

Is There a Metropolitan Bias? The relationship between poverty and city size in a selection of developing countries

*Céline Ferré, Francisco H.G. Ferreira, and Peter Lanjouw*¹

This paper provides evidence from eight developing countries of an inverse relationship between poverty and city size. Poverty is both more widespread and deeper in very small and small towns than in large or very large cities. This basic pattern is generally robust to the choice of poverty line. The paper shows, further, that for all eight countries, a majority of the urban poor live in medium, small or very small towns. Moreover, it is shown that the greater incidence and severity of consumption poverty in smaller towns is generally compounded by similarly greater deprivation in terms of access to basic infrastructure services, such as electricity, heating gas, sewerage and solid waste disposal. We illustrate for one country – Morocco – that inequality within large cities is not driven by a severe dichotomy between slum dwellers and others. Robustness checks are performed to assess whether the findings in the paper hinge on a specific definition of “urban area”; are driven by differences in the cost of living across city-size categories; by reliance on an income-based concept of well-being; or by the application of small-area estimation techniques for estimating poverty rates at the town and city level. JEL Codes: I32, O18, R12

In the late 1970s and in the 1980s, there was much discussion of “urban bias” in development circles. Following [Lipton \(1977\)](#), development economists increasingly recognized a widespread tendency among (almost always urban-based) governments to pursue policies that – explicitly or implicitly – taxed agriculture and transferred resources to industry and other urban activities.

1. Celine Ferré is an independent consultant based in Amsterdam, the Netherlands; her email address is celinef@gmail.com. Francisco Ferreira is a lead economist in the Development Research Group at the World Bank and a research fellow at IZA; his email address is fferreira@worldbank.org. Peter Lanjouw (corresponding author) is the manager of the Poverty and Inequality Group at the Development Research Group at the World Bank; his email address is planjouw@worldbank.org. We are grateful to Johan Mistiaen for setting us off on this project. We are also much indebted to Victoria Fazio, Philippe George Leite, Ericka Rascón, Martin Ravallion, and Timothy Thomas for advice and numerous contributions. Marianne Fay and participants at a World Bank Workshop on Urban Poverty in June 2007, the World Bank/InWent Development Policy Forum meeting in Berlin, 2007, and the LCSPP Seminar in November 2010, provided useful comments. These are the views of the authors and they should not be attributed to the World Bank.

THE WORLD BANK ECONOMIC REVIEW, VOL. 26, NO. 3, pp. 351–382
Advance Access Publication February 14, 2012

doi:10.1093/wber/lhs007

© The Author 2012. Published by Oxford University Press on behalf of the International Bank for Reconstruction and Development / THE WORLD BANK. All rights reserved. For permissions, please e-mail: journals.permissions@oup.com

The motivation was not exclusively urban self-interest. There was a widespread belief, based on the influential early views of [Rosenstein-Rodan \(1943\)](#), [Prebisch \(1950\)](#), and others, that development was to a large extent synonymous with industrialization – and that industrialization inevitably implied urbanization. As markets could not solely be relied upon to allocate resources to that most dynamic sector, government was required to provide a “big push” to help economies along the righteous path of urban growth.

Against that view, Lipton and his followers argued that urban bias implied a “sacrifice of efficient and equitable growth to rapid urban advance” (p.310). By distorting relative prices and the “intersectoral terms of trade”, such policies induced an inefficient allocation of resources that could lead to perpetually infant industries, at the expense of farmers, many of whom were the poorest people in the land.² That was a time when an estimated 80-90% of the world’s poor lived in rural areas, and an important part of the argument against urban bias was that, in addition to distorting the allocation of capital and other resources, these policies were also anti-poor.³

In 2012, the situation is somewhat different. Urbanization has proceeded apace in the last quarter century such that the world’s urban population is now as large as its rural population. Extreme poverty remains a predominantly rural phenomenon, with some 75% of those who subsist on expenditures below \$1-a-day still residing in rural areas in 2002, even when higher cost of living in urban areas is taken into account.⁴ But urban poverty has been falling more slowly than rural poverty – in part because urbanization has been a key driver behind rural poverty reduction, but some of those who migrate to urban areas remain poor. Urban poverty therefore accounts for a growing share of global poverty: [Ravallion et al. \(2007\)](#) estimate that the urban share of total extreme poverty rose from 19% in 1993 to 25% in 2002. In some regions, like Latin America (76.2%); Eastern Europe and Central Asia (63.5%) and Middle-East and North Africa (55.8%), urban poverty is already dominant.

Poverty is expected to continue to urbanize, in the sense that the share of the total number of poor who live in urban areas is expected to continue to grow (with some exceptions, notably in Eastern Europe). Some expect that the urban share of \$1-a-day (\$2-a-day) poverty may reach 40% (51%) around 2030. Urban poverty is also thought to be accompanied by a different set of characteristics and challenges, including health and sanitation problems in urban slums, unemployment, and a greater incidence of violent crime. In response, strategies to fight urban poverty – and its specific peculiarities – are growing in importance, both at the national and at the international level.

2. A classic study by [Bates \(1981\)](#) documented the use of price regulation and marketing boards in Ghana, Nigeria and Zambia to extract surplus from farmers to the benefit of urban food consumers.

3. [Ravallion et al. \(2007\)](#) produced arguably the first global poverty statistics that cover the majority of the world’s population and disaggregate between urban and rural areas. They estimate that the urban share of the world’s extreme poor in 1993 was 19%.

4. See [Ravallion et al. \(2007\)](#).

Yet urban poverty is far from a homogeneous phenomenon, even within a single country. It is often remarked that poverty is spatially heterogeneous. Usually this is stated with reference to a marked rural-urban dichotomy in measured poverty. But there is also considerable spatial heterogeneity *among* urban areas, and one important dimension of that heterogeneity is across city sizes. In Brazil, for instance, while most anecdotal discussion of urban poverty focuses on the sprawling slums of Rio de Janeiro or São Paulo, over 50% of the country's urban poor live in towns with fewer than 50,000 inhabitants. Only around 10% live in cities with populations greater than a million. In Kazakhstan, the incidence of poverty in smaller towns is six times larger than in Almaty. And there are large differences in access to local public goods and services too: in Morocco, average access to sewerage is over 80% in cities greater than a million, but less than 50% in the smallest towns.

A greater understanding of how poverty – both in terms of incomes or consumption expenditures and in terms of access to public services – varies across different types of cities should help inform the discussion of appropriate poverty reduction strategies in most countries. Yet, the evidence base needed for this disaggregated analysis is seldom available, since household surveys – on which most poverty assessments are based – are seldom representative at the level of any but the largest metropolitan areas in the developing world. They are certainly not representative for smaller towns and cities, and information is not usually disaggregated along these lines.⁵

In this paper, we draw upon insights generated by small area poverty estimation (based on the combination of welfare estimates from household surveys with “sample” sizes from National Censuses) to investigate the relationship between poverty and city size in eight developing countries, namely Albania, Brazil, Kazakhstan, Kenya, Mexico, Morocco, Thailand and Sri Lanka. We find substantial variation in the incidence and depth of consumption poverty across city sizes in seven of the eight countries. For all seven countries where the data permits some kind of disaggregation of the incidence of public service access, there is also considerable variation across city sizes. In all cases, poverty is lowest and service availability is greatest in the largest cities – precisely those where governments, the middle-classes, opinion-makers and airports are disproportionately located.

At a minimum this “poverty gradient” across city sizes needs to be borne in mind whenever considering options and priorities for addressing urban poverty. More speculatively, this evidence might lead us to ask whether,

5. It is rare for household surveys to include identifiers for the specific town or city within which the survey respondent resides – unless the city constitutes a specific stratum. In those settings where such identifiers are included, one can estimate urban poverty rates for different city size intervals (see for example, the poverty profile of Brazil by Ferreira, Lanjouw and Neri (2003) and of India by World Bank (2011)). While city size intervals constructed from survey data are able to reveal differential poverty outcomes across the urban spectrum, they do not convey the differences that may exist between individual towns and cities (see further below).

alongside Lipton's original urban bias, there exists also a "metropolitan bias" in the allocation of resources (including policy attention) to larger cities, at the expense of smaller towns, where most of the urban poor are located.

There are a number of caveats which require that our results be treated with care. First, although our samples within countries are representative, our *sample of countries* is not. Although these eight countries are located in all six regions into which the World Bank routinely divides the developing world, they are not random draws.⁶ They are countries where there was an early interest in (and the data required for) constructing a poverty map. Second, we use national, rather than international, poverty lines. This has the advantage that poverty is measured in the terms which each particular country's residents feel is appropriate. But it has the disadvantage that poverty does not mean the same living standard across the eight countries. Third, we do not apply a uniform definition of "urban area" across our eight countries. Instead, we rely on the definition of urban settlements that each respective country reports in its census documents. Such administrative definitions are likely to vary across countries, and may not correspond closely to economic definitions of towns and cities (linked to population density).⁷ This could have a bearing on our results. We probe the robustness of our conclusions, in the context of Brazil, by re-defining urban settlements so as to take explicit account of population density.

A fourth caveat concerns our inability to systematically adjust for cost-of-living differences between cities. These differences might be expected to be smaller than those between urban and rural areas, but they may still matter, and we report a second robustness test with respect to cost-of-living differences in the only country in our sample for which data permits it, namely Brazil.

Fifth, we do not gauge our findings to variation in equivalence scales – which would matter if family sizes and composition varied systematically across city sizes. Sixth, we assess poverty in a fairly restrictive way: focusing on the share of the population with incomes or consumption levels below the poverty line. It is widely acknowledged that poverty can be viewed more broadly, reflecting multiple dimensions of wellbeing. We seek to mitigate this concern by reporting the association between city size and access to various publicly-provided services and, in one instance, by looking separately at a health outcome (child malnutrition). Nevertheless, it should be acknowledged that the patterns observed in these spaces need not be repeated when other dimensions of poverty are considered.

6. The World Bank divides the developing world into sub-Saharan Africa (AFR), East Asia and the Pacific (EAP), Eastern Europe and Central Asia (ECA), Latin America and the Caribbean (LAC), the Middle-East and North Africa (MENA); and South Asia (SAR).

