"Beyond GDP:" welfare within countries across time

Emilio Depetris-Chauvin* — Universidad Católica de Chile
Alejandro Molnar* — The World Bank
Forhad Shilpi* — The World Bank

August 25, 2026

Abstract

This paper investigates the temporal and spatial evolution of welfare within six developing countries — Malawi, South Africa, Bangladesh, Cambodia, Mexico, and Peru — by applying the Jones and Klenow (2016) expected lifetime utility framework to a consistently defined urban–rural gradient. Using microdata from household surveys and censuses, we construct a welfare metric across three spatial dimensions: subnational administrative units (such as regions or provinces), areas classified by varying degrees of urbanization — including those that have recently urbanized — and towns/cities and their hinterlands. Our results reveal pronounced within-country disparities: in most cases, welfare levels in the poorest regions and rural areas are less than half those of the most prosperous regions and large urban agglomerations. Decomposition analyses show that differences in consumption and life expectancy are the primary drivers of these gaps, partially mitigated by lower inequality in rural and smaller urban areas. Despite substantial initial disparities, we document strong convergence patterns. Regions, urban areas, and hinterlands with lower initial welfare experience faster subsequent growth, driven in part by narrowing gaps in life expectancy. However, convergence remains uneven across countries and across urban typologies, with rural and newly urban settlements in some contexts continuing to fall behind.

JEL Classification: D63, I31, O18, O47, R11, R13

Keywords: Subnational welfare, rural-urban gradient, welfare convergence, consumption, leisure, life expectancy, inequality, expected lifetime utility, developing countries

*We are thankful to Catalina Banfi, Sebastián Gómez Roca, Ana Gorrini, Matías Harari, Pedro Pablo Skorin, Anika Tasnim and Nicolás Zacarías for excellent research assistance, and to Harold Coulombe and Mariel Siravegna for technical assistance with NSO data. We gratefully acknowledge financial assistance from the World Bank's Knowledge for Change Program on *Economic development through a spatial lens*. The views expressed here are those of the authors alone and should not be attributed to the World Bank and its affiliates. Email: aimolnar@worldbank.org

1 Introduction

In many poorer countries, where poverty is starkly visible, traveling from the capital city to remote villages can feel like stepping back in time — from the modern era to centuries past. The differences in per capita consumption – which is the most widely used indicator of poverty and well-being – are large between urban and rural areas. In Sub-Saharan Africa, for example, the incidence of poverty in rural areas is more than twice that in urban areas (World Bank, 2024a). However, income and consumption per capita provide only an incomplete measure of living standards (Stiglitz *et al.*, 2009; Jorgenson, 2018; Stiglitz *et al.*, 2018). Factors such as leisure, life expectancy, inequality, and other aspects of quality of life are not fully captured by these monetary indicators, yet they can substantially shape welfare within a country. Although existing literature documents cross-country differences in welfare and a smaller set of studies examines subnational variation using broader measures of well-being, evidence on welfare differences along the urban–rural gradient in developing countries remains limited. This paper addresses this gap by applying a theoretically grounded expected lifetime utility framework developed by Jones and Klenow (2016) (henceforth JK 2016) to estimate welfare differences across urban and rural areas in six countries: Malawi, South Africa, Bangladesh, Cambodia, Mexico, and Peru.

To illustrate the basic insight behind our welfare measure, suppose we want to compare living standards between Uitkomsdal and Somerset West in South Africa. The estimated welfare measure for Uitkomsdal would indicate the proportion of consumption in Somerset West, given its values for life expectancy, leisure and inequality in both consumption and leisure, that would deliver the same expected utility as the values in Uitkomsdal. A value of less than one for Uitkomsdal would mean that the average person there is less well-off than in Somerset West.

We examine the geographical and temporal evolution of consumption-equivalent welfare across above mentioned six countries with varying levels of economic development. To benchmark our analysis against available evidence on regional disparities within a country, we start by exploring welfare variations at the level of administrative regions or provinces. Our analysis then digs deeper into welfare differences across urban-rural gradients at three levels: rural and urban areas, newly (i.e., urbanized within the last 10 years) and established urban areas, and urban areas of different sizes and their hinterlands. We utilize a definition of urbanization based on globally consistent satellite data from the Global Human Settlement Layer (GHSL) and its Settlement Model (SMOD) component (Schiavina *et al.*, 2023b). While for most countries, the urban classification supplied by National Statistics Offices (NSOs) correlates strongly with the globally consistent definition, we find significant differences between the two for Malawi, South Africa, Cambodia, and Bangladesh. Statistics agencies in these countries have either not kept up with rapid urbanization or have changed their classifications over time due to administrative rather than functional reasons. Finally, utilizing the same settlement data for finer classification, we also compute welfare measures for about 300 urban areas (UAs) and their hinterlands; for more than 250 UAs, we were able to construct time-varying welfare measures. Several key findings emerge from our analysis across different levels of subnational analysis.

First, we find considerable heterogeneity in welfare across regions within a country. The difference between the richest and poorest regions is large. Welfare in the poorest region is less than half of that in wealthiest region in three out of six countries, with South Africa in 1993 displaying the highest welfare inequality.[1] Differences in (demographically-adjusted) consumption are the main component explaining welfare disparities across regions. The correlation between our welfare measure and output (GDP) per capita varies across countries ranging from as high as 0.86 in South Africa in 2023 to as low as 0.4 in Peru in 2007.[2] We find that output per capita tends to underestimate (overestimate) welfare in poorer (wealthier) regions, where inequality is lower (higher).

Second, welfare levels in rural areas are universally lower than in urban areas in all six countries – the urban-rural gaps are as large as regional gaps in many countries. Rural welfare as a percent of urban welfare ranges from 93 percent (i.e., Bangladesh in 2005) to 43 percent (i.e., Peru in 2007). Within urban areas, established areas fare much better than newly urbanized areas in most countries. We find an increasing welfare gradient with respect to sizes of urban areas: welfare in larger cities is higher than smaller ones. The decomposition analysis suggests that differences in consumption and life expectancy widen the welfare gaps between urban and rural areas. Lower inequalities in consumption and leisure in smaller urban areas and rural areas, on the other hand, narrow these gaps. Yet, in most cases, the first two factors predominate, resulting in higher welfare differences than consumption differences. For instance, the welfare difference between Uitkomsdal and Somerset West in 1993 was even greater than the difference due to consumption, primarily due to the additional gap in life expectancy.

Third, a strong pattern of within-country convergence of welfare over time is detected at the regional and urban areas levels. At the regional level, the distribution of

[1] 39 percent in South Africa (in 1993), 46 percent in Cambodia (2009) and 49 percent in Peru (2007). These gaps are smaller in the two poorest countries in our sample (Bangladesh and Malawi).

[2] Regional GDP data are available for Mexico, Peru and South Africa only.

welfare has become more compact over time. The convergence across regions stands in contrast to the stalling regional convergence in per capita output (GDP) in the most recent decade reported in Chatterjee *et al.* (2025).[3] Yet, the differences between the wealthiest and most lagging regions in some cases remain substantial: welfare in the poorest region (Tonle Sap) is only half of the wealthiest region (Phnom Penh) in Cambodia in 2019. Apart from consumption, a major source of regional convergence is the narrowing of regional gaps in life expectancy.

For rural-urban and recently urbanized vs. established urban areas, the pattern of convergence or divergence varies by country. While the relative welfare of rural areas and newly urbanized areas have marginally improved in some countries over time — mainly due to a narrowing gap in life expectancy — others in Cambodia (rural) and Malawi (recently urbanized) have experienced a relative decline.

At the UA level, welfare grew in all countries except Bangladesh. UAs with lower initial welfare experienced the highest subsequent growth rates suggesting a strong pattern of convergence: a one-log-point higher initial welfare position is associated with 4-5-percentage-point lower growth in welfare over our sample period. Further, this qualitative pattern of convergence holds for every sub-component of welfare. Consumption growth is the dominant margin driving welfare gains in most countries, while larger urban areas grow modestly faster, consistent with mild agglomeration effects. Cross-country heterogeneity is substantial: Malawi and Peru exhibit the strongest welfare gains, South Africa and Cambodia more modest improvements, and Bangladesh remains the clear outlier, with negative average welfare growth driven primarily by a sharp decline in demographically adjusted consumption, only partly offset by improvements in inequality and life expectancy. While UAs tend to exhibit higher levels of welfare than their hinterlands, this gap has been closing over time as well. Overall, the evidence suggests that convergence is broad-based across welfare dimensions, but differences across countries are shaped chiefly by variation in consumption growth and, to a lesser extent, mortality improvement. Nonetheless, for some countries (e.g., Malawi and Cambodia), the city–periphery welfare gap remains substantial.

One potential concern with our analysis is that the large within-country welfare gaps that we find could be due to not accounting for some aspect of quality of life. One such important missing aspect is the value of place-based amenities. However, for amenity differences to compensate for our measured welfare gaps would imply much worse

[3] Their sample of countries is tilted to high and upper middle income countries in North and South America, and Europe. There are only three African and nine Asian countries in their sample of 39 countries. Estimation sample periods are also different from our study.

values for amenities in large urban areas than their hinterlands. The available evidence is not supportive of this view. Though the gaps are narrowing over time, residents of large urban areas typically enjoy superior access to education, healthcare, clean water, and electricity, and other physical infrastructure (Charron and Wadhwa, 2025). Rural areas, on the other hand, may outperform urban areas in terms of air quality, natural beauty and noise levels. Yet, on balance, amenities are better in urban areas as reflected in higher housing prices in large UAs than in smaller UAs and rural areas. In terms of public goods, crime, and pollution, Gollin *et al.* (2021) find that the quality of amenities is at least as high in cities as in rural areas in 20 Sub-Saharan African countries. Though amenity differences are not directly accounted for in our welfare measures, they are unlikely to explain the welfare gaps we find across regions and urban gradients.

In our welfare analysis, individuals are observed in their current residences during the surveys. Welfare differences across areas can and do induce migration – rural to urban migration is the predominant type of migration in developing countries. Yet, observed migration rates across regions/areas in developing countries are about half of those in the United States (World Bank, 2024b). The estimated costs of migration in developing countries are also more than double those in the US (Bryan and Morten, 2019; Lagakos, 2020; Selod and Shilpi, 2021; World Bank, 2024b). The higher migration costs in developing countries are due to a lack of infrastructure and information, credit constraints, language barriers, and migration restrictions (Selod and Shilpi, 2021). And for a highly mobile country such as the US, Falcettoni and Nygaard (2023) find that incorporating migration in the analysis does not change their overall findings on the welfare differences and convergence patterns across US states. Higher migration costs can arguably sustain the persistence of the large welfare differences we find for six developing countries.

This paper makes contributions in three areas. Following the publication of JK (2016), several studies used their methods to examine the cross-country differences in living standards using mostly macro data (Bannister and Mourmouras, 2017; Gallardo-Albarrán, 2019). Brouillette *et al.* (2025) study racial differences in welfare in the US. This approach has been extended to investigate within-country regional differences in living standards in two recent studies. Falcettoni and Nygaard (2023) find that the dispersion in welfare levels across US states exceeded the corresponding dispersion in real per capita GDP with little evidence of welfare convergence over time. Using data from Nigeria, Okoye and Pandey (2024) find that welfare differences across provinces are four times greater than those suggested by consumption alone due to differences in life expectancy and inequality. To the best of our knowledge, this paper is the first one to apply the JK (2016) methodology to examine the welfare differences across urbanization gradients.

Evidence on the dispersions in welfare and its evolution over time is of great importance to policymakers in both developed and developing countries. But providing evidence on these disparities along the urban-rural gradient in developing countries has been challenging due to data considerations. While GDP data are available at the regional level in some countries (three out of six countries in our sample), they are simply unavailable by urban classification. Existing studies on urban-rural disparities use poverty or wealth indices – measured using per capita consumption expenditure in former case (e.g., Nakamura *et al.*, 2023) and a set of household level characteristics in the latter case (Yeh *et al.*, 2020). Our empirical extension of the JK (2016) approach in this paper provides a broader measure of welfare disparities using the standard utility accounting framework.

Several measures of poverty – such as Human Development Index (HDI; UNDP, 2013) and Poverty-Adjusted Life-Expectancy (PALE) index (Baland *et al.*, 2021) incorporate life expectancy but not inequality of consumption and leisure into welfare computation. The mash-up indices such as HDI and multi-dimensional deprivation index (Alkire *et al.*, 2015) incorporate wider indicators of well-being, but these measures are sensitive to choices and respective weights of those indicators which remain arbitrary (Ravallion, 2012; Döpke *et al.*, 2017; Decerf, 2024; Okoye and Pandey, 2024).[4] In contrast, the sub-components of welfare in JK (2016) are aggregated using the standard approach to welfare accounting in economics.

Our analysis shows that this broader approach is replicable in countries where micro data from household surveys and censuses are available – a point also emphasized by Okoye and Pandey (2024). As micro data has become widely available across developing countries in recent years, the methodological approach in JK (2016) provides a consistent alternative to measuring broader welfare in data-constrained settings. When combined with recent advances in small areas estimation using machine learning, this approach can be extended to estimation of welfare at finer geographical levels (Zheng *et al.*, 2026).

Policymakers often struggle with persistent regional and urban rural gaps. Our results show not only which areas within a country are falling behind, but also whether they are catching up to better-off areas. The overall picture is encouraging: we find convergence across regions and along the urban-rural gradient. Our results also suggest ways to speed up this process. Two major drivers of disparity are per capita consumption and life expectancy. Lowering barriers to migration by improving

[4]Recent studies on Europe, however, derived the weights from individual preferences obtained from attitude surveys (Boarini *et al.*, 2022), but such stated-preference data are not available for most developing countries.

infrastructure, access to information, and financial services can support moves from lagging (often rural) areas to leading areas (large UAs) and reduce welfare differences. By contrast, place-based policies work best when they raise welfare without distorting migration decisions (Kline and Moretti, 2014). To the extent that health services do not induce migration arbitrage (Schwartz and Sommers, 2014), our finding that life expectancy explains a meaningful share of spatial welfare gaps suggests that locally provided health care can be a powerful tool for welfare convergence.

The rest of the paper is organized as follows. Section 2 develops the welfare accounting framework. Section 3 describes the framework's calibration to national and sub-national data. Section 4 describes the results; we use breakouts to summarize key facts and remarks thought this section. Section 5 concludes.

2 Theory

In this section, we summarize the main methodology presented by Jones and Klenow (2016) and incorporate the necessary adjustments to guide our within-country analysis. Let us assume an individual who obtains flow utility u from consumption C and leisure ℓ. She will face a lottery determining her life in a specific location, unaware whether she will be rich or poor, work few or many hours, be healthy or not. Her expected lifetime utility will be given by

$$U = E \sum_{a=1}^{100} \beta^a u(C_a, \ell_a) S(a)$$

where β is the usual discount factor. Note that the expectation operator applies to her uncertainty regarding consumption and leisure. Furthermore, over her lifetime (i.e., 100 years), she will draw from cross-sectional distributions (across age groups a) of consumption, leisure, and survival rates. Let $S(a)$ represent the probability of an individual surviving to age a. We define a consumption-equivalent measure λ as follows:

$$U_i(\lambda) = E_i \sum_{a=1}^{100} \beta^a u(\lambda C_{ai}, \ell_{ai}) S_i(a)$$

Note that λ indicates the factor by which we must adjust the individual's consumption to make her indifferent between living her life as a random person in our reference location r and living in another location i. Thus, it satisfies:

$$U_r(\lambda_i) = U_i(1)$$

To take the theory to the microdata, several assumptions are needed to implement the procedures required to compute the desired welfare measure λ. Recall that, for a particular year, the representative individual will draw from cross-sectional distributions of consumption, leisure, and survival rates. While it seems plausible that this individual can interpret observed survival rates and leisure for older age groups as good indications of her levels of survival and leisure years ahead, it is also plausible that consumption will grow over time along a balanced growth path. Therefore, we assume that leisure and survival rates remain constant whereas consumption grows at a constant rate g over time. We assume that the levels of consumption and leisure we observe in the sample are representative of the possible values an individual could observe for a particular age group a in a given year. Survival rates by age groups in different locations are estimated by using census data and different methodological strategies. Since we have individual-level data on consumption, leisure, location, and age; we have a triplet $\{j, a, i\}$ representing an individual j of age a in location i. Behind the veil of ignorance, expected utility for an individual in location i is

$$U_i = \sum_{a=1}^{100} \beta^a S_a^i \sum_{j=1}^{N_a^i} \bar{\omega}_{ja}^i u(c_{ja}^i e^{ga}, \ell_{ja}^i)$$

Note that, following the original work by Jones and Klenow (2016), flow utility for each individual in our data is going to be weighted within her age group (of size N_a^i) by the factor $\bar{\omega}_{ja}^i$ which is normalized to sum to 1 within each age group (i.e., $\bar{\omega}_{ja}^i \equiv \omega_{ja}^i / \sum_{j=1}^{N_a^i} \omega_{ja}^i$). Then, each weighted average of utility flows is adjusted by the survival rates specific to each location i for each age group a (i.e., S_a^i).

The next step requires specifying the flow utility function for the individual. For our main analysis, we follow Jones and Klenow (2016) and assume the following function which is additive on (log) consumption and leisure, and includes an intercept.

$$u(C, \ell) = \bar{u} + \log C + \nu(\ell) \tag{1}$$

Recall that the consumption-equivalent welfare metric λ_i is implicitly defined by $U_r(\lambda_i) = U_i(1)$. Given our assumption regarding the specification of the flow utility, we can write

$$U_r(\lambda_i) = \sum_{a=1}^{100} \beta^a S_a^r[\bar{u} + \log \lambda_i C_{ar} + \nu(\ell_{ar})] = \sum_{a=1}^{100} \beta^a S_a^r[u_a^r + \log \lambda_i] \tag{2}$$

where,

$$u_a^r \equiv \bar{u} + ga + \sum_{j=1}^{N_a^r} \bar{\omega}_{ja}^r \left[\log(c_{ja}^r) + \nu(\ell_{ja}^r)\right] \tag{3}$$

Note that the term ga in equation 3 is the result of assuming exponential growth in consumption (at rate g) in a logarithmic function to which a normalized weight summing to 1 within each age group is applied. Using equation 2 we solve for $log(\lambda_i)$ that equates expected utility in the region of reference r and location i, namely $U_r(\lambda_i C_{ar}, \ell_{ar}) = U_i(C_{ai}, \ell_{ai})$

$$\sum_{a=1}^{100} \beta^a S_a^r [u_a^r + \log \lambda_i] = \sum_{a=1}^{100} \beta^a S_a^i u_a^i$$

Given the assumption of separability in the flow utility function, the result is straightforward:

$$\log(\lambda_i) = \left[\frac{1}{\sum_{a=1}^{100} \beta^a S_a^r}\right] \sum_{a=1}^{100} \beta^a \left[(S_a^i - S_a^r)u_a^i + S_a^r(u_a^i - u_a^r)\right]$$

The denominator in the first bracket has a simple interpretation in the case of a fully patient individual (i.e., $\beta = 1$): it represents life expectancy at birth in our region of reference r. Therefore, the higher the life expectancy at birth in the region of reference (i.e., where the individual lives), the lower the consumption-equivalent measure of the standard of living. Further,

$$\log(\lambda_i) = \sum_{a=1}^{100} \left[\frac{\beta^a (S_a^i - S_a^r)}{\sum_{a=1}^{100} \beta^a S_a^r} u_a^i + \frac{\beta^a S_a^r}{\sum_{a=1}^{100} \beta^a S_a^r}(u_a^i - u_a^r)\right]$$

To ease notation, let lower case survival rates (in levels and differences) be normalized by the sum of survival rates in our region of reference $s_a^r \equiv \frac{\beta^a S_a^r}{\sum_{a=1}^{100} \beta^a S_a^r}$ and $\Delta s_a^i \equiv \frac{\beta^a (S_a^i - S_a^r)}{\sum_{a=1}^{100} \beta^a S_a^r}$. Therefore, we can then write the log of the consumption-equivalent welfare as

$$\log(\lambda_i) = \sum_{a=1}^{100} \Delta s_a^i u_a^i + \sum_{a=1}^{100} s_a^r (u_a^i - u_a^r) \tag{4}$$

Now consider a particular specification for the flow utility. By replacing equation 1 as well as its equivalent for a region i in the second term of the equation 4 we get

$$\log(\lambda_i) = \sum_{a=1}^{100} \Delta s_a^i u_a^i$$

$$+ \sum_{a=1}^{100} s_a^r \sum_{j=1}^{N_a^i} \bar{\omega}_{ja}^i \log(c_{ja}^i e^{ga})$$

$$+ \sum_{a=1}^{100} s_a^r \sum_{j=1}^{N_a^i} \bar{\omega}_{ja}^i v(\ell_{ja}^i)$$

$$- \sum_{a=1}^{100} s_a^r \sum_{j=1}^{N_a^r} \bar{\omega}_{ja}^r \log(c_{ja}^r e^{ga})$$

$$- \sum_{a=1}^{100} s_a^r \sum_{j=1}^{N_a^r} \bar{\omega}_{ja}^r v(\ell_{ja}^r)$$

We define demographically-adjusted average utility from consumption and utility from leisure

$$E \log c_i = \sum_{a=1}^{100} s_a^r \sum_{j=1}^{N_a^i} \bar{\omega}_{ja}^i \log(c_{ja}^i e^{ga})$$

$$Ev(\ell_i) = \sum_{a=1}^{100} s_a^r \sum_{j=1}^{N_a^i} \bar{\omega}_{ja}^i v(\ell_{ja}^i)$$

The welfare metric can be written as [5]

$$\log(\lambda_i) = \sum_{a=1}^{100} \Delta s_a^i u_a^i + E \log c_i + Ev(\ell_i) - E \log c_r - Ev(\ell_r) \tag{5}$$

We now define demographically-adjusted average consumption and leisure as $\bar{c}_i \equiv \sum_{a=1}^{100} s_a^r \sum_{j=1}^{N_a^i} \bar{\omega}_{ja}^i c_{ja}^i e^{ga}$ and $\bar{\ell}_i \equiv \sum_{a=1}^{100} s_a^r \sum_{j=1}^{N_a^i} \bar{\omega}_{ja}^i \ell_{ja}^i$. Then, if we add and subtract $\bar{c}_i$, $\bar{\ell}_i$, $\bar{c}_r$, and $\bar{\ell}_r$ in equation 5 we get [6]

[5]Note that $\sum_{a=1}^{100} s_a^r = \sum_{j=1}^{N_a^i} \bar{\omega}_{ja}^i = \sum_{j=1}^{N_a^r} \bar{\omega}_{ja}^r = 1$. Therefore, the growth rate of consumption g is going to be irrelevant in the last part of equation 5 because g will cancel out when adding the two terms $E \log c_i$ and $-E \log c_r$. It may, however, affect the consumption-equivalent measure $log(\lambda)$ through the flow utility in the first term, $\sum_a \Delta s_a^i u_a^i$, of equation 5.

[6]We slightly depart from the original analysis and computation of the main welfare metric (i.e., the consumption-equivalent measure λ) in Jones and Klenow (2016) for two main reasons. First, since we are interested in within-country differences in welfare, instead of computing a welfare measure with respect to the United States, we focus on a measure relative to a reference region r within each country (e.g., the

$$
\begin{aligned}
\log(\lambda_i) = \sum_a \Delta s_a^i u_a^i \\
+ \log \bar{c}_i - \log \bar{c}_r \\
+ \nu(\bar{\ell}_i) - \nu(\bar{\ell}_r) \\
+ E \log c_i - \log \bar{c}_i - (E \log c_r - \log \bar{c}_r) \\
+ E\nu(\ell_i) - \nu(\bar{\ell}_i) - (E\nu(\ell_r) - \nu(\bar{\ell}_r))
\end{aligned}
$$

The decomposition of the welfare measure can be summarized as

$$
\log(\lambda_i) = \Delta \text{LE} + \Delta \text{C} + \Delta \ell + \Delta \text{C Inequality} + \Delta \text{l Inequality} \tag{6}
$$

We next describe each component of the decomposition in equation 6. The first component, $\Delta \text{LE} = \sum_a \Delta s_a^i u_a^i$ with $\Delta s_a^i \equiv \frac{\beta^a (S_a^i - S_a^r)}{\sum_a \beta^a S_a^r}$, captures the effect of differences in survival rates. Note that this term will be negative as long as the reference region presents, in average, lower age-specific mortality rates (or higher life expectancy at birth under the assumption of individuals with no impatience whatsoever). The incidence of these differences in mortality is then weighted by the flow utility in location i -or how much a year of life is worth in terms of utility. The interpretation of the second term, $\Delta \text{C} = \log \bar{c}_i - \log \bar{c}_r$, is much simpler since it refers to the relative difference in (demographically-adjusted) average consumption between locations i and r. Since our reference location tends to be the wealthiest within a country, we expect this component to be negative in most cases. The fourth component, $\Delta \text{l} = \nu(\bar{\ell}_i) - \nu(\bar{\ell}_r)$, denotes the contribution of differences in leisure to the welfare measure. Finally, the last two components, ΔC Inequality and Δl Inequality, represent the contribution of consumption inequality and leisure inequality, respectively.

To conclude, and before moving on to the details of the microdata, a set of additional assumptions about parameters and relevant functional forms for utility is necessary. Specifically, we first need to calibrate the intercept in the flow utility $\bar{u}$, the growth rate

capital of the country, the wealthiest region, or urban areas). Therefore, the simple statistic we compute will indicate what proportion of consumption in the reference region would provide the same expected utility as the value in our region of analysis. In most cases we select the region with the highest income as the reference; adjustment in such cases will require a reduction in consumption. Second, to avoid relying on detailed income data within our country of analysis, we do not compute the welfare as welfare relative to income as in Jones and Klenow (2016). In their main analysis, Jones and Klenow (2016) focus on an equivalent variation relative to income. That is, Jones and Klenow (2016) focus on $\log \frac{\lambda_i}{\tilde{y}_i}$; where $\tilde{y}_i$ is the ratio between country i income and US income. This will simply imply that the component of consumption in the welfare decomposition (i.e., ΔC below) cannot be interpreted as the share of consumption in GDP.

of consumption g, and the discount factor β. In their original work Jones and Klenow (2016) assume a common annual growth rate of consumption of 2 percent. Once one assumes a growth rate of consumption, as well as we assume an interest rate and a mortality discount, an impatience rate (and thus a discount) can be computed. For instance, with an interest rate of 4 percent, growth of consumption of 2 percent, and mortality discount of 1 percent one can come with a 1 percent impatience rate (thus $\beta = 0.99$). To calibrate $\bar{u}$ we need to set a value of life for the country of study. Jones and Klenow (2016) assume a value of life of \$6 million, which leads to $\bar{u} = 5$. However, it has been argued that the implied value of life would be substantially lower in poor countries. We discuss below the details of the calibration for $\bar{u}$. Finally, we need to specify the baseline utility function from leisure. Following Jones and Klenow (2016), we assume a utility function with an implied constant Frisch elasticity, that is:

$$v(\ell) = -\frac{\theta\epsilon}{1+\epsilon}(1-\ell)^{\frac{1+\epsilon}{\epsilon}}$$

Note that this functional form requires information on two parameters: (1) the Frisch elasticity (i.e., ϵ, which determines how flexible the labor supply is in response to wage fluctuations) and (2) the utility weight on leisure (i.e., θ). While Jones and Klenow (2016) set in their main analysis a Frisch elasticity of 1 which implies that the disutility from working increases with the square of the number of hours worked, there is no clear consensus in the literature on its value. We will thus investigate the sensitivity of our results to this choice (i.e., in our main analysis we will set $\epsilon = 1$). Once we have a value of ϵ, we compute the value of θ using the first-order condition for the labor-leisure decision. Given the functional form assumed for the utility function in our main analysis, the first-order condition will deliver $\theta = w(1-\tau)(1-\ell)^{-\frac{1}{\epsilon}}/c$. Therefore, some additional data are required, specifically the marginal tax rate, the ratio of wage to consumption, and average leisure, for which we follow the parametrization of Jones and Klenow (2016) and arrive to a value of 14.2 for θ.

3 Microdata and country-specific calibration

This section describes our main sources of microdata and discusses both the operational strategies used to compute our measures and the country-specific calibrations.

3.1 Consumption and leisure

Our measure of consumption captures household expenditures on non-durable goods and services, including both monetary purchases and in-kind items. It comprises spending on food (including self-produced food), housing (including rent and imputed rent for homeowners), utilities, personal care, clothing, health, education, transportation, and other everyday goods and services. Expenditures on durable goods—such as furniture, appliances, tools, and vehicles—are excluded to avoid issues related to asset accumulation and depreciation. The resulting consumption aggregate reflects actual material well-being over a given period, rather than investment in long-lived assets. In addition to household consumption, we incorporate government expenditure data — obtained from the Penn World Table — to account for publicly provided goods and services that contribute to welfare.

