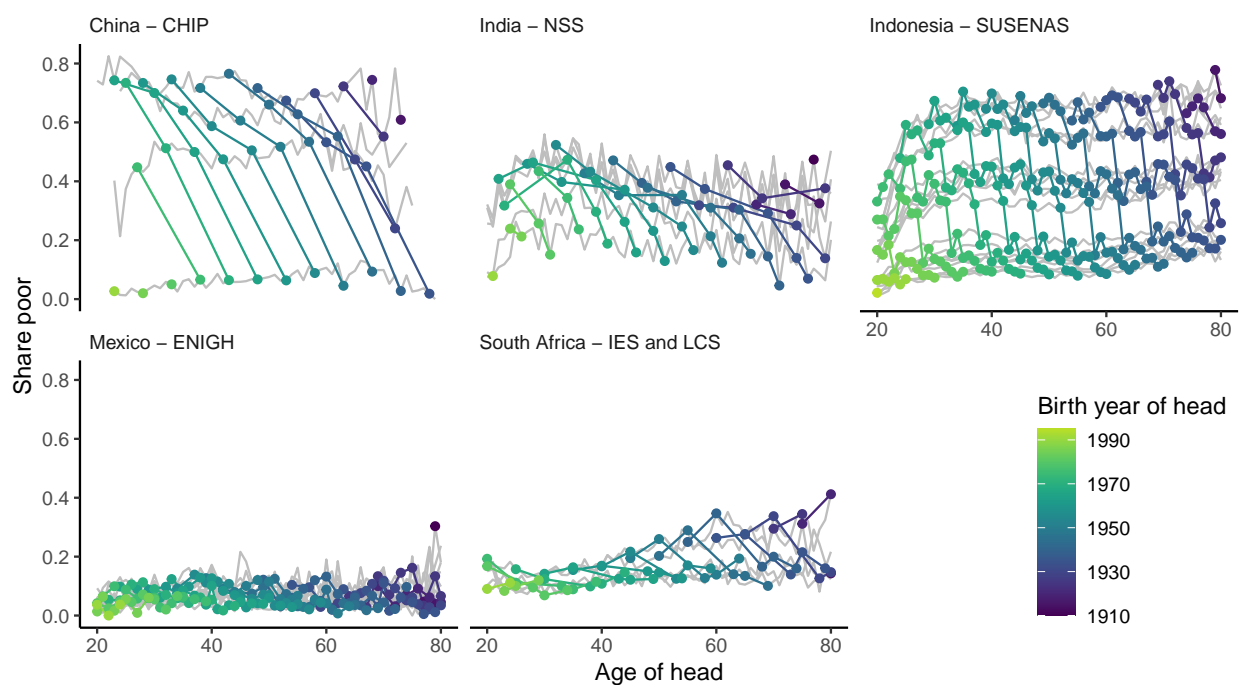
- Pande, Rohini, Charity Troyer Moore, and Erin K Fletcher**, “Women and work in India: Descriptive evidence and a review of potential policies,” CID Working Paper 339 2017.
- Popkin, Barry M, Shufa Du, Fengying Zhai, and Bing Zhang**, “Cohort Profile: The China Health and Nutrition Survey—monitoring 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.
- 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.
- Rosling, Hans, Ola Rosling, and Anna Rosling Rönnlund**, *Factfulness: Ten Reasons We’re Wrong about the World—and Why Things are Better Than You Think*, New York: Flatiron Books, 2018.
- Schoar, Antoinette**, “The Divide between Subsistence and Transformational Entrepreneurship,” *Innovation Policy and the Economy*, 2010, 10, 57–81.
- Shukla, Rajesh**, *How India Earns, Spends and Saves: Unmasking the Real India*, New Delhi: Sage, 2010.
- 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**, “Poverty and Inequality Platform Methodology Handbook,” Available at <https://datanalytics.worldbank.org/PIP-Methodology/> 2023.

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



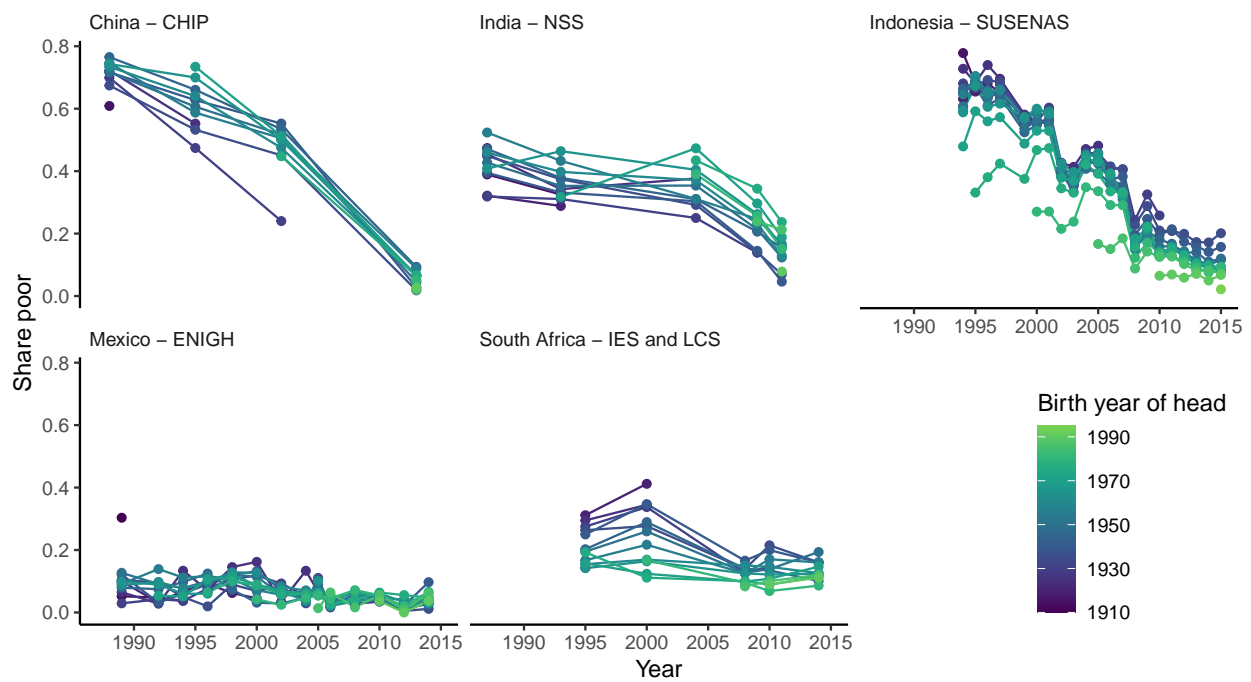
Note: Repeated cross-sectional data. A household is defined as poor if consumption or income is below \$2.15 per person per day. The “all” series represent the share of all individuals living in poor households. The “heads” series represent the share of household heads living in poor households.

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



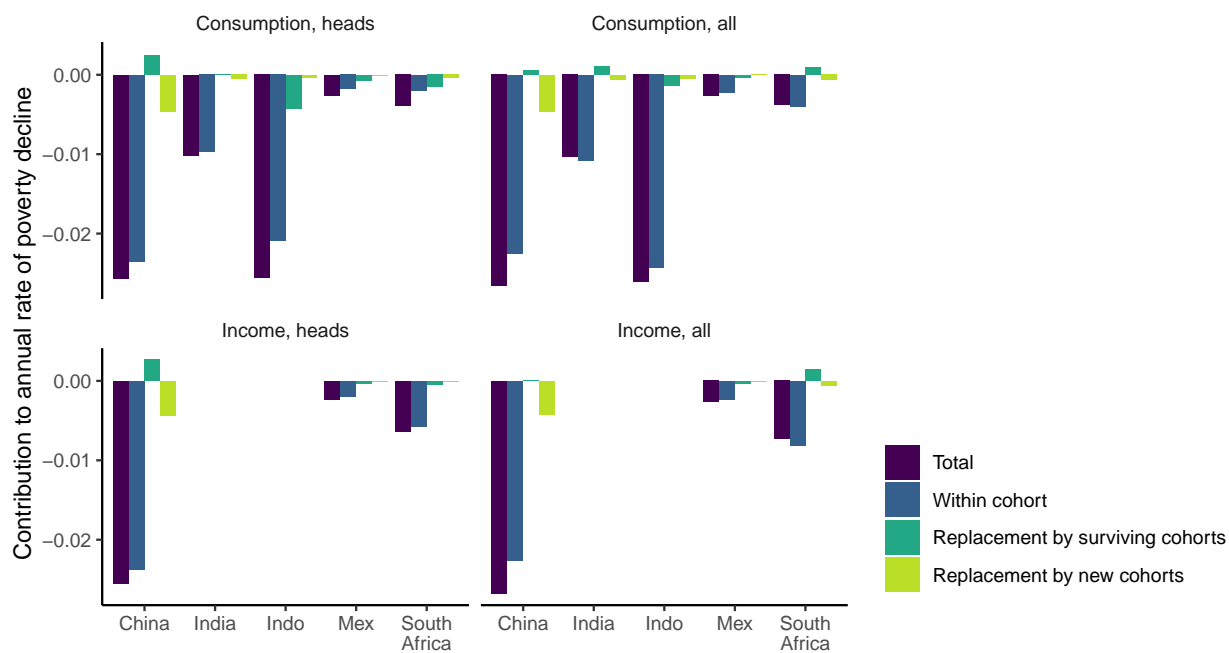
Note: Repeated cross-sectional data. The grey curves are cross-sectional age profiles. The connected scatterplots illustrate poverty over the lifecycle as experienced by cohorts born in years ending in 0 and 5. Households are assigned their heads' demographic characteristics and are classified as poor if consumption per capita is below \$2.15 per day.

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. Households are assigned their heads' demographic characteristics and are classified as poor if consumption per capita is below \$2.15 per day. The figure is different from Figure 2 because cohort poverty rates are projected on time rather than age.

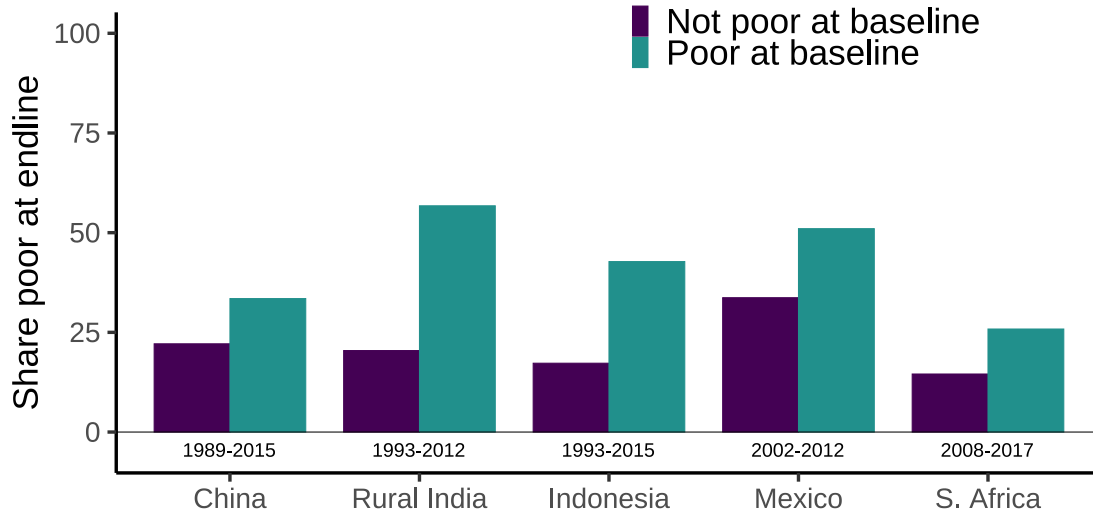
Figure 4: Cohort decomposition of poverty decline



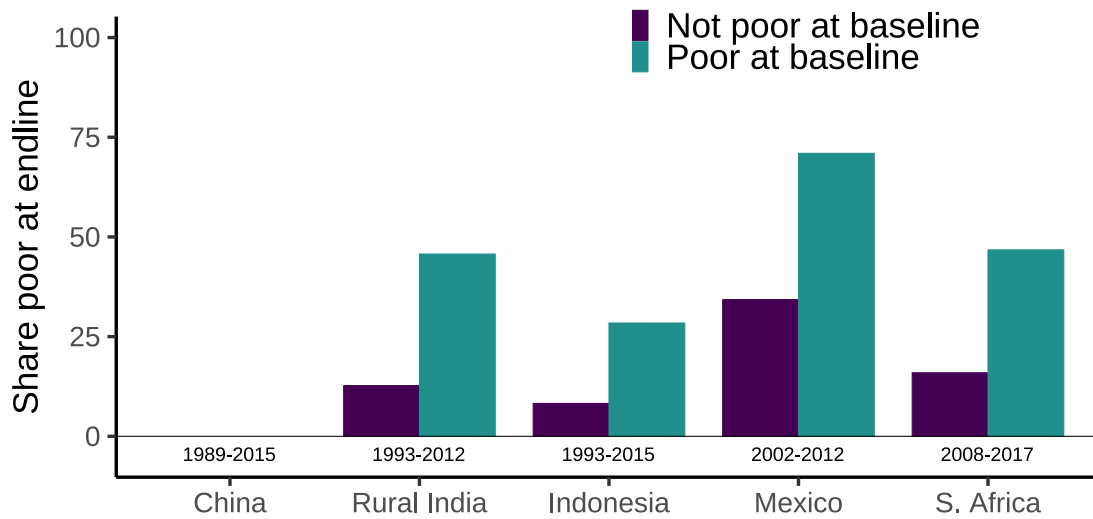
Note: Repeated cross-sectional data. The figure implements the cohort-year decomposition in Equation (3) separately for each country.