7. As noted by Lanjouw and Lanjouw (2001), definitions of rural and urban locality vary widely across countries. For example, in Malawi and Zimbabwe settlements with populations of 3000 and 2500 inhabitants, respectively, are defined as urban while in Taiwan a settlement with less than 250,000 inhabitants is designated as rural.

Each of the limitations of the analysis presented in this paper points to the need for additional research. This paper simply documents the existence of systematic differences in the breadth and depth of poverty (and access to services) across city sizes in eight geographically diverse developing countries.

The paper is structured as follows. Section I provides an overview of the poverty mapping methodology which was used in each of the six countries, in order to generate reliable poverty estimates for every urban area captured in the population census. Section II describes the data sources to which this method was applied, in each country. Section III presents the consumption poverty profiles by city size in each country. Section IV turns to the evidence on access to publicly provided services across city sizes. Section V looks at differences in poverty (and inequality) within specific cities in Morocco, focusing in particular on poverty and inequality differences between slums and non-slum areas in larger towns. Section VI subjects our findings to some robustness checks. We first look to data from Brazil, in order to gauge the sensitivity of our findings to changes in the definition of urban areas and city sizes. We also use Brazilian data to probe for evidence of spatial cost of living variation across city-sizes. A subsequent robustness check is applied to the case of Mexico, to examine whether an alternative dimension of deprivation, namely child malnutrition, exhibits the same gradient across city size categories as income poverty. In a final robustness check we show that, in India, a poverty-city size gradient can be observed both directly from survey data and from small area estimation techniques. We conclude that the inverse relationship is thus not driven exclusively by our reliance on a particular estimation method. Section VII offers tentative conclusions and discusses some of the questions that this descriptive paper raises for further research into urban poverty.

I. METHODOLOGY

The economic analysis of the distribution of living standards in developing countries relies almost entirely on household surveys. If their samples are selected appropriately, these surveys can collect detailed information from a relatively small number of households (perhaps 0.1% of the country's total population), and yet generate information that is representative of the population as a whole. The law of large numbers ensures that the uncertainty about the population which results from sampling (the 'sampling error') becomes very small at sample sizes that are still cost effective. This enables researchers the world over to ask detailed questions from small groups of people, at a fraction of the cost that would be required if entire populations needed to be polled.

But there is one drawback: the samples that are designed to be representative of large populations are not, in general, representative of specific non-random sub-divisions of that population. Indeed, the typical (nationally representative) household survey is not representative of sub-national units, such as states, provinces or districts. There are exceptions, mostly in large countries, such as

China, India or Brazil. But even in those countries, the problem is simply shifted down one level: living standards will vary enormously across different localities of (or towns and cities in) the states of Uttar Pradesh or Minas Gerais; but the Indian *National Sample Survey* and the Brazilian *Pesquisa Nacional por Amostra de Domicílios* are not representative at those levels.

A number of small-area estimation techniques have been developed to seek to address this missing data problem. In this paper, we rely on applications of the “poverty mapping” approach developed by Elbers, Lanjouw and Lanjouw (2002, 2003). This approach typically involves a household survey and a population census as data sources. First, the survey data are used to estimate a prediction model for either consumption or incomes. The selection of explanatory variables is restricted to those variables that can also be found in the census (or some other large dataset) or in a tertiary dataset that can be linked to both the census and survey. The parameter estimates are then applied to the census data, expenditures are predicted, and poverty (and other welfare) statistics are derived. The key assumption is that the models estimated from the survey data apply to census observations.

Let W be a welfare indicator based on the distribution of a household level variable of interest, y_b . Using a detailed household survey sample, we estimate the joint distribution of y_b and observed correlates x_b . By restricting the explanatory variables to those that also occur at the household level in the population census, parameter estimates from this “first stage” model can be used to generate the distribution of y_b for any target population in the census conditional on its observed characteristics and, in turn, the conditional distribution of W . Elbers et al. (2002, 2003) study the precision of the resulting estimates of W and demonstrate that prediction errors will fall (or at least not rise) with the number of households in the target population, and will also be affected by the properties of the first stage models, in particular the precision of parameter estimates. A general rule of thumb is that welfare estimates obtained on this basis will be estimated fairly precisely as long as the target population comprises at least 1,000-5,000 households.

The first-stage estimation is carried out using household survey data.⁸ The empirical models of household consumption allow for an intra-cluster correlation in the disturbances (see Elbers, Lanjouw and Lanjouw, 2002, 2003, Elbers, Lanjouw and Leite, 2008, and Demombynes et al., 2007, for more details). Failing to take account of spatial correlation in the disturbances would result in underestimated standard errors in the final poverty estimates (Tarozzi and Deaton, 2009). Different models are estimated for each region and the specifications include census mean variables and other aggregate level variables in order to capture latent cluster-level effects. All regressions are estimated with household

8. These surveys are stratified at the region or state level, as well as for rural and urban areas. Within each region there are further levels of stratification, and also clustering. At the final level, a small number of households (a cluster) are randomly selected from a census enumeration area.

weights and with parsimonious specifications to be cautious about overfitting. Heteroskedasticity is also modeled in the household-specific part of the residual.

Parameter estimates from all the first-stage models are then taken, in the second stage, to the population census. Since predicted household-level per capita consumption in the census is a function not only of the parameter estimates from the first stage consumption models estimated in the survey, but also of the precision of these estimates and of those parameters describing the disturbance terms in the consumption models, we do not produce just one predicted consumption level per household in the census. Rather, a reasonably large number of predicted expenditures are simulated for each household (typically around 100 simulations). The full set of simulated household-level per capita expenditures are then used to calculate the welfare estimates for each target population. Demombynes, Elbers, Lanjouw and Lanjouw (2007) describe a variety of simulation approaches that are available and document that these all yield closely similar welfare estimates. Validation studies of the poverty mapping methodology remain rare; in settings where one can rigorously check the method, it is likely that it was not needed in the first place. However, a few such studies have been conducted and have yielded encouraging findings (see for example Demombynes, et al., 2007, and Elbers et al., 2008). We examine in Section VI whether there are grounds for suspecting that our broad findings concerning the relationship between urban poverty and city size are due to our employment of small-area estimates of poverty as opposed to direct measures.

II. DATA

Poverty-mapping exercises based on the methodology just described have now been conducted in a number of countries. We have selected eight of these countries, on the basis of the availability of the micro-data files and of regional coverage, for analysis in this paper. Table 1 lists the total urban population (at the Census year) in the eight countries, both in absolute numbers and as a share of the total. The sample includes a wide variety of countries, from the relatively small (e.g. Albania) to the relatively large (e.g. Brazil), and from the predominantly rural (e.g. Sri Lanka) to the highly urbanized (e.g. Brazil). The Table also indicates which household survey was used for the estimation of the household expenditure model, including year and sample size. The year of the nearest available population Census, which was used to generate the small-area welfare estimates, is also included.⁹

9. Further details about the poverty maps analyzed here can be found in respectively (INSTAT, 2004) for Albania, IBGE (2003) for Brazil, Kenya Central Bureau of Statistics (2003) for Kenya, López-Calva et al. (2005) for Mexico, Haut Commissariat au Plan (2005) for Morocco, Healy and Jitsuchon (2007) for Thailand, Department of Census and Statistics (2005) for Sri Lanka. In the case of Kazakhstan, the poverty map for that country was produced on a pilot basis in collaboration with the Agency of Statistics of the Republic of Kazakhstan. The results of this exercise have not been placed in the public domain.

TABLE 1: Data Sources

	Albania	Brazil	Kazakh-stan	Kenya	Mexico	Morocco	Thailand	Sri Lanka
Urban Population	1.3m	125m	8.2m	5.0m	54.5m	12.7m	18.5m	2.2m
Urban Population share	0.42	0.83	0.57	0.19	0.60	0.51	0.31	0.12
Census Year	2001	2000	1999	1999	2000	1994	2000	2001
Survey Year	2002	2002-3	2001	1997	2000	1998	2000	2002
Survey Name	LSMS	POF	HBS	WMS III	ENIGH	ENNVN	SES	HIES
Survey Sample Size	3,600	48,470	11,883	10,874	10,108	5,184	24,747	20,100
Poverty Line ¹	ALL 4,891	BRL 100	KZT 3,157	KES 2,648	PES 768	DHS 283	BAH 1370	LKR 1423
Poverty Line (2005 PPP\$)	85	83	63	147	126	57	88	45
Equivalence Scale	No	No	No	Yes	No	No	No	No

¹ All poverty lines are monthly, and displayed using national currencies. International acronyms apply.

Note: LSMS: Living Standard Measurement Survey; POF: Pesquisa de Orçamentos Familiares; HBS: Household and Budget Survey; WMS: Welfare Monitoring Survey; ENIGH: Encuesta Nacional de Ingresos y Gastos de los Hogares; ENNVN: Enquête Nationale sur les Niveaux de Vie des Ménages; SES: Socio-Economic Survey; HIES: Household Income and Expenditure Survey.

Source: The above surveys and National Censuses. Consumer price indices and PPP exchange rates used for poverty line conversions come from PovcalNet.

Finally, the Table also lists the poverty line used in each country, both in national currency (in survey year prices) and US dollars (in 2006 prices and at PPP exchange rates)¹⁰. As noted in the introduction, we have opted to use national poverty lines, which better capture the meaning of poverty in each specific country. This has the drawback that poverty measures are not defined with reference to comparable standards of living across countries. The alternative of imposing a constant poverty line across countries, however, would have an even greater disadvantage. Had we selected a low internationally comparable poverty line, such as \$1-a-day, we would be comparing traces of poverty driven largely by measurement error and transitory shocks in the richer countries (such as Albania and Brazil) with real poverty in Kenya and Sri Lanka. Had we instead selected a higher line, like those used in Albania, Brazil or Morocco, we would be comparing “reasonable” poverty incidences in the richer countries, with the bulk of the population in the poorer countries. Since this paper is largely about the *relative* extent of poverty in larger and smaller towns, the absolute level of the poverty line is of limited importance. We do, nevertheless, examine the sensitivity of our results to varying the poverty line in some of the countries, in the next section.