We underscore two caveats with consumption measurement. First, household consumption expenditures are subject to spatial variation in the cost of living. Urban areas – and particularly large established cities – systematically exhibit higher costs of living, primarily driven by housing and non-tradable services. Consequently, nominal consumption disparities may overstate real utility differences over space (Ferré *et al.*, 2012). JK (2016) applied their welfare framework at the national level, where spatial variation in the cost of living may have overstated the role of consumption inequality (ΔC Inequality). Since our approach aggregates consumption at the subnational level, any bias from spatial price disparities is likelier to show up in the consumption total (ΔC), rather than the inequality component. Second, relying on household survey microdata to compute the consumption inequality component may under-represent the extremes of the distribution. Household surveys may fail to capture higher income households, or their income may be top-coded, although the effect of selection and top-coding may be greater on income than expenditure statistics (Hurst *et al.*, 2014).

Using individual level survey data, in Figure 1, we display life-cycle consumption for a selected group of regions and years, i.e., the evolution of consumption by age group a (of size N_a^i)) in region i. For each region-country-year, we present four different consumption series.[7]

[7] All consumption series are expressed in 2017 PPP-adjusted US dollars.

First, the gray line shows the evolution of average consumption by age group:[8]

$$c_a^i = \sum_{j=1}^{N_a^i} \bar{\omega}_{ja}^i c_{ja}^i$$

The blue line shows the evolution of demographically-adjusted average consumption when adjusted by the survival rates specific to each location i for each age group a:

$$\bar{c}_a^i = s_a^i \sum_{j=1}^{N_a^i} \bar{\omega}_{ja}^i c_{ja}^i e^{ga}$$

with $s_a^i \equiv \frac{\beta^a S_a^i}{\sum_{a=1}^{100} \beta^a S_a^i}$. This means that the survival adjustment also incorporates impatience and the potential growth g in consumption. Recall we assume that the representative individual will draw from cross-sectional distributions of consumption and survival rates and that these statistics in the sample are representative of the possible values an individual could observe for a particular age group a.

The red line shows the evolution of the demographically-adjusted average consumption in region i when using the age-specific survival probabilities of our reference region r. This is the consumption measure we use in the computation of our welfare measure. Finally, to compare relative positions within the consumption distribution across regions, the gray line depicts the national demographically-adjusted average consumption, i.e., the age-specific average consumption computed using all individuals and adjusted with the national survival curve.

Consistent with previous literature for developed countries (see, for instance, Attanasio and Browning for the United Kingdom) and some developing countries (see, for example, Bairoliya *et al.* for India), consumption over the life cycle appears to follow an inverted-U shape during the usual working years (i.e., from age 20 to 60). [9]

We construct individual-level leisure from reported labor supply in each country's main household survey. Because the reference period for labor supply differs across surveys, we first harmonize all measures into annual hours worked. Depending on the questionnaire, annual hours are obtained by annualizing reported weekly hours or by

[8]Recall that each individual j's consumption in our data is going to be weighted within her age group by the factor $\bar{\omega}_{ja}^i$ which is normalized to sum to 1 within each age group. That is,

$$\bar{\omega}_{ja}^i \equiv \frac{\omega_{ja}^i}{\sum_{j=1}^{N_a^i} \omega_{ja}^i}$$

[9]To make the comparison with previous evidence in the literature more accurate, one should compute consumption relative to the head of the household of a given age (e.g., 25-year-old men).

Figure 1: **Average Consumption by Age and Selected Region-Year**

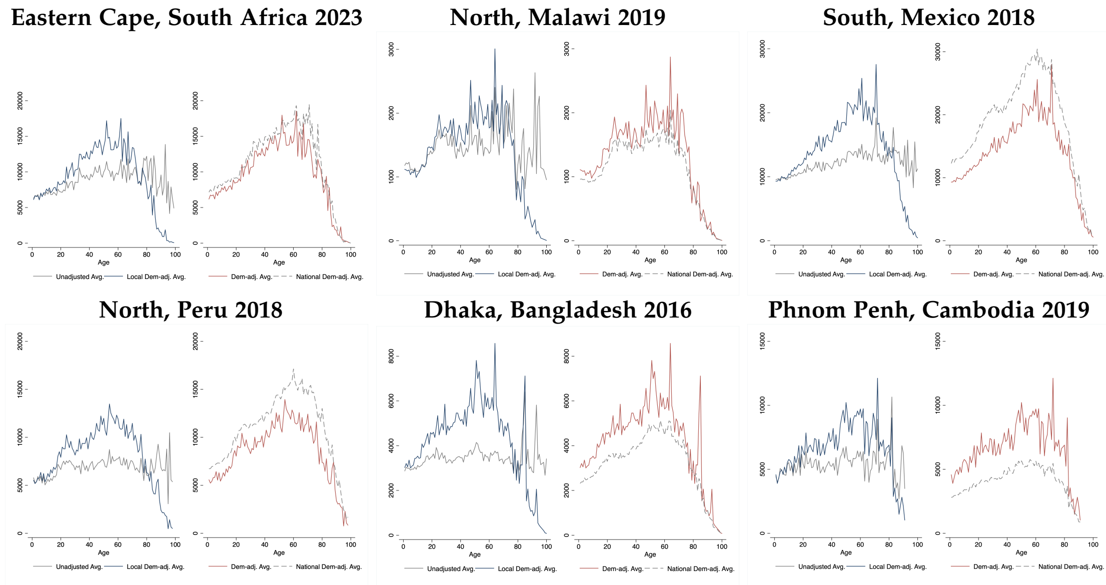

combining reported information on the intensity and duration of work, such as hours, days, weeks, or months worked. Total hours include primary and secondary jobs and, where applicable, other economic activities. Missing work hours are treated as zero.

We assume individuals have 16 discretionary hours per day after accounting for sleep, resulting in 5,840 hours per year. For each individual, leisure is calculated as the share of annual discretionary time not allocated to work. Specifically, we subtract annual hours worked from the annual time endowment and express the remainder as a proportion of that endowment. For individuals below the minimum working-age threshold[10], we assume zero hours worked and assign them a leisure share of one.

An important caveat to the JK (2016) framework for valuing leisure is that non-work hours may reflect unemployment, underemployment or seasonal agricultural lulls. While the first-order effect of unemployment should be captured by lower household earnings and therefore consumption, non-work hours may still not reflect utility-enhancing leisure (e.g., such time may include job-search). Macroeconomic evidence underscores this distinction: adults in low-income countries generally work substantially more hours than those in high-income countries and, within poorer countries, hours worked tend to decrease with individual wages (Bick *et al.*, 2018). To the extent that the JK (2016) approach conflates the time and psychological costs related to unemployment with the utility of leisure it may overstate the positive welfare

[10]We consider individuals aged 15 and older in Peru; 16 and older in South Africa; 6 and older in Bangladesh; 5 and older in Cambodia and Malawi; and 12 and older in Mexico.

contribution of the leisure component in regions with severe unemployment.

For further details on the datasets used and country-specific methodological choices, refer to the Appendix.

3.2 Survival curves

Key inputs for our analysis are survival curves at low levels of aggregation. Since such data is usually unavailable at subnational levels, we estimate survival curves using subnational life expectancy or under-five mortality estimates. To estimate survival curves for all ages between 0 and 100 based on a single statistic, we apply MATCH application from MortPak – the United Nations software package for mortality measurement. MortPak allows the estimation of life tables based on predetermined table patterns and a single input statistic (life expectancy or under-five mortality, in our case). MortPak provides a survival curve in intervals of five years. To expand it for all ages between 0 and 100, we follow Jones and Klenow (2016) and interpolate those results for missing years.

For subnational aggregation, we rely on subnational life expectancy estimates from the Global Data Lab.[11] When a subnational level is too granular, we construct a higher-level life table by computing a population-weighted average of the constituent survival curves. When a subnational level is too coarse, we assume that the survival curves of its subdivisions are equal to the curves of the larger aggregated unit and we propagate the resulting uncertainty to the higher level. Urban–rural differences in survival are derived from under-five mortality estimates based on DHS data, which are derived from NSO urbanization classifications; where necessary, these estimates are aligned with the life-table framework and any required adjustments are documented. When year-specific estimates are unavailable, we interpolate to the target year. For Mexico, where DHS-based urban–rural estimates are unavailable, urban and rural survival curves are constructed from population-weighted averages of state-level survival-rate estimates. To ensure national consistency, the final urban–rural survival curves are adjusted so that their implicit life expectancies match the national life expectancy. Our reliance on a single urban or rural survival curve precludes life expectancy from contributing to variation in welfare between or within urban areas (see Bilal *et al.*, 2019, on variation in life expectancies within six Latin American cities). Note that the implicit life expectancy can be approximated by summing the probabilities of surviving from age 0 to 100. Sometimes, this implicit estimate is higher or lower than the national life expectancy. To

[11]See blogs.worldbank.org/en/opendata/global-data-lab-resource-subnational-development-indicators-household-surveys

correct for this, we compute the ratio between the national and implicit life expectancies and multiply the entire curve by this factor. Because the implicit life expectancy is a linear functional of the age-specific survival probabilities, this scaling yields an exact match to the national life expectancy.

For each country, we select the life table pattern that best matches the national survival curves reported by the WHO in 2013. We then apply MortPak using this selected table and national-level life expectancy figures. Table 1 presents the chosen table for each country. As an illustration, Figure 2 shows the national survival curves provided by WHO data, vis-à-vis the estimates obtained using MortPak, for each country in a selected year. The simple impression of the curves indicates a strong agreement for Mexico, Bangladesh, and, to some extent, Peru and Malawi. There are some discrepancies for people aged over 60 in Cambodia, and a somewhat weaker alignment for South Africa.

Table 1: **Survival Tables Employed to Estimate Survival Curves**

Country	Survival Table
South Africa	New Coale-Demeny North
Malawi	New Coale-Demeny North
Peru	New United Nations Latin American
Mexico	New United Nations Latin American
Bangladesh	New United Nations South Asian
Cambodia	New Coale-Demeny North

Notes: MATCH application from MortPak is applied to subnational life expectancy data (or to under-five mortality estimates for urban–rural comparisons) to estimate subnational survival curves.

3.3 Value of statistical life

To compute the Value of a Statistical Life (VSL) - required to calibrate $\bar{u}$, we follow the methodology proposed by Wijnen *et al.* (2025), which estimates VSL as a function of per capita gross national income (GNI). Specifically, for low and middle income countries, the VSL is given by the expression:

$$VSL = 0.404 \cdot \left(\frac{Y}{5{,}726} \right)^{1.2} \tag{7}$$

where VSL is expressed in millions of USD (2020 prices) and Y denotes the country's GNI per capita, also in USD at 2020 prices.

Figure 2: **Comparison National Survival Curves: WHO vs Mortpack**

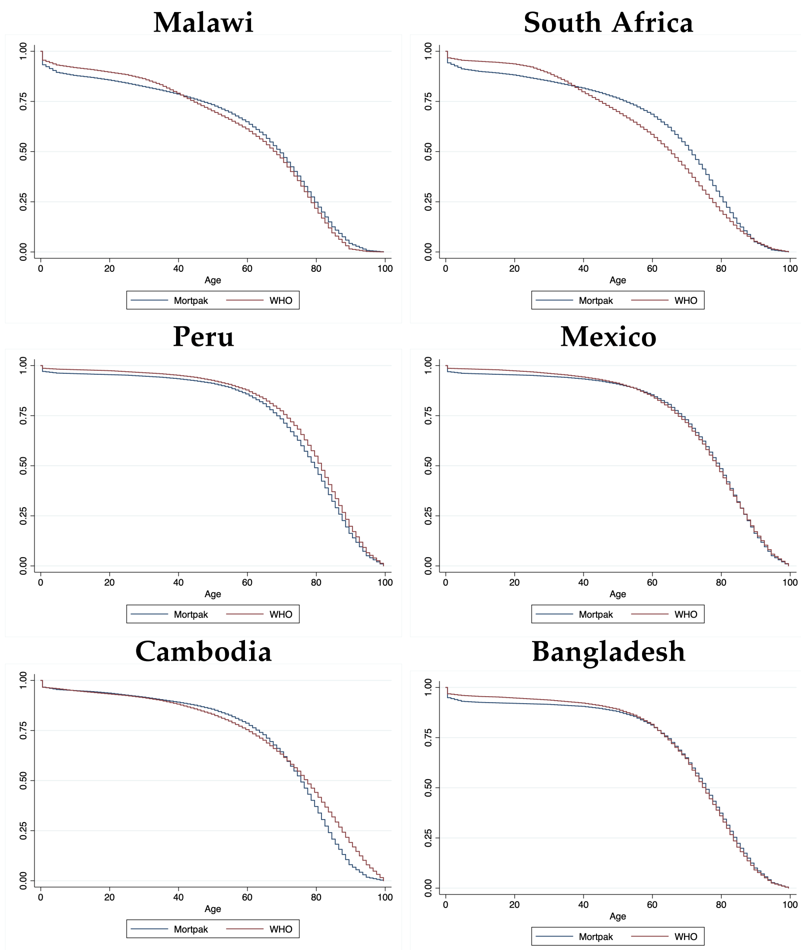

We construct this measure for each country-year in our dataset using GNI per capita figures based on the Atlas method, as reported by the World Bank. Since this GNI data is expressed in current USD, we deflate it to constant 2020 USD using the GDP deflator series, also provided by the World Bank. Finally, to express the VSL in constant 2007 USD, we apply the same deflator to convert 2020 values to 2007 price levels. Table 2 presents the value of statistical life for each country-year in our sample.

3.4 Intercept of the utility function ($\bar{u}$)

An important component of the utility function is the constant term $\bar{u}$, which we calibrate separately for each country-year in our dataset. To determine its value, we equate the present value of the lifetime utility flow—computed under the assumption $\bar{u} = 0$ —for an individual aged 40, with the utility-equivalent value of the VSL.

In practical terms, we proceed as follows. For each country-year, we use the corresponding microdata to compute key empirical moments aggregated by age. Based on these, we calculate the discounted sum of the utility flow from age 40 to 100, assuming a utility function with $\bar{u} = 0$. This provides the present value of lifetime utility, measured in units of "utils," for a 40-year-old individual.

Next, we compute the VSL in the same utility units using the same functional form. The difference between the utility-equivalent of the VSL and the present value of the utility flow defines $\bar{u}$: that is, the value that must be added uniformly to each period's utility so that the discounted lifetime sum matches the VSL in utility terms. Table 2 presents the value of $\bar{u}$ for each country-year in our sample.

Table 2: **Value of a Statistical Life**

Country	Year	VSL	$\bar{u}$
Bangladesh	2005	0.03	4.47
Bangladesh	2016	0.07	3.68
Cambodia	2009	0.04	3.77
Cambodia	2019	0.10	4.30
Malawi	2004	0.02	4.93
Malawi	2019	0.02	4.37
Mexico	2002	0.68	3.28
Mexico	2018	0.63	2.34
Peru	2007	0.21	3.71
Peru	2018	0.40	3.39
South Africa	1993	0.32	3.42
South Africa	2023	0.32	2.53

Notes: VSLs are expressed in million USD (2007).

3.5 Urban - rural classification

We assign an urbanization label to each household that is distinct from the urban-rural status assigned by national statistical offices. Our approach is based on the Degree of Urbanization methodology implemented by the Settlement Model (SMOD) of the Global Human Settlement Layer (GHSL; see Schiavina *et al.*, 2023b). The SMOD uses remote-sensed built-up density and census population data to categorize the planet's land surface into settlement typologies every five years and at a 1 square km resolution. The methodology provides a globally consistent classification for urbanization that is recommended by the UN Statistical Commission, given that official urban/rural definitions vary widely across countries (and may vary across survey rounds within the same country; see Ellis and Roberts (2016) on issues measuring urbanization in South Asia). Our approach also assigns a time-varying urbanization status to fixed locations, allowing us to track changes in status or construct derived metrics such as a rural household's distance to the nearest urban area.

Each SMOD raster cell is assigned a code that reflects a particular type of settlement morphology. We group the codes for urban centers (30), dense and semi-dense urban clusters (23 and 22) and suburban and peri-urban areas (21) as "urban." Conversely, codes 10 through 13 correspond to rural conditions — ranging from low-density settlements to predominantly uninhabited or sparsely populated zones.

The spatial link between survey data and SMOD rasters is implemented using geospatial units from country- and year-specific surveys, typically at the enumeration area level matched to a census round. The details on spatial units and the assignment process employed for each country, year and dataset are provided in the online appendix, but assignment tends to follow two major cases, depending on the nature of the household's geospatial identifier. For point data (available in Peru, Malawi 2019 and for a subset of spatial units in Mexico and Cambodia) the assignment is a simple point-in-polygon test. Each point maps onto a single SMOD raster cell, and receives a unique SMOD code. For polygons, we implement a population-weighted overlay method. First, we intersect the shape with the SMOD raster for urbanization, obtaining a partition of each shape by urbanization class.[12] We then intersect these partitions with a second global data product from GHSL, the 100 meter resolution population layer (POP, see Schiavina *et al.*, 2023a), which we accumulate to the share of the population within each geospatial unit living under each SMOD urbanization class. The shares are applied to each household ID defining an urbanization share weight that applies to each household and urbanization class. When aggregating across urbanization categories, as described below, we multiply standard survey weights by these urbanization share weights to obtain statistics of the population by urbanization status.[13] This aggregation rests on two assumptions. First, that the population in each dataset we use (e.g., a consumer or labor survey) was surveyed within each enumeration area unit in proportion to our urbanization share weights. This condition would be satisfied if households are surveyed in proportion to measured population density within each spatial unit. We cannot test this assumption, which would require knowledge on how each survey

[12] NSOs assign a single urbanization category to each geospatial unit in their census frame, but these do not necessarily coincide, in whole or in part, with GHSL SMOD classifications. We discuss further in Section 3.5.1 below.

[13] As an example, all 10 households in an enumeration area with a 90% urban population share will be weighted 9/10 as urban households and 1/10 as rural households, and further re-weighted by survey weights.

was fielded at the sub-enumeration area level (a form of NSO-specific institutional knowledge that is usually not disclosed, and at which NSOs do not usually keep spatial data). However, we know that the impact of this assumption is small because most geospatial units are assigned to a fully urban or rural status (as shown in Figure 4, and discussed in Section 3.5.1 below). Second, our approach assumes that the survey-specific household weights that each NSO designs to ensure representativeness across urban and rural enumeration areas maintain representativeness across our urbanization classes. This assumption is not testable, but we provide a sensitivity analysis of our main results across alternative methods for classifying urbanization. GHSL provides urbanization and population data every 5 years. We apply linear interpolation at the unit level to assign urbanization status to surveys fielded on off years.

3.5.1 Comparison with urban-rural classifications from NSOs

In this section we compare our classification of urban-rural status with that from national statistical offices (NSOs). For each survey we calculate the share of households classified as rural or urban within each geographical unit by the NSO and by our method, described above. In Figure 3 we aggregate both types of urban-rural population shares at the admin 1 level for the most recent survey in each country.[14] The degree of correspondence varies substantially across countries, depending on two main factors: the granularity of the aggregation unit and the frequency with which the measure provided by the NSO is updated. The correspondence is strongest for Mexican states in 2018, and to some extent, for Peruvian provinces in the same year. For South African provinces in 2023, the correspondence is strong in highly urbanized provinces but weaker in the rest of the country. Cambodian provinces show a moderate degree of agreement (reflected in the clustering of units along the 45-degree line) for about half of the sample. Districts in both Malawi (2019) and Bangladesh (2016) show a weak correspondence between the two approaches: the NSOs in both countries classify the majority of the population in almost all level 1 units as highly rural, whereas we find higher (and in the case of Malawi, more nuanced) urbanization weights. In Malawi there appears to be a more continuous distribution of the urbanization rate across districts. Our classification finds uniformly higher urbanization rates for Bangladesh, consistent with "hidden" urbanization in the country's official urbanization rates, as documented in Ellis and Roberts (2016).[15]

[14] Figure A.1 shows the same analysis for the oldest survey from each country. A comparison suggests that the correspondence between the two classifications has improved over time.

[15] Urbanization rates have been undercounted in Bangladesh since the 2011 census, when the Bangladesh Bureau of Statistics changed its classification for urbanization to a strictly administrative definition,

Figure 3: **Correspondence Between Urbanization Rates (Most Recent Survey): NSO vs Our Definition**

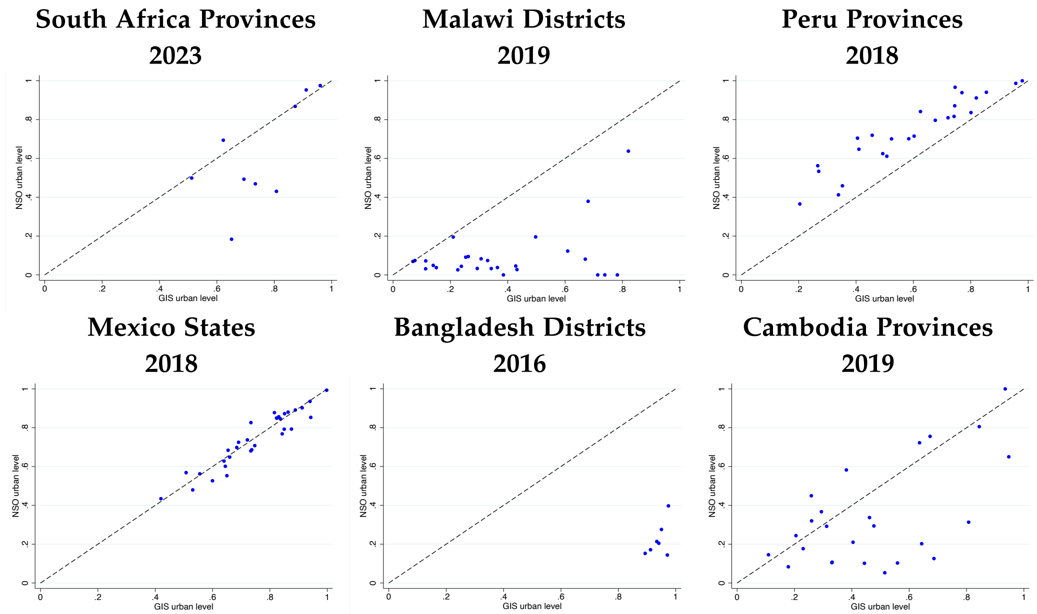

Instead of aggregating our data to administrative districts, in Figure 4 we overlap both definitions using the most granular geospatial unit available for each country in the most recent survey. As expected, the correspondence between the two approaches is much higher at a finer level of granularity (often small enumeration areas or geographical coordinates). Indeed in all the country-years we observed that between 78% to 96% of locations classified as urban by NSO are 100% urban under our SMOD classification. Nevertheless, in some countries there are still clear disparities in how each approach defines a location as urban. Take for instance the cases of Malawi and Cambodia, both in 2019: almost 30% of the enumeration areas that are classified as rural by their NSOs are urban under SMOD classification. The discrepancies are much more evident in the case of Bangladesh, where more than 70% of the smallest geospatial unit at which we can place survey respondents (i.e., Mouzas, a unit larger than enumeration areas) are classified as rural by NSO but are more than 90% urban under SMOD classification. Further, very few administrative units in Bangladesh are classified as rural under SMOD classification: this suggests that despite the issues with NSO classification the GHSL-based classification may also not be suitable in this country. GHSL urbanization classifications are a function of population density thresholds, and

removing the "urban" classification from all Upazila headquarters and growth centers that had not been legally incorporated as Pourashavas (municipal corporations).

the global definitions used in the GHSL may be too inclusive for Bangladesh, the country with the highest population density in the world (excluding micro-states).

Figure 4: **GIS-based urban share by NSO classification**

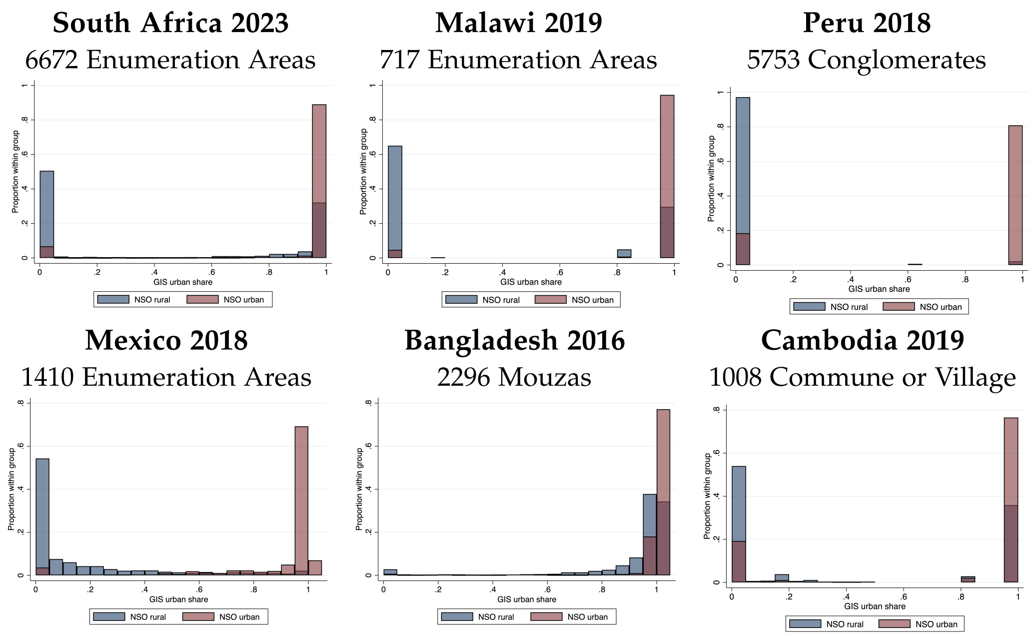

4 Within-country variation in welfare

In this section we describe our welfare results, first by regional administrative units, then by aggregates of a nation's urban vs. rural areas, then with a focus on areas in urban transition, and lastly reporting on statistics at the level of individually labeled urban areas.

Each sub-section contains many descriptive results that are specific to the individual countries in our study. We use "fact" and "remark" break-outs throughout to summarize our key findings, and to highlight relevant caveats.

4.1 Regional differences

4.1.1 South Africa

Table 3 presents the welfare measures and their decompositions for the 9 regions of South Africa (using Gauteng as reference) for 1993 (Panel A) and 2023 (Panel B).[16] The

[16]By the mid-1990s, Gauteng's regional GDP constituted approximately one-third of the economy of South Africa, followed by KwaZulu-Natal (less than 20 percent) and Western Cape (approximately 15

results presented in Table 3 reveal several important takeaways, but one fact stands out: South Africa is a dual society, with the Western Cape and Gauteng regions being significantly better off compared to the rest of the country (see also Figure A.3 for visual inspection of geographical variation in welfare). Indeed, on average, welfare in the two most advanced regions was double that of the rest of the country in 1993. Most of this gap is explained by inequality in consumption between regions. Interestingly, more equal consumption within regions helps to mitigate the negative effect on relative welfare due to differences in average consumption across regions (with the exception of Free State). To a lesser extent, higher levels of leisure in the less advanced regions also helps to reduce the gap in welfare.[17] Another contributor to the welfare gaps between regions is differences in life expectancy. For instance, in regions such as Eastern Cape and KwaZulu-Natal the loss from lower life expectancy (which is weighted by the value of flow utility so we can interpret it as the utility lost from living one year less) subtracts approximately 13 log points to welfare in 1993. Finally, leisure inequality, if anything, helps to reduce the welfare gap explained by consumption and life expectancy for most of the poorer regions.

Results in Panel B suggest that the gaps between regions have decreased after thirty years but remain significant. Furthermore, almost all regions seem to have improved relative to Gauteng, with the exception of the Western Cape, which was marginally better relative to Gauteng in 1993 and exhibited the same levels of welfare in 2023. As was the case in 1993, lower levels of (demographically-adjusted) average consumption and lower life expectancy continue to be the main factors contributing to lower welfare in the poorest regions of South Africa in 2023. Nonetheless, several regions have experienced a substantial improvement in life expectancy in the last thirty years.

In Table A.1 in the Appendix, we perform a set of additional robustness checks. While the conclusions regarding the relative position of each region in South Africa's welfare distribution and the existence of a large gap remain virtually unchanged, we document a series of additional results.