Figure 5: Poverty transition probabilities

(a) Income

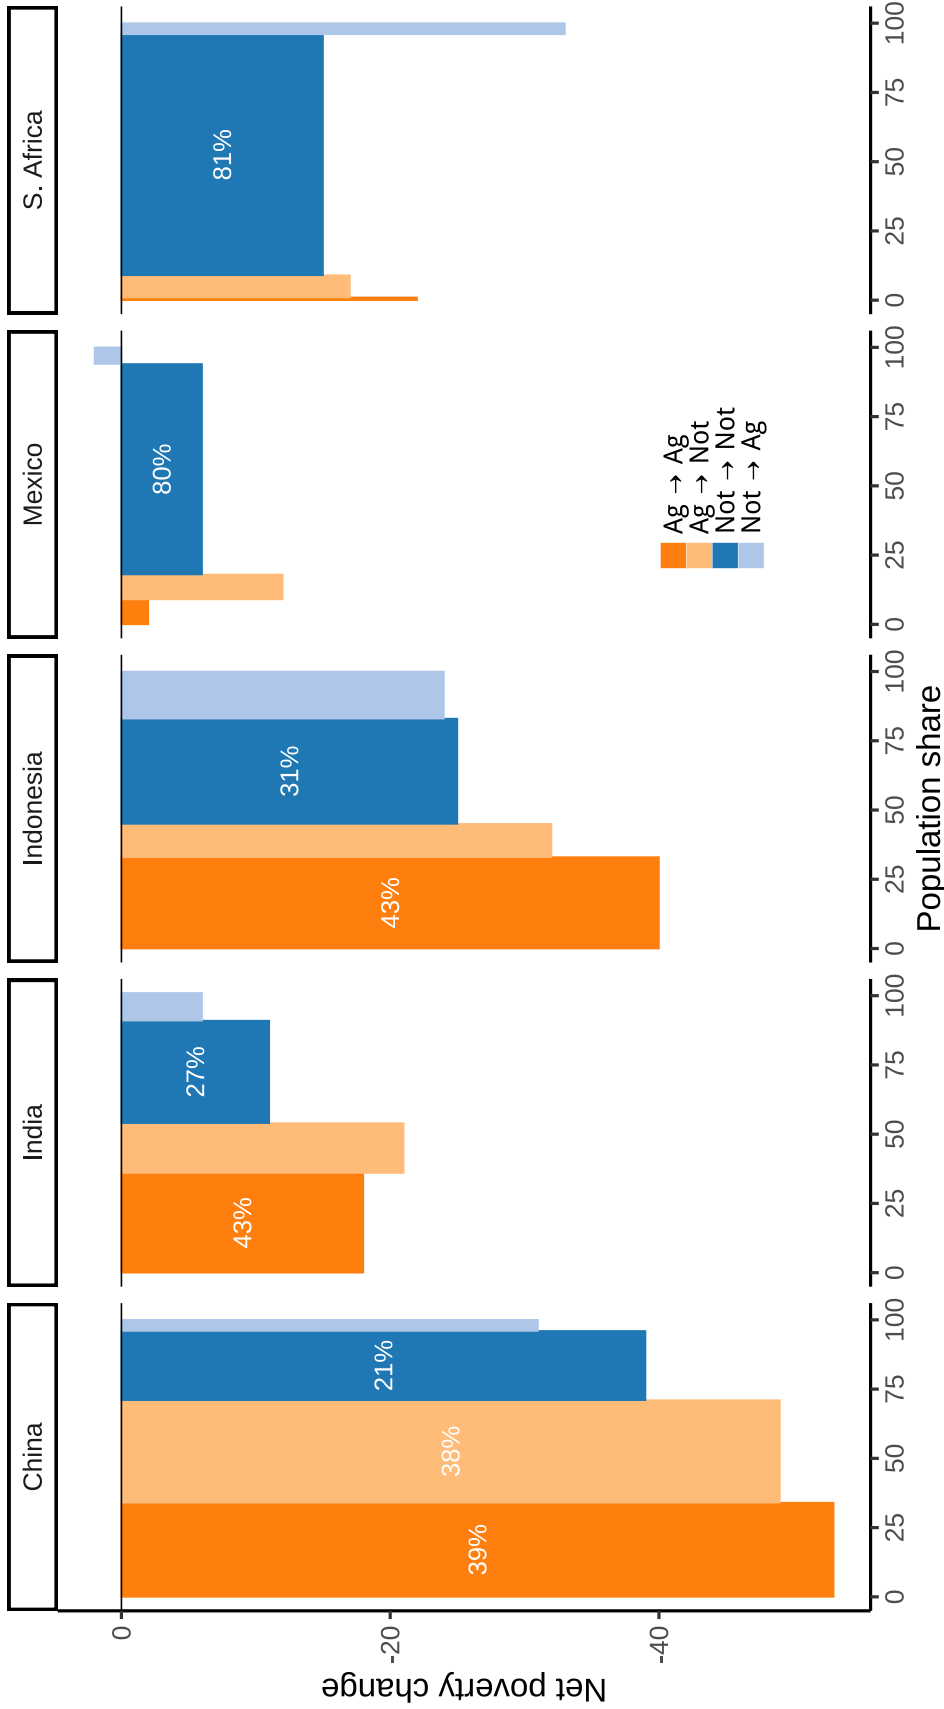


(b) Consumption



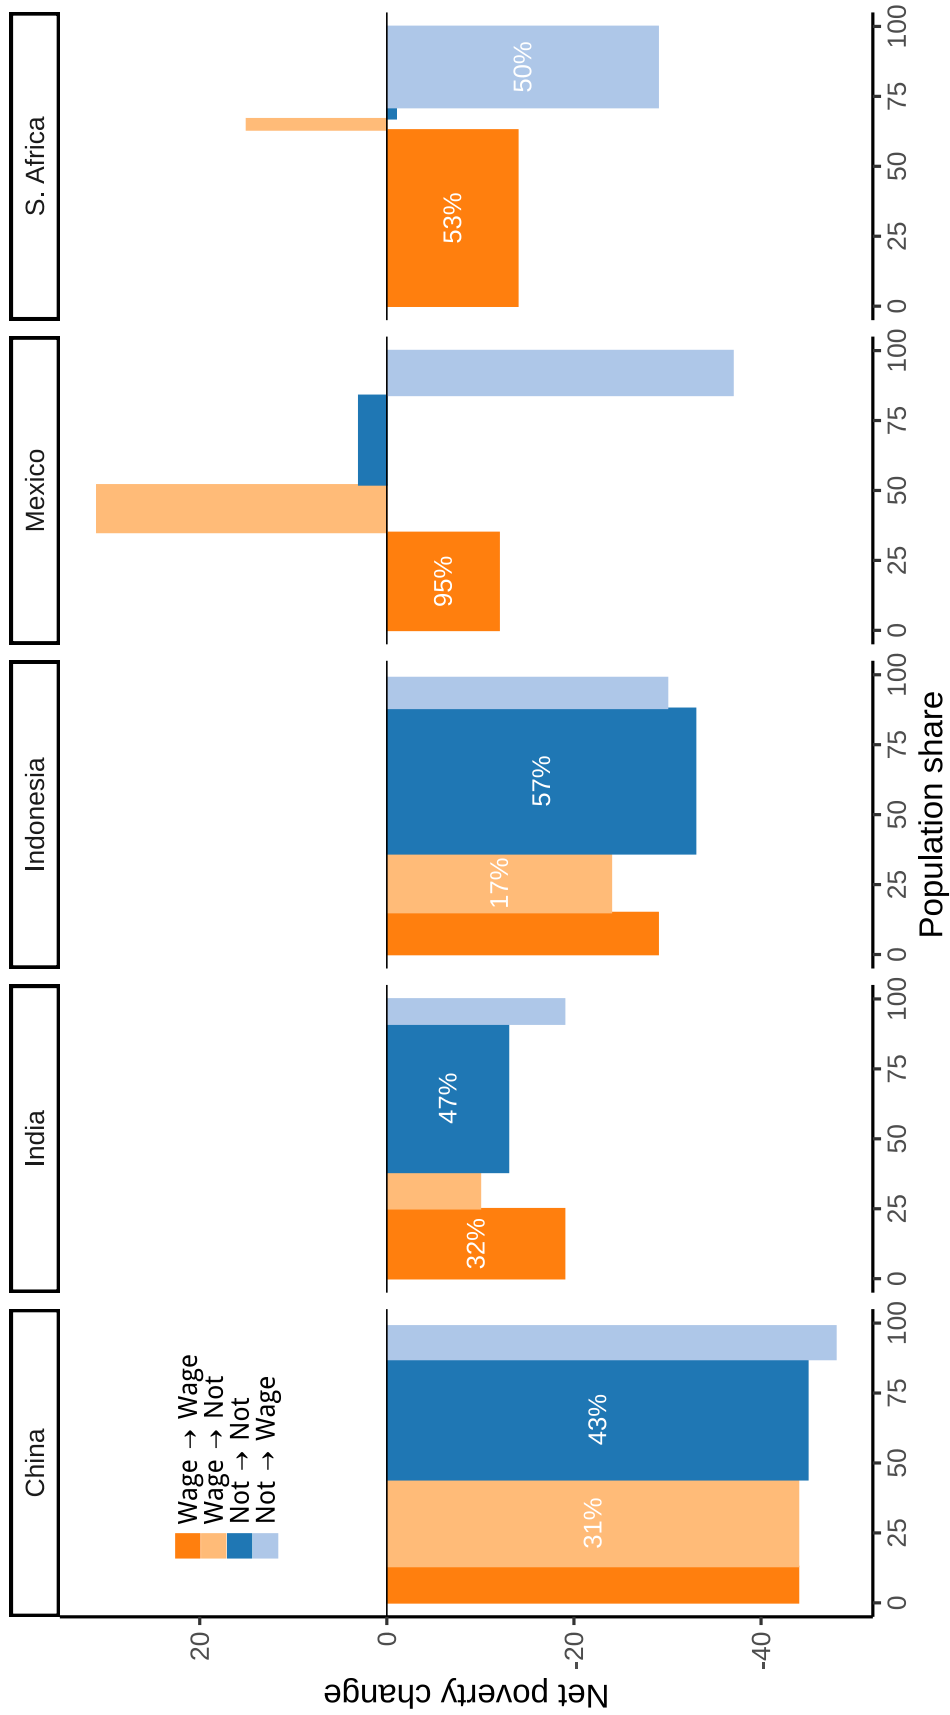
This figure reports transition probabilities between the first and last round of each panel survey: specifically, the probability that a household was poor in the final round conditional on being poor or not poor, respectively, in the initial round. For comparability with Figures A.11a and A.11b the 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