Further caveats relate to the fact that we make no attempt to apply a uniform definition of urban, or to systematically correct for differences in the cost-of-living across different urban categories. In some settings, these differences may be substantial, and future research should attempt to take them into account.¹¹ In addition, with the exception of Kenya, we have used consumption expenditure per capita as the individual welfare indicator throughout. If there are substantive differences in family size or composition across different urban categories, one might like to investigate the robustness of the results with respect to different assumptions regarding equivalence scales. Note that with respect to both cost of living differences and equivalence scales, our findings will be sensitive to systematic differences between large cities and smaller towns. We are not as vulnerable, here, to differences that might exist between urban areas, generically, and rural areas. It is an important empirical question just how much variation there is between cities of different sizes in terms of prices, consumption patterns, and demographic characteristics.

III. CONSUMPTION POVERTY BY CITY SIZE

Table 2 presents our estimates of the three standard FGT poverty measures (as well as population shares and the share of the poor) for each country as a

10. Each poverty line is per capita per month.

11. Section VI reports on a robustness check indicating that our findings for Brazil are not overturned after refining the definition of urban we employ and correcting by cost of living differences across city-size categories.

TABLE 2: Poverty measures and shares for different city sizes in eight countries

	Population share ¹	FGT0	FGT1	FGT2	Share of the Poor ²
<i>Albania</i>		0.25			
<i>Rural</i>		0.30			
Urban	0.42	0.18	0.04	0.01	0.31
M	0.15	0.18	0.04	0.02	0.11
S	0.13	0.18	0.04	0.01	0.09
XS	0.14	0.20	0.05	0.02	0.11
<i>Brazil</i>		0.22			
<i>Rural</i>		0.37			
Urban	0.83	0.19	0.07	0.04	0.72
XL	0.22	0.09	0.03	0.01	0.09
L	0.07	0.17	0.06	0.03	0.06
M	0.24	0.15	0.05	0.03	0.17
S	0.01	0.19	0.07	0.04	0.01
XS	0.28	0.30	0.11	0.06	0.39
<i>Kazakhstan</i>		0.18			
<i>Rural</i>		0.23			
Urban	0.57	0.14	0.04	0.01	0.43
XL	0.08	0.03	0.01	0.00	0.01
M	0.29	0.13	0.04	0.01	0.21
S	0.05	0.18	0.05	0.02	0.05
XS	0.15	0.19	0.05	0.02	0.15
<i>Kenya</i>		0.51			
<i>Rural</i>		0.52			
Urban	0.19	0.47	0.17	-	0.16
XL	0.07	0.44	0.14	-	0.06
L	0.02	0.44	0.16	-	0.02
M	0.03	0.46	0.17	-	0.03
S	0.02	0.55	0.22	-	0.02
XS	0.04	0.49	0.21	-	0.04
<i>Mexico</i>		0.32			
<i>Rural</i>		0.52			
Urban	0.60	0.19	0.06	0.03	0.39
XL	0.27	0.18	0.06	0.03	0.16
L	0.13	0.14	0.04	0.02	0.06
M	0.11	0.19	0.05	0.03	0.07
S	0.04	0.25	0.07	0.04	0.03
XS	0.06	0.31	0.09	0.05	0.07
<i>Morocco</i>		0.17			
<i>Rural</i>		0.23			
Urban	0.51	0.11	0.03	0.01	0.34
XL	0.12	0.04	0.01	0.00	0.03
L	0.09	0.14	0.04	0.02	0.07
M	0.27	0.13	0.03	0.01	0.20
S	0.03	0.16	0.04	0.02	0.03

(Continued)

TABLE 2: Continued

	Population share ¹	FGT0	FGT1	FGT2	Share of the Poor ²
XS	0.01	0.12	0.03	0.01	0.01
<i>Sri Lanka</i>		0.23			
<i>Rural</i>		0.25			
Urban	0.12	0.09	0.02	0.01	0.05
L	0.03	0.08	0.02	0.01	0.01
M	0.03	0.07	0.02	0.01	0.01
S	0.02	0.09	0.03	0.01	0.01
XS	0.04	0.12	0.03	0.00	0.02
<i>Thailand</i>		0.14			
<i>Rural</i>		0.17			
Urban	0.31	0.08	-	-	0.17
XL	0.12	0.02	-	-	0.01
M	0.03	0.04	-	-	0.01
S	0.02	0.09	-	-	0.01
XS	0.14	0.14	-	-	0.13

¹ Proportion of the population living in each category: urban, XL,L,M,S,XS.

² Proportion of the country’s poor living in each category: urban, XL,L,M,S,XS.

XL: >1,000, L: 500-1,000, M: 100-500, S: 50-100, XS: <50 (*000 inhabitants).

Source: Authors’ analysis based on data described in the text.

whole, and then for their urban areas, first as an urban aggregate, and then disaggregated into five size categories: towns smaller than 50,000 (“very small” or XS); between 50,000 and 100,000 (“small” or S); between 100,000 and 500,000 (“medium” or M); between 500,000 and 1 million (“large” or L) and above 1 million (“metropolitan areas” or XL). Two countries have no metropolitan areas: Albania and Sri Lanka. Two countries have no large cities: Albania and Kazakhstan.

In all eight countries, both poverty incidence (FGT(0)) and depth (FGT(1)) are highest in either the very small (Albania, Brazil, Kazakhstan, Mexico, Sri Lanka, Thailand) or the small (Kenya and Morocco) categories.¹² This pattern is particularly pronounced in the larger, more urbanized countries of Brazil, Kazakhstan, Mexico, Morocco, and Thailand where FGT(0) in the very small cities is up to six times larger than in the metropolitan areas, and often only slightly lower than in rural areas. In these countries, the more distribution-sensitive poverty measures paint a similar picture: FGT(2) is six times larger for very small towns than for metropolitan areas in Brazil; FGT(1) is five times larger in Kazakhstan.

12. FGT(α) denotes a member of the Foster, Greer and Thorbecke (1984) family of poverty indices, with parameter α .

In the other three countries – Albania (heavily urbanized, but small in total population and area), Kenya and Sri Lanka (predominantly rural) – the pattern is less pronounced, but it is still present. In fact, the coefficient on population in a simple OLS regression of poverty on city size is negative in all cases, and significant at the 10% level in five (four) out of eight cases for FGT(0) (FGT(1)). See Table 3.¹³

The inverse relation between poverty and city size can also be discerned in Figure 1, which presents the distribution of poverty incidence *within* each size category, by means of box-plots. The box-plots indicate that there is much greater variance in poverty rates among smaller towns, as one might expect from their sheer number. But the *median* poverty rate falls markedly and consistently with city size in Brazil and Kazakhstan. It also falls in Albania, although less markedly. In Mexico, the gradient is clear across all city-size categories except for the metropolitan areas for which the median poverty rate is slightly higher than for the large city-size category (but below all other categories). In Morocco, the negative correlation detected in Table 3 is driven by much lower poverty in metropolitan areas, with no clear pattern among the other size categories. In Sri Lanka and Kenya, the relationship owes to greater poverty incidence in small and very small towns, with no clear pattern among medium and larger towns. Similarly in Thailand it is noteworthy that the overwhelming majority of urban centers belong to the extra small category, with metropolitan Bangkok representing the one very large exception. (Chiang Mai, Thailand's second largest city, had a population of only just over 280,000 in 2008). The median poverty rate in Thailand's smallest towns is markedly higher than in all other city-size categories.

These patterns are also clearly visible in Figure 2, which presents the non-parametric regressions of FGT(0) on the logarithm of city size for each country. Here again, it is least visible in Kenya and Sri Lanka. In Morocco, as we have seen, the negative relationship is driven by markedly lower poverty in Casablanca. With the exception of Kenya and Mexico, metropolitan poverty incidence is less than half of the average urban poverty in every country in our sample that has at least one metropolitan area.

The data underpinning Figures 1 and 2 are of interest not only in providing evidence of a poverty gradient with respect to city size in our eight countries, but also in documenting the great heterogeneity in poverty across towns within a given city size category. Thus, while in Brazil or Thailand the median poverty rate amongst towns in the XS town-size category is clearly higher than in the other categories, there evidently are very small towns that also enjoy very low poverty rates (as well as small towns with near universal poverty). Indeed, in

13. These regression coefficients are presented as illustrative of correlations only. City size is clearly endogenous, and there are evidently many omitted variables, so no inference of causality is possible. Some countries do not display the full set of regressions for lack of data (the Kenyan census for instance being very short, no information is available on infrastructure access).