First, our baseline case considered a Frisch elasticity of 1 for calibration, which implied that the disutility from working increased with the square of the number of hours worked. Given that the rest of South Africa tends to have a higher level of leisure, it is clear from the component $Ev(\ell_i) - Ev(\ell_r)$ in equation 5 that a higher (or

percent). Northern Cape made the smallest contribution, accounting for about 2 percent of the economy of South Africa. Although Gauteng is the smallest region in terms of land area, it is the most populous province in South Africa, with more than a quarter of the national population.

[17] As noted in Section 3.1, this specific welfare component must be interpreted with caution. In regions such as Eastern Cape and KwaZulu-Natal this positive leisure component captures a reality of high, structural unemployment rather than a strictly voluntary substitution toward leisure (Bick *et al.*, 2018).

Table 3: **Welfare Across Regions in South Africa**

	Welfare λ	log λ	Decomposition: Life Exp.	Cons.	Leis.	Cons. Ineq.	Leis. Ineq.	LE
Gauteng	100.0	0.000	0.000	0.000	0.000	0.000	0.000	65.1 (67.3)
Panel A: 1993								
Western Cape	113.5	0.126	0.194	-0.099	0.001	0.022	0.008	70.9
KwaZulu-Natal	55.1	-0.596	-0.125	-0.658	0.054	0.096	0.037	60.4
Northern Cape	53.2	-0.630	-0.040	-1.047	0.076	0.306	0.075	63.4
Mpumalanga	51.6	-0.662	-0.074	-0.777	0.010	0.186	-0.008	62.0
North West	51.3	-0.668	-0.005	-0.814	0.011	0.157	-0.017	64.8
Free State	50.7	-0.679	-0.019	-0.557	0.014	-0.111	-0.006	64.3
Limpopo	47.5	-0.744	-0.022	-1.019	0.077	0.167	0.053	64.1
Eastern Cape	44.1	-0.819	-0.131	-1.057	0.093	0.161	0.115	59.5
Panel B: 2023								
Western Cape	100.2	0.002	0.013	0.137	-0.035	-0.090	-0.022	67.8
Northern Cape	79.2	-0.233	-0.017	-0.337	0.018	0.073	0.029	66.5
Free State	75.3	-0.283	-0.062	-0.344	0.016	0.089	0.017	64.4
KwaZulu-Natal	73.3	-0.311	0.013	-0.408	0.018	0.054	0.012	67.9
Eastern Cape	72.9	-0.316	-0.062	-0.441	0.030	0.121	0.035	64.3
Limpopo	70.9	-0.345	0.049	-0.584	0.014	0.174	0.002	69.8
Mpumalanga	67.5	-0.392	-0.082	-0.492	0.010	0.167	0.005	63.1
North West	66.0	-0.415	-0.066	-0.485	0.028	0.080	0.028	64.0

Notes: Please refer to Section 3 for details on the assumptions required in the calibration and computation of welfare. Last column presents life expectancy of each region in the survey year (for the region of reference, the first year is reported, and the last year is shown in parentheses).

lower) disutility from working (i.e., setting the Frisch elasticity ϵ equal to 2 and 0.5, respectively) increases (or decreases) welfare in the less developed regions. Particularly for the Eastern Cape, Northern Cape, KwaZulu-Natal, and Limpopo, moving from ϵ equal to 0.5 to 2 can result in differences in welfare ranging from 15 to 32 log points in 1993. The potential differences in welfare from departing from a Frisch elasticity of 1 can still be significant, but of a smaller magnitude in 2023. The largest increase in welfare, resulting from raising the Frisch elasticity to 2—specifically in the case of Eastern Cape in 1993—represents a gain of 21 percent in welfare.

Second, decreasing patience by reducing β from 0.99 to 0.96 does not substantially impact welfare gaps. Recall that one of the main channels through which impatience affects our welfare measures is through the life expectancy component, particularly via its term Δs_a^i in equation 5. Indeed, lower life expectancies in the less developed regions of South Africa contribute positively to the regional welfare gaps, as noted in the discussion of Table 3 above. By reducing β, the loss from lower life expectancy (which is weighted by the value of flow utility) becomes smaller, so relative welfare should improve for the lagging regions and worsen (relatively) for the more advanced regions. Nonetheless, we do not observe substantial changes in our welfare measures when assuming more impatient individuals.

Third, despite the role of life expectancy as a key contributor to the welfare gap, reducing (or increasing) the value of statistical life by 25 percent— in our baseline case, we use \$0.32 million in both years for baseline VSLs (in 2007 PPP-adjusted US dollars)— has little effect on improving (or decreasing) the relative welfare in the rest of the country compared to Gauteng.

Fourth, in our baseline analysis, household consumption is evenly distributed among household members. This assumption does not significantly impact the results. As shown in Table A.1, setting individual consumption equal to household consumption divided by the square root of the number of individuals in the household, or using the OECD-modified scale (which accounts for adult equivalency in consumption), does not appear to have a meaningful effect on our welfare measures for any region relative to Gauteng.

Fifth, increasing the threshold age at which we begin evaluating expected utility has a negligible effect on welfare measures. We considered evaluating expected utility only for individuals at or above three age thresholds: 2, 25, and 40 years. We do not find major differences with respect to our baseline analysis.

Finally, the availability of regional GDP data for South Africa allows us to visually assess whether welfare and income tell a similar story. Panel A of Figure 5 depicts the high correlation (coefficient of 0.86) between GDP per capita and consumption-equivalent welfare in both sample years (both measures in log scales). Since there are no clear departures from the 45-degree line in the same figure, GDP per capita appears to be a good indicator of welfare across regions in South Africa. Nevertheless, regions with average GDP per capita tend to locate below the 45-degree line, suggesting that welfare levels in these regions are marginally lower than GDP per capita would indicate. In Panel B of Figure 5, we plot the ratio of welfare to per capita GDP across the regions of South Africa to provide a closer look at the deviations of each region's welfare in relation to GDP per capita. The regions below the line have welfare measures that are typically 80 to 90 percent below their incomes whereas regions above the line have measures of up to 20 percent above their relative income. Note that North West, Northern Cape, Free State, and Mpumalanga are regions for which the consumption-equivalent welfare measure suggests lower welfare than GDP per capita would indicate.

Figure 5: **Welfare and Income Across Provinces in South Africa 2023**

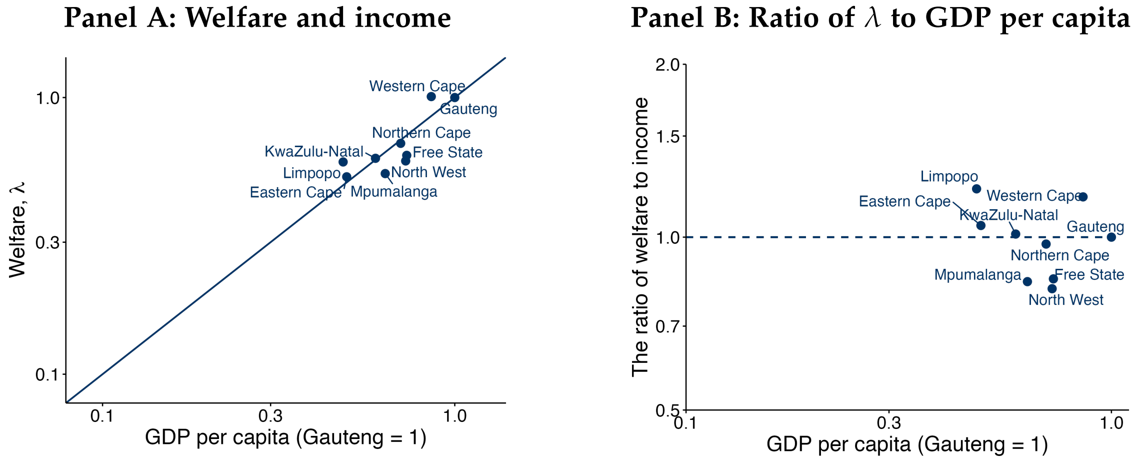

4.1.2 Malawi

We now focus on another African, yet poorer, country: Malawi. Table 4 presents the welfare measures and their decompositions for the 3 regions of Malawi (using Central Region as reference) for 2004 (Panel A) and 2019 (Panel B).[18] Several results are worth mentioning. First, the differences in welfare between regions in Malawi are not very substantial (although the inequality in welfare is somewhat more pronounced when we compare districts; see Figure A.4). Second, there is a complex dynamic in the evolution of welfare over time and across the regions of Malawi. In 15 years, the Central region became the most disadvantaged in terms of relative welfare, while the Northern region became the most advanced, and the Southern region achieved a percentage growth of over 30 percent in its well-being compared to the Central region (as can be deduced from the two panels in Table 4).

Regarding the differences in welfare observed in 2004, it is important to note that most disparities in well-being between the Central Region and the rest of Malawi were due to differences in consumption. These disparities were partially offset by better indicators of consumption inequality in the Northern and Southern regions. The Northern Region almost entirely compensated for these differences in consumption with a larger life expectancy's component and, to a lesser extent, more equal consumption. Indeed, differences in life expectancy are key. Life expectancy was 54.8 years in the North compared to 50.2 years in the Central Region. This difference alone contributes 12.6 log points to relative welfare in 2004 for the North. It is also important to note that

[18] The Central region has the largest population (marginally larger than that of the Northern region) and is the location of Malawi's capital, Lilongwe.

the relatively poor performance of the South, due to lower average consumption values, is further compounded by greater inequality in leisure.

The decomposition of welfare in 2019 shows that the Northern Region is marginally behind in terms of consumption inequality but performs much better (relative to the Central Region) in consumption, life expectancy, and both measures of leisure. The relative gain in the consumption components is substantial (i.e., 30 log points) and explains most of the relative improvement of the North. Surprisingly, the Southern Region has surpassed the Central Region in almost all welfare components except for life expectancy (and marginally in terms of leisure inequality). Importantly, (demographically-adjusted) average consumption contributes 9.2 log points to welfare in the South.

Table 4: **Welfare Across Regions in Malawi**

	Welfare λ	log λ	Decomposition					LE
			Life Exp.	Cons.	Leis.	Cons. Ineq.	Leis. Ineq.	
Central	100.0	0.000	0.000	0.000	0.000	0.000	0.000	50.2 (64.7)
Panel A: 2004								
North	100.5	0.005	0.126	-0.183	0.007	0.052	0.003	54.8
South	80.7	-0.215	0.000	-0.195	-0.011	0.019	-0.028	50.2
Panel B: 2019								
North	141.3	0.346	0.044	0.299	0.002	-0.002	0.003	66.8
South	108.7	0.083	-0.015	0.092	0.003	0.004	-0.002	63.9

Notes: Please refer to Section 3 for details on the assumptions required in the calibration and computation of welfare. Last column presents life expectancy of each region in the survey year (for the region of reference, the first year is reported, and the last year is shown in parentheses).

As mentioned above, differences in welfare appear to undergo significant shifts over time. Consider, for example, the comparison between the Central and Southern regions. In 2004, the consumption-equivalent welfare of the South was around 80 percent of that of the Central Region. This indicates that we would need to reduce an individual's consumption by approximately 20 percent to make our representative agent indifferent between living as a random person in our reference location (i.e., Central Malawi) and living in the South. Fifteen years later, in 2019, the welfare gap between these two regions has reversed, with the Southern Region now nearly 8 log points ahead of Central Region. Most of this shift is explained by changes in consumption patterns. Finally, it is important to note that, while differences in leisure do exist, they do not seem to significantly affect the differences in welfare between the regions of Malawi.

We next zoom in on differences in welfare within regions. The two panels of Figure A.4 convey the same message: the differences in welfare between districts are more pronounced than between regions. This should not be surprising given the disparities that exist between districts in Malawi regarding levels of urbanization and, in particular, mortality and consumption. However, it is important to highlight a significant caveat of the analysis when focusing on smaller units; survival rate calculations (which are particularly important for differences in the life expectancy component) tend to be noisier for districts with low population counts.[19]. Unfortunately, there are no district or regional GDP data available in Malawi to compare the relationship between GDP and welfare.

In Table A.2 in the Appendix, we present the robustness of our findings under alternative assumptions in the computation of welfare. Following Jones and Klenow (2016), we recompute our welfare measures by considering different specifications for the utility function, the growth rate of consumption, discounting, the distribution of consumption across household members, the starting age group for calculating expected utility, and two parameters in the calibration: the value of statistical life needed to determine the intercept in the utility function, and the value of the Frisch elasticity.

The main takeaway is that the results in Table 4 remain quite robust. The only exception occurs when we evaluate expected utility exclusively for individuals aged 25 and older, for those aged 40 and older, and when we assume no growth in consumption and no discounting for impatience. In these cases, welfare in the North remains below our reference region (i.e., the Central region), with drops in welfare (relative to the benchmark case) ranging between 3.8 and 6.9 log points.

4.1.3 Peru

We now study the evolution of welfare and its decompositions in Peru. Table 5 shows welfare across Peruvian regions with West Region (composed by Metropolitan Lima and the province of Ancash) as reference for 2007 (Panel A) and 2018 (Panel B). Regional differences in welfare in Peru are greater than those in Malawi but are of similar magnitude to those in South Africa. In 2007, welfare in the rest of Peru was, on average, 57 percent of that in the West Region — equivalent to roughly 60 log points lower welfare on average, or, expressed in consumption units, a 43 percent reduction relative to West. The components that explain these gaps in welfare are, in order of importance, consumption, life expectancy, and leisure. On the other hand, inequalities in consumption and leisure in 2007 help mitigate these gaps and contribute to welfare, ranging from 7 log points in the North Region to 16 log points in the East Region.

[19]It appears that our approach to compute survival rates tends to overestimate life expectancy in low population areas.

It is important to note that differences in average consumption are, by far, the main contributor to differences in welfare across regions.

While all the aforementioned main conclusions for 2007 hold for 2018 (Panel B), there is a reduction in the difference in welfare. Indeed, welfare in the rest of Peru was, on average, about 72 percent of that in the West Region in 2018. The improvements in the East, Northeast, and Central regions are noteworthy, as they reduced their welfare gap compared to the West Region by between 24 and 32 log points over 11 years.

The reduction in the welfare gap is mainly explained by a decrease in the differences in average consumption between regions (although the differences are still high) as well as a relative improvement in the life expectancy component (although it still contributes to the welfare gap since life expectancy remains highest in the West Region).

Table 5: **Welfare Across Regions in Peru**

			Decomposition					
	Welfare λ	$\log \lambda$	Life Exp.	Cons.	Leis.	Cons. Ineq.	Leis. Ineq.	LE
West	100.0	0.000	0.000	0.000	0.000	0.000	0.000	76.1
Panel A: 2007								
South	68.2	-0.383	-0.087	-0.369	-0.016	0.064	0.024	72.1
North	58.8	-0.531	-0.127	-0.466	-0.009	0.014	0.057	70.0
North East	55.3	-0.593	-0.054	-0.571	-0.047	0.043	0.036	73.3
East	52.5	-0.644	-0.036	-0.714	-0.067	0.080	0.092	74.1
Central	49.6	-0.702	-0.092	-0.705	-0.030	0.045	0.080	71.1
Panel B: 2018								
South	77.7	-0.252	-0.056	-0.189	-0.041	0.054	-0.020	75.2
North East	75.1	-0.286	-0.021	-0.308	-0.019	0.041	0.022	76.6
East	72.5	-0.322	-0.009	-0.374	-0.054	0.061	0.054	77.1
North	71.7	-0.332	-0.082	-0.291	-0.011	0.033	0.018	74.1
Central	63.0	-0.462	-0.079	-0.444	-0.044	0.059	0.046	74.0

Notes: Please refer to Section 3 for details on the assumptions required in the calibration and computation of welfare. Last column presents life expectancy of each region in the survey year (for the region of reference, the first year is reported, and the last year is shown in parentheses).

In Table A.3 in the Appendix, we present the robustness of our findings in Table 5 to alternative assumptions in the computation of welfare for Peru. The main takeaway from Table A.3 is that the documented results are robust to different assumptions. The within-country average deviation from our benchmark case ranges from 0.1 to 2.2 log points in 2007 and from 0.2 to 2.6 log points in 2018. There are, however, some assumptions that result in larger deviations from our benchmark computations. In particular, if we evaluate expected utility exclusively for individuals aged 40 and older, we obtain welfare measures that are smaller than our benchmark in between 1 and 11 log points in 2007 (7 logs points in average) and between 4 and 11 log points in 2018. Increasing the discount rate can also have a significant impact on our measure, resulting in higher welfare measures between 2 to 7 log points regardless of the year we look at. Finally, setting a higher disutility from working (i.e., Frisch elasticity ϵ equal to 2) decreases welfare between 1.4 and 7.3 log points relative to the benchmark.

Figure A.5 illustrates the geographical distribution of welfare across Peruvian departments (of which there are 25) for the years 2007 (Panel A) and 2018 (Panel B). The magnitudes of the welfare differences at the departmental level are consistent with those documented in Table 5 at the regional level, albeit more pronounced. For example, take Huancavelica, the department with the lowest level of consumption-equivalent welfare in 2007, which has a welfare level that is 33 percent of Lima's. Moreover, in the same year, 10 Peruvian departments had levels of welfare that represented less than half of Lima's. Another noteworthy point in Panel A is that the departments on the Peruvian coast exhibit the highest levels of well-being. Panel B shows that after 11 years, the distribution of our welfare measure exhibits lower variance. Some departments have experienced a significant increase in relative welfare between 2007 and 2018. This process of convergence in welfare is particularly notable for Moquegua and Huancavelica. In the case of Moquegua, its welfare gap with Lima is close to closing in 2018, with an improvement of 35 log points. For Huancavelica, there is an improvement in relative welfare of 53 log points, mainly due to increases in average consumption and improvements in consumption inequality. Additionally, the only department with a relative decline in its welfare measure during our analysis period is Arequipa.

Data for regional GDP is also available for Peru. Consequently, we calculate GDP per capita at the departmental level and correlate it with our measures of welfare for 2007 and 2018 (Panel A of Figure 6). We also plot the ratio of welfare to per capita GDP across Peruvian departments (Panel B of Figure 6). Two main patterns emerge from Figure 6. First, the correlation between welfare and income per capita is much weaker in Peru than in South Africa, with correlation coefficients of 0.4 in 2007 and 0.5 in 2018. Notably, while most departments are positioned away from the 45-degree line, the degree of deviation has decreased over time. Second, there is a negative correlation between the ratio of welfare to income and GDP per capita, implying that income per capita tends to underestimate (overestimate) welfare in less (more) developed departments.

4.1.4 Mexico

Now, we turn our attention to another Latin American country—Mexico—which has a larger economy than Peru. We focus on eight regions, based on the geographic

Figure 6: **Welfare and Income Across Departments in Peru (2007 and 2018)**

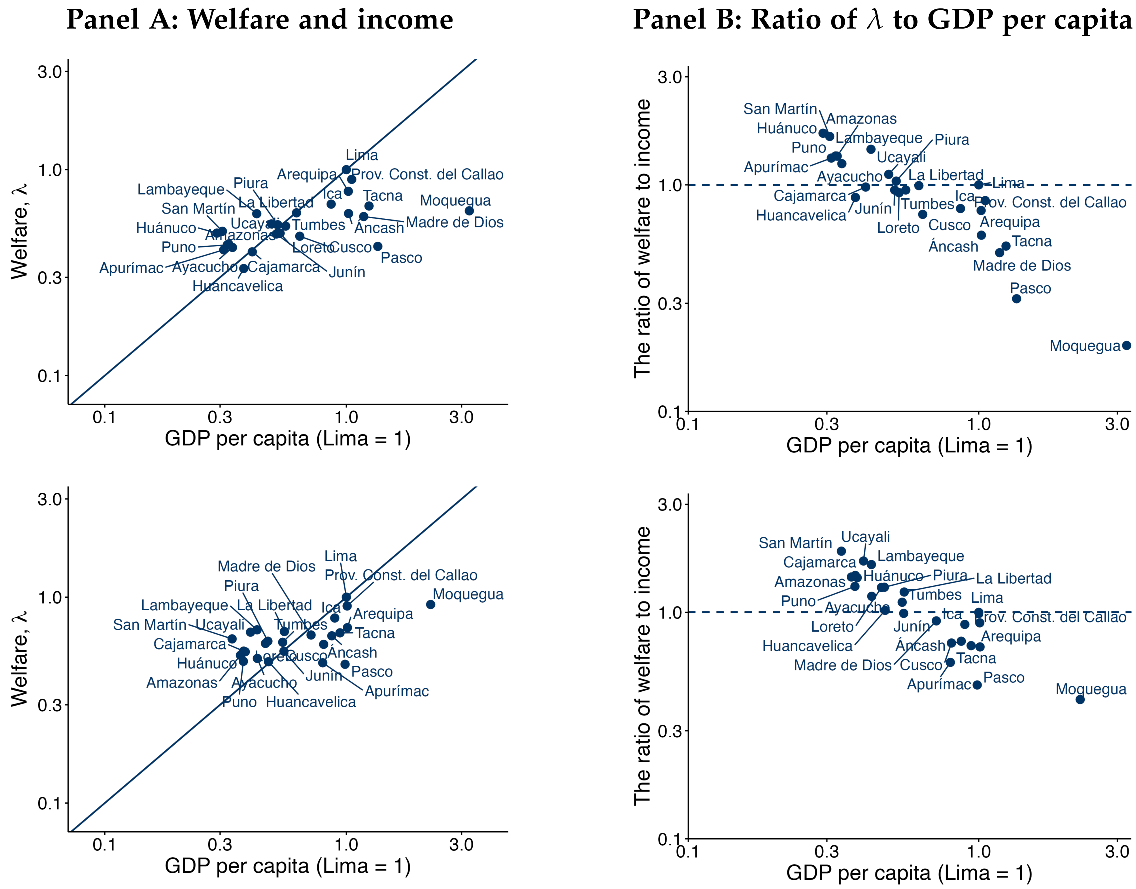

distribution proposed by Bassols Batalla (1991), using the East-Central region (home to Mexico City) as a reference.[20] Table 6 depicts our welfare measures and its decomposition across Mexican regions for the years 2002 (Panel A) and 2018 (Panel B). Several observations can be drawn from an analysis of this table. First, unlike Peru, the distribution of welfare across regions is much more continuous and has a lower variance in Mexico. Although in 2002 there was a region (i.e., the South) significantly lagging behind the rest of Mexico, its gap relative to the East-Central region decreased considerably. Second, almost every region improves its relative welfare with respect to our region of reference. In fact, by 2018, five regions were located only between 4 and

[20]Bassols Batalla (1991)'s taxonomy of Mexican regions is the following (listing states in parentheses): Northwest (Baja California, Baja California Sur, Nayarit, Sinaloa, and Sonora), North (Coahuila, Chihuahua, Durango, San Luis Potosí, and Zacatecas), Northeast (Nuevo León and Tamaulipas), West-Central (Aguascalientes, Colima, Guanajuato, Jalisco, and Michoacán), East-Central (Ciudad de México, Hidalgo, México, Puebla, Querétaro, Tlaxcala, and Morelos), South (Chiapas, Guerrero, and Oaxaca), East (Tabasco and Veracruz), and Yucatán Peninsula (Campeche, Quintana Roo, and Yucatán).

13 log points away from our reference region, with two of these regions (i.e., the two northern regions of Mexico) showing levels of welfare above East-Central.

The bulk of the convergence observed in most regions can be attributed to improvements in consumption. For instance, while all but one region exhibited lower levels of leisure in 2018, they all also showed improvements in consumption inequality, helping to reduce the welfare gap. On the other hand, regional differences in leisure inequality do not appear to significantly impact overall welfare levels. The Yucatán Peninsula stands out for having reduced its relative welfare gap by 24 log points over sixteen years, largely due to a substantial relative improvement in average consumption. Similarly, the Northwest region performs relatively better, primarily because of its higher life expectancy and, in particular, its higher levels of demographically-adjusted average consumption.

Table 6: **Welfare Across Regions in Mexico**

			Decomposition					
	Welfare λ	log λ	Life Exp.	Cons.	Leis.	Cons. Ineq.	Leis. Ineq.	LE
East-Central	100.0	0.000	0.000	0.000	0.000	0.000	0.000	75.2 (74.3)
Panel A: 2002								
Northwest	120.9	0.190	0.056	0.166	0.004	-0.029	-0.007	76.9
Northeast	100.2	0.002	0.055	-0.052	0.002	-0.008	0.004	76.9
West-Central	92.2	-0.081	-0.009	-0.018	-0.007	-0.035	-0.012	74.9
North	91.1	-0.093	-0.044	-0.073	0.020	-0.006	0.010	73.7
Yucatán Peninsula	76.3	-0.271	0.002	-0.158	-0.033	-0.057	-0.026	75.3
East	74.2	-0.299	-0.040	-0.140	0.000	-0.102	-0.017	73.8
South	59.8	-0.515	-0.041	-0.375	-0.037	-0.033	-0.028	73.7
Panel B: 2018								
Northwest	113.4	0.126	0.022	0.091	-0.007	0.012	0.008	75.1
Northeast	109.1	0.087	0.021	0.035	0.014	-0.004	0.021	75.1
West-Central	99.8	-0.002	0.005	0.006	-0.024	0.015	-0.005	74.5
Yucatán Peninsula	96.9	-0.031	0.012	0.001	-0.036	0.021	-0.030	74.8
North	95.9	-0.042	0.002	-0.050	-0.008	0.004	0.009	74.4
East	83.4	-0.181	0.005	-0.203	-0.011	0.028	0.000	74.5
South	72.6	-0.320	-0.013	-0.343	-0.014	0.039	0.011	73.7

Notes: Please refer to Section 3 for details on the assumptions required in the calibration and computation of welfare. Last column presents life expectancy of each region in the survey year (for the region of reference, the first year is reported, and the last year is shown in parentheses).

In Table A.4 in the Appendix, we present the robustness of our findings in Table 6 to alternative assumptions in the computation of welfare for Mexico. Our results are indeed highly robust to changes in the main assumption. No region exhibits (average) welfare differences exceeding 0.9 log points compared to the benchmark calculation. Furthermore, the maximum deviation from the benchmark is only 6 log points, observed for the Yucatán Peninsula in 2018 when setting the Frisch elasticity ϵ to 2.

Figure 7: **Welfare and Income Across States in Mexico 2018**

Panel A: Welfare and Income **Panel B: Ratio of λ to GDP per capita**

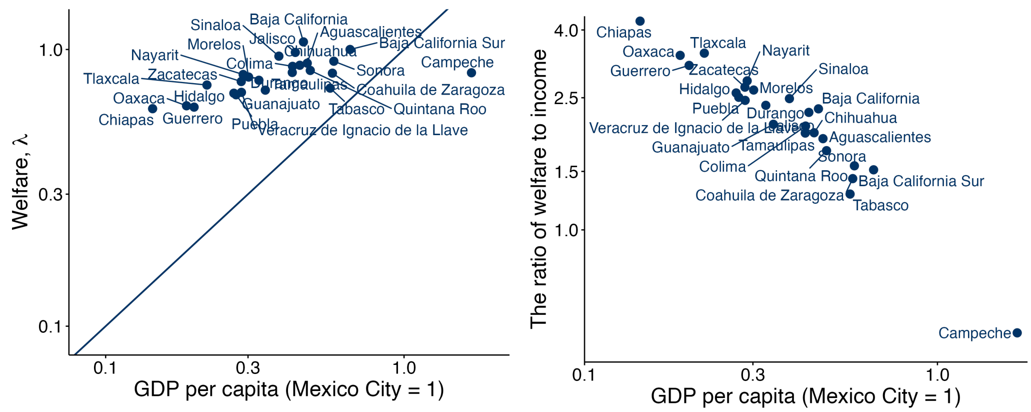

In Figure A.6 illustrates the geographical distribution of welfare across Mexican states (of which there are 32) for the years 2002 (Panel A) and 2018 (Panel B). From Panel A, it can be observed that the distribution of welfare among states in 2002 was more unequal than the distribution observed among regions for the same period. For instance, the consumption-equivalent welfare in Guerrero was roughly one-third of that observed in Mexico City. When analyzing the results at the state level for the year 2018, the conclusion is similar to the case of regions discussed based on Panel B of Table 6: there is a process of reduction in the variance of welfare across states. Although most of the variation in welfare is explained by consumption, differences in life expectancy across states also play a role, reducing welfare by up to 15 log points (e.g., Colima and Coahuila).

In Figure 7, we compare our measure of welfare with GDP per capita at the state level (Panel A). We also plot the ratio of welfare to per capita GDP as a function of relative GDP per capita (Panel B of Figure 7). Three main patterns emerge from this figure. First, the correlation between welfare and income per capita is similar to the one observed across Peruvian regions, with a correlation coefficient of 0.45. Second, although most states are positioned away from the 45-degree line, the rankings based on the two measures appear to be well preserved. Third, there is a strong negative correlation between the ratio of welfare to income and GDP per capita, suggesting that income per capita tends to underestimate (overestimate) welfare in less (more) developed Mexican states.