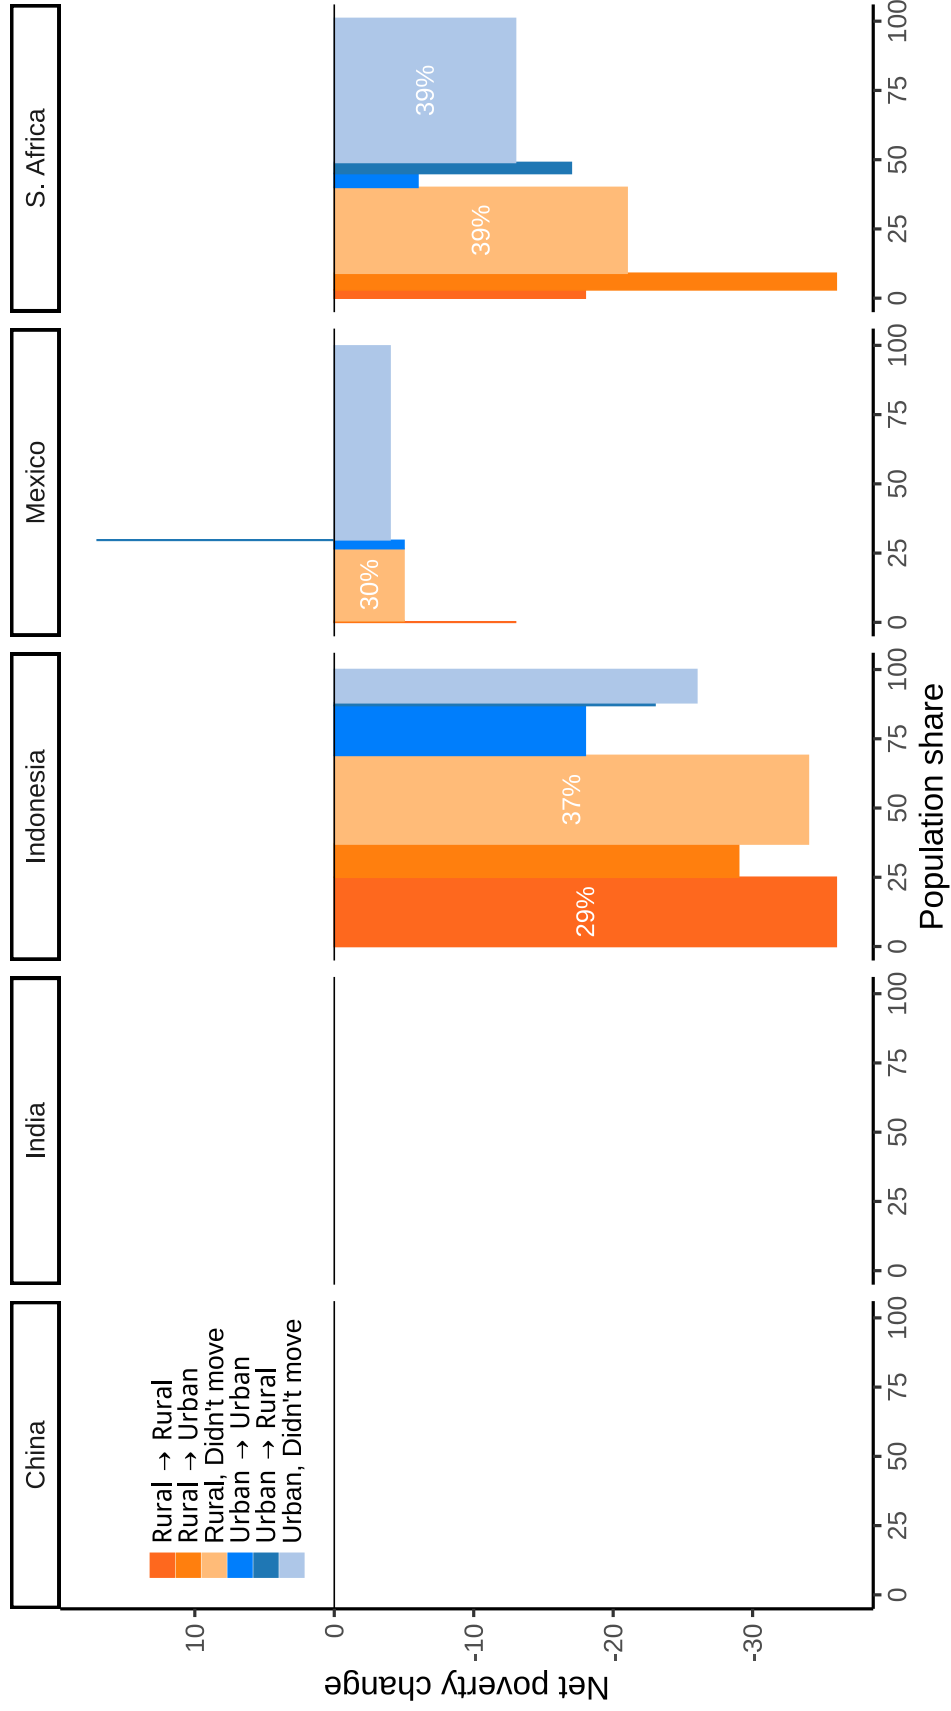
This figure reports changes in poverty headcounts between the first and last round of the panel surveys, broken down by groups defined based on whether the household's primary source of income was in the agricultural or non-agricultural sector in each round. For example, the bar in the first row represents the net change in the number of poor households, as a share of the total population, among households who earned their income primarily in the non-agricultural sector in both survey rounds. The table at right provides further details as follows: "Pop. share" is the share of the population in each transition group (and so sums to 1 within each country); "Net exit rate" is equal to one minus the share poor within that transition group in the last round divided by the share poor within that group in the first round; "Component" is the product of these two figures (and hence the quantity displayed in the bar chart); and "Effect share" is the ratio of the component to the within-country sum of components, i.e. to the total change in the headcount poverty rate. The sample consists of households observed in both the first and last waves of each survey. See Appendix B.7.1 for details on the classification of households into sectors.

Figure 7: Occupational choice and poverty status



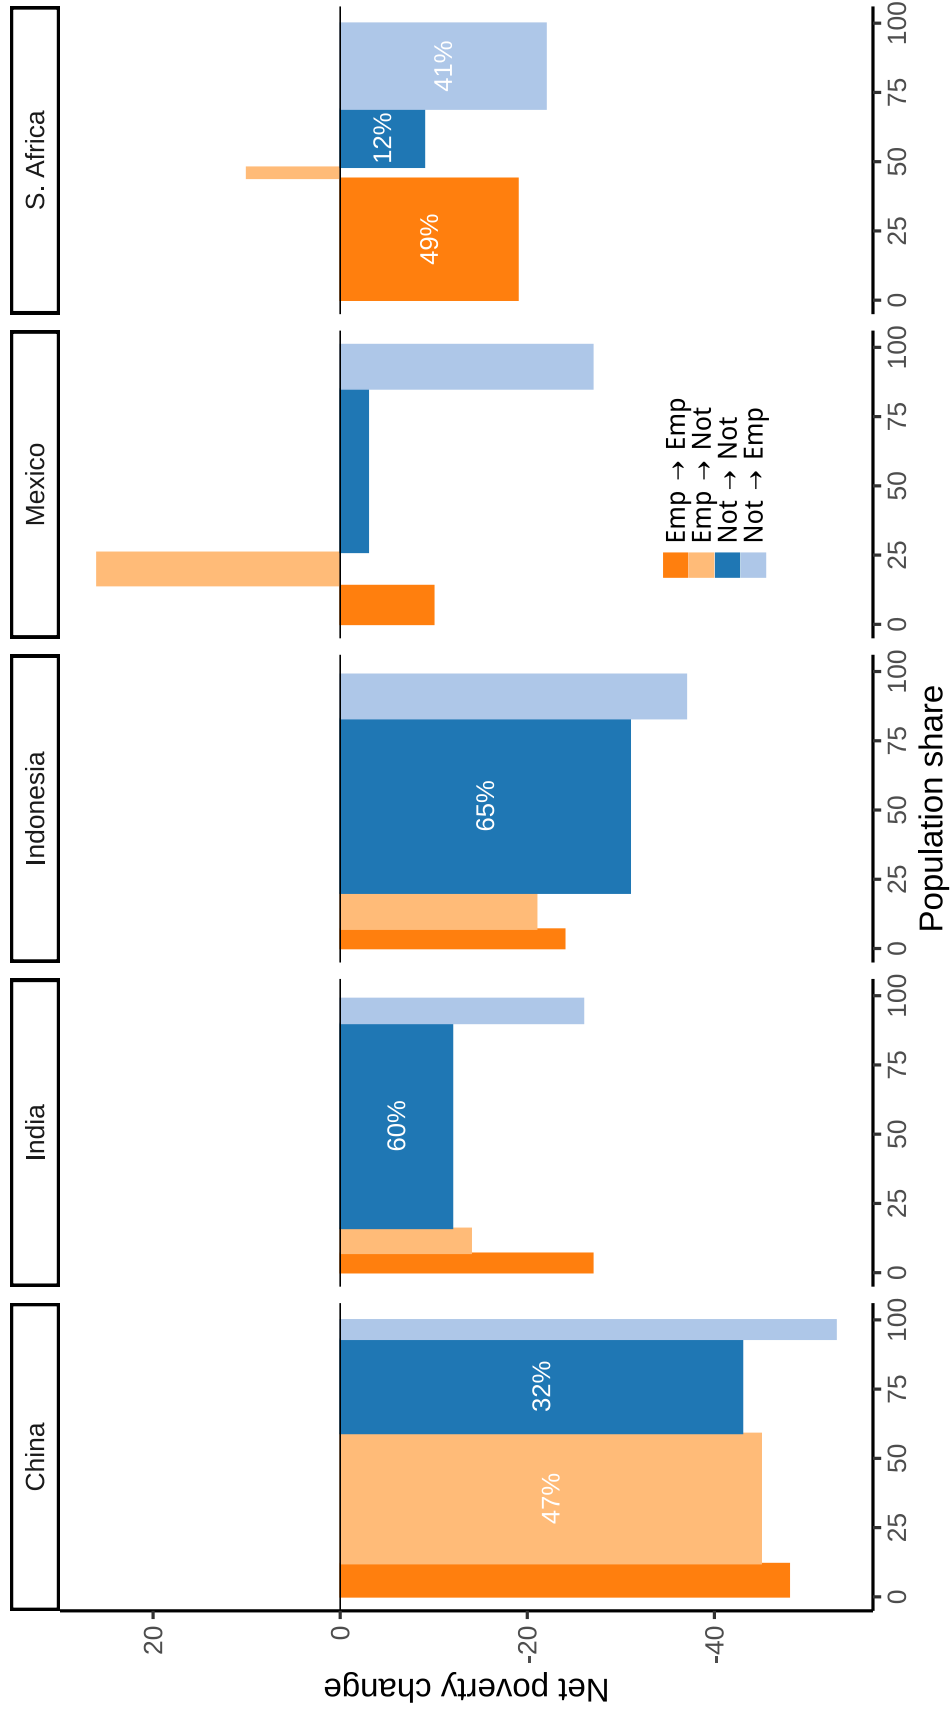
This figure reports changes in the poverty rate broken down as in Figure 6, but with groups defined based on whether or not the household earned its income primarily from wage employment. See Appendix B.7.3 for details on this classification. Note that households in which no member reported being in the labor force are grouped here with others that were not primarily wage-employed; Figure A.14 reports a more detailed breakdown that further distinguishes between these groups.

Figure 8: 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.7.2 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.7.4 for details on this classification.

Table 1: Data sources

Country	Longitudinal surveys		Repeated cross-section surveys	
	Name	Description	Name	Description
China	China Health & Nutrition Survey (CHNS)	10 waves, 1989-2015, 8/22 provinces	Chinese HH Income Project (CHIP)	4 waves, 1988-2013, national
India	India Human Development Survey (IHDS)	3 waves, 1993-2012, national*	National Sample Survey (NSS), thick rounds	5 waves, 1988-2012, national**
Indonesia	Indonesia Family Life Survey (IFLS)	5 waves, 1993-2015, 13/27 provinces	National Socio-economic Survey (SUSENAS)	21 waves, 1994-2015, national†
Mexico	Mexico Family Life Survey (MxFLS)	3 waves, 2002-2012, national	National Survey of HH Income & Spending (ENIGH)	17 waves, 1984-2014, national
South Africa	National Income Dynamics Study (NIDS)	5 waves, 2008-2017, national	Income/Expenditure & Living Conditions Surveys (IES & LCS)	5 waves, 1995-2015, national‡

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	Income poor		Consumption poor		Housing poor	
	At endline	Every wave	At endline	Every wave	At endline	Every wave
China	28%	4%			32%	20%
India	53%	42%				
Indonesia	37%	16%	20%	9%	16%	10%
Mexico	51%	37%	69%	57%	69%	72%
South Africa	29%	13%	35%	19%	33%	23%

This table reports the share of initially poor households who are 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.