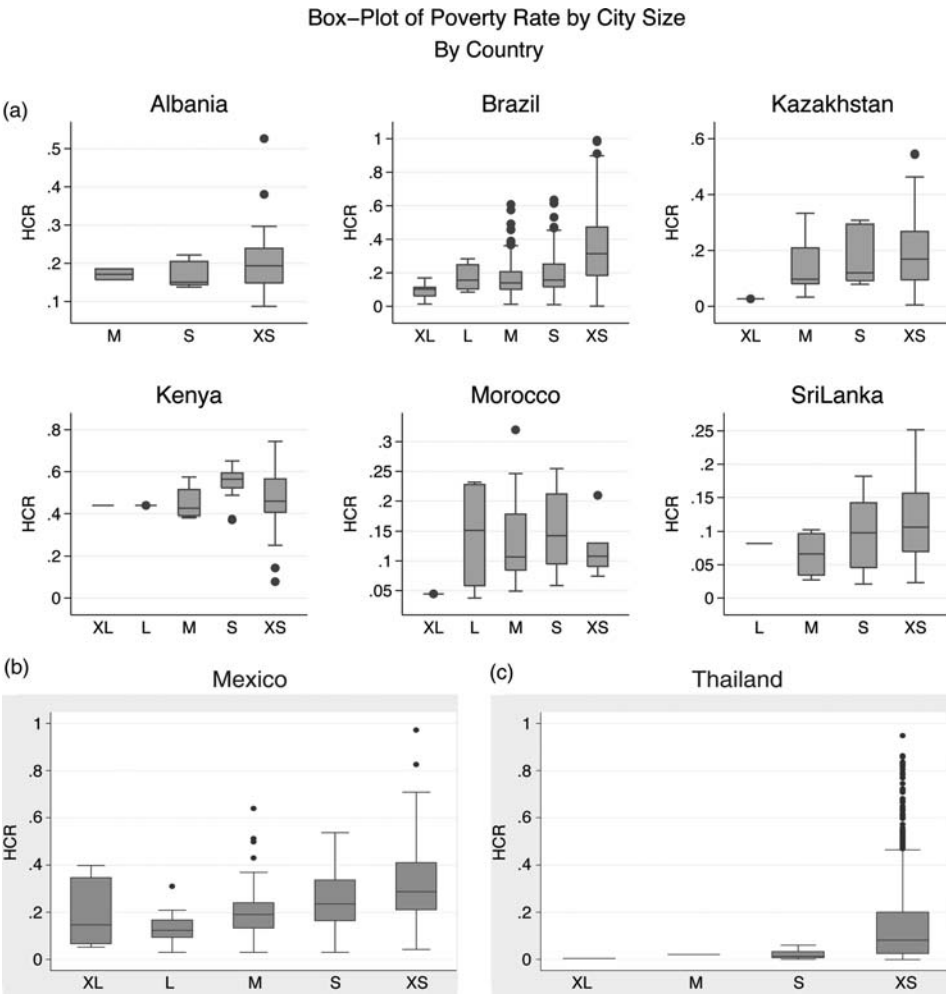
TABLE 3: Simple regressions of poverty indicators on city size, OLS

Country	Dependent variable:									
	FGT0		FGT1		Water		Electricity		Sewer	
	coeff	p-value	coeff	p-value	coeff	p-value	coeff	p-value	coeff	p-value
<i>Albania</i>	-0.17	0.21	-0.04	0.27	-	-	-	-	-	-
<i>Brazil</i>	-0.10	0.03	-0.04	0.03	0.03	0.02	0.01	0.03	0.03	0.04
<i>Kazakhstan</i>	-0.20	0.00	-0.06	0.00	-	-	0.00	0.06	0.92	0.02
<i>Kenya</i>	-0.02	0.73	-0.04	0.38	-	-	-	-	-	-
<i>Mexico</i>	-0.02	0.06	-0.004	0.28	0.007	0.24	0.002	0.07	0.007	0.16
<i>Morocco</i>	-0.03	0.01	-0.01	0.02	0.06	0.01	0.05	0.02	0.08	0.23
<i>Thailand</i>	-0.04	0.15	-0.01	0.22	1.00	0.04	-	-	0.06	0.00
<i>Sri Lanka</i>	-0.08	0.08	-0.03	0.05	0.40	0.23	-0.03	0.54	-	-

Explanatory variable: city size in million of inhabitants ('000,000). All poverty indicators take values between 0 and 1.

Source: Authors' analysis based on data described in the text.

FIGURE 1.



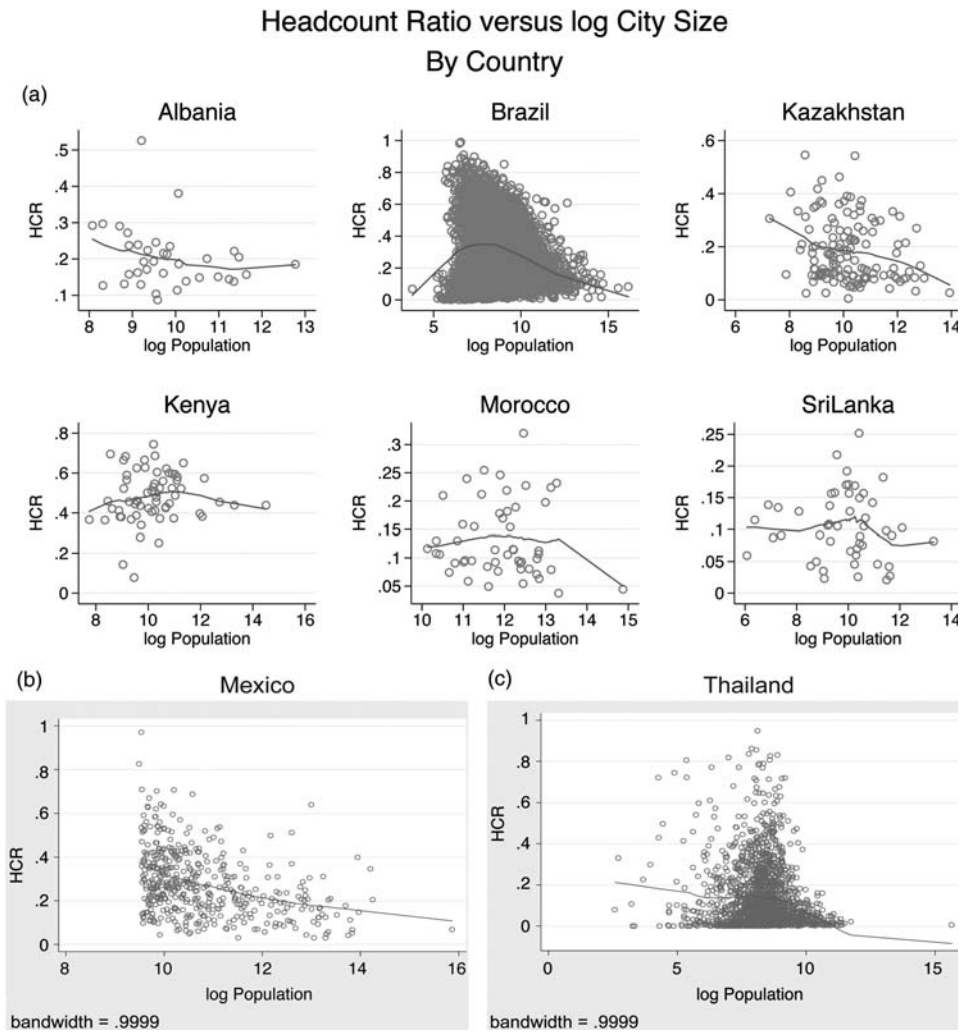
Source: Authors' calculations based on the household surveys and censuses listed in Table 1.

Note: XL: >1,000, L: 500–1,000, M: 100–500, S: 50–100, XS: <50 (thousands inhabitants).

Albania the evidence indicates that the lowest estimated urban poverty rates are found amongst towns in the XS category. It is thus important to bear in mind that while on average poverty rates in smaller towns tends to be higher than in medium and large cities, this is far from a general rule.

To investigate the robustness of the inverse poverty-city size relationship to variations in the poverty line, we plotted the cumulative distribution function

FIGURE 2.




Source: Authors' calculations based on the household surveys and censuses listed in Table 1.

Note: fitted with Lowess regression – bandwidth = 1.

separately by size category for each country.¹⁴ Poverty is always higher in the smallest towns (XS), for any poverty line, in Albania, Brazil, Sri-Lanka and Thailand (up to the 90th percentile). It is generally lowest for metropolitan areas in the vicinity of the national poverty lines, but this ranking is not everywhere robust to larger changes in the poverty line. Figures 3 and 4 illustrate two polar cases: Brazil and Morocco. Figure 3 shows that metropolitan areas

14. Again with the exception of Kenya, for which we do not have the disaggregated poverty mapping data.

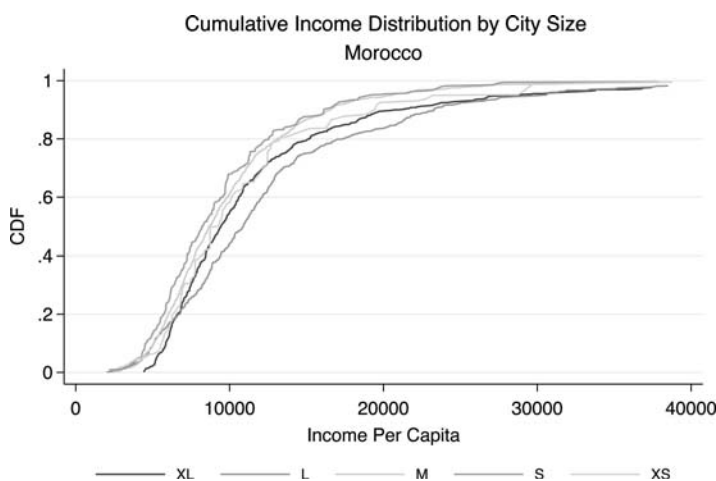
FIGURE 3.



Source: Brazilian poverty map, constructed from the POF survey and the 2000 census.

Note: XL: >1,000, L: 500–1,000, M: 100–500, S: 50–100, XS: <50 (thousands inhabitants).

FIGURE 4.



Source: Moroccan poverty map, constructed from the ENNVN survey and the 1994 census.

Note: XL: >1,000, L: 500–1,000, M: 100–500, S: 50–100, XS: <50 (thousands inhabitants).

first-order stochastically dominate all other size categories in Brazil: poverty is lower in these large cities than in any other type of town, by any poverty line. Conversely, very small towns are first-order stochastically dominated by every other size grouping: poverty is higher in this size category than in any other, by any poverty line.

A very different picture (in terms of dominance relationships) is that of Morocco, shown in Figure 4. The poverty ranking between metropolitan areas and large towns which is observed at the country's poverty line of Dhs 3,400 reverses at higher poverty lines (above Dhs 8,000). Similarly, there is no dominance relationship among very small, small and medium towns in Morocco: their cumulative distribution functions cross many times. Even in Morocco, however, which displays the largest number of cumulative distribution function crossings in our sample, there is still one broad regularity: taken as a group, large and very large cities (L, XL) do provide a lower envelope for the smaller towns (XS, S, M). There is no strict stochastic dominance but it is evident that, for almost every poverty line one could think of, poverty is lower in the group of larger cities than in other urban settings.