4.1.5 Bangladesh

Let us now turn to Asia. In this section, we examine the evolution of welfare and its decompositions in Bangladesh. Table 7 presents welfare across Bangladeshi regions, using Dhaka as the reference for 2005 (Panel A) and 2016 (Panel B). Bangladesh is an interesting case, and two key findings can be quickly observed from the two panels in the table. First, given Bangladesh's high levels of income inequality, our welfare measures suggest less inequality across regions, similar to what is observed in Mexico. Second, this equitable distribution has improved over time — that is, welfare among regions appears to be converging. Figure A.7 shows the geographic distribution of welfare across Bangladeshi regions.

Based on Panel A of Table 7, in 2005, the most lagging region of Bangladesh had a consumption-equivalent welfare level that was 59 percent of Dhaka's. As with the analysis for other countries, most of the welfare differences within Bangladesh are due to variations in consumption. By 2005, the consumption gap between regions and Dhaka ranged from 10 to 50 log points.

Eleven years later, in 2016 (Panel B of Table 7), we observe that most regions experienced a relative improvement in welfare compared to Dhaka, although Mymensingh and Khulna deteriorated. In fact, now Chattogram is the region with the highest relative welfare. Despite having a lower life expectancy and average consumption than Dhaka, Chattogram's higher leisure levels along with lower inequality in both leisure and consumption , make it the region with the highest overall welfare.

The most lagging region, Rangpur, experienced the largest relative improvement in welfare of approximately 23 log points. This improvement is primarily explained by a reduction of leisure inequality. At the same time, Mymensingh experienced a deterioration in relative welfare over time, reflecting a larger shortfall in the consumption component. Notably, the differences in the consumption component remains similar to the one in 2005 and can contribute as much as 48 log points to welfare underperformance relative to Dhaka. Additionally, the life expectancy component accounts for at most 7.6 log points of this gap (Sylhet). Both leisure levels and inequalities in consumption and leisure help offset the negative effects of lower consumption and shorter life expectancy. These components can contribute up to 8.1 (Sylhet), 5.3 (Mymensingh), and 13.5 (Sylhet) log points, respectively.

In Table A.5 in the Appendix, we present the robustness of our findings in Table 7 to alternative assumptions in the computation of welfare for Bangladeshi regions. Overall, the main results hold. For 2005, the within-country deviation from the benchmark tend to be on average very small, whereas for 2016, the average deviation ranges from -3 to 0

Table 7: **Welfare Across Regions in Bangladesh**

	Welfare λ	$\log \lambda$	Decomposition: Life Exp.	Cons.	Leis.	Cons. Ineq.	Leis. Ineq.	LE
Dhaka	100.0	0.000	0.000	0.000	0.000	0.000	0.000	69.6 (72.6)
Panel A: 2005								
Chattogram	94.7	-0.054	-0.060	-0.096	0.035	0.011	0.057	66.7
Khulna	90.9	-0.096	0.003	-0.226	0.042	0.011	0.075	69.7
Mymensingh	82.8	-0.189	-0.099	-0.209	0.031	-0.001	0.089	64.5
Barisal	81.6	-0.203	-0.017	-0.337	0.045	0.028	0.078	68.7
Sylhet	77.9	-0.249	-0.121	-0.125	-0.002	-0.001	0.000	63.3
Rajshahi	75.8	-0.277	-0.018	-0.402	0.020	0.056	0.067	68.6
Rangpur	59.4	-0.522	-0.011	-0.498	-0.040	0.052	-0.024	68.8
Panel B: 2016								
Chattogram	108.2	0.079	-0.040	-0.090	0.074	0.027	0.108	70.6
Sylhet	90.9	-0.096	-0.076	-0.276	0.081	0.040	0.135	68.7
Khulna	89.8	-0.107	0.002	-0.341	0.060	0.050	0.121	72.7
Barisal	88.0	-0.128	-0.016	-0.326	0.070	0.041	0.104	71.7
Rajshahi	87.3	-0.136	-0.025	-0.315	0.054	0.032	0.117	71.2
Mymensingh	75.5	-0.281	-0.041	-0.401	0.034	0.053	0.073	70.1
Rangpur	74.6	-0.293	-0.008	-0.477	0.049	0.024	0.119	72.1

Notes: Please refer to Section 3 for details on the assumptions required in the calibration and computation of welfare. Last column presents life expectancy of each region in the survey year (for the region of reference, the first year is reported, and the last year is shown in parentheses).

log points. Changes to our assumptions regarding the value of the Frisch elasticity generate the largest deviations from the benchmark. Specifically, setting the Frisch elasticity ϵ to 0.5 results in an average deviation of -9 log points, while setting ϵ to 2 leads to an average deviation of 10 log points, both in 2016.

4.1.6 Cambodia

We now turn our attention to another Asian country: Cambodia. Table 8 presents the welfare measures and their decompositions for the five regions of Cambodia (using Phnom Penh as the reference) for 2009 (Panel A) and 2019 (Panel B). Several results are worth highlighting. First, there are significant welfare disparities between Phnom Penh and the rest of the country. Specifically, in 2009, welfare differences ranged from 62 to 79 log points. Although there was a marginal improvement over time, these disparities remain quite large ten years later, in 2019. Second, while the rest of the country exhibits better relative measures in components such as consumption inequality, leisure, and leisure inequality —particularly in 2009 — the differences in life expectancy and, especially, in average consumption are of much larger magnitude. For instance, the consumption component explains between 73 and 91 log points of the welfare gap in 2009 and between 46 and 77 log points in 2019. Note that the difference in the life expectancy component remains high but has decreased by an average of 7 log points during our period of analysis.

Figure A.8 illustrates the geographical distribution of welfare across Cambodia's 25 provinces for the years 2009 (Panel A) and 2019 (Panel B). The findings are very similar to those obtained when analyzing regions in Table 8. Specifically, there are significant disparities in welfare across Cambodian provinces. The most lagging province, Ratanakiri, had welfare levels that were approximately 25 and 30 percent of those in our reference province in 2009 and 2019, respectively. Additionally, by 2019, the second most advanced region, Kep, had welfare levels that were 35 log points below those in Phnom Penh. Finally, all provinces exhibit levels of both life expectancy and consumption components below those of our reference province.

The majority of our conclusions remain qualitatively robust to departures from our benchmark assumptions. However, there are some differences in magnitude, depending on certain assumptions. In 2009, the average within-country deviation from the benchmark ranges between 3.5 and 5 log points. By 2019, this variance decreases, ranging from 2.1 to 3.2 log points. Perhaps the most notable difference arises from the assumption regarding the minimum age threshold used to evaluate expected utility—particularly ages 25 and 40 and above—and whether we assign children adult leisure within the household. Altering these assumptions can lead to average differences of around 15 log points above the benchmark in 2009 (around 9 log points in 2019). Likewise, setting the Frisch elasticity ϵ to 2 results in an average deviation of 12 log points above our measures in the benchmark for the year 2009. This is consistent with the relative importance of increased leisure outside Phnom Penh in mitigating the negative impact on relative welfare from lower consumption and life expectancy.

4.1.7 Summary

From our discussion above we generalize the following findings on welfare across regions:

Fact 1. With the notable exception of Bangladesh, we observe a process of convergence between regions. As a result, the regional distribution of welfare appears to be becoming more compact in our sample of countries. However, the differences between the wealthiest and most lagging regions remain substantial in some cases, potentially reaching over 60 log points in the most recent year analyzed. Nonetheless, even across the most spatially unequal countries we never find within-country differences that are of a comparable magnitude to those reported by Jones and Klenow (2016) between the

Table 8: **Welfare Across Regions in Cambodia**

	Welfare λ	$\log \lambda$	Decomposition					LE
			Life Exp.	Cons.	Leis.	Cons. Ineq.	Leis. Ineq.	
Phnom Penh	100.0	0.000	0.000	0.000	0.000	0.000	0.000	76.7 (74.0)
Panel A: 2009								
Coast	53.7	-0.621	-0.149	-0.726	0.087	0.011	0.157	67.9
Tonle Sap	49.9	-0.695	-0.136	-0.843	0.107	0.017	0.160	68.1
Plain	49.3	-0.708	-0.151	-0.810	0.073	0.046	0.133	67.2
Plateau/Mountain	45.6	-0.786	-0.186	-0.908	0.101	0.034	0.173	64.5
Panel B: 2019								
Coast	61.3	-0.489	-0.091	-0.463	-0.001	0.029	0.037	70.1
Plain	52.8	-0.638	-0.075	-0.683	0.009	0.060	0.051	70.7
Plateau/Mountain	50.6	-0.682	-0.086	-0.711	0.003	0.048	0.063	70.0
Tonle Sap	49.9	-0.695	-0.064	-0.773	0.016	0.070	0.057	71.0

Notes: Please refer to Section 3 for details on the assumptions required in the calibration and computation of welfare. Last column presents life expectancy of each region in the survey year (for the region of reference, the first year is reported, and the last year is shown in parentheses).

wealthiest and poorest countries.[21]

Fact 2. Differences in (demographically-adjusted) average consumption tend to be the main component explaining welfare disparities across regions. Differences in average consumption are partially offset by more equitable consumption distributions in the lagging regions.

Fact 3. With the exception of our Latin American sample, higher levels of leisure in lagging regions tend to help narrow the welfare gap.

Fact 4. While differences in life expectancy remain responsible for differences in welfare between the wealthiest and most lagging regions, they have become less important over time.

Fact 5. In countries where we have measures of gross domestic product at the regional level, we find varying degrees of correlation (though positive) between our welfare computations and output per capita, ranging from as high as 0.86 in South Africa in 2023 to as low as 0.4 in Peru in 2007. Further, we find a negative correlation between the ratio of welfare to output and GDP per capita, implying that output per capita tends to underestimate (overestimate) welfare in less (more) developed regions.

4.2 Urban and rural areas

Having documented the heterogeneity in welfare across administrative regions in six countries, we now describe the patterns for welfare differences between the rural and

[21] E.g., the difference in national welfare between the US and Malawi is close to 500 log points.

urban areas in these countries. To do this, we aggregate our consumption-equivalent measure of welfare over the urban and rural classification for households, as explained in Section 3.5. We collect the results from this aggregation in Table 9. Before describing results for each country in detail, we outline some basic facts.

Fact 6. Rural areas, in the aggregate, exhibit lower levels of welfare compared to urban areas. These differences range from 7 log points (Bangladesh, 2005) to nearly 84 log points (Peru, 2007).

Remark 7. There is no single, overarching trend on rural convergence. In our sample of six countries we find two countries (South Africa and Peru) where the welfare of rural areas has caught up with urban areas, one where rural areas have fallen further behind (Cambodia), and three where the rural gap has remained relatively constant (Malawi, Mexico and Bangladesh).

Fact 8. Differences in consumption, followed by life expectancy, are the primary drivers of welfare disparities between urban and rural areas.

As a caveat, spatial price disparities may account for some of the gap between urban and rural areas. To the extent that higher costs of living in urban areas do not reflect the pricing of urban amenities the differences due to consumption should be read as an upper bound.

Fact 9. Consumption and leisure inequality tend to be lower in rural areas, which mitigates the overall welfare gap.

We find that differences in average leisure have a small effect in shrinking the urban-rural welfare gap, with the exception of Peru. Although this fact may be explained by unemployment, underemployment, or seasonal employment in rural settings, such causes do negatively affect our welfare measure through lower consumption. However, as discussed in Section 3.1, non-work hours that follow from structural unemployment may produce psychological costs rather than utility-enhancing leisure (Aghion *et al.*, 2016).

Remark 10. Bangladesh stands out among the six countries, as the welfare gap between urban and rural areas is substantially smaller than in other cases — primarily because differences in consumption are less pronounced.

We now discuss country-specific results, with relevant caveats that arise from the 13 robustness exercises we perform for the headline welfare statistics for each country and period. In **South Africa**, the results presented in Panel A of Table 9 show that rural areas are between 28 and 66 log points below the welfare level of urban areas (note that these differences are smaller in magnitude to those between regions). However, there was a substantial improvement of approximately 37 log points between 1993 and 2023. Although there are still disparities in life expectancy between rural and urban areas in 2023, the gap is closing. In 1993, life expectancy in urban areas was 10 years higher than in rural areas. By 2023, this difference had substantially narrowed. The improvement for rural areas is mainly driven by a substantial improvement in the components of life expectancy and, especially, in consumption, which delivers a 30-log-point gain.

Robustness. Table A.7 presents the calculations of λ under alternative assumptions. The main takeaway for South Africa is that the above results are robust across multiple specifications and data decisions. The only exception is the case where we assign children the leisure of their parents; in 2023, this raises the relative welfare of rural areas by a substantial 14-15 log points. Note that using the official classification for the urban status of surveyed households does not affect the results qualitatively (i.e., welfare is lower in rural areas but has improved over time). However, it does impact the magnitudes, as the calculations based on the official classification yield lower values of consumption equivalent welfare for rural areas in both survey years.

Panel B of Table 9 shows the results for **Malawi**. As we noted in the discussion of regional differences, the disparity between urban and rural areas in Malawi is less than that observed in South Africa. Although rural areas in Malawi exhibit lower levels of consumption equivalent welfare compared to urban areas, these differences are roughly 23-25 log points. In other words, in 2004 and 2019, the consumption-equivalent welfare of rural areas was approximately 78 percent and 79 percent of that of urban areas, respectively. Most of this difference in welfare is primarily due to differences in demographically-adjusted average consumption and, to a lesser extent, to the combination of losses from lower life expectancy and lower levels of flow utilities. This urban-rural difference would have been even greater if there had been no better indicators of inequality in consumption and leisure observed in rural areas of Malawi.

Robustness. Table A.7 shows that variations in most of our core assumptions do not produce significant differences in the rural welfare gaps for Malawi. There are only two exceptions worth highlighting. First, evaluating expected utility only for individuals aged 40 and above can result in a welfare measure that is up to 13 log points higher in 2004. Second, we find considerably lower rural welfare indicators when using the official urban and rural definitions from Malawi. In 2004, the measure would decrease from 78 to 62.1, while in 2019, it would drop from 79.3 to 57.6. These differences are consistent with the more encompassing definition of urbanization in the GHSL degree of urbanization measures for Malawi, and higher welfare (closer to the urban benchmark)

Table 9: **Welfare Comparison Between Urban and Rural Areas by Country**

			Decomposition					
	Welfare λ	log λ	Life Exp.	Cons.	Leis.	Cons. Ineq.	Leis. Ineq.	LE
Urban	100.0	0.000	0.000	0.000	0.000	0.000	0.000	-
Panel A: Rural areas in South Africa								
1993	52.0	-0.655	-0.217	-0.670	0.043	0.154	0.036	57.8
2023	75.4	-0.282	-0.057	-0.371	0.024	0.099	0.024	64.0
Panel B: Rural areas in Malawi								
2004	78.0	-0.249	-0.144	-0.282	0.013	0.106	0.058	49.2
2019	79.3	-0.232	-0.052	-0.279	0.006	0.056	0.037	63.2
Panel C: Rural areas in Peru								
2007	43.0	-0.844	-0.169	-0.767	-0.085	0.067	0.110	70.3
2018	64.7	-0.435	-0.063	-0.460	-0.047	0.056	0.078	75.1
Panel D: Rural areas in Mexico								
2002	71.7	-0.333	-0.024	-0.249	0.000	-0.052	-0.008	74.4
2018	73.8	-0.303	-0.004	-0.362	-0.006	0.057	0.011	74.3
Panel E: Rural areas in Bangladesh								
2005	92.9	-0.074	-0.027	-0.179	0.020	0.051	0.062	66.5
2016	90.8	-0.097	-0.014	-0.190	0.028	0.008	0.071	70.5
Panel F: Rural areas in Cambodia								
2009	73.0	-0.315	-0.114	-0.382	0.023	0.089	0.068	65.2
2019	64.7	-0.435	-0.189	-0.365	0.013	0.067	0.040	67.0

in the areas that are on the margin between classifications.

The differences in welfare between rural and urban areas in **Peru** are of a similar magnitude to those already documented in Table 5 between Lima and the poorest regions of Peru. As shown in Panel C of Table 9 welfare in rural areas was 43 percent of that in urban areas in 2007. This translates to a difference of almost 84 log points in welfare, explained primarily by average consumption (77 log points), life expectancy (17 log points) and leisure (9 log points). This difference exists despite better measures of inequality in both consumption (which adds 7 log points to welfare) and leisure (which adds 11 log points). Although these components also explain the difference in welfare between rural and urban areas in 2018 in similar proportions, it is important to note a substantial catch-up in the welfare of rural areas, where relative welfare improved by approximately 41 log points between 2007 and 2018.

Robustness. The results in Table A.7 show that our calculations are robust to alternative assumptions. In particular, the results using the official classification of urban-rural status are very similar in the case of Peru.

In Panel D of Table 9, we present our results for **Mexico**. Differences in welfare between rural and urban areas of Mexico are not as pronounced as in Peru. While rural areas indeed have lower levels of consumption-adjusted welfare, the gap was 33 log points in 2002 and 30 log points in 2018. Most of this difference is explained by gaps in average consumption. Differences in life expectancy are negligible for welfare under our method.

Robustness. The results in Table A.7 show that our calculations are robust to alternative assumptions. The only exception is when using the official classification of urban-rural status in 2002. In this case, rural welfares based on NSO definitions were 10.0 log points higher in 2002 and 1.7 log points lower in 2018.

Panel E of Table 9 shows the results for **Bangladesh**. As mentioned above, the difference in welfare between rural and urban areas is quite small. In 2005, rural areas were about 7 log points behind urban areas. These gaps remained small over the next eleven years, widening slightly to just under 10 log points. The difference in life expectancy is smaller, and although the differences in consumption reached around 20 log points in 2016, the better leisure indicators (both in terms of mean value and variance) have helped keep the welfare gap below 10 log points. **Robustness**. It is important to note that these striking results regarding the small welfare gap observed between urban and rural areas in Bangladesh are not an artifact of our definition of rural and urban areas. Welfare gaps for rural areas are comparable when computed using NSO urbanization definitions (see Table A.7). Although NSO welfare measures display a larger rural gap in 2005 and an apparent catch-up in 2016, these results are spurious: the Bangladesh Bureau of Statistics assigned urban status more strictly starting with the 2011 Census, effectively selecting higher consumption areas into a rural status in the 2016 consumption survey. This case highlights the importance of urbanization classifications that are globally comparable and consistent over time.

Finally, the differences in welfare between rural and urban areas of **Cambodia** are presented in Panel F of Table 9. We observe that relative welfare in rural areas deteriorated between 2009 and 2019, with the welfare gap increasing from a negative 30 log points to 43 log points, despite reductions in consumption differences (which remain substantial). The primary reason for the relative decline in welfare in rural Cambodia is a failure to keep pace with life expectancy improvements in urban areas. In fact, the life expectancy component — which links mortality differences with flow utility — went from accounting for 11 log points of rural areas' welfare gap in 2009 to 19 log points in 2019. Additionally, while both leisure components (average and variance) and consumption inequality help mitigate the rural gap, their relative importance has decreased over time.

Robustness. Results for Cambodia are robust to alternative assumptions and data decisions (see Table A.7), but gaps are larger and change less over time using NSO urbanization definitions.

Summing up the robustness exercises across all countries we find that, while NSO definitions of urban-rural status can differ significantly from ours (in particular, for Malawi, South Africa, Cambodia and Bangladesh) the qualitative patterns for rural welfare gaps are mostly robust across classifications. In every case the gap is larger using NSO definition (except Mexico in 2002), however:

Remark 11. Rural welfare gaps are higher using NSO classifications in all cases except Mexico in 2002.

NSO definitions for urbanization may be too coarse, for instance when adhering to administrative rather than functional classifications, or may not keep pace with recent urbanization trends. A change enacted by the Bangladesh Bureau of Statistics for the 2011 Census compounded both issues, reverting to a strictly administrative classification during an urbanization process.

4.3 Urban transitions

In this section we explore the interaction between welfare and the dynamics of urbanization. Specifically, we partition our urban classification of households in each consumption or labor survey based on whether their location was classified as urban within the previous 10 years. We refer to areas that have been urban for 10 or more years as "established," otherwise the areas are in "urban transition." We set the established urban areas as our location of reference for the computation of λ in this section. Table 10 below shows the results for the 12 country-wave pairs. A final clarification is necessary before starting the analysis: as discussed in Section 3.2, we cannot measure survival curves within or between urban areas. We assume no differences in mortality rates between established and "in transition" urban areas. As a result, our estimates of the welfare gap for areas in transition are likely biased, with the size and direction of the bias depending on unobserved differences in mortality rates (or the implied life expectancy at birth). Such potential differences may depend on differential access to health care, which may also depend on the prevalence of informal urban settlements and the provision of public goods and transportation to such areas (Bilal *et al.*, 2019). Table A.8 presents summary statistics for consumption, leisure and life expectancy for these three types of areas in the six countries and twelve wave-years analyzed. Before discussing each country in detail, we present some basic facts derived from Table 10.

Fact 12. Areas in urban transition enjoy higher welfare than rural areas (with the exception of Mexico in 2018), and lower welfare than established urban areas.

Fact 13. The welfare gap between areas in urban transition and established urban areas has remained relatively constant over time, with the exception of a relative improvement in South Africa and a deterioration in Malawi.

Note that "established" areas include places that would have been classified as in-transition 10 years earlier. The urban areas that are in transition in the latter period for each country are typically expanding out from the urban core, and Fact 13 reflects change in welfare at what constitutes the margin of urban areas at each point in time.

Note that we do not expect large changes in the relative welfare of rural areas in this section, compared to the last. Rural households are the same in both analyses. Changes in lambda observed between the tables are solely due to the change in the reference point: all urban areas in the previous section, and established urban areas here.

Welfare for areas in urban transition in South Africa as of 1993 were positioned roughly in the middle between established urban areas and rural zones, with approximately 30 log points of distance from each extreme.[22] The lag relative to established areas was entirely explained by lower average consumption, despite better inequality indicators for both leisure and consumption, as well as higher levels of leisure. Despite having higher average consumption than rural areas, these newly urbanized zones had worse inequality and lower levels of leisure. Thirty years later, the areas recently urbanized by 2023 had significantly narrowed the gap with respect to areas urbanized 10 years prior, primarily due to higher average consumption levels.

The trend for welfare in newly transitioned urban areas was different for the other five countries in our study. In Malawi, the gap for urban areas in transition widened by 19 log points between 2004 and 2019, coming to parity with traditionally rural zones by 2019. Average consumption remains the main factor behind the gap between established and in transition urban areas, while the contribution from lower inequality (in both consumption and leisure) has decreased in relative importance.

In Peru, the welfare gap for urban areas in transition decreased slightly between 2007 and 2018. Rural areas caught up substantially to areas in urban transition, with a welfare gap declining from 57 log points in 2007 to 20 log points by 2018. Most of this reduction is explained by a relative improvement in average consumption; by 2018, average consumption in rural areas was about 80% of that in newly urbanized areas.

[22] Note that townships such as Soweto in Johannesburg and Khayelitsha in Cape Town are classified as established by 1993.

Several patterns emerge when analyzing the evolution of relative welfare for new urban and rural areas in Mexico between 2002 and 2018. First, there is no evidence of divergence over time. Rural areas experience a slight improvement of about 4 log points, while newly urbanized areas remain broadly stable. Although there was lagging demographically-adjusted average consumption, there was also decline in the standard deviation of consumption, which translate into a decline in overall welfare. Second, welfare for both rural and newly urban areas remains between 20 and 24 log points below established urban areas [23]. Third, rural life expectancy was not a major source of the welfare gap in 2002 (2 log points), and the difference in life expectancy had almost vanished by 2018 (around 74 years for rural and urban areas in 2018) so we do not find the life expectancy component to be accounting for much of the variation in welfare in Mexico.

The results for Bangladesh show a slight deterioration in the relative welfare of newly urbanized and rural areas over time. Notably, by 2005, the difference in welfare among rural, in-transition, and established urban areas appeared small because lower average consumption was partly offset by the two leisure components and consumption inequality. In addition, rural areas were approximately 1.5 log points behind, due to a combination of lower life expectancy and average consumption. Eleven years later, in 2016, both areas were still only between 8 and 10 log points below established urban areas, primarily due to lagging average consumption, despite higher leisure and lower inequality in both consumption and leisure.

Finally, the results for Cambodia show that areas in urban transition lagged established urban areas by 15 log points in 2009 and 14 points in 2019. Meanwhile, rural areas experienced a relative decline between 2009 and 2019, with the gap compared to established urban areas increasing by more than 10 log points. This lag in welfare is explained by lower demographically-adjusted average consumption and shorter life expectancy, despite improved leisure (both in terms of average and standard deviation) and lower consumption inequality.

4.4 Welfare at the level of individual urban areas

We now turn to tracking welfare at the level of individual urban areas. Our approach is to partition our prior urban classifications, applying identifiers to some of the

[23]The case of Mexico illustrates how demographic adjustment can substantially influence welfare measurements. According to Table A.8, the average consumption in 2018 for historically urban and newly urbanized areas was $16,230 and $16,224, respectively, in 2017 PPP-adjusted US dollars. Consequently, their difference in logs is less than 1 point. However, the demographically-adjusted average consumption (before adding per capita government consumption) for the same years was $20,521 and $14,554, respectively, in 2017 PPP-adjusted US dollars, which results in a difference of approximately 34 log points.

Table 10: **Welfare Comparison Between Urban Transitions**

Group	Country	Year	Welfare λ	log λ	Decomposition Life Exp.	Cons.	Leis.	Cons. Ineq.	Leis. Ineq.	LE
Urban Transition	ZAF	1993	73.6	-0.306	0.000	-0.442	0.025	0.093	0.017	67.3
Rural	ZAF	1993	51.1	-0.672	-0.217	-0.690	0.044	0.155	0.036	57.8
Urban Transition	ZAF	2023	83.9	-0.176	0.000	-0.247	0.001	0.073	-0.003	66.7
Rural	ZAF	2023	75.1	-0.286	-0.057	-0.372	0.023	0.096	0.023	64.0
Urban Transition	MWI	2004	92.2	-0.082	0.000	-0.224	-0.003	0.091	0.056	55.0
Rural	MWI	2004	76.6	-0.266	-0.144	-0.327	0.013	0.123	0.070	49.2
Urban Transition	MWI	2019	76.0	-0.275	0.000	-0.378	0.002	0.065	0.036	66.3
Rural	MWI	2019	74.8	-0.291	-0.053	-0.349	0.007	0.060	0.044	63.2
Urban Transition	PER	2007	75.6	-0.280	0.000	-0.366	-0.089	0.049	0.125	80.4
Rural	PER	2007	42.7	-0.851	-0.169	-0.773	-0.086	0.066	0.111	70.3
Urban Transition	PER	2018	78.4	-0.243	0.000	-0.238	-0.093	0.065	0.022	77.8
Rural	PER	2018	64.5	-0.439	-0.063	-0.463	-0.047	0.056	0.078	75.1
Urban Transition	MEX	2002	78.9	-0.238	0.000	-0.209	-0.003	-0.017	-0.008	75.2
Rural	MEX	2002	78.7	-0.239	-0.023	-0.148	-0.007	-0.054	-0.007	74.4
Urban Transition	MEX	2018	79.5	-0.230	0.000	-0.282	-0.013	0.076	-0.010	74.5
Rural	MEX	2018	82.0	-0.198	-0.004	-0.239	-0.012	0.055	0.002	74.3
Urban Transition	BGD	2005	93.9	-0.063	0.000	-0.156	0.007	0.050	0.036	67.9
Rural	BGD	2005	92.5	-0.078	-0.027	-0.188	0.021	0.053	0.063	66.5
Urban Transition	BGD	2016	91.9	-0.084	0.000	-0.195	0.019	0.047	0.044	71.2
Rural	BGD	2016	90.3	-0.102	-0.014	-0.197	0.028	0.008	0.073	70.5
Urban Transition	KHM	2009	86.5	-0.145	0.000	-0.366	0.037	0.104	0.079	73.1
Rural	KHM	2009	71.7	-0.332	-0.114	-0.413	0.028	0.091	0.075	65.2
Urban Transition	KHM	2019	86.6	-0.144	0.000	-0.302	0.000	0.072	0.085	75.3
Rural	KHM	2019	64.2	-0.443	-0.192	-0.377	0.018	0.065	0.044	67.0

Notes: These results were obtained using the rural/urban points or share by population raster. Urban transition areas are defined by a transition to urbanization over the previous 10 years.

contiguous urban areas and then tracking welfare at the level of such identifiers.