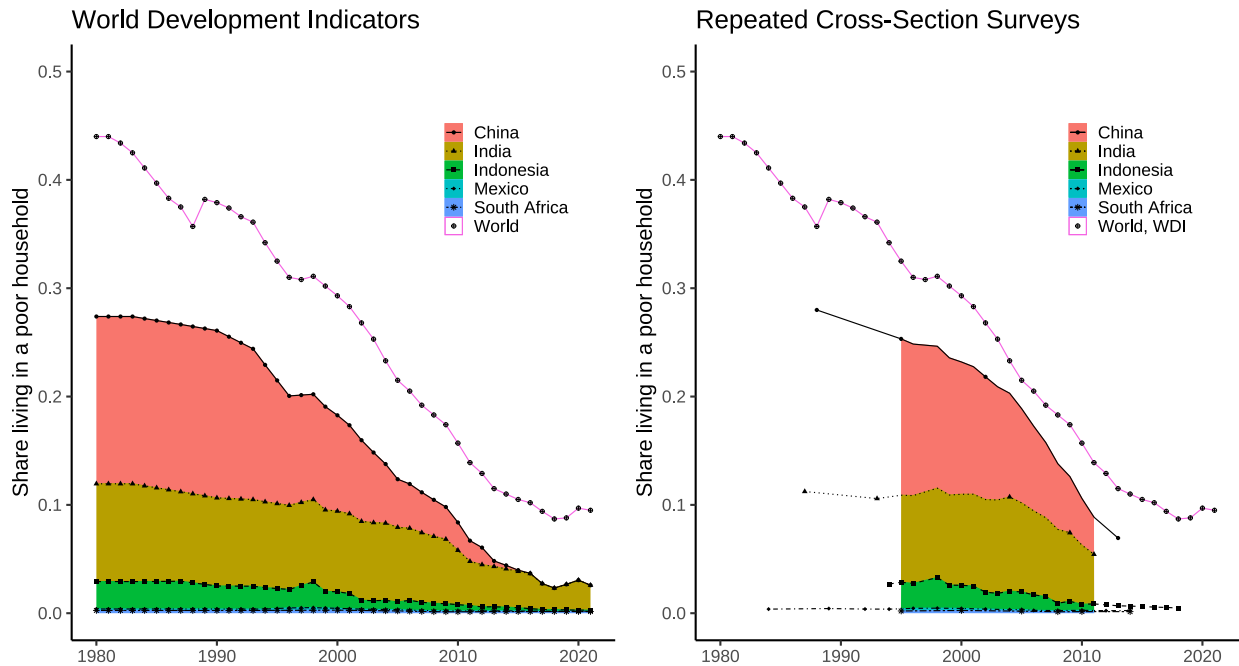
Table 3: Livelihood changes among households whose poverty status changed

	China		India		Indonesia		Mexico		South Africa	
	Exits	Entries	Exits	Entries	Exits	Entries	Exits	Entries	Exits	Entries
Share of sample	53	5	26	9	42	6	23	19	18	24
Any change	40	34	30	29	28	34	14	12	13	11
Ag to not	37	27	21	14	13	7	10	6	7	8
Not to ag	3	7	8	14	14	28	4	6	7	3
Any change	44	43	22	25	31	40	39	43	39	28
Self to wage	13	8	10	7	13	10	33	6	36	20
Wage to self	31	35	12	17	18	30	6	37	2	8
Any change					12	14	1	1	7	8
Rural to urban					11	13	1	1	5	5
Urban to rural					1	1	0	0	3	4
Any change	55	56	22	15	30	25	29	29	36	37
LFP to not	47	55	9	10	11	15	4	22	2	10
Not to LFP	8	1	12	5	19	10	25	7	35	27
Ag	-26	-9	9	12	-36	-21	3	-1	-1	-2
Wage	42	1	-9	-5	-5	-15	20	-23	33	-2
Self Employment	-23	-16	10	15	8	19	-7	1	-5	1
Capital			0	0					2	0
Transfers	4	36	-1	-10	-3	-3	-45	24	-30	1
Male	6	-12	-7	-11	2	-31	35	-27	23	1
Female	4	-20	0	5	6	3	12	2	9	-2
Unattributed	-10	32	7	6	-8	28	-47	25	-32	1

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 33% 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 20%). Blank entries imply the absence of data. For example, only the IHDS and NIDS have questions about 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.7, respectively.

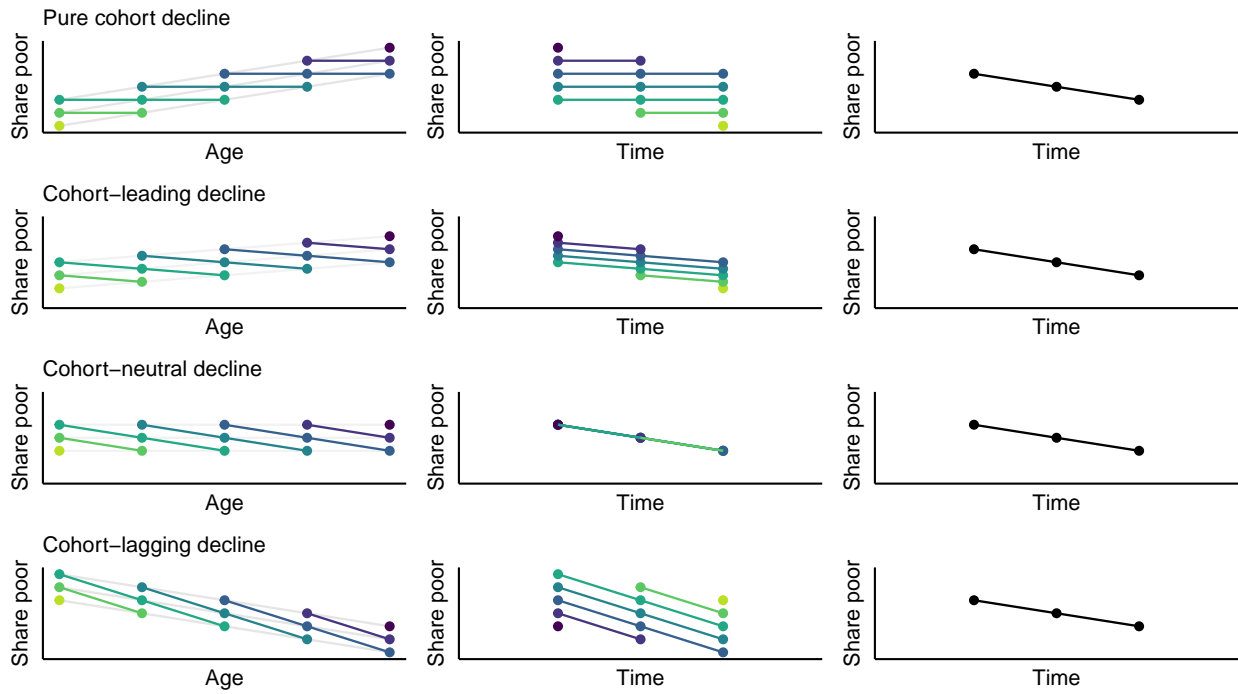
A Additional exhibits

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



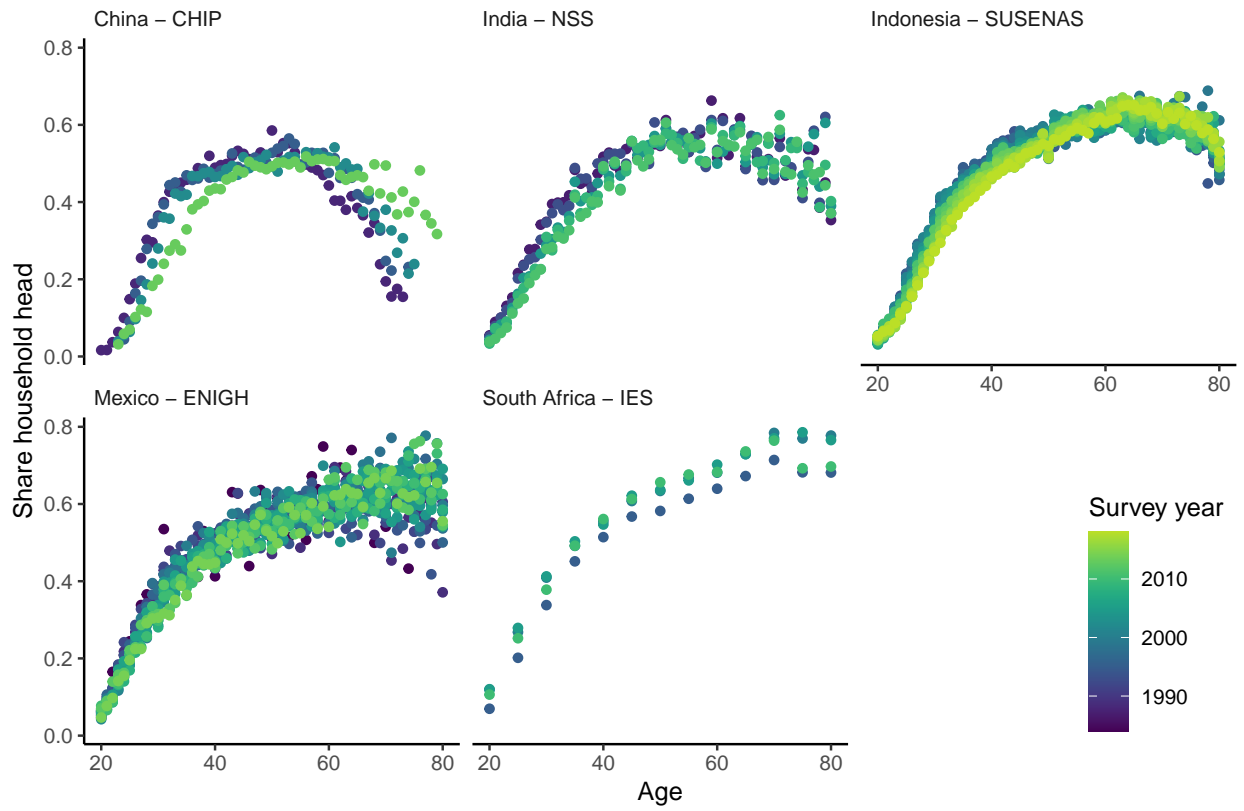
Note: Series are stacked in increasing order of their contribution to the world poverty headcount. Poverty is defined as living in a household with consumption per capita below \$2.15 per day. The left panel plots World Bank estimates; the right panel replaces the World Bank's country-specific estimates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

Figure A.2: Typology of cohort poverty decline



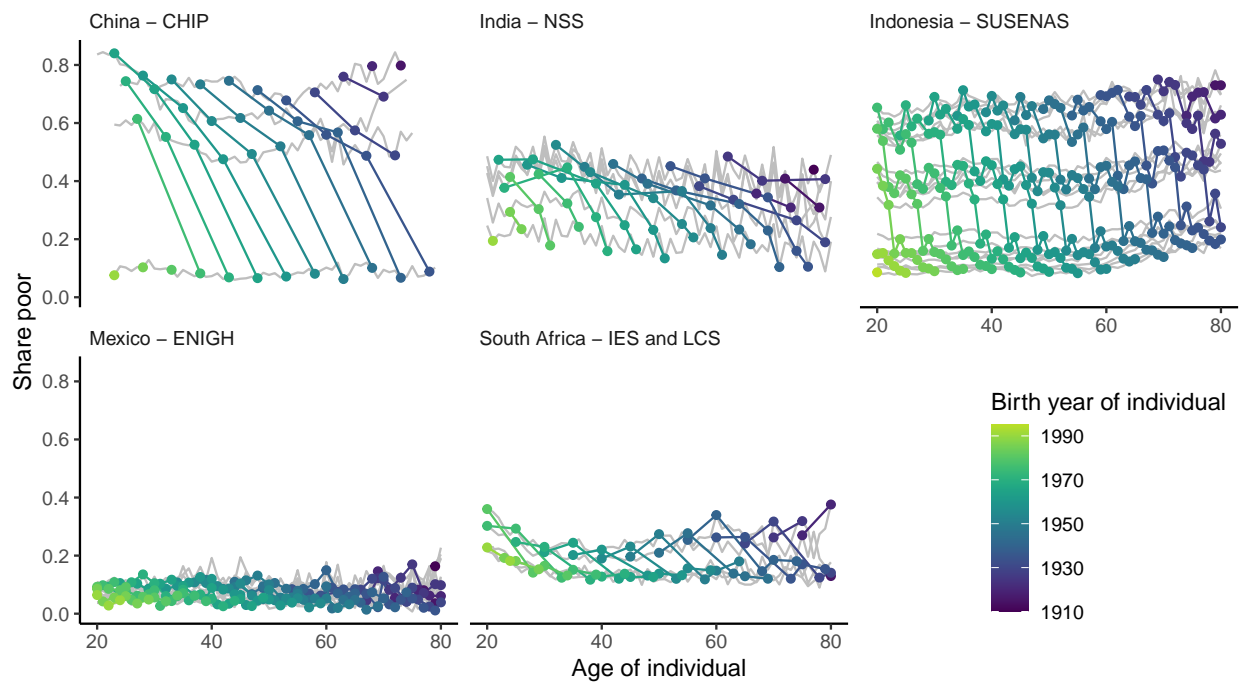
Note: Four types of cohort poverty decline, all consistent with the same aggregate poverty decline. The first column plots cohort poverty rates over the lifecycle, with lighter colors indicating later cohorts. Grey curves correspond to cross-sectional poverty-age profiles. The second column plots cohort 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