It is possible, of course, that poverty is both more widespread and deeper in smaller towns, but that population is so concentrated in large cities that the bulk of the poor live there. If this were the case, greater attention to (and resources for) metropolitan poverty might be justified on the basis that the *share of poverty* is greatest there. But Table 2 shows that this is nowhere the case. In fact, the share of the poor is lower than the population share in every country that has a metropolitan area: the difference is relatively small in Kenya, but very substantial elsewhere. In Brazil, although 22% of the population live in cities greater than 1 million, only 9% of the country's poor do. In Kazakhstan, 14% of the population lives in Almaty, but only 3% of the poor. In Mexico, 27% of the population resides in Mexico City and the other very large metropolitan areas of the country, but only 16% of the poor live in these conurbations. In Morocco, 12% of the population lives in Casablanca but only 3% of the poor.

At the other end of the size distribution, a majority of the country's urban poor live in small or very small towns in four of our eight countries: Albania, Brazil, Sri Lanka and Thailand. If we add medium towns to the list, this rises to seven of the eight countries, including Kenya. And even in the case of Mexico, where the population weight of metropolitan areas is particularly large, the share of the urban poor in medium or smaller sized cities exceeds 40%.

IV. ACCESS TO SERVICES BY CITY SIZE

Even though people are poorer in smaller towns than in large cities and even though a greater number of the poor live in those smaller towns, one might think that, due to the higher population densities in metropolitan areas, per-capita availability of publically provided basic services was lower there. This does not appear to be the case, however. Table 4 presents the proportion of households with access to various basic infrastructure services by city size, in seven of our eight countries.¹⁵

15. Kenya is once again omitted for data reasons.

TABLE 4: Access to services for different city sizes in seven developing countries

	Water	Electricity	Sewer	Gas	Garbage	Fridge	Electric Heat
<i>Albania</i>							
Urban						0.88	0.62
M						0.91	0.68
S						0.87	0.57
XS						0.87	0.60
<i>Brazil</i>							
Urban	0.96	0.99	0.92		0.86		
XL	0.98	1.00	0.94		0.92		
L	0.97	1.00	0.91		0.89		
M	0.97	1.00	0.93		0.91		
S	0.96	0.99	0.94		0.89		
XS	0.92	0.98	0.90		0.76		
<i>Kazakhstan</i>							
Urban		1.00	0.68	0.55			
XL		1.00	0.73	0.81			
M		1.00	0.80	0.62			
S		1.00	0.67	0.36			
XS		1.00	0.40	0.31			
<i>Mexico</i>							
Urban	0.93	0.99	0.94			0.83	
XL	0.95	0.99	0.97			0.84	
L	0.93	0.99	0.93			0.87	
M	0.92	0.98	0.93			0.81	
S	0.91	0.98	0.91			0.78	
XS	0.89	0.98	0.90			0.76	
<i>Morocco</i>							
Urban	0.77	0.82	0.87				
XL	0.84	0.87	0.87				
L	0.86	0.87	0.80				
M	0.71	0.79	0.91				
S	0.73	0.78	0.91				
XS	0.75	0.78	0.45				
<i>Thailand</i>							
Urban	0.65		0.16			0.86	
XL	0.87		0.29			0.88	
M	0.76		0.23			0.90	
S	0.61		0.14			0.87	
XS	0.50		0.07			0.84	
<i>Sri Lanka</i>							
Urban	0.57	0.89					
L	0.53	0.86					
M	0.68	0.90					
S	0.60	0.92					
XS	0.51	0.89					

XL: >1,000, L: 500-1,000, M: 100-500, S: 50-100, XS: <50 ('000 inhabitants).

Source: Authors' analysis based on data described in the text.

Access to piped water is generally quite high in Brazil, but it declines from 98% in metropolitan areas, to 92% in very small towns. In Mexico and Thailand the comparable figures are 95% to 89%, and 87% to 50%, respectively. In Morocco and Sri Lanka, the picture is less clear. In Morocco, as for income poverty, access to piped water is higher in the two largest size categories (L, XL), than in the other three (M, S, XS). In Sri Lanka, there is an inverted U curve, with access lowest in large and very small towns. Similar patterns hold in each of these countries with respect to access to electricity, although overall access rates tend to be higher. Access increases monotonically with city size in Brazil and Mexico; it is higher in L and XL cities than in S, XS and M towns in Morocco (but with no clear pattern within these two blocks), and it follows an inverted U in Sri Lanka.

Access to networked sanitation and sewerage facilities is on average scarcer than piped water or electricity in most developing countries. And in our sample of countries, there is also a clear positive association between city size and access to networked sewerage services. In all five countries that report data on this service (Brazil, Kazakhstan, Mexico, Morocco and Thailand), very small towns have the lowest access rates – in two cases just barely half the rates observed in larger towns. Interestingly, however, in both Kazakhstan and Morocco, medium-sized towns report higher access rates than metropolitan areas.

Access to piped natural gas is an important infrastructure service in Kazakhstan (for cooking and heating). Access is clearly and monotonically increasing with city size. The differences are quite sizable, with 81% of connected households in Almaty, but only 31% in very small towns. A similar pattern attains for electric heating apparatus in Albania. Access to organized solid waste disposal (garbage collection) is only reported for Brazil, where it is once again highest in metropolitan areas, and lowest in very small towns.¹⁶

V. LOOKING WITHIN CITIES: THE CASE OF MOROCCO

A further plausible argument for focusing one's poverty-reduction efforts on metropolitan areas might be that – even if poverty is less widespread or intense there; even if a smaller share of the poor live there; and even if they already enjoy superior access to services – these very large urban centers are deeply divided between rich and poor. If relative incomes matter for well-being, then the stark contrast between the crowded and steep hillsides of *Rocinha* and the neighboring verdant gardens of *Gávea* in Rio de Janeiro may be so inherently objectionable as to raise the priority that should be accorded to fighting poverty in large cities.

16. Although the relationship for intermediate size categories is not monotonic, and there is very little difference between large, medium and small towns in this respect.

There may well be something to the argument that stark local inequalities may have greater costs than geographically diffuse inequality. There is some evidence that relative incomes in one's vicinity do affect well-being directly (Luttmer, 2005), and that local inequality may lead to increased property violence (Demombynes and Özler, 2005). But there is much less evidence that inequality is indeed so much greater in metropolitan areas than in smaller towns. Although this is a popular notion, it is one for which very limited statistical backing exists – in large part for previously mentioned reasons: household surveys are not representative at the level of smaller towns, and so we know very little about local inequality in them. It may be that the accumulation of anecdotal evidence of large inequalities in developing country metropolises is itself simply another reflection of metropolitan bias: Journalists and photographers, like most economists and policy analysts, prefer to visit Casablanca than Figuig¹⁷, and Rio de Janeiro than Bertolinia.¹⁸

To shed some additional light on this matter, we now turn to some evidence from Morocco. Table 5 presents FGT(0) and three inequality measures (the Gini coefficient and the two Theil indices) for each of the five largest cities in the country, as well as the aggregate inequality for three city size categories. Overall intra-city inequality does not appear to be positively correlated with city size in this small sample, but this is not the main point. Taking advantage of the fine spatial disaggregation made possible by a poverty map, we calculated inequality measures for various individual neighborhoods within each of these cities. We further classified these neighborhoods into slums and non-slums¹⁹. We then decomposed the two Theil indices of inequality²⁰ for each of the five cities, into a component due to inequality *within* each of the two groups of neighborhoods, and a component *between* the two. For the argument that within-city inequality is egregiously large in metropolitan areas and large cities to hold (in Morocco), it would be necessary (but not sufficient) that the between-group shares reported in the last two columns of Table 5 be substantial. In the event, it appears that most inequality in the five largest cities in Morocco is not due to some great divide between slum areas and other parts of the town. Inequality appears to be considerably more widely dispersed *within* these two broad groups.

17. Casablanca is the biggest agglomeration of Morocco (2.9 million inhabitants), Figuig is a small town in L'Oriental (49,000).

18. Bertolinia is a small town in the Brazilian state of Piauí, with fewer than 40,000 inhabitants.

19. A district (smallest level of disaggregation after the census track) was considered as a slum if less than 10% of the population had access to water and less than 10% of the population had access to electricity.

20. GE(0), or mean log deviation, is the Theil-L index. GE(1) is the Theil-T index. Both are perfectly decomposable into within- and between-group components, in the sense that the decomposition has no residual. GE(0) weighs within-group inequalities by population shares, while GE(1) weighs them by incomes shares.

TABLE 5: The role of metropolitan slums: Poverty and inequality decompositions within five cities in Morocco

Morocco	All Urban Areas					2 groups: slums and non-slums ¹			
	Population	FGT0	GINI	GE0 ²	GE1 ²	W0 ³	W1 ³	B0 ³	B1 ³
Casablanca	2,875,326	0.05	0.39	0.25	0.29	0.24	0.29	0.01	0.01
Rabat	604,680	0.04	0.41	0.27	0.30	0.27	0.29	0.01	0.01
Salé	575,600	0.12	0.38	0.24	0.26	0.22	0.25	0.02	0.02
Marrakech	503,802	0.07	0.37	0.23	0.23	0.23	0.23	0.00	0.00
Fès	501,592	0.23	0.47	0.38	0.41	0.38	0.41	0.00	0.00
200,000-500,000	4,930,980	0.12	0.36	0.21	0.23	0.21	0.22	0.00	0.00
<200,000	2,736,390	0.15	0.22	0.13	0.14	0.13	0.11	0.00	0.00
<i>All urban</i> ⁴	<i>12,728,370</i>	<i>0.11</i>	<i>0.37</i>	<i>0.23</i>	<i>0.25</i>	<i>0.23</i>	<i>0.24</i>	<i>0.01</i>	<i>0.01</i>

¹ A slum is a district (smallest disaggregation above the census tract) where less than 10% of the population has access to water and less than 10% has access to electricity.

² GE0, or mean log deviation, is the Theil-L index. GE1 is the Theil-T index. Both are perfectly decomposable into within- and between-group components, in the sense that the decomposition has no residual. GE0 weighs within-group inequalities by population shares, while GE1 weighs them by incomes shares.