We make use of the "Functional Urban Areas" (FUA) layer from European Commission. Joint Research Centre. (2019), a global dataset of urban footprints. Our goal is not to make use of the boundaries in that dataset to delineate cities, as this is a concept open to substantial ambiguity (see Dong *et al.*, 2024, for a discussion). We simply use the FUA layer to attach an identifier to households that we had previously classified as urban and that fall within each FUA boundary. Rural households that fall within a FUA boundary remain classified as rural, and any urban household not within any boundary are excluded from the analysis in this section. Since FUAs use an arguably "expansive" definition of urban boundaries, the first of these cases is more common than the second. We also define "hinterlands," which are the rural areas that are closest to a particular urban area in straight-line distance.

Out of 321 labeled urban areas (UAs) in our sample of countries we are able to

Table 11: **Welfare Comparison Between Urban and Rural Areas Relative to Urban Mexico 2002**

			Decomposition					
Area	Year	Welfare λ	log λ	Life Exp.	Cons.	Leis.	Cons. Ineq.	Leis. Ineq.
Urban MEX	2002	100.0	0.000	0.000	0.000	0.000	0.000	0.000
Panel A: South Africa								
Rural	1993	22.5	-1.491	-0.356	-1.316	0.087	0.067	0.028
Urban	1993	40.6	-0.900	-0.205	-0.643	0.044	-0.090	-0.008
Rural	2023	50.1	-0.691	-0.201	-0.709	0.095	0.075	0.050
Urban	2023	64.6	-0.437	-0.175	-0.332	0.072	-0.028	0.026
Panel B: Malawi								
Rural	2004	4.0	-3.222	-0.497	-3.016	-0.011	0.239	0.062
Urban	2004	4.8	-3.033	-0.404	-2.757	-0.029	0.151	0.007
Rural	2019	6.6	-2.720	-0.181	-2.914	0.076	0.178	0.121
Urban	2019	8.1	-2.513	-0.150	-2.644	0.069	0.128	0.084
Panel C: Peru								
Rural	2007	17.8	-1.728	-0.078	-1.566	-0.203	0.198	-0.078
Urban	2007	42.8	-0.849	0.138	-0.810	-0.124	0.138	-0.190
Rural	2018	40.2	-0.912	0.002	-0.933	-0.171	0.259	-0.070
Urban	2018	62.4	-0.471	0.076	-0.480	-0.128	0.209	-0.148
Panel D: Bangladesh								
Rural	2005	7.1	-2.649	-0.152	-2.794	0.036	0.268	-0.006
Urban	2005	7.6	-2.581	-0.133	-2.613	0.019	0.214	-0.068
Rural	2016	17.2	-1.759	-0.081	-1.953	0.059	0.204	0.012
Urban	2016	18.9	-1.668	-0.072	-1.765	0.031	0.197	-0.059
Panel E: Cambodia								
Rural	2009	11.5	-2.159	-0.157	-2.247	-0.037	0.257	0.025
Urban	2009	15.8	-1.848	-0.047	-1.862	-0.058	0.165	-0.045
Rural	2019	11.5	-2.164	-0.171	-1.969	-0.177	0.248	-0.096
Urban	2019	17.6	-1.737	0.014	-1.602	-0.192	0.179	-0.137

Notes: These results were obtained using the rural/urban points or share by population raster. All welfare values are normalized to urban Mexico in 2002 = 100.

construct welfare measures for 257 across two survey years.[24] The dataset forms an unbalanced panel of 578 observations. [25]

To compute our welfare measures at the UA level we use each country's largest city as the reference location. The urban areas of reference are Dhaka, Johannesburg,

[24]The number of identified UA-year observations varies across countries: 248 in Mexico, 122 in Bangladesh, 102 in South Africa, 81 in Peru, 16 in Cambodia, and only 9 in Malawi.

[25]We use two population-count measures at the UA level. The first is time-invariant and reflects population size in 2015 as measured by European Commission. Joint Research Centre. (2019), while the second is time-varying.

Lilongwe, Lima, Mexico City, and Phnom Penh. Our computations also require survival rates specific to each UA and its hinterland. To the best of our knowledge, such data is not publicly available. We calibrate UA- and hinterland-specific survival rates by reallocating our data on regional life expectancy into urban and rural components to match the national urban–rural gap. For each country-year, we compute the national urban–rural difference (e.g., in Peru 2018, 77.8 versus 75.1 implies a 3.6% urban advantage). Each UA is assigned to its region and the region's life expectancy is split into urban and rural values such that (a) the regional mean stays the same, and (b) the urban–rural gap matches the national gap. For example, a North Peru region life expectancy of 70.0 decomposes to about 71.3 (urban) and 68.8 (rural), preserving the 3.6% national gap. These urban and rural statistics feed MortPak MATCH to produce life-table survival curves, using country-specific model life-table patterns. UAs within the same region and country-year share the same urban and rural survival curves (one for the urban area, and one for the hinterland).

Figure 8 depicts kernel density estimates for our welfare computations and their decomposition across six panels. Each panel pools all available data for our sample of 12 country-years and plots kernel density estimates. Several patterns emerge: First, as our urban areas of reference (i.e., largest cities in each country) are among the richest urban areas in our sample, the overall distribution of welfare (Panel A) and, particularly, its consumption component (Panel B) have negative means (a negative value of our welfare measure implies lagging behind the reference urban area) and negative skewness. Second, consumption inequality (Panel E) tends to be much larger in our UAs of reference as the mass of the distribution of this welfare component is concentrated on the left of the figure (i.e., right-skewed) and presents a positive mean. Third, while the mean of the leisure component is around zero (Panel D), its kernel density exhibits substantial negative skewness due to a fat left tail, indicating that a considerable number of UAs are substantially worse in terms of leisure, relative to their national reference. Fourth, the life expectancy component (Panel C) does not appear unimodal, even though its largest peak is near zero.

We now analyze the distribution of our main welfare calculation (λ) for each country separately. Figure 9 shows kernel estimations for λ. Consistent with a previous discussion at the regional level, South Africa (yellow line) has the largest variance (i.e., 0.23) and the lowest mean (-0.4) and mode (-0.39) in welfare. However, it is important to note that this is mostly driven by the welfare distribution in 1993 (see discussion below). Peru's density function (blue line) also stands out for being more compact, with the smallest variance (0.04), but it has a substantially negative mean value (-0.38). Notably, almost every UA in Peru, in terms of welfare, lags behind the capital city,

Figure 8: **Distribution of Welfare and its Components Across Urban Areas**

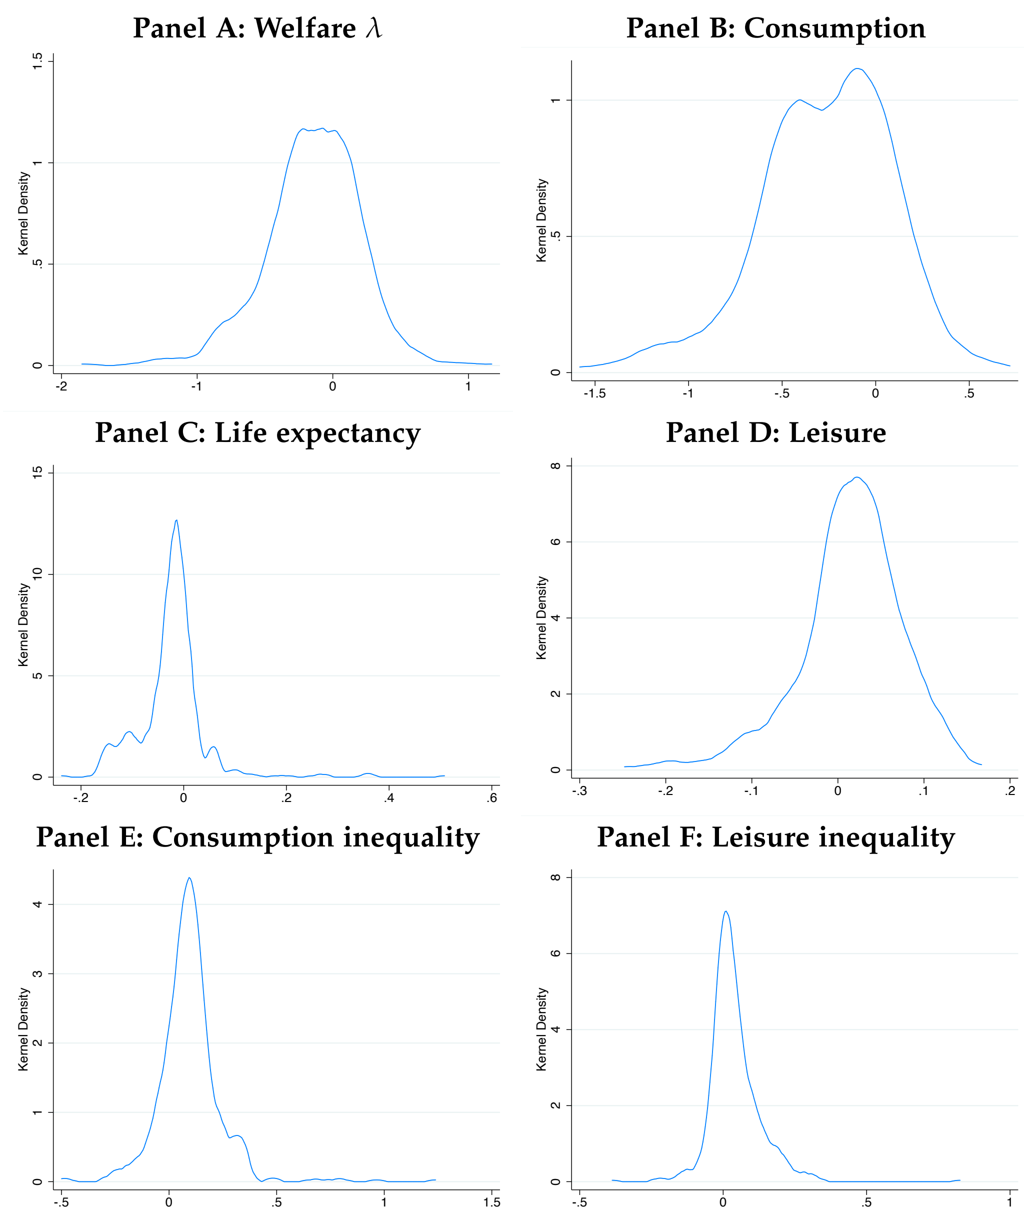

Notes: Each panel depicts a kernel density estimate of welfare and its decompositions from equation 6. The blue densities combine measures across all UAs in available years, while red densities use measures residualized from country × year fixed effects.

Lima. Cambodia's density function (purple line) appears to be bimodal, but, similar to Malawi (green line), our estimates rely on a very small number of UAs, preventing us from being conclusive. Finally, the welfare density functions for Mexico (red line) and Bangladesh (orange line) indicate the tightest distributions of welfare across cities, with both relatively small variances and more than 20% of the UAs performing better than their respective references.

Figure 9: **Distribution of Welfare Across UAs by Country**

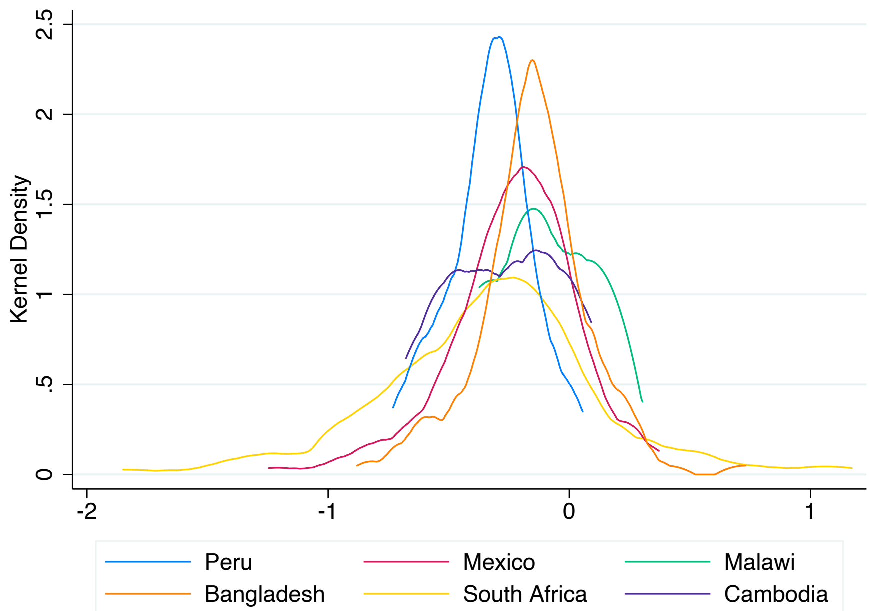

Notes: Each panel shows a kernel density estimate of welfare by country, pooling the two available years of data for each case. Welfare for each UA is measured relative to the largest city in the country during the corresponding year. The reference urban areas (with years of available data in parentheses) are: Mexico City (2002 and 2018), Dhaka (2005 and 2016), Johannesburg (1993 and 2023), Lilongwe (2004 and 2019), Lima (2007 and 2018), and Phnom Penh (2009 and 2019).

In Figure 10, we plot kernel density functions for welfare in each country-year separately.[26] The previous country-level analysis masked some within-country heterogeneity over time. Consider, for instance, the case of South Africa: its welfare distribution shifted toward zero and became more compact from 1993 to 2023, leading to an increase in kurtosis. A similar pattern can be observed in Peru from 2007 to 2018, which also exhibited a fatter right tail in the later period—evidence that some UAs have caught up with Lima. Mexico (from 2002 to 2018) and Bangladesh (from 2005 to 2016)

[26]We omitted Malawi and Cambodia from this analysis due to a small number of observations in each country-year sample.

Figure 10: **Distribution for Welfare Across UAs, Within Countries and Over Time**

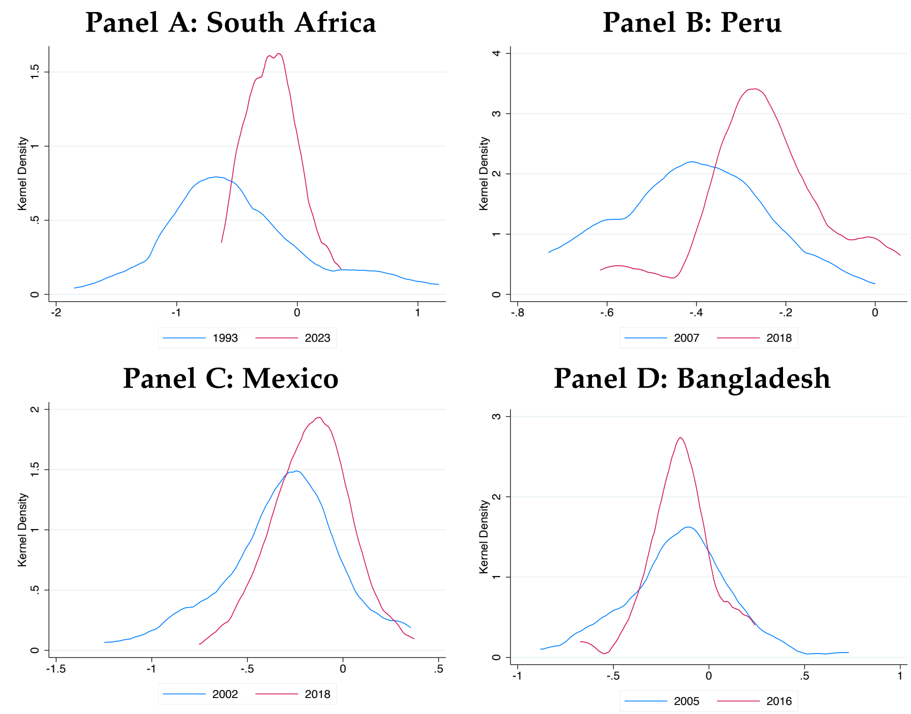

Notes: Each panel shows a kernel density estimate of welfare for each country and year. The blue kernel represents the distribution of UA welfare in the first available year of data, while the red kernel represents the second year. Malawi and Cambodia are omitted from the analysis due to the small number of observations in their country–year samples.

display similar dynamics, with their welfare distributions across urban areas becoming more compact with shorter left tails over time.

We now describe the relationship between population size and welfare within countries across UAs. We estimate two variations of a simple log-log model. The first model exploits within country-year variation:

$$W_{i(c)t} = \alpha_0 + \alpha_1 Log(Population)_i + \gamma_{ct} + \varepsilon_{it} \tag{8}$$

where $i(c)t$ indicates UA i located in country c during sampling year t; $W_{i(c)t}$ is the welfare-related outcome of interest, e.g., our main measure λ or its components of consumption, leisure, and life expectancy. $Log(Population)_i$ refers to our time-invariant measure of an UA's population size in year 2015. We also estimate a second specification in which population is allowed to vary over time, using $log(Population)_{it}$ measured in the corresponding year for each UA. γ_{ct} is our set of country-years fixed effects. We cluster standard errors at the UA level.

Table 12: **Population and Welfare at the UA level**

Panel A	Consumption-Equivalent Welfare $\log(\lambda_i)$			Consumption Average $\log \bar{c}_i - \log \bar{c}_r$			Consumption Inequality $(E[\log c_i] - \log \bar{c}_i) - (E[\log c_r] - \log \bar{c}_r)$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Log of Population	0.049***	0.043***	0.040***	0.061***	0.064***	0.063***	-0.031***	-0.024***	-0.024***
	(0.011)	(0.010)	(0.010)	(0.012)	(0.012)	(0.011)	(0.005)	(0.005)	(0.005)
Adj. R^2	0.034	0.136	0.133	0.043	0.182	0.181	0.063	0.318	0.319
Panel B	**Life Expectancy $\sum_a \Delta s_a^i u_a^i$**			**Leisure Average $\nu(\bar{\ell}_i) - \nu(\bar{\ell}_r)$**			**Leisure Inequality $\big(E[\nu(\ell_i)] - \nu(\bar{\ell}_i)\big) - \big(E[\nu(\ell_r)] - \nu(\bar{\ell}_r)\big)$**		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Log of Population	0.004**	0.001	0.000	0.019***	0.011***	0.011***	-0.004	-0.010***	-0.011***
	(0.002)	(0.002)	(0.002)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.003)
Adj. R^2	0.005	0.255	0.254	0.042	0.208	0.208	-0.000	0.271	0.273
Time-Varying Population	N	N	Y	N	N	Y	N	N	Y
Country×Year FE	N	Y	Y	N	Y	Y	N	Y	Y
Obs.	578	578	578	578	578	578	578	578	578
N. UA	321	321	321	321	321	321	321	321	321

Notes: Standard errors clustered at the UA level. ***p<0.01, **p<0.05, *p<0.1. All components come from the welfare decomposition in equation 6. When population is indicated as time-invariant, 2015 figures are used. When time-varying, population represents counts in the year corresponding to each welfare measure.

Results in columns 1 to 3 of Panel A in Table 12 suggest that larger UAs (as measured by 2015 population) tend to have higher welfare. In column 1 we set $\gamma_{ct} = 0$ and document a statistically significant link between our consumption-equivalent welfare measure and population size in 2015. Adding country-year fixed effects improves model fit, but the point estimate for α_1 remains similar and statistically significant: a 1 percent increase in 2015 population is associated with a 0.04 percent increase in urban area welfare. The result in column 3 of the same panel suggests qualitatively and quantitatively similar findings when using time-varying population data for a restricted set of UAs for which population counts were available in both sampling years. Most of the welfare–population relationship is mediated by the link between population size and average consumption. As shown in columns 4 to 6, the underlying consumption to population elasticities range from 0.061 to 0.064. Results in columns 7 to 9 show that larger UAs tend to have more unequal consumption.

Turning to Panel B of Table 12, several results are worth noting. Results for life expectancy are included for completeness, but given our inability to measure city-specific life expectancy these are driven by our national rescaling of regional differences. A statistical correlation between population size and life expectancy holds in the cross section, but vanishes once we include country-year fixed effects. This makes sense because administrative regions contain urban areas of multiple sizes, all of which we need to assign the same life expectancy. Therefore, the apparently precise finding of no correlation between population and life expectancy should not be given a conclusive interpretation. Finally, people in larger UAs appear to work more hours and face greater inequality in leisure hours.[27]

So far we have investigated a linear relationship between population size and welfare. We next test whether non-linearities in population are relevant for explaining differences in welfare across UAs. To do so, we modify equation 8 replacing $Log(Population)_i$ with a set of indicators for quintiles of the population distribution in 2015:

$$W_{i(c)t} = \alpha_0 + \sum_{\tau=2}^{5} \alpha_\tau Pop_i^\tau + \gamma_{ct} + \varepsilon_{it} \tag{9}$$

where τ indexes the population quintile. In Figure A.9, we present the point estimates from equation 9 using our six different welfare measures. Although our previous results are more evident and gain statistical significance in large UAs (i.e., UAs with an average population size of 2.4 million inhabitants), there does not appear to be substantial

[27] As Figure A.10 shows, these results remain similar for different alternative samplings in which we omit different time periods (Panel A) or exclude one year (Panel B) or one country (Panel C) at a time.

nonlinearity in the statistical relationship between welfare and population size.

We now analyze whether the data provide any evidence of within-country (across UAs), conditional or unconditional, convergence in welfare. To do so, we estimate different versions of the following equation:

$$\Delta W_{i(c)t} = \alpha_0 + \alpha_1 W_{i0} + \alpha_2 Log(Population)_{i0} + \delta_c + \varepsilon_i \tag{10}$$

where $\Delta W_{i(c)t}$ is the growth rate of our welfare measures between the two available periods.[28] Two clarifications are necessary regarding this implicit growth rate. First, since the two data periods (and the distance between these periods) vary by country, our dependent variable is either the annual compound growth rate or the average growth rate of welfare.[29] Second, to allow within-country comparisons over time, all welfare measures are computed using each country's largest city in the first period as the baseline. The variable W_{i0} represents welfare (i.e., log(λ)) in the first period where $Log(Population)_{i0}$ accounts for an UA's population in year 2015. Finally, δ_c is a set of country fixed effects. A negative estimate of α_1 would indicate that lagging UAs experienced faster growth in welfare, implying convergence. When we set $\delta_c = \alpha_2 = 0$, a negative α_1 would indicate unconditional convergence, whereas conditional convergence in welfare exists when the negative coefficient persists after controlling for ("conditioning on") the other covariates.[30] Figure 11 shows evidence consistent with both unconditional (Panel A) and conditional (Panel B) convergence in our main welfare measure (i.e., Consumption-Equivalent Welfare). The negative relationship between the initial levels of welfare and growth thereafter seems to be large across urban areas. Including country fixed effects and population counts does not alter the previous result, but improves the precision of the estimate.

Table 13 provides point estimates and confirms the results using different specifications of equation 10. Columns 1–4 use the annual compound growth rate of

[28] The multiplicative normalization to the reference UA's first-period welfare cancels out in the time-difference ($\Delta W_{i(c)t}$), so growth rates are absolute changes in consumption-equivalent welfare for UA_i, not changes relative to the largest UA of the country.

[29] We compute annual compound growth rates as $\Delta W_{i(c)t} = ((W_{it}/W_{it_0})^{\frac{1}{t-t_0}} - 1) \times 100$. Since some components of the welfare decomposition within an UA can change sign over time (this issue is not present when using our main welfare measure), computing an annual compound growth rate would be mathematically impossible if the number of years (i.e., compounding period) within periods is an even number. To avoid this problem we instead compute average growth rates, as $\Delta W_{i(c)t} = ((W_{it}/W_{it_0}) - 1) \times \frac{1}{t-t_0}) \times 100$.

[30] The seminal papers in the empirical growth literature (Mankiw *et al.*, 1992; Barro, 1991) called this result "conditional convergence," because it reflects the convergence of regions after one controls for differences in the determinants of steady states. Therefore, it is important to clarify that a rigorous conclusion about the existence or absence of conditional convergence would require a theoretical model that identifies the variables that need to be included to account for welfare differences at steady states.

Figure 11: **Convergence in Consumption-Equivalent Welfare Across Urban Areas**

Panel A: Unconditional **Panel B: Conditional**

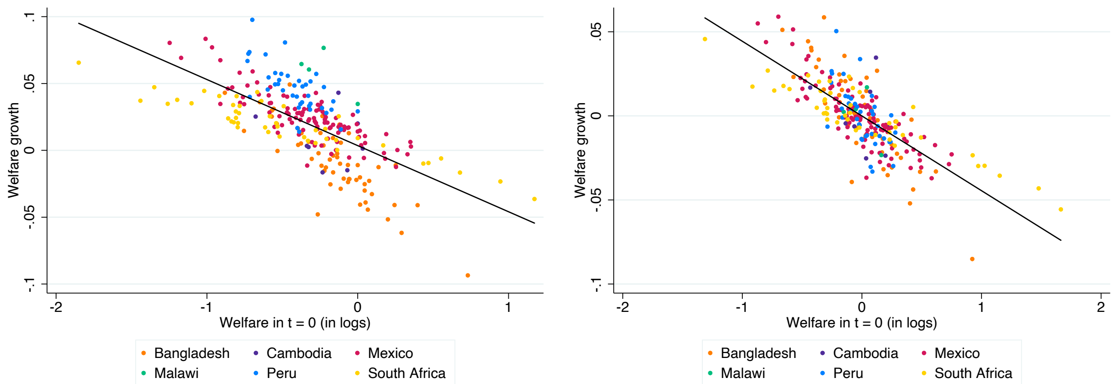

Notes: These panels show the unconditional (Panel A) and conditional (Panel B) empirical relationship between within-UA welfare growth and initial level of welfare as represented in equation 10.

log(λ); column 5 uses the implied arithmetic average growth rate. Column 1 has no controls. Columns 2–5 add country fixed effects; columns 3–5 also condition on the (log) population of the UA in 2015; column 4 additionally weights observations by population. Several findings are worth highlighting.

The coefficient on initial welfare is negative and highly significant in every column, ranging from -0.049 (column 1) to -0.043 (column 2), and -0.044 once both country fixed effects and population are included. A one-log-point higher initial welfare position relative to the country's largest UA is associated with roughly 4–5 percentage points lower annual welfare growth thereafter. Read in the relative-welfare framework, this means that UAs sitting further below their country's frontier in the first period closed a substantial portion of that gap by the second period. The implied half-life of within-country welfare differentials is on the order of log 2 / 0.045 $\approx$ 15 years.

The convergence coefficient barely moves across specifications. Comparing column 1 and column 2 shows that most of the convergence we observe is within country rather than driven by cross-country differences in average growth; country fixed effects sharpen precision without changing the slope. Column 4 confirms that the result is not driven by a few small or large UAs: population-weighting yields a very similar coefficient (-0.047).

Second, larger UAs grow modestly faster. The coefficient on log population is positive and significant. A one-log-point increase in initial population (an e-fold rise) is associated with about a 0.2 percentage point higher annual welfare growth, or roughly 10% of the sample's mean welfare growth rate. This is consistent with mild agglomeration effects in welfare growth, and it is the channel through which controlling for size matters most for the Bangladesh fixed effect discussed below.

Third, there is a pronounced cross-country heterogeneity in average welfare growth. Country dummies (with no intercept) capture each country's average annual welfare growth rate conditional on a UA's initial welfare position and (where controlled) its population. Importantly, since ΔW is absolute, these dummies should be read as country-specific average growth in consumption-equivalent welfare itself, not as shifts in the country's gap with respect to its reference city. Malawi and Peru thus stand out as cases of broad-based welfare gains across the urban hierarchy; South Africa and Cambodia show modest improvement; Mexico is intermediate. Bangladesh is the only country in which the conditional average growth rate of consumption-equivalent welfare is negative. Note also that the Bangladesh coefficient moves from -0.016 (column 2) to -0.038 (column 3) once log population enters, because Dhaka — the largest UA in our sample — has positive welfare growth that is partly attributed to size in column 3. With size partialled out, the residual country-specific component for Bangladesh is more negative.

A natural follow-up question is whether the negative Bangladesh dummy reflects the country's largest city "pulling away" from the rest. It does not. Whether the largest UA in Bangladesh is gaining ground on, or losing ground to, the rest of the urban system is captured by α_1, the convergence slope (which is estimated on within-country variation once country fixed effects are included). Since $\alpha_1 < 0$ in Bangladesh as in every other country, smaller and poorer UAs closed the gap on Dhaka over the sample. The negative Bangladesh dummy instead says that, on average and after holding initial welfare position and population fixed, absolute consumption-equivalent welfare deteriorated across Bangladeshi UAs during the period covered — a result that the decomposition in Table 14 attributes to the consumption and inequality components, partially offset by life-expectancy gains.