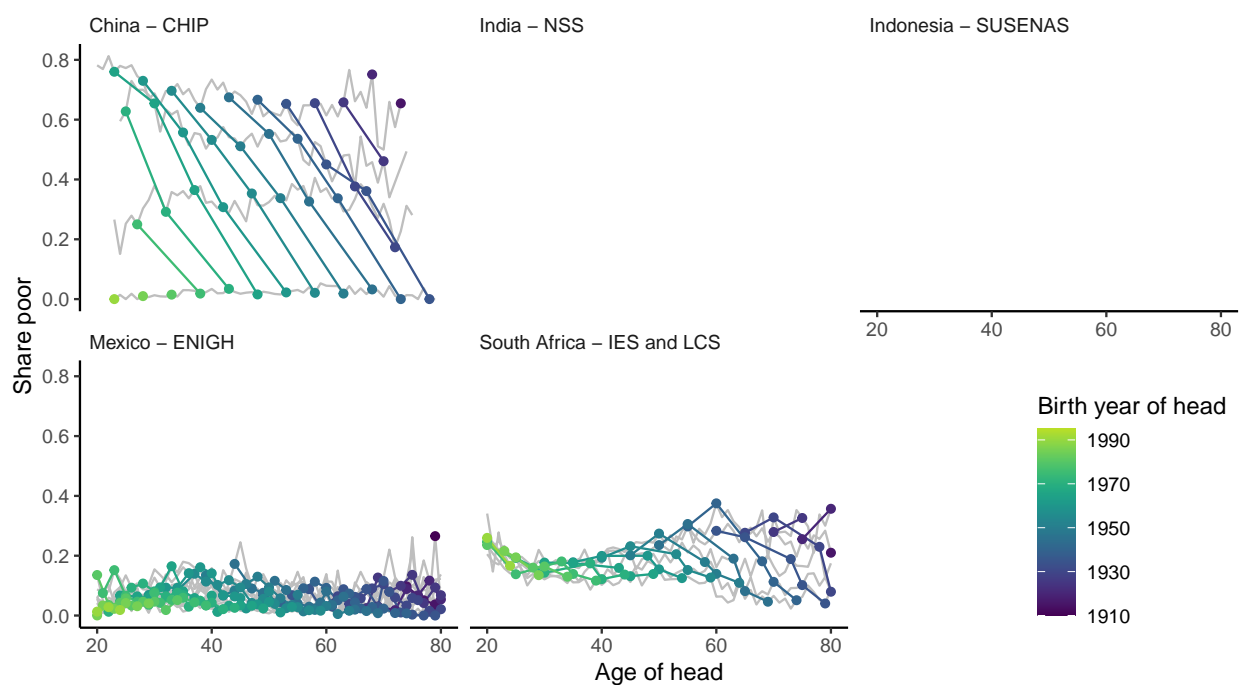
Note: Repeated cross-sectional data. Share of individuals who are household heads by age and year. In South Africa, age is categorized in five-year intervals; all other countries use single-year intervals.

Figure A.4: Consumption poverty over age of individual, by year and cohort



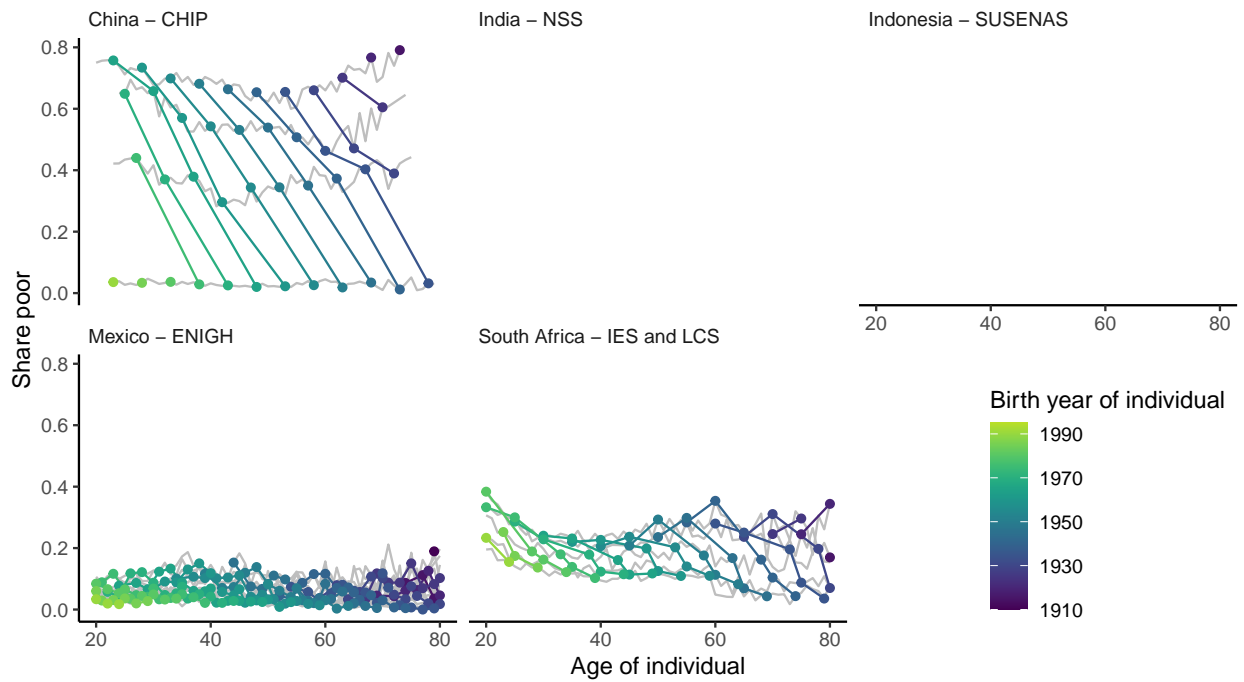
Note: Counterpart to Figure 2 using data on consumption poverty for all individuals. Repeated cross-sectional data. The grey curves are cross-sectional age profiles. The connected scatterplots illustrate poverty over the lifecycle as experienced by cohorts born in years ending in 0 and 5.

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



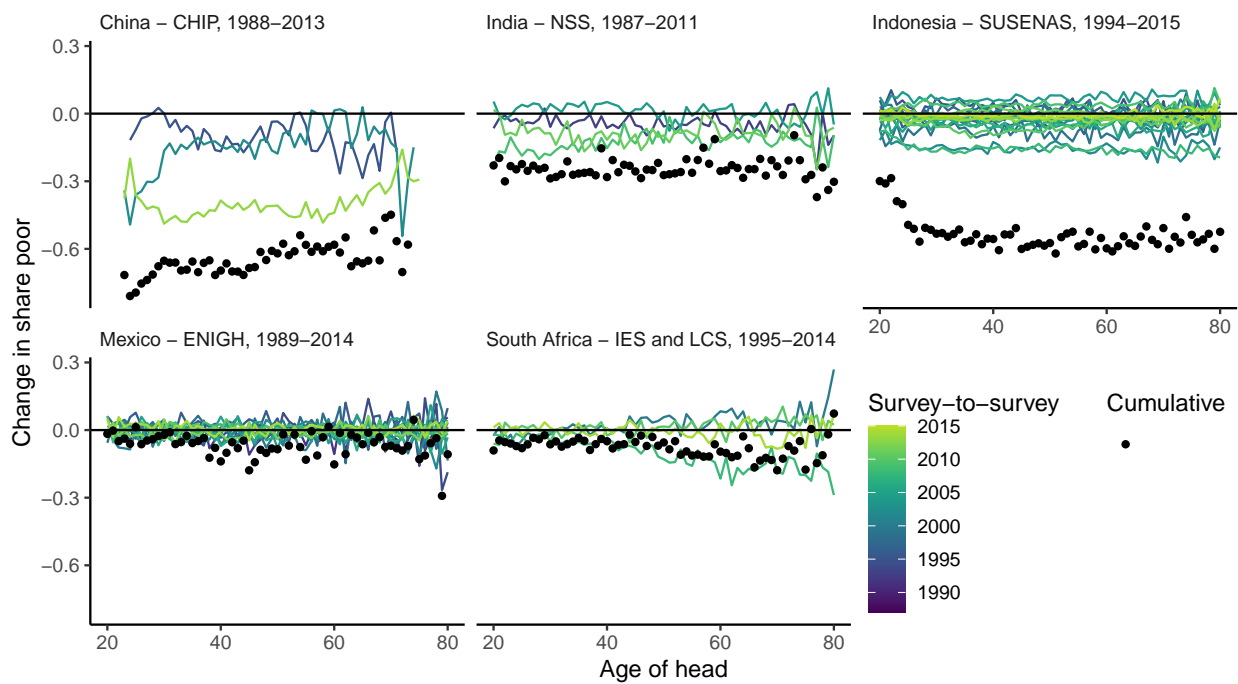
Note: Counterpart to Figure 2 using data on income poverty for household heads. Repeated cross-sectional data. The grey curves are cross-sectional age profiles. The connected scatterplots illustrate poverty over the lifecycle as experienced by cohorts born in years ending in 0 and 5.

Figure A.6: Income poverty over age of individual, by year and cohort



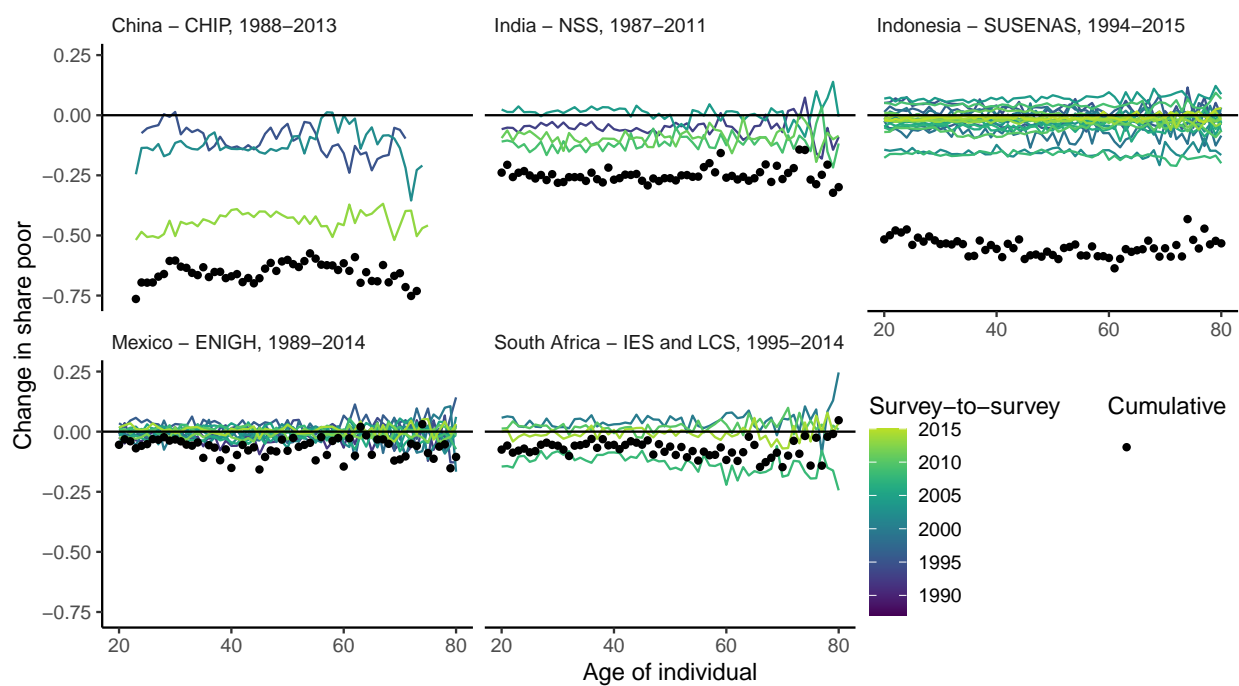
Note: Counterpart to Figure 2 using data on income poverty for all individuals. Repeated cross-sectional data. The grey curves are cross-sectional age profiles. The connected scatterplots illustrate poverty over the lifecycle as experienced by cohorts born in years ending in 0 and 5.