³ W0 and W1 display within-group inequality associated with the GE0 and GE1 measures respectively; B0 and B1 display the corresponding between-group inequality component.

⁴ Each index presented here was computed at the city level and then aggregated into each category (all urban, etc).

Source: Authors' calculations based on the ENNM: Enquête Nationale sur les Niveaux de Vie des Ménages.

VI. FOUR ROBUSTNESS CHECKS

We noted at the outset several important caveats attached to the broad findings in this paper. While we are unable to subject all of them to exhaustive scrutiny, we attempt, in this section, to probe several with a view to gauging the robustness of our findings. We assess, first, our reliance on official, country-specific, definitions of what constitutes an urban setting in our eight countries. In most countries, the designation of an area as urban is based on several criteria, combining both administrative functions and population characteristics. As a starting point, many countries designate as urban any area, irrespective of population size or density, that happens to have been declared a municipality or corporation by state law, or that has a Cantonment Board or notified town area committee. In general such administrative criteria are combined with additional criteria such as a minimum population (e.g. 5000 inhabitants), a minimum density of population (e.g. at least 400 persons per sq. km), and/or a minimum degree of economic diversification out of agriculture (e.g. at least 75% of male working population engaged in non-agricultural pursuits). As mentioned in the introduction, we take in this paper the definition of urban

TABLE 6: Poverty and City Size in Urban Brazil: Modified Definition of Towns and Cities based on Contiguous Census Tracts

City Size	Total Population (000's)	FGT0	Share of Urban Population	Share of Urban Poor
25,000-100,000	21,261	0.22	0.20	0.29
100,000-500,000	22,384	0.16	0.21	0.22
> 500,000	61,655	0.12	0.59	0.49
All Urban	105,300	0.15	1.00	1.00

Source: Thomas (2008).

employed by each respective country, and do not try to impose a uniform definition across the countries in our sample. The question thus arises whether our finding of a town size–poverty gradient would survive the adoption of an alternative definition of urban. In particular, it is of interest to know whether using a more data driven, and less administrative, definition would affect our conclusions.

To explore this concern we take the Brazil census data and derive our own urban area definitions for that country.²¹ We start with census tracts (groupings of roughly 100 geographically contiguous households) that have been designated as either rural or urban by IBGE, the Brazilian National Statistical Office. We then assemble these census tracts into towns and cities by joining contiguous urban census tracts together. In some cases these form huge units. For example, a coastal area near Rio de Janeiro is 10 km wide and over 170 km long. Compared to the GRUMP (Global Rural Urban Mapping Project at the Columbia University Center for International Earth Science Information Network) settlement area database, this approach would appear to join together several distinct cities into megacities. Because these seem too large, we cut them up using Brazilian micro-regions (officially recognized agglomeration of municípios, but subunits of states) which divides up these “super agglomerations” into areas that closely approximate what GRUMP considers independent cities in Brazil, together with their immediate metropolitan areas. While the definition of a town or city, under this approach, is not tied strictly to a population density criterion (probably the most economically logical way of defining a town) it does employ a single, uniform, criterion and is likely to correlate reasonably well with a purely density-based definition.

Table 6 recalculates poverty across city size categories when cities have been defined in this alternative fashion. We also drop from consideration, here, the large subset of very small towns with less than 25,000 inhabitants²². One might

21. We draw here on work undertaken by Timothy Thomas in support of the project described in this paper (Thomas, 2008).

22. There are nearly 10,000 such settlements in Brazil.

argue that such very small towns are hard to distinguish clearly from rural areas, particularly where those rural areas are reasonably densely populated.²³ We examine whether our city size – poverty gradient for Brazil survives the exclusion of these very smallest towns. We see from Table 6 that amongst all urban centers with a population of at least 25,000 inhabitants, poverty in small towns (with a population of 25,000 to 100,000) is markedly higher than in the medium sized cities (100,000-500,000) which in turn record greater poverty than the large cities of 500,000 inhabitants or more. In terms of poverty shares, our earlier findings also survive: while the inhabitants of Brazil's largest cities represent 59% of the urban population (as per our new definition), they represent just 50% of the urban poor. Conversely, small towns of 25,000 to 100,000 inhabitants account for just 20% of the urban population but nearly 30% of the urban poor.

Our second robustness check assesses the possibility that there may be important cost of living differences between urban settlements of different sizes. The findings reported above have not attempted to adjust for such cost-of-living differences, because spatial price indices across city-size categories are not generally available. It is well recognized in the literature, however, that observed differences in poverty rates between urban and rural areas can be significantly attenuated once one corrects for the fact that the cost of living in urban areas may be much higher than in rural areas (generally because of the higher cost of food, housing and transport services). The possibility exists that our broad findings of greater poverty in small towns than in metropolitan areas might also be driven, at least in part, by our failure to allow for a higher cost of living in metropolitan areas.

While household survey datasets are not generally large enough in sample size to permit the construction of a cost-of-living index across different city-size categories, our survey data for Brazil constitute an important exception. We are able to draw on the 2002 POF data (see Table 1) to construct a cost-of-living index across the broad city size categories employed in this paper, and can check whether our findings for Brazil, reported in previous sections, are robust to this correction.

There are many ways in which spatial price indices can be constructed. We follow here the approach applied by Ferreira, Lanjouw and Neri (2003) to the construction of a *regional* price index (that distinguished also between urban and rural areas) using 1996 PPV data for Brazil. This approach was subsequently applied in World Bank (2007) to produce a regional price index based on the 2002 POF data, and is based on unit-value information provided in the

23. Nevertheless, the large number of small towns in developing countries, and the high poverty rates that characterize them (Table 2), raises the interesting question of the role of such towns in the classic rural-urban debates in development economics (for a recent overview, see Christiaensen, Demery and Kuhl, 2011). Lanjouw and Murgai (2009) suggest, in the context of India, that growth in such small towns has promoted rural non-farm diversification which, in turn, contributes to rural poverty reduction by both providing new employment opportunities for the poor and by raising agricultural wages via agricultural labor market tightening (see also World Bank, 2011).

TABLE 7: Spatial Price Indices across City Size Categories in Urban Brazil
Laspeyres Price Indices Based on the Cost of Food and Housing

Region	City Size Category	Laspeyres Price Index
North	Large (>500,000)	0.94
	Medium (100,000-500,000)	0.75
	Small (<100,000)	0.68
North-East	Large (>500,000)	0.66
	Medium (100,000-500,000)	0.72
	Small (<100,000)	0.60
South-East	Large (>500,000)	0.55
	Medium (100,000-500,000)	0.49
	Small (<100,000)	0.84
São Paulo		1.00
South	Large (>500,000)	0.76
	Medium (100,000-500,000)	0.65
	Small (<100,000)	0.62
Center-West	Large (>500,000)	0.86
	Medium (100,000-500,000)	0.80
	Small (<100,000)	0.64

Source: Authors' calculations based on the POF: Pesquisa de Orçamentos Familiares.

POF survey on food items, as well as a hedonic model of rent. We re-apply the method here, but focus solely on urban areas and construct a Laspeyres price index that captures price differences across city-size categories (and regions). Because of the limited sample size of even the unusually large POF household survey we employ a three-way city size breakdown, distinguishing between cities larger than 500,000 persons, large towns with a population between 100,000 and 500,000, and towns with fewer than 100,000 inhabitants. For reasons of data availability our index captures only food and housing price differentials.²⁴ The reference basket employed in our Laspeyres index is based on the consumption patterns of the second quintile of the national urban per capita consumption distribution.

Table 7 presents our Laspeyres spatial price index based on the cost of food and housing in urban Brazil. Relative to the reference region of metropolitan São Paulo, the cost of living in other regions and city size categories of Brazil is generally lower. In the regions of the North, South and Center West there is evidence of lower cost of living as towns become smaller. This pattern is less clear-cut in the North East, where the cost of living appears particularly high in large towns, and the South East, where those living in small towns appear to face the highest cost of living in the region (although still well below metropolitan São Paulo).

24. As noted above, transport costs are also likely to vary substantially across city size categories. Among the poor, these costs typically account for a lower share of expenditures than food and housing. Nevertheless, the sensitivity of our results to spatial variations in the cost of living that include transport costs remains an issue for future work.

TABLE 8: Poverty measures for different city sizes in Brazil Checking for Robustness to Cost of Living Differences

	Population share ¹	FGT0 (nominal expenditure)	FGT0 (real expenditure)
Urban	0.83	0.19	0.18
XL	0.22	0.09	0.06
L	0.07	0.17	0.10
M	0.24	0.15	0.10
S	0.01	0.19	0.11
XS	0.28	0.30	0.19

XL: >1,000, L: 500-1,000, M: 100-500, S: 50-100, XS: <50 ('000 inhabitants).

Source: Authors' analysis based on data described in the text.

We next take the price indices reported in Table 7 and apply them to the small-area based estimates of per capita consumption for each household in the population census. We then re-calculate poverty rates across region and city-size categories. Does this cost of living adjustment overturn our conclusion that urban poverty in smaller towns is significantly higher than in the large and metropolitan areas? Table 8 indicates that it does attenuate the “gradient” between poverty and city size somewhat, but is far from sufficient to negate or overturn our broad finding. In Brazil, it remains the case that the incidence of poverty in the smallest towns is roughly three times higher than in Metropolitan centers.

In a third robustness check we investigate whether the finding of a negative gradient between poverty and city-size is somehow an artifact of the focus, in this paper, on income poverty as opposed to a broader conceptualization of deprivation. While the broad pattern of higher poverty and lower access to services in small towns was found to be quite robust across our eight countries, an important potential caveat to this assessment concerns health outcomes. It has been suggested in the literature that health outcome indicators in large cities in the developing world may lag behind those in smaller towns. For example, [Chattopadhyay and Roy \(2005\)](#) demonstrate that a variety of indicators of child mortality are more pronounced in the large cities of India than in towns and medium sized cities. This study finds that while infant mortality amongst the wealthiest classes in large cities are particularly low, infant mortality rates amongst the poorest classes are quite pronounced – and indeed are higher than amongst the poorer segments in small and medium sized towns. These are suggestive findings and may be related to the particularly unhealthy living conditions in over-crowded slum areas of large cities. However, evidence on health outcomes across city sizes categories remains scarce and there does not appear to be a broad consensus in the literature on the relatively higher health risks in large cities. For example, [Kapadia-Kundu and Kanitkar \(2002\)](#) argue, also with reference to India, that urban public health services generally place greater emphasis on mega-cities and metro-centers, to the relative neglect of smaller cities and towns.