To conclude, the goodness of fit rises sharply from column 1 ($R^2 = 0.46$) to column 2 ($R^2 = 0.76$) once country fixed effects are added, confirming that country-level shifts account for most of the residual variance in welfare growth left unexplained by initial welfare alone. Adding population and population-weighting raises R^2 modestly (to 0.77 and 0.80 respectively). Column 5, which replaces the compound growth rate with the simple-average annual growth rate, returns essentially identical point estimates — as expected for the moderate growth rates and short time horizons in our sample.

Table 14 re-estimates equation 10 with each of the five components of the welfare decomposition in equation 6 as the dependent variable. Each component is expressed in the same log-units of consumption-equivalent welfare as $\log(\lambda_i)$ itself, and each

Table 13: **Convergence in consumption-equivalent welfare across UAs**

	Dependent Variable: Growth in Consumption-Equivalent Welfare $\log(\lambda_i)$				
	(1)	(2)	(3)	(4)	(5)
Initial Welfare	-0.049***	-0.043***	-0.044***	-0.048***	-0.044***
	(0.004)	(0.003)	(0.003)	(0.004)	(0.003)
Log of Population			0.002**	0.002**	0.002**
			(0.001)	(0.001)	(0.001)
KHM		0.014**	0.016**	0.028***	0.016**
		(0.007)	(0.007)	(0.005)	(0.007)
MEX		0.026***	0.027***	0.027***	0.027***
		(0.003)	(0.003)	(0.003)	(0.003)
MWI		0.065***	0.065***	0.057***	0.064***
		(0.006)	(0.007)	(0.005)	(0.007)
PER		0.044***	0.045***	0.040***	0.045***
		(0.003)	(0.003)	(0.002)	(0.003)
ZAF		0.014***	0.015***	0.016***	0.015***
		(0.003)	(0.003)	(0.003)	(0.003)
BGD		-0.016***	-0.038***	-0.040***	-0.040***
		(0.002)	(0.009)	(0.010)	(0.009)
R^2	0.464	0.761	0.765	0.796	0.765
Country FE	N	Y	Y	Y	Y
Population weights	N	N	N	Y	N
Obs.	257	257	257	257	257

Notes: Coefficients with robust standard errors in parentheses. ***p<0.01, **p<0.05, *p<0.1. Columns (1)–(4) use the compound annual growth in welfare as dependent variable. Column (5) uses the mean annual growth in welfare across periods.

component's *level* is defined relative to the country's largest UA in the first period — so initial conditions enter the right-hand side as relative positions, while the dependent variable (the mean annual change in the component) is, exactly as in Table 13, an *absolute* change in that component over time. Before reading the table, three sign conventions are worth bearing in mind. The consumption component $\Delta C = \log \bar{c}_i - \log \bar{c}_r$ is negative whenever UA i has lower demographically-adjusted average consumption than the reference, and its growth coincides with the growth rate of $\bar{c}_i$ since the reference cancels. The two inequality components, ΔC Inequality and $\Delta \ell$ Inequality, are differences in (negative) Jensen gaps relative to the reference: positive *levels* indicate *less* inequality than in the reference UA, and positive *growth* means that inequality fell in UA i. The life-expectancy component $\Delta\text{LE} = \sum_a \Delta s_a^i u_a^i$ weighs survival differentials by flow utility; because urban and rural survival curves are defined at the country level in our framework, the within-country across-UA variation in this term operates only through the flow-utility weight u_a^i, and its growth captures the joint movement of survival differentials and flow utility over time. All specifications include country fixed effects and the (log) UA population in 2015, exactly as in column 3 of Table 13.

Several findings are worth highlighting from Table 14. First, convergence is widespread across all five components. The coefficient on the initial condition is negative and highly significant in every column, ranging from -0.038 (life expectancy) to -0.065 (leisure inequality). UAs with a higher starting point for a component – i.e., closer to or above the reference UA's first-period value – grow that component more slowly thereafter. The point estimate for the consumption component (-0.043) is essentially identical to the convergence coefficient on overall welfare in Table 13 (also -0.044), as one would expect given that consumption is the largest contributor to $\log(\lambda)$. Importantly, the convergence read off these regressions is *within country* in every component: the country fixed effects absorb cross-country differences in average growth, so the slope tracks how UAs catch up to the leader in each country.

Second, the role of population is component-specific. Larger UAs grow average consumption faster ($+0.003$, significant): a one-log-point increase in initial population raises annual consumption growth by about 0.3 percentage points, somewhat larger than the population effect on overall welfare in Table 13. Population has no detectable effect on the life-expectancy, average-leisure, or leisure-inequality components – consistent with the fact that, in our framework, survival probabilities do not vary across UAs of the same urban category within a country and the leisure margin is averaged at the UA level. The point estimate on consumption inequality is negative (-0.002) but not statistically significant; if anything it points to a tendency for the consumption-inequality contribution to fall faster in larger UAs (i.e., growing consumption inequality in big cities), consistent with well-known patterns of urban primacy in inequality, but precision precludes a stronger statement.

Third, country dummies read very differently component by component, and they are the natural place to ask *which margins of welfare moved* in each country over the sample window. Because the dependent variable is an absolute change in the component, these coefficients should be read as country-specific average annual changes in the welfare contribution of consumption, life expectancy, leisure, and the two inequality terms, conditional on initial relative position and population. Consumption growth (column 1) is the central margin: five of the six countries record positive and highly significant consumption-component growth – Peru (+5.6%) and Malawi/Cambodia (+3.9%) leading, Mexico (+3.2%) intermediate, and South Africa (+1.8%) lagging. Bangladesh is the conspicuous exception: at −7.3% per year (conditional on size and starting position), the decline in average demographically-adjusted consumption in Bangladeshi UAs accounts for the bulk of the negative welfare growth documented in Table 13. The ordering of countries in column 1 mirrors – but amplifies – the ordering of overall welfare growth, confirming consumption as the dominant lever.

Fourth, the life-expectancy component (column 3) adds rather little to growth differences across UAs within countries, except for Malawi. Malawi displays a large and highly significant positive coefficient (+2.3% per year), consistent with the well-known post-AIDS recovery in mortality. South Africa, Mexico, Peru and Cambodia show coefficients close to zero (and not statistically distinguishable from it), and Bangladesh shows a small but significant positive contribution (+0.2%). The high R^2 of this regression (0.92) is largely driven by the strong convergence coefficient and by survival being modelled at the urban/rural country level: this column captures most of the variation in the component without much country-by-country nuance.

Fifth, the inequality and leisure terms reveal more subtle, sometimes offsetting, contributions. On consumption inequality (column 2), Bangladesh (+2.5%) and South Africa (+1.0%) show statistically significant welfare-improving movements (i.e., falling consumption inequality), while Cambodia (−0.4%) shows a mild deterioration. Notably, the large positive entry for Bangladesh partially offsets the heavy negative consumption-average shock in column 1: a non-trivial part of the distribution of consumption became more compressed, even as the level fell. On average leisure (column 4), every country except Malawi shows a small negative coefficient – Cambodia is the largest at −1.7% – indicating that average leisure utility declined over the sample, presumably driven by rising work hours; this margin works *against* welfare growth in those countries. On leisure inequality (column 5), the pattern reverses across countries: Bangladesh again stands out with a welfare-improving fall in leisure inequality (+1.4%), while the other countries register small but significant *increases* in leisure inequality (coefficients between −0.3% and −1.1%).

Finally, putting the country dummies of Table 14 alongside those of Table 13 makes the components-of-welfare narrative concrete. Malawi (+6.5% overall) is driven jointly by strong consumption growth and the largest life-expectancy contribution in the sample, with negligible movement in the leisure or inequality terms. Peru (+4.5%) is consumption-led, reinforced by a small but statistically significant fall in consumption inequality and dragged only modestly by a small decline in average leisure and a small rise in leisure inequality. Mexico (+2.7%) is essentially a consumption story, with marginally rising consumption inequality, slight leisure declines, and small leisure-inequality increases acting as offsets. Cambodia (+1.6%) shows consumption growth tempered by the largest decline in average leisure and rising consumption- and leisure-inequality contributions. South Africa (+1.5%) records modest consumption growth, partially aided by lower consumption inequality, but is the smallest "good news" story among the five countries with positive overall welfare growth. Bangladesh (−3.8%) is dominated by a substantial collapse in demographically-adjusted consumption (−7.3%), partially – but not fully – offset by the largest measured reductions in consumption and leisure inequality and a small life-expectancy contribution. The R^2 values, in turn, rank the components by how concentrated their variation is around a common within-country dynamic: highest for life expectancy (0.92) and consumption (0.73), and lower for the leisure components (0.48 and 0.57), where genuine UA-level idiosyncrasies remain that initial conditions, country, and size do not absorb.

Read alongside Table 13, Table 14 therefore sharpens the welfare-accounting story in two ways. Convergence in consumption-equivalent welfare is not the artifact of a single component: every margin – consumption, life expectancy, leisure, consumption inequality, and leisure inequality – exhibits robust within-country convergence, with UAs starting below the reference catching up across the board. At the same time, the cross-country heterogeneity uncovered in Table 13 reflects very different combinations of these margins: consumption-led growth almost everywhere (Bangladesh being the exception, where consumption fell sharply), an exceptional life-expectancy contribution in Malawi, mild leisure declines in most countries, and a mixed but quantitatively secondary role for the inequality terms. Because each component is expressed in the same units of consumption-equivalent welfare relative to the country's reference UA, the contributions can be read additively, and they reconcile with the overall welfare growth rates of Table 13 – by construction rather than as a separate empirical claim.

Table 14: **Decomposition of mean growth in welfare components across UAs**

	Dependent Variable: Mean Annual Growth in Welfare Components				
	Consumption Average	Consumption Inequality	Life Expectancy	Leisure Average	Leisure Inequality
	$\log \bar{c}_i - \log \bar{c}_r$	$(E[\log c_i] - \log \bar{c}_i) - (E[\log c_r] - \log \bar{c}_r)$	$\sum_a \Delta s_a^i u_a^i$	$\nu(\bar{\ell}_i) - \nu(\bar{\ell}_r)$	$\big(E[\nu(\ell_i)] - \nu(\bar{\ell}_i)\big) - \big(E[\nu(\ell_r)] - \nu(\bar{\ell}_r)\big)$
	(1)	(2)	(3)	(4)	(5)
Initial Condition	-0.043***	-0.057***	-0.038***	-0.054***	-0.065***
	(0.003)	(0.004)	(0.002)	(0.007)	(0.006)
Log of population	0.003***	-0.002	0.000	0.000	-0.000
	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
KHM	0.039***	-0.004**	0.006	-0.017***	-0.011***
	(0.010)	(0.002)	(0.000)	(0.005)	(0.001)
MEX	0.032***	0.004***	-0.004	-0.002**	-0.007***
	(0.003)	(0.001)	(0.000)	(0.001)	(0.001)
MWI	0.039***	0.002	0.023***	0.000	-0.003***
	(0.007)	(0.002)	(0.001)	(0.001)	(0.001)
PER	0.056***	0.004***	-0.008	-0.005***	-0.004***
	(0.003)	(0.001)	(0.000)	(0.001)	(0.001)
ZAF	0.018***	0.010***	-0.005	-0.005***	-0.004***
	(0.004)	(0.002)	(0.000)	(0.001)	(0.001)
BGD	-0.073***	0.025***	0.002**	-0.002	0.014***
	(0.011)	(0.005)	(0.001)	(0.003)	(0.004)
Mean Dep. Variable	0.014	0.003	-0.003	0.000	0.003
Mean Initial Condition	-0.402	0.081	0.023	0.000	0.017
R^2	0.728	0.656	0.923	0.479	0.574
Country FE	Y	Y	Y	Y	Y
Obs.	257	257	257	257	257

Notes: Coefficients with robust standard errors in parentheses. ***p<0.01, **p<0.05, *p<0.1. Initial condition refers to the level of the welfare component in the first period of the available data.

4.4.1 Welfare in urban areas and their hinterlands

We now compare the welfare of each urban area with the welfare of its corresponding hinterland, and ask whether UA–hinterland gaps narrowed over time. To allow within-country comparisons, welfare levels are normalized by the main city of each country in the base year. As discussed above, the survival rates assigned to each UA correspond to the region's urban category, while those assigned to the hinterland correspond to the region's rural category. We were able to construct relative welfare measures for 236 UA–hinterland pairs across our 12 country–year observations, and a balanced panel of 161 pairs for which both waves are available.

Before turning to the regression analysis, Figure 12 shows the distributions of hinterland consumption in 2017 PPP-adjusted US dollars (Panel A) and the ratio of hinterland to UA consumption (Panel B), by five distance rings (0–10 km, 10–30 km, 30–50 km, 50–100 km, and above 100 km). Consumption tends to be lower on average in close hinterlands but with lower variance than in distant hinterlands; the consumption ratio to the closest UA, on the other hand, tends to be closer to one for close hinterlands, meaning they are similar to the urbanized area on average (although in approximately 80% of the cases close hinterlands are still worse off than their matched urban area). Figure 13 reveals substantial heterogeneity in UA–hinterland welfare across countries and years. In this figure, each point represents a single UA and its hinterland. Within a country, blue and red points identify the first and last year in our data for that country, respectively. A scatter plot concentrated around the 45-degree line would indicate that the welfare gap between UAs and their rural hinterlands is not substantial. Points below the line indicate rural hinterlands that are worse off than their matched UA, and points that are up and to the right indicate an urban-hinterland pair with a better relative welfare status than the country's reference city. Additionally, a scatterplot with a slope that is less than one (as we find in all countries) and more than zero (everywhere but South Africa in 1993 and Cambodia in 2009) indicates that the rural hinterlands of higher welfare cities in turn have higher welfare, but rural welfare grows less than proportionately so that the urban-hinterland gap is higher for higher welfare cities.

To test the gap-closing pattern formally, we estimate a β-convergence regression at the pair level. Let pair p denote a unique UA–hinterland pair (i, h_i) in country c, observed at a baseline year t_0 and a later year t_1. Define the welfare gap of pair p in year t as

$$G_{p,t} \equiv W_{i(c)t} - W_{h_i(c)t},$$

and its temporal change as $\Delta G_p \equiv G_{p,t_1} - G_{p,t_0}$. Pairs whose UA–hinterland gap narrows over time have $\Delta G_p < 0$. We estimate

$$\Delta G_p \;=\; \beta_0 \;+\; \beta_1\, G_{p,t_0} \;+\; \beta_2 \log(\text{Population})_i \;+\; \delta_c \;+\; \varepsilon_p, \tag{11}$$

where δ_c are country fixed effects and the unit of observation is the UA–hinterland pair (one observation per pair, $N = 161$).[31] The coefficient of interest is β_1. A negative β_1 indicates *β-convergence of UA–hinterland gaps*: pairs that started further apart closed their gap by more between t_0 and t_1. The coefficient β_2 then asks whether, conditional on the initial gap, larger UAs converged faster or slower with their own hinterlands. Standard errors are clustered at the pair level.[32] We re-estimate equation (11) component-by-component using each entry of the welfare decomposition in equation (6) as the dependent variable (and the same component's initial gap as the regressor), exactly parallel to the structure of Table 14.

Several findings stand out from Table 15. First, every component of welfare shows strong β-convergence between UAs and their hinterlands. The slope on the initial gap is negative and highly significant in all six columns, ranging from -0.655 (average leisure) to -1.023 (leisure inequality), with the overall welfare gap converging at -0.961 and demographically-adjusted consumption at -0.879. Pairs that began the sample with a larger gap in any given component closed nearly all of that gap by the final wave. The point estimate for life expectancy (-0.949) is very close to the overall-welfare slope. We note, however, that the within-country across-UA variation in the life-expectancy component is, by methodology, limited: UAs in the same region share the same urban survival curve, so the LE component's convergence in this table is driven primarily by cross-region rather than cross-UA differences within region.

Second, larger UAs do *not* see faster hinterland convergence. The coefficient on log population in the overall welfare-gap column is essentially zero. For the consumption component the coefficient is positive but imprecisely estimated; for average leisure it is positive and marginally significant (0.018, $p < 0.1$). The one component for which size matters meaningfully is consumption inequality (-0.031, $p < 0.01$): conditional on the initial consumption-inequality gap, larger UAs saw their consumption-inequality gap with the hinterland narrow more strongly. This is consistent with rising consumption inequality at the top of the urban hierarchy bringing UA-level inequality closer to (typically higher) hinterland inequality, but the sign is the opposite of what one would

[31] Because the unit of observation is the pair-level temporal difference, country fixed effects and country-year fixed effects coincide in this specification: each pair has a single $t_0 \to t_1$ window pinned down by its country.

[32] This specification is the panel-differenced analog of the convergence regression we run across UAs in equation (10).

predict if larger UAs were systematically pulling away from their rural surroundings; the population coefficient on the welfare gap itself does not support that "divergence-with-size" narrative.

Third, the fit varies sharply across components in a way that is informative about which margins of the UA–hinterland gap move coherently within countries. The adjusted R^2 is 0.81 for overall welfare, 0.92 for life expectancy, and 0.72 for average consumption – all reflecting tight within-country convergence patterns – but only 0.20 for average leisure and 0.36–0.41 for the two inequality components. The leisure margin in particular leaves substantial idiosyncratic pair-level variation that the initial gap, size, and country do not absorb.

Two caveats on the magnitudes are worth mentioning. First, slopes close to -1 deserve a Galton's-fallacy caveat: classical measurement error in G_{p,t_0} biases $\hat{\beta}_1$ toward -1, and with only two survey waves per pair we cannot instrument the initial gap. Two pieces of evidence make us read the convergence as substantive, rather than purely mechanical: (i) the slope is fairly uniform across components with very different measurement-error structures (life expectancy is constructed from administrative life-table inputs, consumption from household survey microdata), and (ii) the R^2 of the welfare-gap regression is high (0.81), so the initial gap genuinely explains most of the variation in subsequent gap-closing rather than the slope being an attenuation artifact. Second, the dependent variable in equation (11) is the *total* change in the gap, not the annualized change, so the magnitudes in Table 15 are not directly comparable to the per-year convergence speeds in Table 13.

In sum, equation (11) isolates the temporal dimension of the UA-hinterland convergence question. Read this way, Table 15 delivers a clean message: across all five welfare components, UA–hinterland gaps narrowed in pairs that began further apart, and they did so at broadly similar rates across components. The size of the UA, conditional on the initial gap, does not predict the speed of overall welfare convergence with the hinterland — the one exception being consumption inequality, where larger UAs saw their inequality gap with the hinterland close faster. Combined with the descriptive patterns of Figure 13 (red dots above blue in five of six countries), the data are consistent with a widespread, component-balanced narrowing of UA–hinterland welfare differences during our sample window, with Cambodia as the only country in which the descriptive picture does not move toward convergence.

Figure 12: **Consumption in hinterlands by distance from urban area**

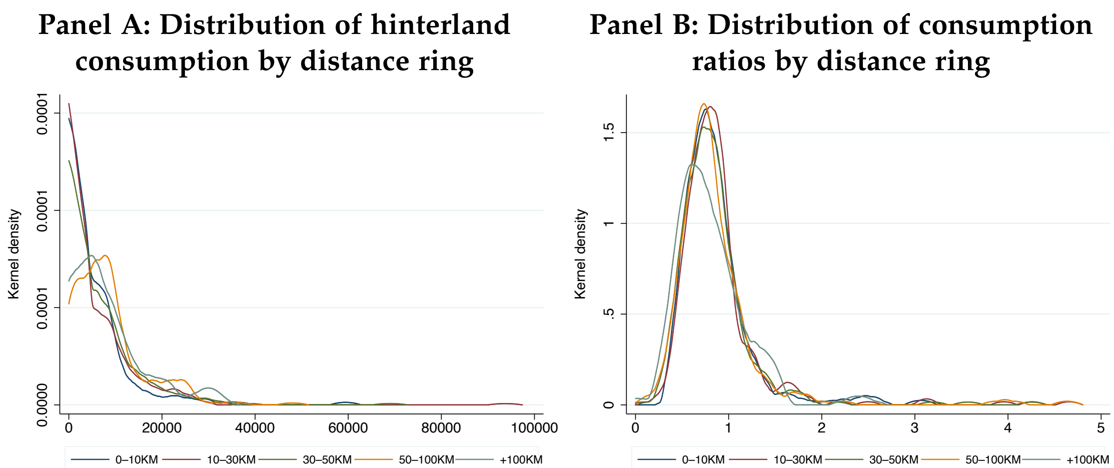

Notes: These panels show the distributions of hinterland consumption in 2017 PPP-adjusted US dollars (Panel A) and its relative value with respect to its closest UA (Panel B) by five distance rings (0–10 km, 10–30 km, 30–50 km, 50–100 km, and above 100 km).

Figure 13: **Urban area vs. hinterland**

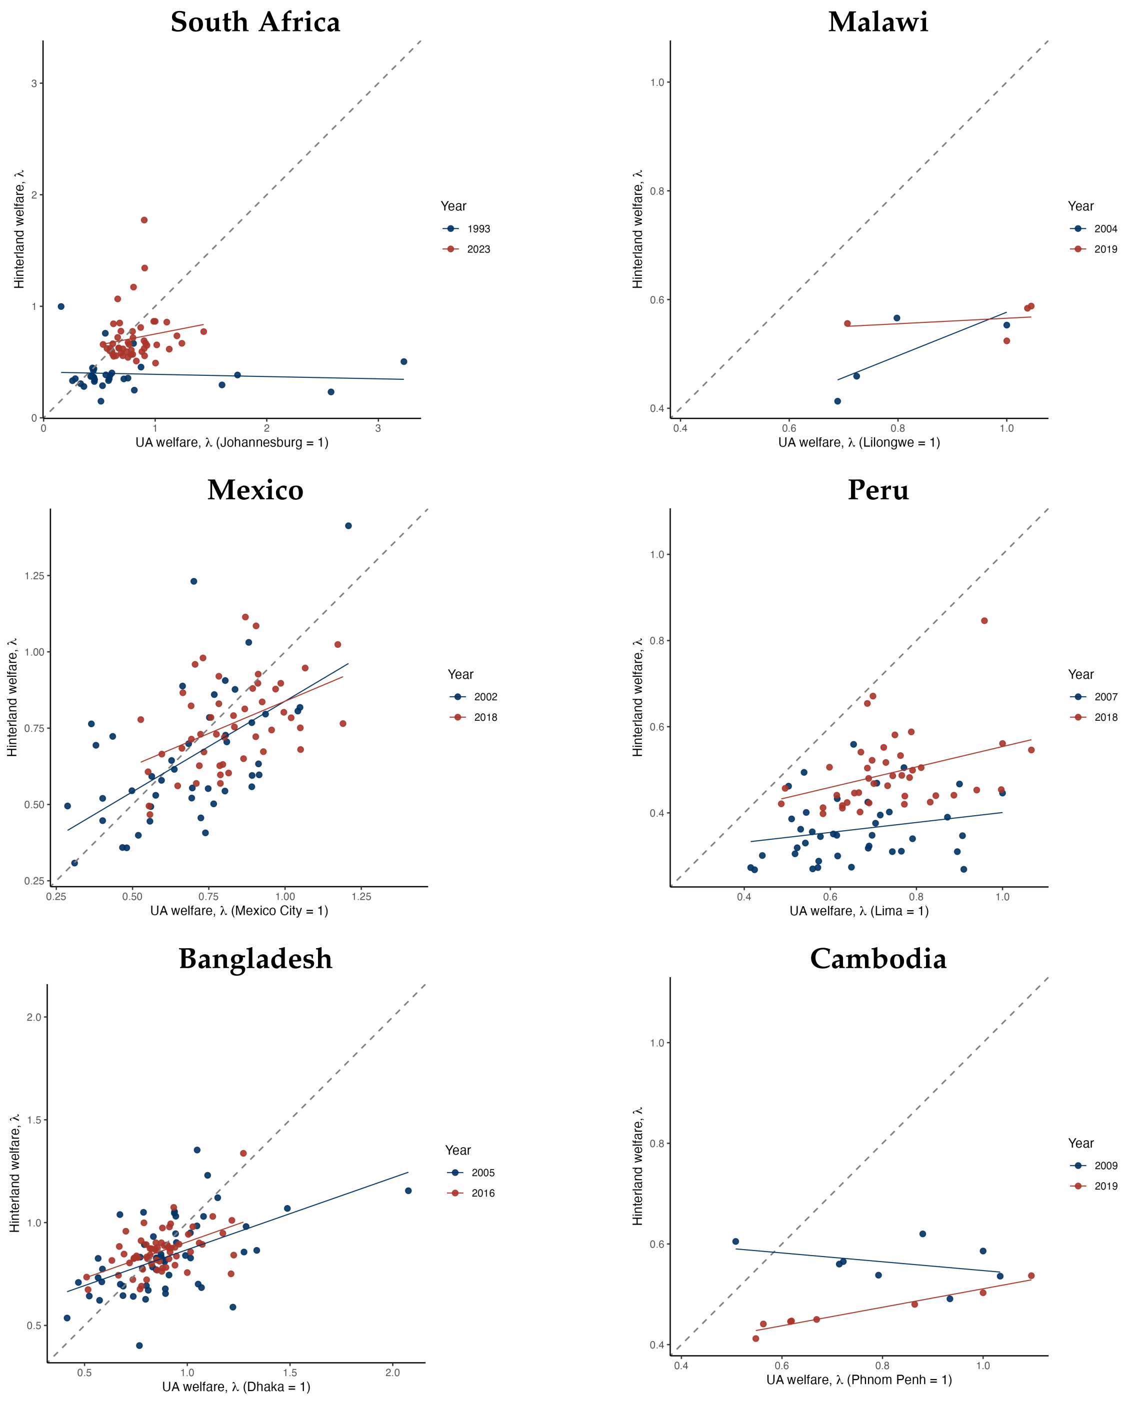

Notes: These panels compare the welfare level of each UA with that of its corresponding hinterland. One scatter plot is presented per country, where each point represents a single UA and its hinterland. Welfare levels are normalized by the main city of each country, allowing for consistent comparisons within country.

Table 15: **Difference in welfare and components between UAs and hinterlands**

	△ Welfare Gap	△ Consumption Gap	△ Cons. inequality Gap
	(1)	(2)	(3)
Initial Gap	-0.961***	-0.879***	-0.713***
	(0.033)	(0.052)	(0.192)
Log of Population	-0.001	0.017	-0.031***
	(0.013)	(0.016)	(0.012)
Adj. R^2	0.812	0.721	0.357
	△ Life expectancy Gap	△ Leisure Gap	△ Leisure inequality Gap
	(4)	(5)	(6)
Initial Gap	-0.949***	-0.655**	-1.023***
	(0.048)	(0.256)	(0.093)
Log of Population	-0.000	0.018*	-0.005
	(0.001)	(0.009)	(0.010)
Adj. R^2	0.922	0.203	0.408
Country-year FE	Y	Y	Y
Obs.	161	161	161
Number of UAs	161	161	161

Notes: Coefficients with standard errors clustered at the UA level in parentheses. All specifications include log population and country–year fixed effects. The hinterland component corresponds to the same welfare component accounted for in the computation of the welfare gap (dependent variable), measured in the hinterland. ***p< 0.01, **p< 0.05, *p< 0.1.

5 Conclusion

This paper extends Jones and Klenow (2016)'s framework for consumption-equivalent welfare to the subnational level, leveraging rich, geocoded household data to shed light on how welfare varies over time and across space within six economies at varying degrees of development. By examining three distinct dimensions —within-country regional disparities, urban–rural contrasts, and the dynamics of urban transitions between urban areas and their rural hinterlands—, we document the empirical distribution of welfare that goes beyond traditional income or poverty measures. The aim is not only to document disparities but to understand their structural drivers and how they evolve in the presence of demographic, spatial, and economic development heterogeneity.

Across these dimensions, our analysis reveals a consistent pattern: welfare converges across regions within countries, even as large gaps persist between the most and least developed areas. This convergence is nuanced by country-specific trajectories, with notable heterogeneity in how population size, urbanization, and the distribution of consumption and leisure and disparities in life expectancies may influence welfare dynamics. Importantly, the subnational focus highlights that aggregated national indicators can obscure meaningful local patterns; in several contexts, areas undergoing urban transitions or experiences of rapid urban growth exhibit distinct welfare paths from established urban centers. These findings have direct implications for policy design, suggesting that targeted, geographically aware interventions — paired with improvements in data quality and harmonization — are essential to promote inclusive growth.