Figure A.7: Changes in consumption poverty by age of household head



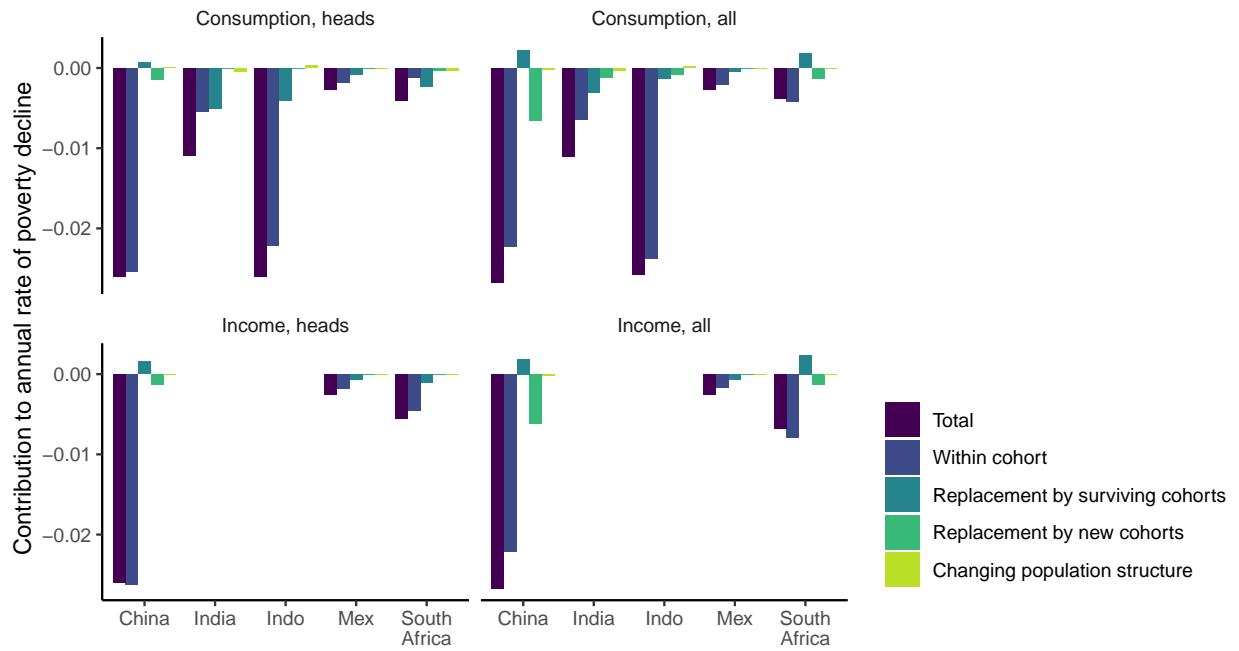
Note: Repeated cross-sectional data. Curves represent changes between adjacent surveys. Scatterplot represents the change from the first to the last survey wave. Households are assigned their heads' demographic characteristics and are classified as poor if consumption per capita is below \$2.15 per day. In South Africa, age and cohort are categorized in five-year intervals; all other countries use single-year intervals.

Figure A.8: Changes in consumption poverty by age of individual



Note: Counterpart to Figure A.7 using data on all individuals instead of only household heads. Repeated cross-sectional data. Curves represent changes between adjacent surveys. Scatterplot represents the change from the first to the last survey wave. Households are assigned their heads' demographic characteristics and are classified as poor if consumption per capita is below \$2.15 per day. In South Africa, age and cohort are categorized in five-year intervals; all other countries use single-year intervals.

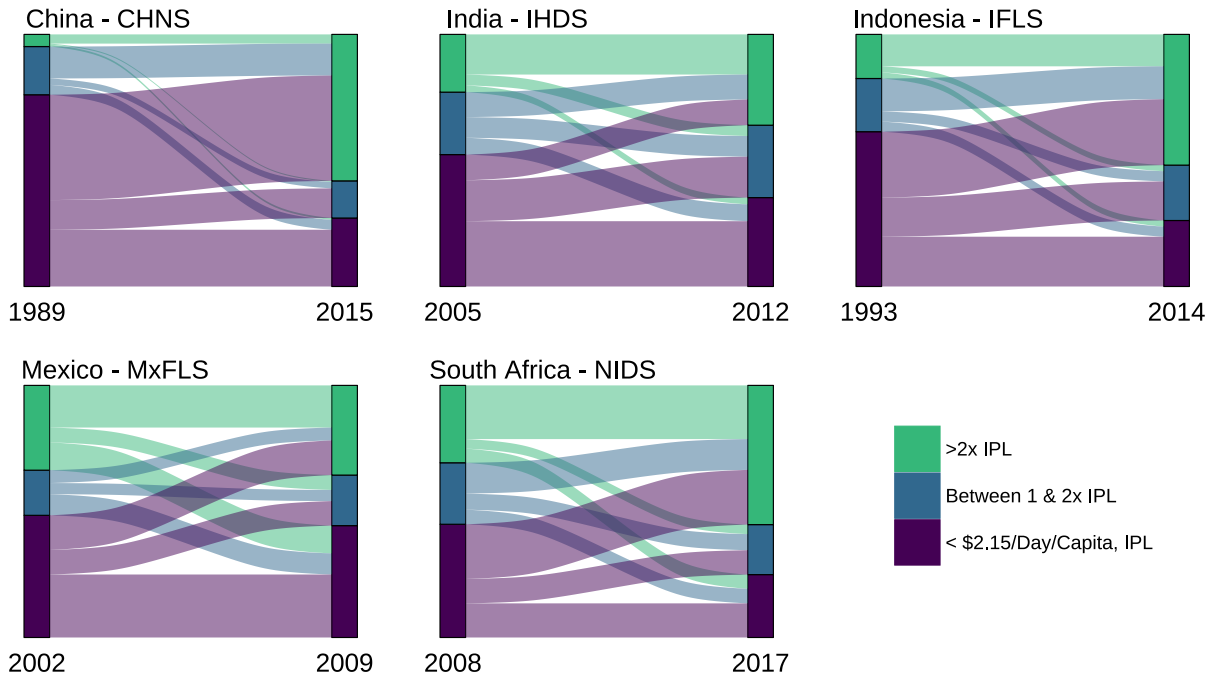
Figure A.9: Cohort decomposition with variable population shares



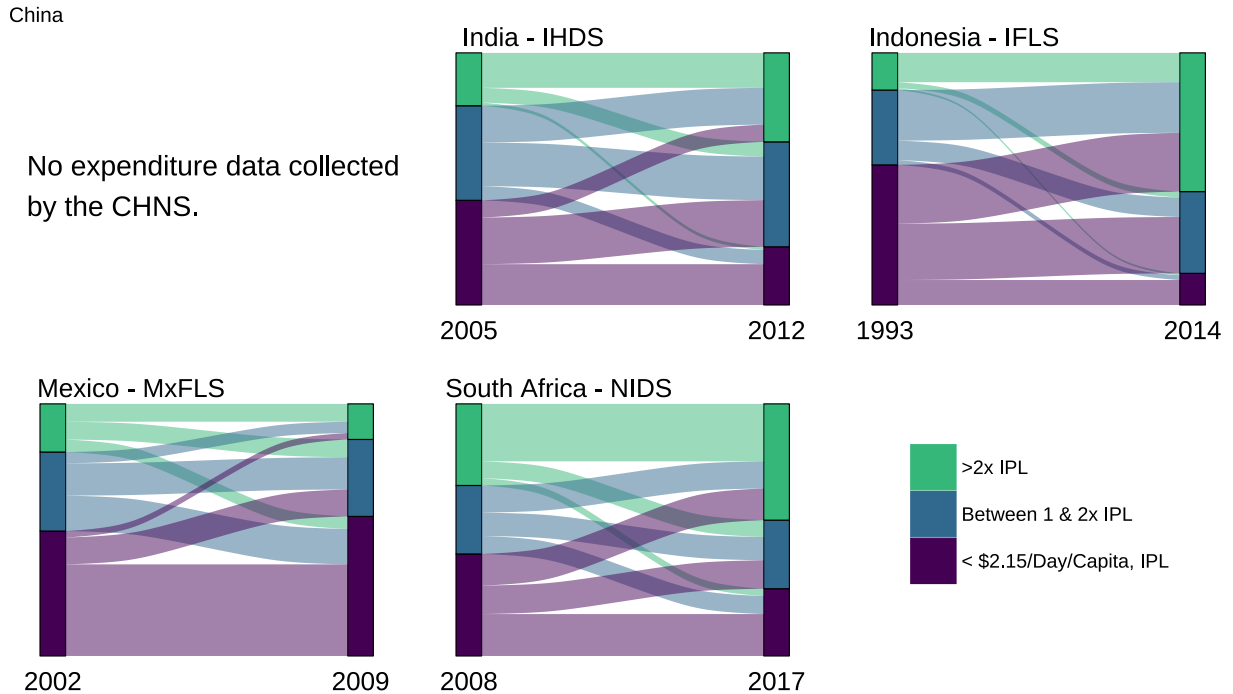
Note: Counterpart to Figure 4 allowing cohorts to have variable population shares. Repeated cross-sectional data.