We probe this concern by examining the gradient of child anthropometric outcomes across city size categories in Mexico. We draw on a small area estimation effort undertaken by Rascón (2010) that parallels the work reported in preceding sections, but focuses on anthropometric outcomes rather than income poverty. Rascón combines the Mexican National Survey of Health and Nutrition 2006 with the Second Count of the Population and Dwellings 2005 in order to apply a variant of the Elbers et al. (2003) small area estimation procedure to the incidence of stunting and underweight amongst children aged 5 and below in Mexico.²⁵ Lanjouw and Rascón (2010) examine the correlation of child health outcomes in urban areas with city size. Table 9 summarizes their results and documents that the incidence of low height for age (stunting) and low weight for age (underweight) amongst children displays a similar gradient across city sizes as we have seen for income poverty. In Mexican cities that are larger than 500,000 inhabitants, the incidence of stunting and of underweight amongst children is 9%. In the case of stunting the incidence rises monotonically as cities decline in size. Amongst the smallest cities (of less than 10,000 inhabitants) the incidence is as high as 16%. In the case of underweight, the incidence also rises, but less markedly: from 9% in the largest cities to 11% in the small cities. The higher incidence of child malnutrition in small towns also translates into more malnourished children: 27% of stunted children in urban areas are found in the largest cities, while 29% are in towns with less than 15,000 inhabitants. Similarly, 27% of underweight urban children reside in the largest cities, but 30% reside in the smallest towns.

In a final robustness check, we ask whether our observed gradient between poverty and city-size is driven by our reliance on small area estimates of poverty for each city rather than direct measures of poverty for such localities.²⁶ We have already noted that direct measures of poverty for individual towns and cities are not generally available in developing country settings. The household surveys that underpin poverty analysis in these countries do not generally cover sufficiently large samples to permit poverty measurement at this detailed level. As was described in Section 2, the small area estimation procedure applied in the present paper combines household survey with unit-record population census data in an effort to circumvent this small sample problem. The approach takes advantage of the full population coverage of the

25. Fujii (forthcoming) adapts the Elbers et al. procedure for the estimation of anthropometric outcomes and applies this methodology to Cambodia. Rascón adapts this procedure further to apply it to Mexican data.

26. Tarozzi and Deaton (2007) have recently expressed a concern that the small area estimation procedure employed by ELL (2002, 2003) may overstate the precision of local level poverty estimates. They base their argument on Monte Carlo simulation results. (See also Molina and Rao, 2010) Elbers, Lanjouw and Leite (2008) examine this issue with data for the state of Minas Gerais in Brazil, and find little evidence in that specific setting for concern. It remains true, though, that the ELL procedure estimates poverty, rather than directly measuring it, and as such there is interest in assessing whether the findings reported in this paper would also hold had poverty been directly measured.

TABLE 9: Child malnutrition estimates for different city sizes in Mexico Small area estimates of malnutrition amongst children under 5 in urban areas

Locality size (inhabitants)	Stunting			Underweight		
	Incidence	Share of Urban Population	Share of National Population	Incidence	Share of Urban Population	Share of National Population
L	0.09	0.27	0.15	0.09	0.27	0.18
M	0.11	0.44	0.24	0.09	0.43	0.28
S	0.16	0.29	0.15	0.11	0.30	0.19

L: >500, M: 15-500, S: 2.5–15 ('000 inhabitants).

Source: Lanjouw and Rascón (2010).

population census and then applies statistical techniques to insert into the census an indicator of per capita expenditure or income for each household. This is necessary because in most developing (and developed) countries, the population census fails to collect detailed income or expenditure data.

India offers an opportunity to probe the contention that our findings are merely an artifact of the methods we have employed. The Indian National Sample Survey Organization (NSSO) fields a very large sample survey every five years with a sample size that is sufficiently large to permit a breakdown of urban areas into city-size categories.²⁷ Table 10 draws on a World Bank study (World Bank, 2011) to illustrate that at the national level for the years 1983, 1993 and 2004/5, National Sample Survey data show a clear gradient in poverty by city size. This gradient holds both at the national level, as well as at the level of individual states.²⁸ A recent study applies the small area estimation methodology used here to estimate poverty at the local level in three states of India in 2004/5 (Gangopadhyay et al., 2010). The study confirms that in West Bengal, Orissa and Andhra Pradesh the poverty-city size gradient observed from NSS data also emerges from estimates derived out of the small-area estimation procedure (Table 11). Thus, at least in India, the finding of an inverse poverty-city size gradient is robust to alternative empirical methods. This provides some support to the claim that the findings reported in preceding sections are not driven by our reliance on small-area estimation techniques.

27. Every five years the NSS fields a “thick round” with a sample size of around 120,000 households, The “thin rounds” fielded in the other years have sample sizes of around 30-40,000 households.

28. World Bank (2011) also shows that the pattern of differential per capita access to public services across city size categories is skewed in India, with small towns faring more poorly than large cities.

TABLE 10: Poverty in India's Small Towns Exceeds Poverty in the Large Cities: Direct Evidence from the NSS.

	1983	1993-94	2004-05
Rural	0.465	0.368	0.281
Urban:	0.423	0.328	0.258
Small towns	0.497	0.434	0.300
Medium towns	0.423	0.315	
Large towns	0.290	0.202	0.147

Notes: Poverty rates based on NSS 1983, 1993 and 2004/5 surveys using Uniform Reference Period consumption and official poverty lines. Small < 50K, Medium 50K-1m, Large \geq 1m.

Source: World Bank (2011).

TABLE 11: Small area estimates reveal high poverty in small towns in three Indian states

City Size	West Bengal				Orissa				Andhra Pradesh			
	No. of towns	Share of Pop	Share of Poor	FGT0	No. of towns	Share of Pop	Share of Poor	FGT0	No. of towns	Share of Pop	Share of Poor	FGT0
XL	1	0.20	0.08	0.05	-	-	-	-	1	0.18	17	0.23
L	1	0.05	0.04	0.12	2	0.21	0.20	0.34	3	0.13	7	0.14
M	54	0.48	0.46	0.13	6	0.22	0.19	0.31	37	0.39	37	0.24
S	28	0.09	0.12	0.17	15	0.19	0.19	0.36	40	0.15	20	0.33
XS	298	0.18	0.31	0.23	121	0.38	0.42	0.39	104	0.15	18	0.31

Note: XL > 1m; L: 500K-1m; M: 100K-500K; S: 50K-100K; XS < 50K.

Source: Gangopadhyay et al. (2010) and World Bank (2010).

VII. CONCLUDING REMARKS

Using highly disaggregated poverty map data from eight countries drawn from all six regions of the developing world, we have shown evidence of a common – although not universal – inverse relationship between poverty and city size. In all countries in our sample, poverty is both more widespread (higher FGT(0)) and deeper (higher FGT(1)) in very small and small towns (those with a population below 100,000) than in large or very large cities (those with a population greater than 0.5 million). Metropolitan poverty, in particular, is considerably lower than poverty in other urban areas in all countries in our sample, except for Kenya and Mexico. Dominance analysis of cumulative distribution functions indicates that the basic pattern is generally robust to the choice of poverty line.

Neither is it true that, because of sheer population size, most poor people in these countries live in large cities. In fact, in all eight countries, a majority of the urban poor lives in medium, small or very small towns. In four of them

(Albania, Brazil, Sri Lanka and Thailand), a majority of the urban poor lives in towns smaller than 100,000 people.

The greater incidence and severity of consumption poverty in smaller towns is compounded by similarly greater deprivation in terms of access to basic infrastructure services, such as electricity, heating gas, sewerage and solid waste disposal. This pattern is not absolute. It does vary by type of service and across countries. Access rates seldom increase strictly monotonically with city size, but they do generally increase, so that for most services and in most countries, large cities and metropolitan areas have higher coverage rates than smaller towns.

Finally, we have also shown for one particular country – Morocco – that inequality within large cities is not driven by a severe dichotomy between slum dwellers and others. The notion of a single cleavage between slum residents and well-to-do burghers as the driver of urban inequality in the developing world appears to be unsubstantiated – at least in this case. Perhaps more important than the highly visible inequalities within our large cities are the less obvious differences between them and smaller urban settlements. In countries like those studied here (with the possible exception of Kenya), poverty is greater and deeper in smaller towns, in both income and (at least some) non-income dimensions.

While the findings reported in this paper are suggestive, we have also pointed to a number of caveats that merit further attention in order to gauge the robustness of the poverty - city size gradient, as well as its applicability in other developing countries. To begin with, while our household samples are representative of the populations within countries, our *sample of countries* is not statistically representative of the developing world as a whole.

In a number of cases, we have been able to perform tentative robustness checks for a subset of these caveats. First, we have indicated that the criteria used to define urban areas vary across countries. It is not uncommon to observe large numbers of small towns designated as urban on the basis of some administrative criterion rather than one linked to population characteristics (e.g. population size, density or economic diversification). We have found, in the case of Brazil, that the inverse relationship between city size and poverty is robust to an alternative, non-administrative definition of urban areas. But it is important to probe more generally whether the gradient we observe in other countries might be a construct of the way cities are defined.

Second, differences in cost of living between cities of different sizes can be significant. Failure to correct for such differences could generate a misleading sense of higher living standards in larger cities. Again in Brazil, we have found that adjusting for cost of living variation attenuates, but does not overturn, our observed city size-poverty gradient.