Looking forward, this study opens several avenues for further research. Extending the analysis to additional countries would test the generalizability of the observed convergence patterns and illuminate how different development trajectories and stages in the urbanization process shape subnational welfare. Incorporating richer measures of well-being — such as health, education (or access to other public goods), and environmental quality — could help us go further beyond GDP in accounting for welfare.[33] Finally, a deeper exploration of causal mechanisms — for example, through the roll-out of infrastructure, governance, or social protection programs — would help translate the observed associations into transferable policy insights aimed at reducing regional disparities and supporting sustainable urban transitions. Our approaches to

[33] Note that this welfare measure does not directly account for access to roads, sewers, electricity, schooling, etc. Some forms of private or public consumption - the latter being evenly imputed in per capita terms to each household in our approach - may reflect these indirectly.

attributing welfare to urban areas provide a promising lens to investigate these issues. Further, adding consistent urbanization measures to NSO definitions is valuable to facilitate comparisons of specifically urban or rural outcomes across countries. To conclude, our framework is useful for understanding spatial trade-offs. While a growing literature rationalizes subnational residential location choices in response to wages (see Redding and Rossi-Hansberg, 2017; Redding, 2024), migrants may trade off broader measures of welfare with subnational components that can be modeled to change over time.

References

Aghion, P., Akcigit, U., Deaton, A. and Roulet, A. (2016). Creative Destruction and Subjective Well-Being. *American Economic Review*, **106** (12), 3869–3897.

Alkire, S., Foster, J., Seth, S., Santos, M. E., Roche, J. M. and Ballon, P. (2015). Chapter 5. The Alkire-Foster counting methodology. In *Multidimensional Poverty Measurement and Analysis*, Oxford University Press, pp. 144–185.

Attanasio, O. and Browning, M. (1993). *Consumption over the Life Cycle and over the Business Cycle*. Tech. Rep. w4453, National Bureau of Economic Research, Cambridge, MA.

Bairoliya, N., Chanda, A., Louisiana State University, Fang, J. and Yang, F. (2025). Household Consumption and Savings over the Life Cycle: The Roles of Demographics and Durables. *Federal Reserve Bank of Dallas, Working Papers*, **2025** (2537).

Baland, J.-M., Cassan, G. and Decerf, B. (2021). The Poverty-Adjusted Life Expectancy index: a consistent aggregation of the quantity and the quality of life.

Bannister, G. J. and Mourmouras, A. (2017). Welfare vs. Income Convergence and Environmental Externalities. *IMF Working Papers*, **2017** (271), iSBN: 9781484331514.

Barro, R. J. (1991). Economic growth in a cross section of countries. *Quarterly Journal of Economics*, **106** (2), 407–443.

Bick, A., Fuchs-Schündeln, N. and Lagakos, D. (2018). How Do Hours Worked Vary with Income? Cross-Country Evidence and Implications. *American Economic Review*, **108** (1), 170–199.

Bilal, U., Alazraqui, M., Caiaffa, W. T., Lopez-Olmedo, N., Martinez-Folgar, K., Miranda, J. J., Rodriguez, D. A., Vives, A. and Diez-Roux, A. V. (2019). Inequalities in life expectancy in six large Latin American cities from the SALURBAL study: an ecological analysis. *The Lancet. Planetary Health*, **3** (12), e503–e510.

Boarini, R., Fleurbaey, M., Murtin, F. and Schreyer, P. (2022). Well-being during the Great Recession: new evidence from a measure of multi-dimensional living standards with heterogeneous preferences*. *The Scandinavian Journal of Economics*, **124** (1), 104–138.

Brouillette, J.-F., Jones, C. I. and Klenow, P. J. (2025). Race and Economic Well-Being in the United States. *American Economic Review: Insights*, **7** (4), 429–446.

Bryan, G. and Morten, M. (2019). The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia. *Journal of Political Economy*, **127** (5), 2229–2268.

Charron, B. and Wadhwa, D. (2025). From water to electricity: The rural-urban infrastructure gap in seven charts.

Chatterjee, S., Giannone, E., Kleineberg, T., Kuno, K. and Looser, L. (2025). Unequal Global Convergence.

Decerf, B. (2024). *Multidimensional Well-Being Measurement Practices: A Review Focused on Improving Global Multidimensional Poverty Indicators*. Tech. rep., The World Bank, Washington, DC.

Dong, L., Duarte, F., Duranton, G., Santi, P., Barthelemy, M., Batty, M., Bettencourt, L., Goodchild, M., Hack, G., Liu, Y., Pumain, D., Shi, W., Verbavatz, V., West, G. B., Yeh, A. G. O. and Ratti, C. (2024). Defining a city — delineating urban areas using cell-phone data. *Nature Cities*, **1** (2), 117–125.

Döpke, J., Knabe, A., Lang, C. and Maschke, P. (2017). Multidimensional Well-being and Regional Disparities in Europe: Well-being and Regional Disparities in Europe. *JCMS: Journal of Common Market Studies*, **55** (5), 1026–1044.

Ellis, P. and Roberts, M. (2016). *Leveraging Urbanization in South Asia: Managing Spatial Transformation for Prosperity and Livability*. Washington, DC: The World Bank.

European Commission. Joint Research Centre. (2019). GHS-FUA R2019A - GHS functional urban areas, derived from GHS-UCDB R2019A, (2015), R2019A.

Falcettoni, E. and Nygaard, V. M. (2023). A Comparison of Living Standards Across the United States of America. *International Economic Review*, **64** (2), 511–542.

Ferré, C., Ferreira, F. H. and Lanjouw, P. (2012). Is There a Metropolitan Bias? The relationship between poverty and city size in a selection of developing countries. *The World Bank Economic Review*, **26** (3), 351–382.

Gallardo-Albarrán, D. (2019). Missed opportunities? Human welfare in Western Europe and the United States, 1913–1950. *Explorations in Economic History*, **72**, 57–73.

Gollin, D., Kirchberger, M. and Lagakos, D. (2021). Do urban wage premia reflect lower amenities? Evidence from Africa. *Journal of Urban Economics*, **121**, 103301.

Hurst, E., Li, G. and Pugsley, B. (2014). Are Household Surveys Like Tax Forms? Evidence from Income Underreporting of the Self-Employed. *The Review of Economics and Statistics*, **96** (1), 19–33.

Jones, C. I. and Klenow, P. J. (2016). Beyond GDP? Welfare across Countries and Time. *American Economic Review*, **106** (9), 2426–2457.

Jorgenson, D. W. (2018). Production and Welfare: Progress in Economic Measurement. *Journal of Economic Literature*, **56** (3), 867–919.

Kline, P. and Moretti, E. (2014). People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs. *Annual Review of Economics*, **6** (Volume 6, 2014), 629–662.

Lagakos, D. (2020). Urban-Rural Gaps in the Developing World: Does Internal Migration Offer Opportunities? *Journal of Economic Perspectives*, **34** (3), 174–192.

Mankiw, N. G., Romer, D. and Weil, D. N. (1992). A contribution to the empirics of economic growth. *Quarterly Journal of Economics*, **107** (2), 409–439.

Nakamura, S., Combes, P.-P., Moellerherm, R., Robert, C., Roberts, M., Stewart, B. and Yakubenko, S. (2023). *Where Is Poverty Concentrated?: New Evidence Based on Internationally Consistent Urban and Poverty Measurements*. World Bank, Washington, DC.

Okoye, D. and Pandey, S. (2024). Welfare comparisons within countries beyond GDP: An application to Nigeria. *Economic Modelling*, **140**, 106831.

RAVALLION, M. (2012). Mashup Indices of Development. *The World Bank Research Observer*, **27** (1), 1–32.

REDDING, S. J. (2024). Spatial Economics.

— and ROSSI-HANSBERG, E. (2017). Quantitative Spatial Economics. p. 40.

SCHIAVINA, M., FREIRE, S. and MACMANUS, K. (2023a). GHS-POP R2023A - GHS population grid multitemporal (1975-2030).

—, MELCHIORRI, M. and PESARESI, M. (2023b). GHS-SMOD R2023A - GHS settlement layers, application of the Degree of Urbanisation methodology (stage I) to GHS-POP R2023A and GHS-BUILT-S R2023A, multitemporal (1975-2030).

SCHWARTZ, A. L. and SOMMERS, B. D. (2014). Moving For Medicaid? Recent Eligibility Expansions Did Not Induce Migration From Other States. *Health Affairs*, **33** (1), 88–94.

SELOD, H. and SHILPI, F. (2021). Rural-urban migration in developing countries: Lessons from the literature. *Regional Science and Urban Economics*, **91**, 103713.

STIGLITZ, J. E., FITOUSSI, J.-P. and DURAND, M. (2018). *Beyond GDP: Measuring What Counts for Economic and Social Performance*. OECD.

STIGLITZ, P. J. E., SEN, P. A. and FITOUSSI, P. J.-P. (2009). Report by the Commission on the Measurement of Economic Performance and Social Progress.

UNDP (2013). *The Rise of the South: Human Progress in a Diverse World*. Human Development Report, New York, NY: United Nations Development Programme.

WIJNEN, W., DAHDAH, S. and PKHIKIDZE, N. (2025). The value of a statistical life in the context of road safety: a new value transfer approach. *Traffic Injury Prevention*, pp. 1–8.

WORLD BANK (2024a). *Poverty, Prosperity, and Planet Report 2024: Pathways Out of the Polycrisis*. Washington, DC: World Bank.

WORLD BANK (2024b). *World Development Report 2024: The Middle-Income Trap*. Washington, DC: World Bank.

YEH, C., PEREZ, A., DRISCOLL, A., AZZARI, G., TANG, Z., LOBELL, D., ERMON, S. and BURKE, M. (2020). Using publicly available satellite imagery and deep learning to understand economic well-being in Africa. *Nature Communications*, **11** (1), 2583.

ZHENG, Z., WU, T., LEE, R., NEWHOUSE, D., KILIC, T., BURKE, M., ERMON, S. and LOBELL, D. B. (2026). Dynamic, high-resolution poverty measurement in data-scarce environments. *Journal of Development Economics*, **179**, 103691.

Table A.1: **Robustness Results for Regional Differences in South Africa**

	Western Cape	Eastern Cape	Northern Cape	Free State	KwaZulu-Natal	North West	Mpumalanga	Limpopo
Panel A: 1993								
Benchmark	113.5	44.1	53.2	50.7	55.1	51.3	51.6	47.5
No discounting/growth	111.1	46.2	53.4	51.1	57.0	50.4	51.3	47.4
Discounting: $\beta = 0.96$	106.3	42.5	53.3	50.2	54.7	52.8	54.7	47.7
Household size (sqrt)	115.5	46.9	55.0	51.0	61.3	51.8	54.8	51.4
OECD Scale	113.6	46.0	55.4	51.7	56.7	52.3	52.8	49.4
Ages 2 and above	108.8	45.4	53.6	50.8	56.7	51.2	52.3	47.7
Ages 25 and above	107.6	50.5	56.2	51.8	59.9	49.8	50.0	48.2
Ages 40 and above	119.9	48.2	54.0	52.2	57.6	47.4	44.8	46.0
Kids get adult leisure	113.4	50.4	58.0	51.6	59.6	52.1	52.4	53.0
Frisch elasticity = 0.5	112.6	38.7	48.1	50.2	51.9	51.2	51.3	43.0
Frisch elasticity = 2.0	114.2	53.3	61.9	51.9	60.5	52.2	52.3	55.7
Value of life 25% smaller	110.6	45.2	53.6	50.9	56.3	51.3	52.3	47.7
Value of life 25% larger	116.4	43.0	52.9	50.5	54.0	51.2	50.9	47.3
Panel B: 2023								
Benchmark	100.2	72.9	79.2	75.3	73.3	66.0	67.5	70.9
No discounting/growth	101.2	75.1	80.1	77.6	72.9	68.9	69.6	67.1
Discounting: $\beta = 0.96$	98.5	71.9	77.8	75.6	72.0	64.6	69.3	73.3
Household size (sqrt)	103.2	75.2	81.5	75.0	78.2	66.2	70.2	74.4
OECD Scale	100.3	74.2	80.2	76.3	73.8	67.1	68.7	71.7
Ages 2 and above	100.1	74.0	79.9	76.4	73.3	67.0	68.6	70.2
Ages 25 and above	100.6	76.5	82.0	76.9	74.4	69.6	67.1	67.3
Ages 40 and above	108.7	75.2	82.4	74.7	76.0	69.3	64.6	64.2
Kids get adult leisure	99.9	92.6	87.8	84.3	77.5	70.2	73.9	80.8
Frisch elasticity = 0.5	104.7	69.3	76.7	73.2	71.3	63.2	66.4	69.7
Frisch elasticity = 2.0	93.4	78.7	82.8	78.6	76.6	70.9	69.5	73.1
Value of life 25% smaller	99.9	74.1	79.5	76.6	73.0	67.3	69.2	69.9
Value of life 25% larger	100.5	71.6	78.8	74.1	73.5	64.8	65.9	71.9

Table A.2: **Robustness Results for Regional Differences in Malawi**

	North	South
Panel A: 2004		
Benchmark	100.5	80.7
No discounting/growth	96.6	80.3
Discounting: $\beta = 0.96$	99.2	81.6
Household size (sqrt)	102.4	78.0
OECD Scale	100.4	80.5
Ages 2 and above	97.3	80.6
Ages 25 and above	93.8	79.0
Ages 40 and above	95.3	79.0
Kids get adult leisure	100.0	79.2
Frisch elasticity = 0.5	98.6	82.4
Frisch elasticity = 2.0	103.4	79.2
Value of life 25% smaller	98.9	80.7
Value of life 25% larger	102.1	80.7
Panel B: 2019		
Benchmark	141.3	108.7
No discounting/growth	140.0	109.7
Discounting: $\beta = 0.96$	138.0	108.0
Household size (sqrt)	142.6	106.4
OECD Scale	140.9	109.5
Ages 2 and above	140.0	109.2
Ages 25 and above	140.2	110.3
Ages 40 and above	141.9	111.9
Kids get adult leisure	140.9	108.7
Frisch elasticity = 0.5	140.7	108.4
Frisch elasticity = 2.0	142.3	109.4
Value of life 25% smaller	140.8	108.8
Value of life 25% larger	141.9	108.5

Table A.3: **Robustness Results for Regional Differences in Peru**

	North	Central	South	East	North East
		Panel A: 2007			
Benchmark	58.8	49.6	68.2	52.5	55.3
No discounting/growth	61.0	50.8	68.4	51.8	55.2
Discounting: $\beta = 0.96$	60.1	50.6	72.1	54.9	57.7
Household size (sqrt)	58.9	49.2	66.3	50.6	56.8
OECD Scale	60.2	51.0	69.2	53.7	56.7
Ages 2 and above	60.2	50.4	69.3	52.8	55.8
Ages 25 and above	62.6	51.7	67.7	52.1	55.4
Ages 40 and above	58.2	47.2	61.4	47.3	50.2
Kids get adult leisure	61.3	51.2	69.4	52.6	56.1
Frisch elasticity = 0.5	58.6	49.6	68.7	53.6	56.2
Frisch elasticity = 2.0	57.5	47.6	66.5	48.8	52.6
Value of life 25% smaller	59.8	50.2	68.9	52.8	55.7
Value of life 25% larger	57.9	48.9	67.5	52.3	54.9
		Panel B: 2018			
Benchmark	71.7	63.0	77.7	72.5	75.1
No discounting/growth	72.7	63.7	77.1	71.3	74.8
Discounting: $\beta = 0.96$	74.2	65.5	83.2	74.7	76.4
Household size (sqrt)	71.4	61.4	75.1	69.8	76.4
OECD Scale	72.9	64.3	78.3	73.2	76.8
Ages 2 and above	72.7	63.8	78.3	72.5	75.4
Ages 25 and above	72.4	63.1	73.8	71.6	75.7
Ages 40 and above	69.1	59.6	69.4	66.3	72.5
Kids get adult leisure	73.8	63.7	76.4	72.1	77.8
Frisch elasticity = 0.5	71.8	63.9	80.6	73.5	75.2
Frisch elasticity = 2.0	70.7	60.0	74.1	68.7	73.9
Value of life 25% smaller	72.5	63.7	78.2	72.6	75.3
Value of life 25% larger	71.0	62.3	77.1	72.4	74.9

Table A.4: **Robustness Results for Regional Differences in Mexico**

	Northwest	North	Northeast	West-Central	South	East	Yucatán Peninsula
Panel A: 2002							
Benchmark	120.9	91.1	100.2	92.2	59.8	74.2	76.3
No discounting/growth	118.9	92.4	99.3	93.3	60.2	74.5	77.1
Discounting: $\beta = 0.96$	120.1	91.2	99.3	91.6	61.8	76.5	75.5
Household size (sqrt)	118.7	89.3	97.4	91.5	59.3	72.4	75.0
OECD Scale	121.6	92.8	100.1	93.5	61.4	75.0	77.3
Ages 2 and above	119.2	91.6	98.7	92.0	60.0	74.2	75.9
Ages 25 and above	120.0	92.6	100.4	94.1	57.7	72.5	76.2
Ages 40 and above	119.0	94.2	103.8	95.2	57.7	72.3	79.5
Kids get adult leisure	122.3	94.0	101.9	94.4	59.9	75.5	75.9
Frisch elasticity = 0.5	120.3	89.3	99.7	93.2	62.1	74.7	79.1
Frisch elasticity = 2.0	122.3	93.8	100.6	91.2	56.8	74.1	72.8
Value of life 25% smaller	119.8	91.7	99.3	92.3	60.2	74.6	76.2
Value of life 25% larger	121.9	90.5	101.0	92.1	59.4	73.7	76.4
Panel B: 2018							
Benchmark	113.4	95.9	109.1	99.8	72.6	83.4	96.9
No discounting/growth	113.2	96.1	109.1	99.5	72.8	83.1	96.0
Discounting: $\beta = 0.96$	112.3	95.6	107.3	100.6	73.3	84.2	98.9
Household size (sqrt)	111.2	94.8	107.3	100.2	72.5	82.1	95.9
OECD Scale	113.4	96.5	109.2	100.3	73.6	84.0	96.5
Ages 2 and above	112.8	95.8	108.7	99.5	72.6	83.2	96.4
Ages 25 and above	113.2	96.1	109.7	98.9	72.5	82.3	93.6
Ages 40 and above	115.9	98.0	112.3	100.5	72.9	83.5	94.9
Kids get adult leisure	112.0	95.6	110.0	99.3	73.1	83.5	94.1
Frisch elasticity = 0.5	113.8	96.2	106.7	102.1	72.9	84.1	101.1
Frisch elasticity = 2.0	112.0	94.4	112.3	95.7	71.2	82.0	91.3
Value of life 25% smaller	113.1	95.8	108.8	99.7	72.7	83.3	96.7
Value of life 25% larger	113.8	95.9	109.4	99.9	72.5	83.5	97.1

Table A.5: **Robustness Results for Regional Differences in Bangladesh**

	Barisal	Chattogram	Khulna	Mymensingh	Rajshahi	Rangpur	Sylhet
Panel A: 2005							
Benchmark	81.6	94.7	90.9	82.8	75.8	59.4	77.9
No discounting/growth	82.7	98.0	91.3	87.4	77.0	59.4	82.8
Discounting: $\beta = 0.96$	79.7	92.2	88.9	80.0	74.2	60.3	77.7
Household size (sqrt)	83.3	101.0	91.2	82.0	73.9	58.6	83.9
OECD Scale	83.5	98.0	88.6	82.8	72.9	58.5	79.0
Ages 2 and above	81.9	96.4	90.8	85.0	76.2	59.3	80.1
Ages 25 and above	86.0	103.3	94.7	93.5	79.0	59.8	85.0
Ages 40 and above	84.8	100.8	90.6	91.3	78.2	56.9	83.9
Kids get adult leisure	84.3	97.5	92.7	86.0	76.5	60.0	79.4
Frisch elasticity = 0.5	76.3	90.5	84.9	78.4	72.9	62.0	78.7
Frisch elasticity = 2.0	88.0	99.6	97.8	86.5	78.1	55.9	76.8
Value of life 25% smaller	81.8	95.3	90.8	83.6	76.0	59.4	78.9
Value of life 25% larger	81.5	94.2	90.9	81.9	75.7	59.3	77.0
Panel B: 2016							
Benchmark	88.0	108.2	89.8	75.5	87.3	74.6	90.9
No discounting/growth	89.2	110.7	89.8	77.3	88.7	75.2	94.6
Discounting: $\beta = 0.96$	85.5	104.8	87.0	74.4	84.7	72.3	87.7
Household size (sqrt)	90.7	115.4	88.5	75.1	86.1	74.2	101.5
OECD Scale	88.0	109.1	86.6	75.6	85.0	73.7	90.7
Ages 2 and above	88.6	109.8	90.2	76.4	88.2	74.9	93.0
Ages 25 and above	91.2	115.5	93.9	78.6	92.1	78.5	99.5
Ages 40 and above	90.9	112.7	89.9	77.8	89.7	77.3	97.0
Kids get adult leisure	90.6	112.9	92.4	77.0	89.5	76.2	95.4
Frisch elasticity = 0.5	79.6	97.7	81.4	71.6	79.9	68.4	80.8
Frisch elasticity = 2.0	98.7	122.0	99.3	79.4	95.1	80.7	103.7
Value of life 25% smaller	88.1	108.7	89.8	75.9	87.6	74.7	91.6
Value of life 25% larger	87.8	107.8	89.8	75.1	87.1	74.5	90.1

Table A.6: **Robustness Results for Regional Differences in Cambodia**

	Plain	Tonle Sap	Coast	Plateau/Mountain
Panel A: 2009				
Benchmark	49.3	49.9	53.7	45.6
No discounting/growth	53.4	53.8	57.8	50.5
Discounting: $\beta = 0.96$	48.7	48.8	53.0	44.3
Household size (sqrt)	47.3	48.1	51.5	44.0
OECD Scale	51.0	51.5	55.5	47.4
Ages 2 and above	50.3	51.0	54.9	46.9
Ages 25 and above	58.1	59.2	62.9	56.7
Ages 40 and above	55.0	54.4	57.6	51.8
Kids get adult leisure	56.4	58.2	61.1	54.3
Frisch elasticity = 0.5	44.2	43.2	47.1	39.5
Frisch elasticity = 2.0	53.8	57.0	60.1	51.4
Value of life 25% smaller	50.0	50.5	54.4	46.4
Value of life 25% larger	48.6	49.3	53.0	44.7
Panel B: 2019				
Benchmark	52.8	49.9	61.3	50.6
No discounting/growth	54.6	51.0	62.6	52.2
Discounting: $\beta = 0.96$	52.1	49.6	62.4	50.1
Household size (sqrt)	54.5	51.5	62.8	52.7
OECD Scale	54.0	51.3	62.6	51.7
Ages 2 and above	53.5	50.5	62.4	51.4
Ages 25 and above	57.7	54.7	64.7	56.5
Ages 40 and above	58.7	52.1	63.0	54.5
Kids get adult leisure	58.3	56.4	67.5	56.9
Frisch elasticity = 0.5	51.4	48.0	60.2	49.0
Frisch elasticity = 2.0	53.5	51.1	61.6	50.9
Value of life 25% smaller	53.4	50.3	62.0	51.2
Value of life 25% larger	52.3	49.5	60.7	50.0

Table A.7: **Robustness results for urban-rural gaps, by country and year**

	MWI		**ZAF**		**PER**		**MEX**		**BGD**		**KHM**	
	2004	2019	1993	2023	2007	2018	2002	2018	2005	2016	2009	2019
Benchmark	78.0	79.3	52.0	75.4	43.0	64.7	71.7	73.8	92.9	90.8	73.0	64.7
No discounting/growth	83.2	82.4	55.8	77.8	45.5	65.1	72.9	73.6	94.1	92.4	77.3	68.0
Discounting: $\beta = 0.96$	76.6	78.4	53.9	75.5	46.4	66.9	70.1	74.6	92.5	87.8	73.4	70.4
Household size (sqrt)	79.6	80.6	52.4	76.8	42.7	63.3	73.5	74.6	92.3	92.1	72.2	65.2
OECD Scale	79.4	80.3	53.5	76.7	44.9	66.4	73.2	74.8	92.8	90.0	73.8	66.0
Ages 2 and above	80.6	80.1	53.8	76.2	44.0	65.2	72.1	73.8	93.7	91.2	74.6	65.9
Ages 25 and above	88.5	82.7	55.2	77.5	44.3	66.1	73.7	73.8	95.9	96.2	79.3	65.7
Ages 40 and above	90.9	87.4	53.5	76.0	40.0	62.0	75.4	73.3	93.8	95.6	76.8	61.4
Kids get adult leisure	80.8	80.6	55.5	85.5	44.2	66.5	73.2	75.5	94.2	92.6	77.0	69.0
Frisch elasticity = 0.5	75.9	77.6	49.9	72.7	44.7	65.1	71.9	73.7	89.2	86.2	70.0	64.1
Frisch elasticity = 2.0	78.9	80.6	55.5	80.0	38.6	61.4	71.7	73.4	96.3	95.6	75.0	64.3
Value of life 25% smaller	79.4	79.7	54.1	76.6	44.0	65.2	72.0	73.9	93.2	90.9	73.9	66.2
Value of life 25% larger	76.6	78.9	49.9	74.2	42.0	64.2	71.4	73.8	92.6	90.7	72.0	63.3
NSO-specific definition	62.1	57.6	42.9	68.3	42.6	64.3	80.7	72.1	81.8	88.0	59.8	60.0

Figure A.1: **Correspondence Between Urbanization Rates (Oldest Survey): NSO vs Our Definition**

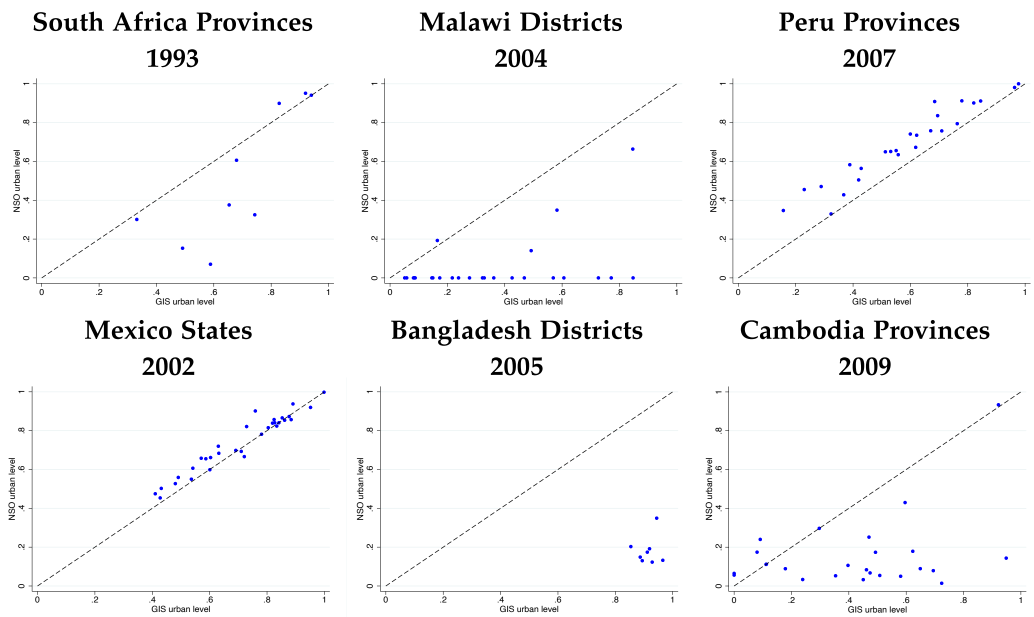

Figure A.2: **Distribution of welfare growth measures across UAs**

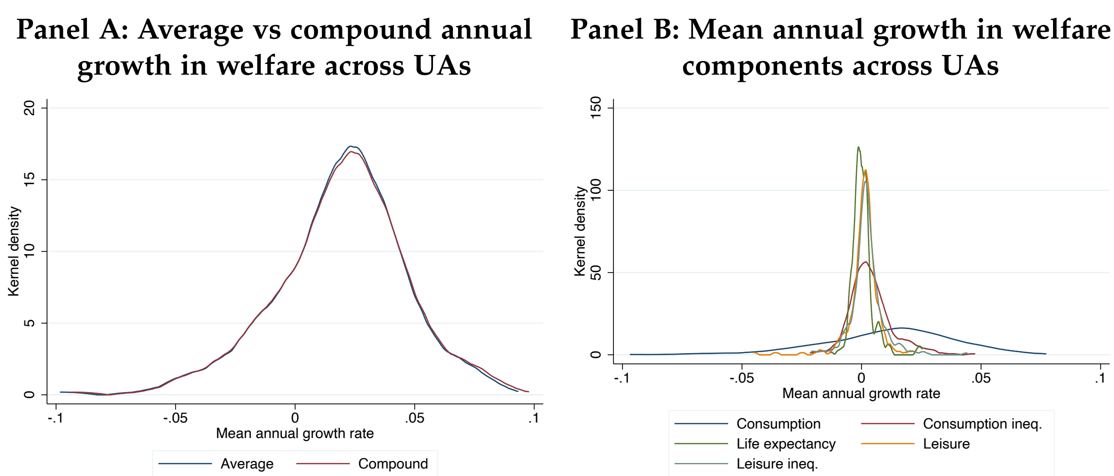

Table A.8: **Life Expectancy, Consumption, and Leisure by Urban Transitions**

Year	Category	LE	Avg. cons.	Avg. cons. adj.	SD cons.	Avg. leis.	SD leis.
			Panel A: South Africa				
1993	Urban	67.3	5,806	8,824	0.731	0.903	0.179
1993	New Urban	67.3	5,112	5,178	0.696	0.910	0.175
1993	Rural	57.8	4,526	3,743	0.672	0.921	0.167
2023	Urban	66.7	6,804	10,302	0.587	0.900	0.177
2023	New Urban	66.7	4,839	7,413	0.501	0.930	0.156
2023	Rural	64.0	5,048	6,212	0.513	0.927	0.159
			Panel B: Malawi				
2004	Urban	55.0	945	844	0.577	0.887	0.171
2004	New Urban	55.0	766	644	0.494	0.892	0.156
2004	Rural	49.2	715	567	0.457	0.891	0.152
2019	Urban	66.3	1,311	1,303	0.633	0.925	0.139
2019	New Urban	66.3	900	860	0.552	0.930	0.121
2019	Rural	63.2	852	889	0.505	0.930	0.117
			Panel C: Peru				
2007	Urban	80.4	6,096	9,611	0.512	0.802	0.250
2007	New Urban	80.4	4,514	5,680	0.506	0.814	0.234
2007	Rural	70.3	2,749	3,873	0.472	0.812	0.213
2018	Urban	77.8	8,949	11,973	0.428	0.796	0.237
2018	New Urban	77.8	7,047	8,713	0.390	0.797	0.241
2018	Rural	75.1	5,235	6,664	0.369	0.800	0.213
			Panel D: Mexico				
2002	Urban	75.2	10,795	15,457	0.653	0.860	0.186
2002	New Urban	75.2	10,430	12,089	0.650	0.862	0.185
2002	Rural	74.4	10,098	12,989	0.651	0.865	0.184
2018	Urban	74.5	16,230	20,521	0.533	0.873	0.181
2018	New Urban	74.5	16,224	14,554	0.535	0.873	0.180
2018	Rural	74.3	16,031	15,354	0.530	0.873	0.180
			Panel E: Bangladesh				
2005	Urban	67.9	1,207	1,273	0.469	0.880	0.205
2005	New Urban	67.9	1,052	1,066	0.410	0.889	0.193
2005	Rural	66.5	1,027	1,032	0.396	0.890	0.191
2016	Urban	71.2	2,691	3,191	0.481	0.882	0.201
2016	New Urban	71.2	2,443	2,582	0.446	0.889	0.191
2016	Rural	70.5	2,374	2,575	0.439	0.890	0.187
			Panel F: Cambodia				
2009	Urban	73.1	2,602	3,040	0.555	0.829	0.206
2009	New Urban	73.1	1,806	2,007	0.428	0.855	0.176
2009	Rural	65.2	1,677	1,906	0.421	0.849	0.179
2019	Urban	75.3	3,426	4,262	0.533	0.770	0.235
2019	New Urban	75.3	2,693	2,977	0.445	0.786	0.226
2019	Rural	67.0	2,413	2,842	0.450	0.789	0.222

Notes: Household consumption values are expressed in per capita terms, in constant 2017 international dollars (PPP-adjusted). Two measures of average consumption are reported: one weighted using survey weights only, and another that additionally applies a demographic adjustment. Standard deviations of consumption are computed over log-transformed values; for leisure they are computed in levels.