Figure A.10: Transitions in living standards

(a) Income

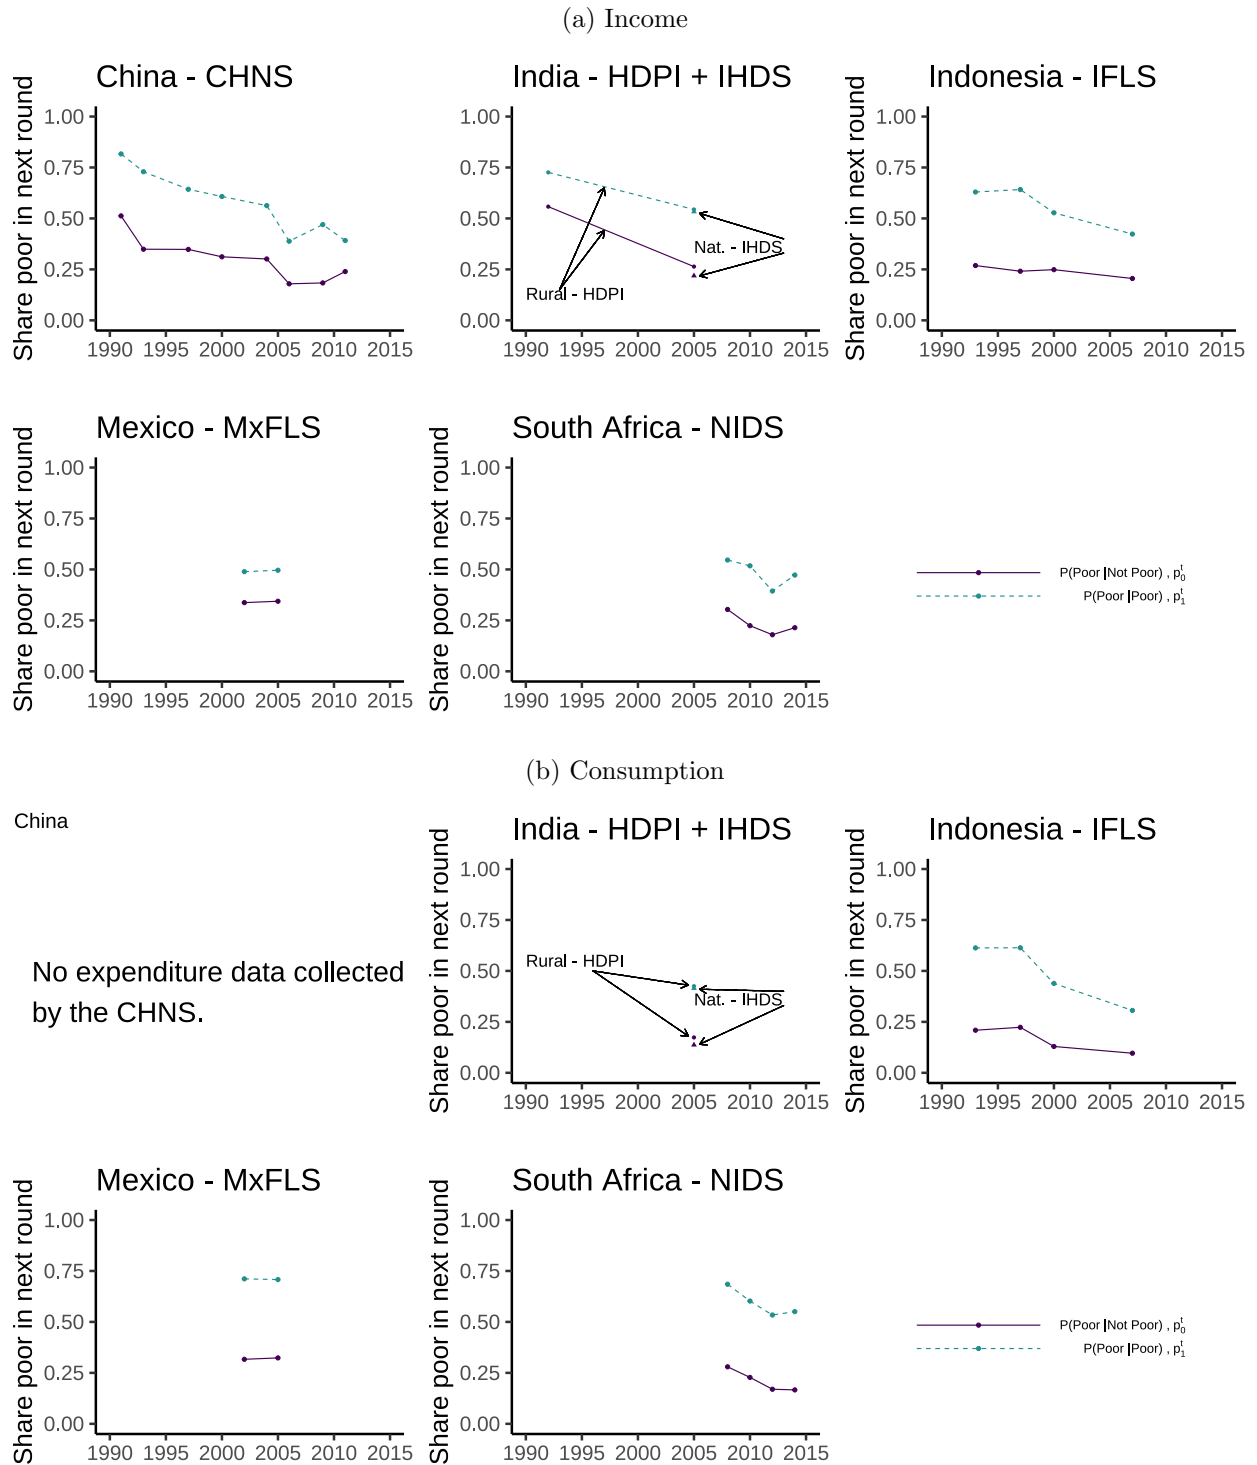


(b) Consumption



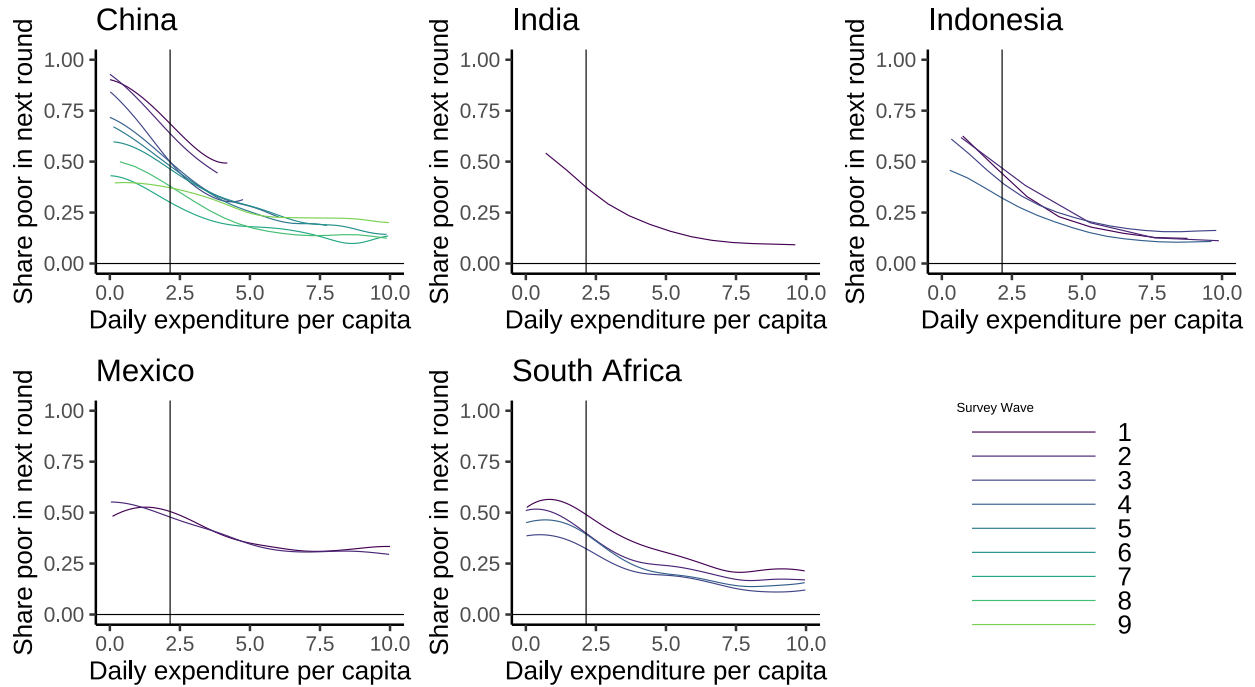
This figure illustrates transitions between different living standards categories between the first and last rounds of each panel survey (conducted in the years denoted at the bottom of each pane). The sample includes households observed in both of those rounds. Living standards are measured using income per capita in the top . 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.

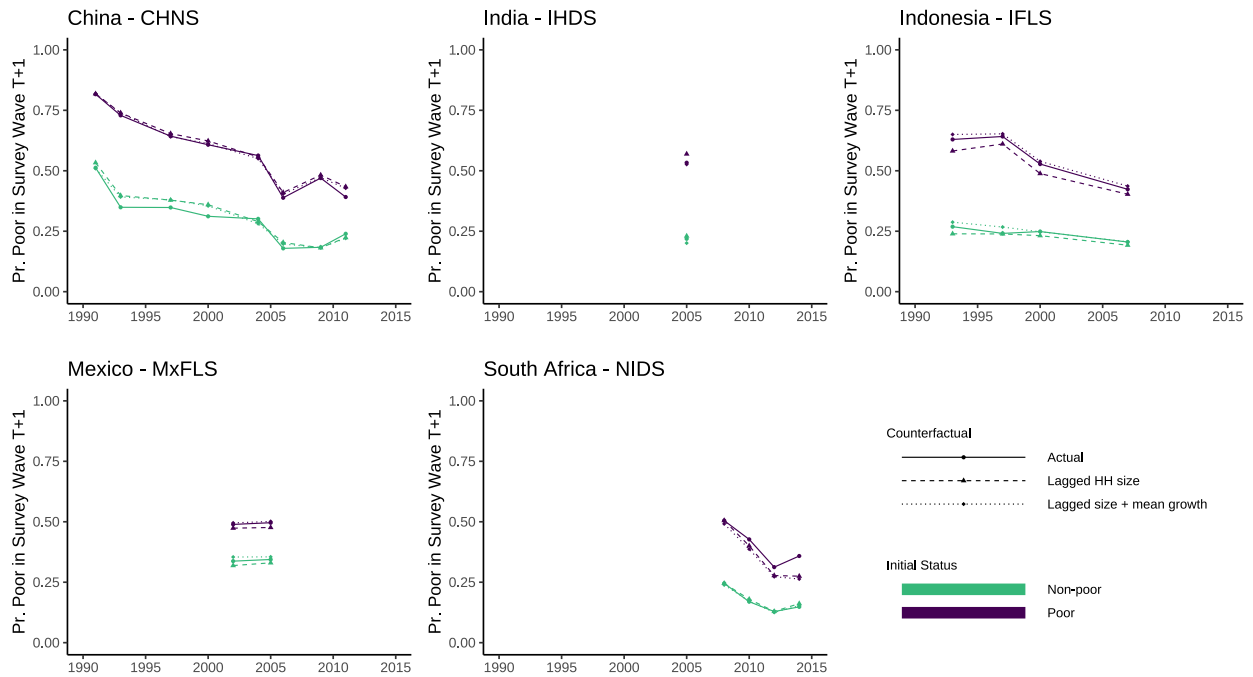
Figure A.12: Transition probabilities as non-parametric functions of initial income



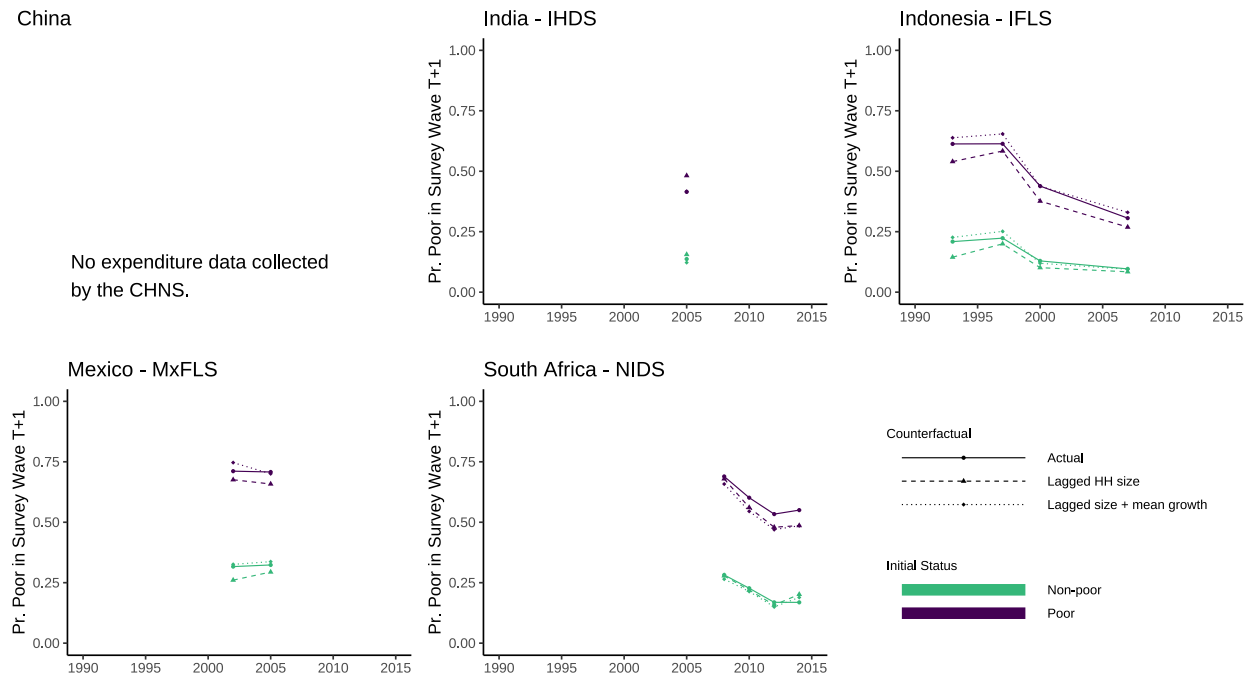
This figure reports the round-by-round probabilities of being poor in the subsequent round of each panel survey, as a smooth function of income in a given round. Results for each survey wave are plotted as a distinct line. Nonparametric fits were estimated using local linear regression applied to the 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: Transition probabilities with counterfactual household sizes

(a) Income

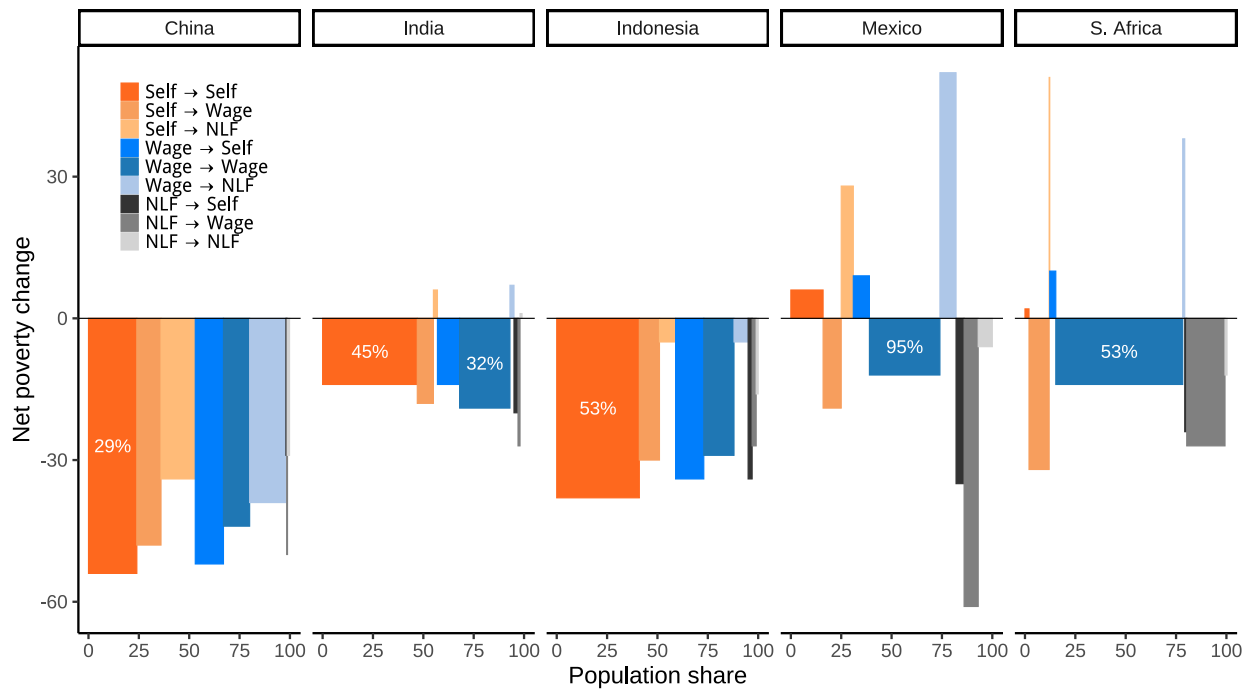


(b) Consumption



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.14: Occupational choice and poverty status, splitting non-workers from non-wage workers