Third, we have noted a contrasting literature suggesting that health outcomes might be worse in large cities than small towns. We have explored this conjecture by examining child nutritional outcomes across Mexican towns and cities, and have found that in this context a gradient for city size-child health mirrors, rather than offsets, the city size-poverty gradient uncovered in this paper. The general

applicability of this finding also merits further attention.²⁹ Finally, we acknowledge that our findings in this paper are based on small area estimates of poverty that combine survey-based prediction models with population census data. It is important to probe the findings of a city size-poverty gradient also with directly measured poverty data for different cities. We have illustrated that, in the case of India, direct measures and small area estimates yield the same conclusion. It is extremely difficult to find survey data that is representative at sufficient levels of disaggregation to measure poverty directly in small towns – which is why we have used poverty mapping data. Nevertheless, to the extent that additional comparisons with direct survey-based measures are possible elsewhere, they are definitely of interest. A full research agenda also presents itself with respect to uncovering possible reasons for the city size-poverty gradient observed in our sample of countries. We have, somewhat provocatively, postulated that higher poverty rates in small towns might be the consequence of a “metropolitan bias” among policy makers, resulting in a disproportionate allocation of resources to large cities. Such a bias would be consistent with the evidence we have presented of differential per-capita availability of a variety of infrastructure and public services across the city-size spectrum. However, there are many other plausible explanations, including the possibility that the cost of infrastructure provision may be lower in larger towns and cities.

More generally, we have noted in Figures 1 and 2, the considerable heterogeneity in poverty outcomes amongst towns and cities within all of the size-categories we have considered. Thus, even amongst small towns there are some with low poverty and others with high poverty rates. Important questions arise as to what explains this variation. The “new” economic geography literature and also the recent *World Development Report*, entitled “Reshaping Economic Geography” devote considerable attention to the mechanisms through which concentration of population and economic activities can generate various kinds of externalities (for example, Krugman 1999, Henderson, Shalizi and Venables, 2001, World Bank, 2009). It is possible that the inverse association between urban poverty and city size reflects primarily urban agglomeration effects.

An additional possibility is that the *location* of towns matters: small towns located near major urban centers may experience low poverty rates while those in remote areas are poorer. There may also be a tendency for towns located near major metropolitan areas to tend to be larger in size.³⁰ Although beyond

29. A potentially interesting exercise is to directly estimate infant mortality rates from census data and to examine differentials of this welfare outcome across the urban city-size spectrum. We leave this exercise to future work.

30. An examination of these questions in the Indian state of West Bengal reveals that when towns are split into three groups – within a 100-kilometer radius of Kolkata, in a radius of 100-200 kilometers, and more than 200 kilometers – poverty sharply rises with distance from Kolkata when the towns are within a 100 kilometer radius (World Bank 2011). The relationship is weaker in the second group, and completely absent in the third, most distant, group. Thus in the first group the agglomeration effect that really matters is the one generated by Kolkata. However, in the second group, and even more strongly so in the third group, a separate agglomeration effect (proxied by the town’s own size) is discernible.

the scope of this paper, it is clear that a great deal of additional work is needed to better understand how the empirical patterns we have uncovered come about, and how they might best be addressed. What is also clear, though, is that wherever a clear city size-poverty gradient holds, it should be acknowledged and confronted in any strategy for urban poverty reduction.

REFERENCES

- Bates, Robert H. 1981. *Markets and States in Tropical Africa*. Berkeley, C.A.: University of California Press.
- Chattopadhyay, A., and T.K. Roy. 2005. "Are urban poor doing better than their rural counterpart in India? A study of fertility, family planning and health." *Demography India* 34(2):299–312.
- Christiaensen, Luc, Lionel Demery, and Jesper Kuhl. 2011. "The (evolving) role of agriculture in poverty reduction – An empirical perspective", *Journal of Development Economics* 96: 239–254.
- Demombynes, Gabriel, and Berk Özler. 2005. "Crime and Local Inequality in South Africa." *Journal of Development Economics* 76 (2):265–92.
- Demombynes, Gabriel, Chris Elbers, Jean O. Lanjouw, and Peter Lanjouw (2007) "How Good a Map? Putting Small Area Estimation to the Test" Policy Research Working Paper No. 4155, The World Bank.
- Department of Census and Statistics, Sri Lanka. 2005. *A Poverty Map for Sri Lanka: Findings and Lessons* (Colombo: Department of Census and Statistics).
- Elbers, Chris, Jean O. Lanjouw, and Peter Lanjouw. 2002. "Micro-Level Estimation of Welfare", Policy Research Working Paper No. 2911, The World Bank.
- . 2003. "Micro-level Estimation of Poverty and Inequality", *Econometrica*, 71 (1), pp. 355–364.
- Elbers, Chris, Peter Lanjouw, Johan Mistiaen, Berk Özler, and Kenneth Simler. 2004. "On the Unequal Inequality of Poor Communities." *World Bank Economic Review* 18(3):401–21.
- Elbers, Chris, Peter Lanjouw, and Phillippe Leite. 2008. 'Brazil within Brazil: Testing the Poverty Mapping Methodology in Minas Gerais' World Bank Policy Research Working Paper No. 4513, The World Bank.
- Ferreira, Francisco, H.G., Peter Lanjouw, and Marcelo Neri. 2003. "A Robust Poverty Profile for Brazil Using Multiple Data Sources", *Revista Brasileira de Economia*, 57 (1), pp. 59–92.
- Foster, James, Joel Greer, and Erik Thorbecke. 1984. "A Class of Decomposable Poverty Measures", *Econometrica* 52 (3): 761–766.
- Fujii, Tomoko. Forthcoming. "Micro-level Estimation of Child Undernutrition Indicators in Cambodia", *World Bank Economic Review*.
- Gangopadhyay, S., P. Lanjouw, T. Vishwanath, and N. Yoshida 2010. *Identifying Pockets of Poverty: Insights from Poverty Mapping Experiments in Andhra Pradesh, Orissa and West Bengal*, forthcoming *Indian Journal of Human Development*.
- Henderseon, J.V., Z Shalii, and A.J. Venables. 2001. "Geography and Development" *Journal of Economic Geography* 1:81–105
- Haut Commissariat au Plan, Maroc. 2005. *Pauvreté, Développement Humain et Développement Social au Maroc: Données Cartographiques et Statistiques d'après le Recensement Général de la Population et de l'Habitat* (Rabat: Haut Commissariat au Plan).
- Healy, A., and S. Jitsuchon 2007. 'Finding the Poor in Thailand' *Journal of Asian Economies* 18(5), 739–759.
- Instituto Brasileiro de Geografia e Estatística, IBGE. 2003. 'Mapa de Pobreza e Desigualdade Municípios Brasileiros' DVD ROM available from IBGE at: http://www.ibge.gov.br/english/presidencia/noticias/noticia_visualiza.php?id_noticia=1293&id_pagina=1

- INSTAT (2004). Poverty and Inequality Mapping in Albania. (Tirana, INSTAT).
- Kapadia-Kundu, N., and T. Kanitkar. 2002. "Primary Healthcare in Urban Slums". *Economic and Political Weekly*, December 21.
- Kenya Central Bureau of Statistics. 2003. *Geographic Dimensions of Well-Being in Kenya* (Regal Press Kenya, Nairobi).
- Krugman, Paul. 1999. 'The Role of Geography in Development' *International Regional Science Review*, 22(2):142–61.
- Lanjouw, Peter, and Jean O. Lanjouw. 2001. 'The Rural Non-Farm Sector: Issues and Evidence from Developing Countries', *Agricultural Economics* 26 (2001) 1–23.
- Lanjouw, Peter, and Rinku Murgai. 2009. 'Agricultural Wage Labour, Non-Farm Employment and Poverty in Rural India', *Agricultural Economics* 40 (2009) 243–263.
- Lanjouw, Peter, and Erika Rascón. 2010. 'Town Size and Malnutrition in Mexico', mimeo, DECRG, the World Bank.
- Lipton, Michael. 1977. *Why Poor People Stay Poor: Urban Bias and World Development*. London: Temple Smith.
- López-Calva, L., A. Melendez, E. Rascón, L. Rodríguez-Chammusy, and M. Székely (2005) 'Poniendo al Ingreso de los Hogares en el Mapa de México' Working Paper EGAP-2005–04, Tecnológico de Monterrey, Campus Ciudad de Mexico.
- Luttmer, Erzo F.P. 2005. "Neighbors As Negatives: Relative Earnings And Well-Being," *Quarterly Journal of Economics*, V120 (3, Aug), 963–1002.
- Molina, I., and J.N.K. Rao (2010) 'Small Areas Estimation of Poverty Indicators' *The Canadian Journal of Statistics* 38 (3): 369–385
- Prebisch, Raul. 1950. "The Economic Development of Latin America and its Principal Problems", CEPAL.
- Ravallion, Martin, Shaohua Chen, and Prem Sangraula. 2007. "New Evidence on the Urbanization of Global Poverty". *Population and Development Review* 33(4), 667–701.
- Rascón, Erika. 2010. "Cracking Malnutrition: an Resolved Challenge for Mexico?: Micro-Level Estimation of Child Malnutrition Indicators", mimeo, Development Economics Research Group, the World Bank.
- Rosenstein-Rodan, Paul. 1943. "Problems of Industrialization in Eastern and Southeastern Europe", *Economic Journal*, 53, pp. 202–12.
- Tarozzi, A., and A. Deaton 2009. 'Using Census and Survey Data to Estimate Poverty and Inequality for Small Areas' *Review of Economics and Statistics* 91(4) 773–792.
- Thomas, T. (2008) 'Is There Evidence of an Urban-Rural Spillover? An Empirical Analysis Using Spatially Explicit Data from Three Countries', mimeo, Development Economics Research Group, the World Bank.
- World Bank (2007) *Measuring Poverty Using Household Consumption* Brazil Poverty Assessment, Poverty Reduction and Economic Management Network, The World Bank.
- (2009) *World Development Report: Reshaping Economic Geography* (Washington D.C.: World Bank).
- (2011) *Perspectives on Poverty in India: Stylized Facts from Survey Data*, India Poverty Assessment, Poverty Reduction and Economic Management Network, the World Bank.