Figure A.3: **Welfare Across Provinces in South Africa**

Panel A: 1993

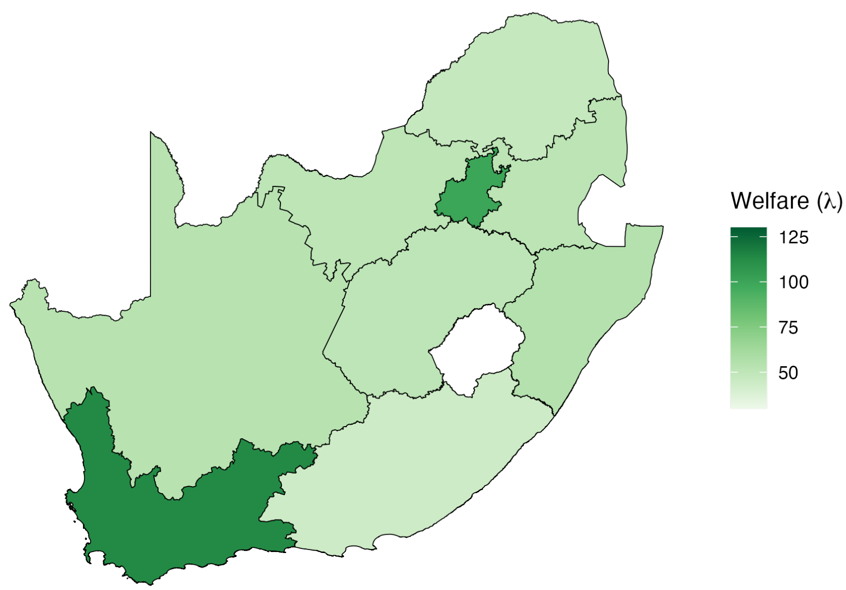

Panel B: 2023

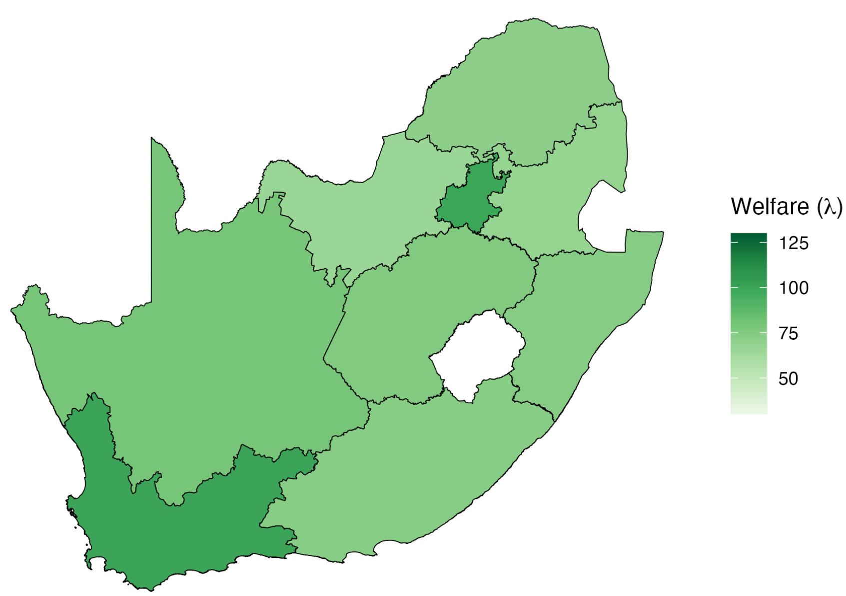

Figure A.4: **Welfare Across Districts in Malawi**

Panel A: 2004

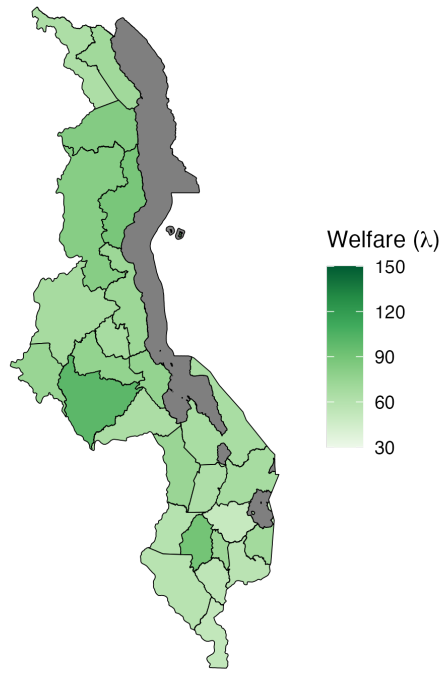

Panel B: 2019

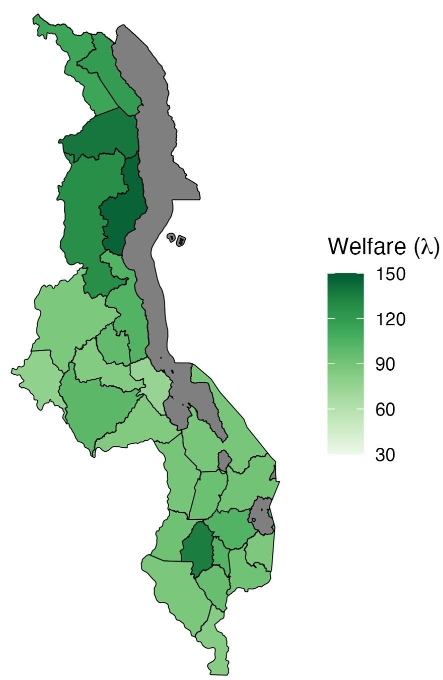

Figure A.5: **Welfare Across Departments in Peru**

Panel A: 2007

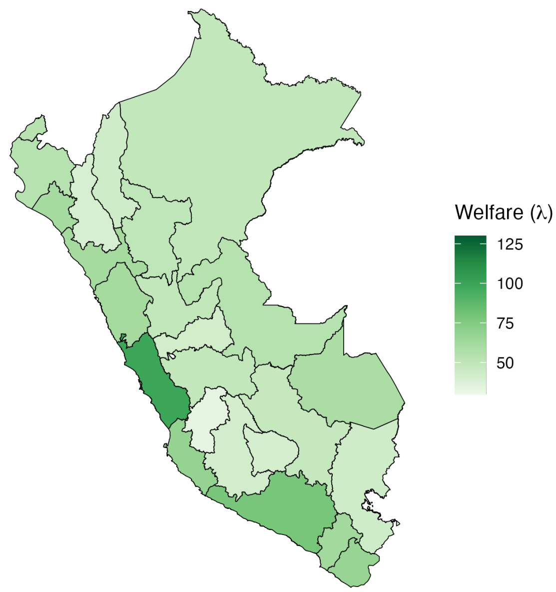

Panel B: 2018

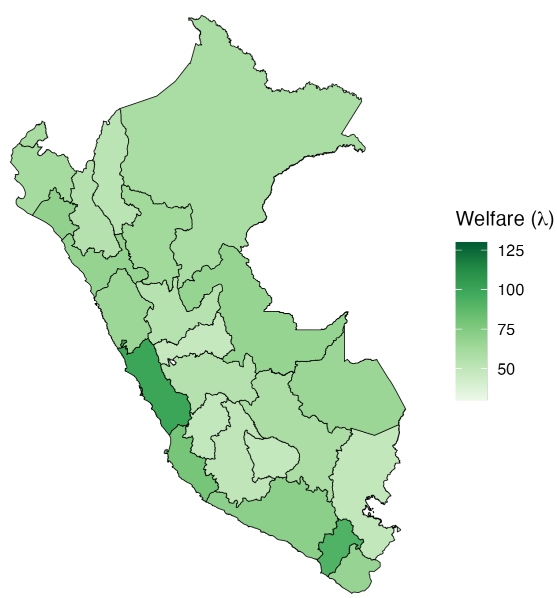

Figure A.6: **Welfare Across States in Mexico**

Panel A: 2002

Panel B: 2018

Figure A.7: **Welfare Across Divisions in Bangladesh**

Panel A: 2005

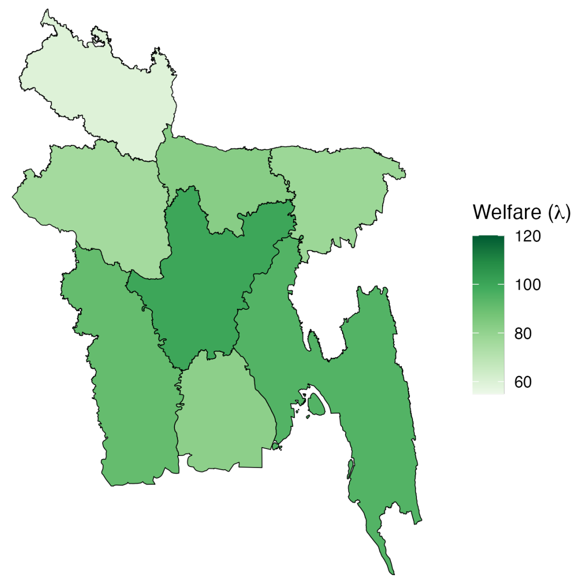

Panel B: 2016

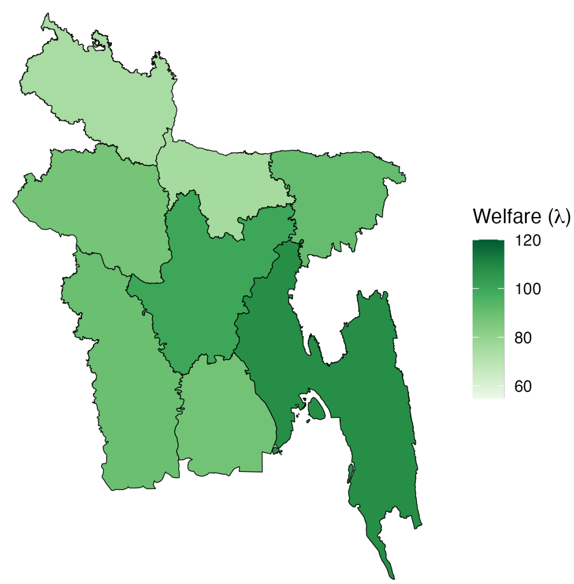

Figure A.8: **Welfare Across States in Cambodia**

Panel A: 2009

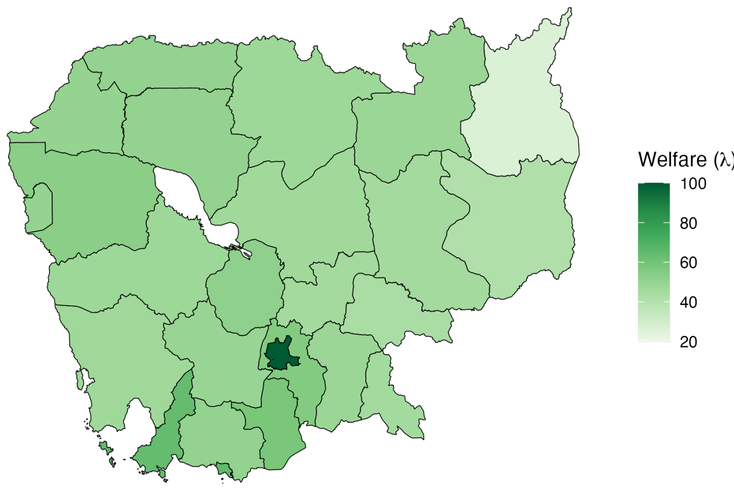

Panel B: 2019

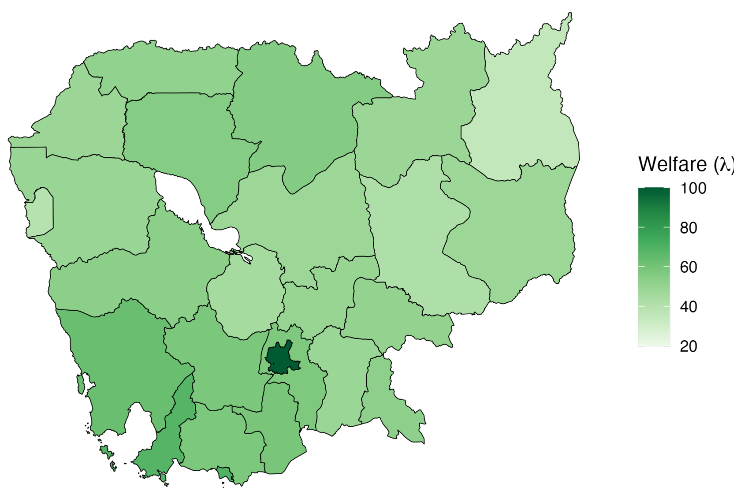

Figure A.9: **Semi-elasticities by Population Size**

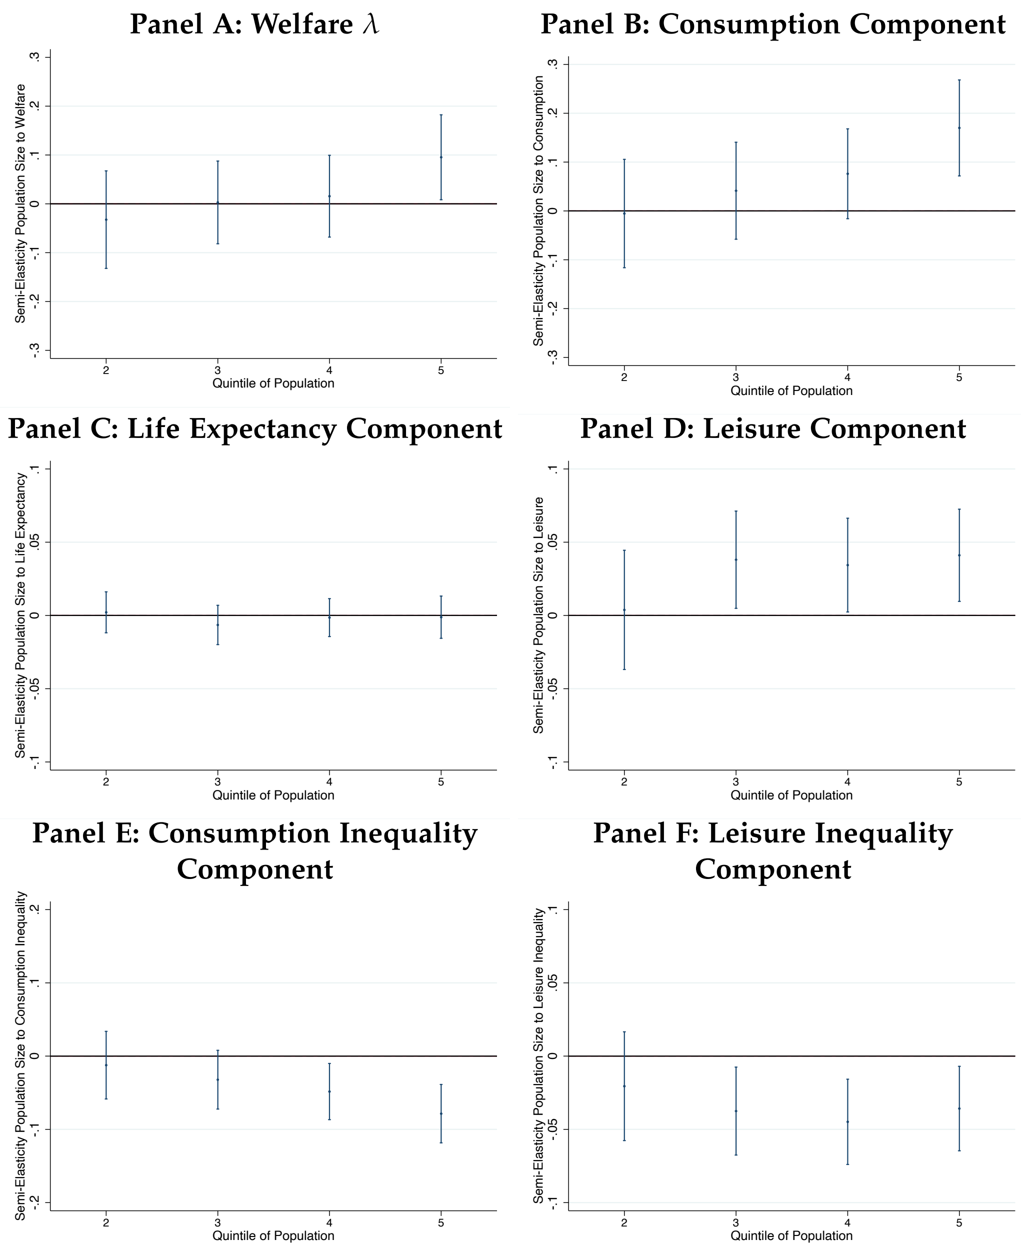

Notes: Each panel shows the point estimates for α_τ in $\sum_{\tau=2}^{5} \alpha_\tau Pop_i^\tau$ as specified in equation 9 where the quintile $\tau = 1$ is the category of reference. Each quintile τ contains between 118 and 120 Functional Urban Areas. Average population sizes in each quintile are 72,777; 116,602; 208,023; 437,804; and 2,630,814 inhabitants.

Figure A.10: **Population-Welfare Elasticities by Sampling**

Panel A: Main Results by Period

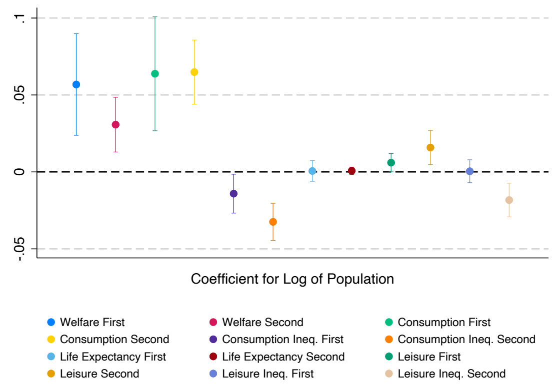

Panel B: Excluding One Year at Time

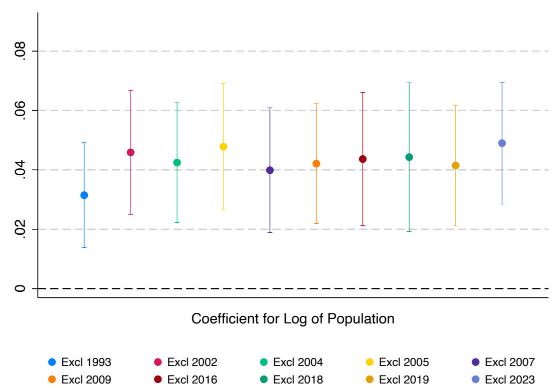

Panel C: Excluding One Country at Time

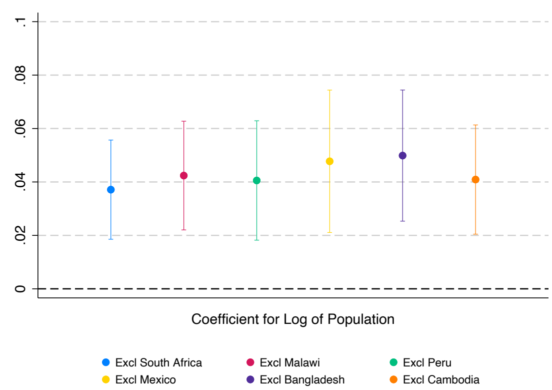

Notes: Each panel shows point estimates for each dependent variable as specified in equation 8 for different alternative samples. In Panel A we estimate equation 8 for the first available survey in each country and for the second one separately. In Panel B we estimate equation 8 when the dependent variable is our main measure of welfare and omit one year in our sample at a time. In Panel C we estimate equation 8 when the dependent variable is our main measure of welfare and omit one country in our sample at a time.

Figure A.11: **Urban Welfare Trajectories**

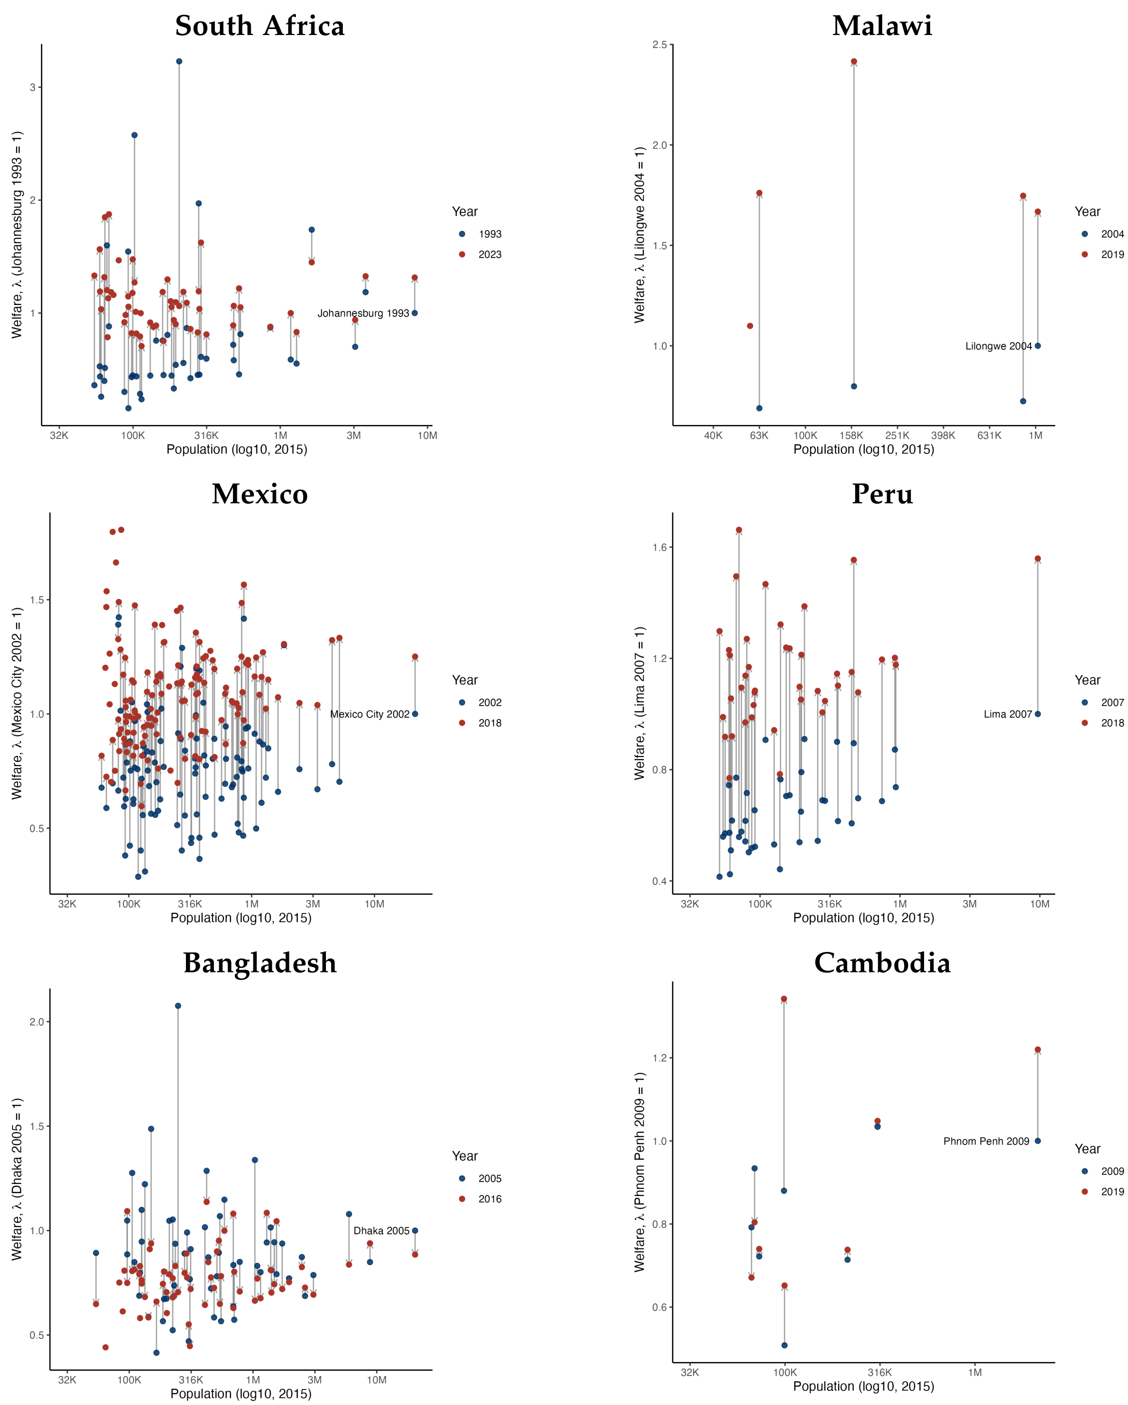

Notes: One scatter plot is presented per country, where each point represents a single UA in a specific year. Welfare levels are normalized by the main city of each country in the base year, allowing for consistent comparisons within-country.