Note: Effects 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: Missingness in panel data sources

Country	Waves	Household-year observations	Unique households present		
			At baseline	At baseline and endline	In all rounds
China	10	44,340	3,795	1,925	1,416
India (National)	2	78,330	41,554	34,639	34,639
India (Rural Only)	3	31,432	10,792	9,848	9,848
Indonesia	5	57,553	7,224	6,019	5,704
Mexico	3	26,265	8,440	7,182	6,824
S. Africa	5	35,948	7,296	5,672	4,894

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.3: Changes in household demographics and changes in poverty status

	Income definition					Expenditure definition				
	China (1)	India (2)	Indonesia (3)	Mexico (4)	S. Africa (5)	China (6)	India (7)	Indonesia (8)	Mexico (9)	S. Africa (10)
Males under 15	-0.192** (0.091)	-0.031 (0.029)	0.031 (0.044)	0.035 (0.068)	0.149 (0.096)		-0.057** (0.027)	-0.031 (0.042)	-0.073 (0.061)	0.033 (0.087)
Males 15-65	-0.269*** (0.085)	-0.401*** (0.024)	-0.293*** (0.043)	-0.320*** (0.059)	-0.101 (0.081)		-0.316*** (0.023)	-0.161*** (0.039)	-0.244*** (0.054)	-0.069 (0.073)
Males over 65	-0.175* (0.101)	-0.177*** (0.036)	-0.127** (0.059)	-0.104 (0.076)	-0.302*** (0.094)		-0.195*** (0.032)	0.070 (0.054)	-0.236*** (0.068)	-0.095 (0.085)
Females 15-65	-0.276*** (0.090)	-0.295*** (0.022)	-0.103** (0.042)	-0.135*** (0.052)	-0.123** (0.059)		-0.285*** (0.021)	-0.052 (0.038)	-0.267*** (0.047)	-0.020 (0.052)
Females over 65	-0.177* (0.104)	-0.142*** (0.033)	0.026 (0.055)	0.030 (0.072)	-0.193** (0.077)		-0.202*** (0.030)	0.011 (0.048)	-0.264*** (0.065)	0.031 (0.068)
Females under 15										
Household Size	0.009 (0.007)	0.014*** (0.001)	0.010*** (0.003)	-0.031*** (0.005)	0.006 (0.008)		0.036*** (0.001)	0.034*** (0.003)	0.00004 (0.004)	0.057*** (0.008)
Constant	-0.461*** (0.021)	-0.149*** (0.003)	-0.336*** (0.009)	-0.021** (0.009)	-0.129*** (0.012)		-0.154*** (0.003)	-0.400*** (0.008)	0.077*** (0.008)	-0.016 (0.011)
Y mean	-0.49	-0.17	-0.36	-0.03	-0.17		-0.18	-0.43	0.06	-0.06
$\vec{\beta} \cdot \vec{X}$	-0.01	-0.02	-0.01	-0.01	-0.04		-0.02	0.00	-0.02	-0.01
Observations	1,888	34,638	5,961	6,761	5,669		34,638	5,961	6,761	5,669
Adjusted R ²	0.010	0.033	0.025	0.017	0.005		0.056	0.034	0.007	0.012

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.4: Structural change: Ravallion & Huppi (1991) comparison

Country	Category	HPF Share	RH Share
China	Ag → Ag	39%	73%
	Ag → Not	38%	53%
	Not → Ag	3%	-38%
	Not → Not	21%	21%
	RH interaction term		-10%
India	Ag → Ag	43%	67%
	Ag → Not	25%	35%
	Not → Ag	4%	-21%
	Not → Not	27%	25%
	RH interaction term		-6%
Indonesia	Ag → Ag	43%	53%
	Ag → Not	12%	-11%
	Not → Ag	13%	7%
	Not → Not	31%	50%
	RH interaction term		1%
Mexico	Ag → Ag	4%	9%
	Ag → Not	19%	33%
	Not → Ag	-2%	-21%
	Not → Not	80%	78%
	RH interaction term		2%
S. Africa	Ag → Ag	2%	15%
	Ag → Not	8%	9%
	Not → Ag	9%	-9%
	Not → Not	81%	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 v.s. non-agricultural.

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 throughout our analysis the definitions of a household, household head, etc. that were employed by the underlying surveys. For reference, these are as follows for the panel sources:

CHNS The CHNS documentation does not explicitly define a concept of household.²⁵ 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 discrepancies between rosters across waves.

IHDS A household is defined as “all those who live under the same roof and share the same kitchen for 6+ months.”

IFLS A household is defined as “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.” The head of the household is defined as “a person among the group of householders who is responsible for satisfying daily necessities of the household or a person who is regarded/assigned as the head of the household.” A householder is defined as “anyone who usually lives in the household, whether she/he is at home during the survey or is temporarily absent. A householder who has been away for 6 or more months, and a householder who has been away for less than 6 months but plans to move out/be away for 6 or more months is not regarded as a householder. A guest who has stayed in the household for 6 or more months or a guest who has stayed in the household for less than 6 months but plans to stay for 6 or more months is regarded as a householder.”

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.” It further states that household members include

- Any person who usually lives in the household, regardless his presence or temporary absence. For example someone on vacation or who has left the household temporarily (for less than one year) for labor reasons is considered a household member.
- A person who has lived in the household for one year or more or who has lived in the household for less than one year but is planning to stay in the household for a year or more is considered a household member.

²⁵See <https://www.cpc.unc.edu/projects/china/about/design/sample>.

- The guests who fulfill the criteria mentioned above and who sleep in the household, share the meals prepared in the household and are free to use the kitchen.
- Domestic servants or any other household workers who fulfill the criteria mentioned above are considered household members.

but do not include

- A person who has not lived in the household for one year or more, or who has lived in the household for less than a year but is planning to stay away for a year or more (since the day of departure) is NOT considered a household member

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.”

These definitions are important for interpreting our measures of poverty since those measures are based on income and consumption per capita, which we calculate by dividing total household income or consumption measures by household size. We include and treat equally all members listed on the household roster when doing so. While the definitions that determine household size are fairly similar across sources, 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.

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 (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 (Series PA.NUS.PRVT.PP). 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. We do not adjust for sub-national variation in price levels as information on regional differences in prices is not consistently available across the countries and years in our samples.

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, 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 that rent is again the majority of this expenditure.
- In the NIDS we start our calculation from partially aggregated subcategories of consumption prepared by the research team. 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 use

the reported market prices to value goods consumed, and impute the local median of market prices where the household did not report one.

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, interest and capital gains received.
- **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. Conceptually these flows typically reflect a mix of labor and capital income, but it is generally not possible to distinguish them convincingly. 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.

We value self-produced goods as described in Section B.3 above, so that the number which enters 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. Note also that 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 / gift; we classify this as self-produced as 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. As is standard in the literature (e.g. Ravallion and Chen, 2007) we do not attempt

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	0.22 (0.21)	0.27 (0.23)	0.3 (0.24)	0.33 (0.25)	0.36 (0.26)

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.

to price publicly provided services such as education for which appropriate reference prices are not available.

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. We attribute capital income and transfer income to the non-agricultural sector. Capital income questions about money earned from dividends and interest are present in both the IHDS and NIDS questionnaires.
- 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 where the survey clearly identifies an enterprise as belonging to or being run by a single person in the household; in cases where this is not possible 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%).²⁶ 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 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.

²⁶<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.

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	1,925	119	1,806	3,612	2,731	(76%)	881	(24%)	55%
India	34,639	255	34,384	68,768	57,529	(84%)	11,239	(16%)	49%
Indonesia	6,019	160	5,859	11,718	9,843	(84%)	1,875	(16%)	43%
Mexico	7,182	1,356	5,826	11,652	8,415	(72%)	3,237	(28%)	22%
S. Africa	5,672	225	5,447	10,894	10,449	(96%)	445	(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.7.1. 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”).

B.7 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 household observed in both the first and last round of each panel.

B.7.1 Classification by sector

We first classify as many households as possible using sectoral classifier codes. The IHDS and NIDS provide a single 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 72% and 96% 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.7.2 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.

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.7.3 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.²⁷ 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 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%, 36%, and 34% 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%, 11%, 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

²⁷In 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 NREGA 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.

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 11% 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.14 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.²⁸

Table B.5 tabulates the resulting final classification of households. Overall, averaging across survey rounds, we classify 43%, 47%, and 11% of households as primarily self-employed, primarily in wage employment, and out of the labor force, respectively.

²⁸Note 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)	16.3%	20.3%	21.6%	30.3%	16.6%
	Inactive, Not Labor Force Age (not 18–70)	20.6%	33.3%	37.1%	38.3%	30.4%
	Unable to classify	0.2%	0.0%	0.0%	0.7%	3.8%
Self-employment	Animal work		8.6%			
	Boss/employer				1.7%	
	Business		4%			2.7%
	Family worker in hh business- no remuneration				1.7%	
	Family worker			5.9%		1.9%
	Farm work		19.9%			
	Paid family worker	2.2%				
	Peasant on plot					
	Self-employed, independent operator with no employees	26%			2.1%	
	Self-employed, owner-manager with employees	2.3%		9%	5.1%	
Self-employed, owner-manager with temporary paid help			0.6%			
Unpaid family worker	0.2%		8.8%			
Worker without remuneration on non-hh business				0.3%		
Wage employment	Agriculture laborer		3.7%			
	Contractor with other people or enterprise	8.2%				
	Government worker			3.1%		
	Non-agricultural laborer		4.4%		16.9%	2.7%
	Nrega work		2.1%			
	Privateworker			13.7%		
	Rural laborer				2.9%	
	Temporary worker	13.1%				
	Unemployed					19.9%
	Wage/salary job		3.7%			22.1%
Works for another person or enterprise as permanent employee	11%					

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%	25%	15%	6%	28%
	Wage employment	30%	22%	18%	29%	54%	31%
	Not in labor force	22%	33%	57%	55%	41%	42%
	Unable to classify	0%					0%
Not labor force age (not 18–70)	Self-employment	7%	11%	5%	2%	1%	5%
	Wage employment	11%	2%	3%	3%	5%	5%
	Not in labor force	81%	87%	92%	95%	94%	90%
	Unable to classify	0%					0%
Age unknown	Self-employment		0%	8%	2%	0%	3%
	Wage employment		0%	6%	4%	0%	3%
	Unable to classify		100%	86%	94%	100%	95%

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%	13.9%	42.7%
Primary wage employment	44.7%	37.8%	37.4%	50.3%	63.0%	46.6%
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 v.s. 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	0.47	0.48	0.46	0.68	0.34	0.84	0.33	0.55
Poor (Consumption)			0.36	0.22	0.39	0.29	0.50	0.48	0.26	0.41
Female head	0.13	0.29	0.11	0.38	0.15	0.50	0.21	0.37	0.43	0.60
Head has primary education	0.50	0.62	0.22	0.31	0.27	0.25	0.43	0.34	0.54	0.35
Age of head	47.9	65.4	49.0	58.1	49.2	57.7	48.5	56.0	46.3	49.2
Household size	4.4	3.1	5.2	3.1	6.2	6.2	5.1	3.9	3.9	3.7

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 in labor force	China	India	Indonesia	Mexico	S. Africa
Yes	Yes	38.9% &	16.2% &	22.6% &	25.6% &	61.9% &
Yes	No	51.7% &	81.7% &	70.2% &	65.2% &	20.6% &
No	Yes	0.2% &	0.1% &	0.2% &	0.1% &	2.5% &
No	No	9.2% &	2% &	7% &	9% &	15% &
Final classification: female outside employment		39.1%	16.3%	22.8%	25.7%	64.4%

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. Households in the fifth row, “unable to classify” are those for which age and gender information is unavailable for all members of the households.

B.7.4 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 five types of households after incorporating data limitations, 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. These account for between 16% and 62% of households across countries and survey rounds. These households we classify 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 18% 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.