COMPARING THE HEALTH AND WEALTH PERFORMANCE OF METROPOLITAN REGIONS

Antoine Paccoud
LSE Cities, London School of Economics
Houghton Street, London WC2A 2AE
United Kingdom
a.m.paccoud@lse.ac.uk

Abstract: This paper presents a methodology to construct comparable estimates of health and wealth performance for 126 metropolitan regions globally that puts spatial comparability on an equal footing with data comparability. It will be used to investigate the relationship between health and wealth performance at the metropolitan level. The point of departure is the construction of a new spatial unit, the Extended Metropolitan Region (EMR), based on sub-national administrative units that proxy the maximum spatial extent of the city under consideration. A standardised ratio of EMR to national performance is then computed in the health and wealth dimensions based on available data in each national context and applied to internationally comparable national-level indicators to obtain an estimate of EMR performance in the health and wealth dimensions. EMRs are then divided into two main types: those which over or underperformed their national contexts to an equal degree, and those which did not.The hypothesis that will be tested is that EMRs where the percentage of private expenditure in total health expenditure is high can have a health performance that matches their wealth performance only if the level of national inequality is low.

Keywords: Comparative, metropolitan regions, inequality, health, wealth

COMPARING THE HEALTH AND WEALTH PERFORMANCE OF METROPOLITAN REGIONS

1. INTRODUCTION

This paper presents a methodology to construct comparable estimates of health and wealth performance for 126 metropolitan regions globally that puts spatial comparability on an equal footing with data comparability. It will be used to investigate the relationship between health and wealth at the metropolitan scale and to test a hypothesis concerning the relation between health and wealth performance at the metropolitan level, the financing national health systems and national income inequality. The point of departure is the construction of a new spatial unit, the Extended Metropolitan Region (EMR), based on sub-national administrative units that proxy the maximum spatial extent of the city under consideration. A standardised ratio of EMR to national performance is then computed in the health and wealth dimensions based on available data in each national context and applied to internationally comparable national-level indicators to obtain an estimate of EMR performance in the health and wealth dimensions. EMRs are then divided into two main types: those which over or underperformed their national contexts to an equal degree, and those which did not.The hypothesis that will be tested is that EMRs where the percentage of private expenditure in total health expenditure is high can have a health performance that matches their wealth performance only if the level of national inequality is low.

2. ESTIMATING METROPOLITAN HEALTH AND WEALTH PERFORMANCE

2.1 A New Spatial Unit

In order to construct comparable estimates of metropolitan health and wealth performance, the first step is to define what is meant by a metropolitan region. This is because comparisons across metropolitan regions are only valid if they represent the same kind of entity: administrative cities with administrative cities, metropolitan regions with metropolitan regions, etc. The difficulty with this in practice is that no two countries administratively organize their territories in the same way: some define metropolitan regions, some do not, some create administrative boundaries around the central area of their cities, and some do not, etc. It is thus not possible to rely on existing spatial units if we want to be able to compare the same type of urban entity in all places.

The use of existing spatial units in international comparisons of cities and metropolitan regions undermines many such efforts. An awareness of the differences between the administratively defined entity (whether it be a city or a metropolitan region) and the reality on the ground is key to approaching spatial comparability. This is because a city's spatial extent, in terms of the urban fabric that constitutes it, has not consideration for administrative boundaries. When studies such as Mercer's Quality of Life index, Kearney's Global City Index, the Economist Intelligence Unit's Green City Indices or Angel et al's Atlas of Global Urban Expansion (2010) compare cities or metropolitan regions, they are in fact comparing different types of spatial units. Even the United Nation's World Urbanisation Prospects database, a key reference in the comparison of metropolitan regions (or urban agglomerations as they refer to them), is forced by missing data to substitute cities for metropolitan regions. This means that the city of Jakarta's population of 10 million is compared, as a metropolitan region, to Tokyo's 37 million, when Indonesian estimates for the Jakarta metropolitan region range from 28 to 34 million inhabitants. Needless to say that any values of indicators that take Jakarta as a 10 million city or a 30 million metropolitan region will be wildly different. Studies which uncritically use the World Urbanisation Prospect database, such as Pricewaterhouse Coopers (2009) and McKinsey (2012), are thus also affected by these inconsistencies. The issue of the spatial comparability of metropolitan regions is thus of critical importance to any attempt at constructing comparable measures of metropolitan performance.

In order to avoid the problems highlighted above, I decided to construct a new spatial unit to be able to compare the maximum spatial extent of metropolitan regions across all world regions. In order to come up with a suitable alternative to existing classifications, it is necessary to step back and look at what is presented to us at a global scale: almost 200 countries, each sub-divided in their own particular way into varying levels of sub-national entities, and cities which are either contained within one unit or are spread out over many. From this perspective, it can be seen that it is the relationship between cities and these administrative divisions which is of crucial importance. Thus, in order to achieve geographical comparability, we must focus on establishing a consistent relationship between city and administrative boundary that is to be sought in the different national contexts. The new spatial unit at the heart of this project, 'the Extended Metropolitan Region' (EMR), depends on satellite imagery to ascertain a city's maximal spatial extent. This information was then used to decide which existing national administrative sub-divisions that are the basis for statistical data collection could be used as a proxy for the city's maximal spatial extent.

In order to make sure that this proxy for the city constructed using administrative units maintained a relatively consistent relationship to its city in different national contexts, I compared the population obtained through the proxy with the population of its urban agglomeration contained in the United Nation's World Urbanisation Prospects database, supplemented by national metropolitan estimates where those of the UN were problematic. An EMR can thus be defined as administrative unit or combination of administrative units which contain(s) the largest spatial extent of the city yet stay(s) within reasonable bounds of the population of the urban agglomeration as defined by the UN (and national statistical institutes where necessary). This ratio is crucial because it affects the estimation procedure that will be detailed below and which relies on assessing the degree to which an EMR over or under performs its national context. The ratio between the EMR population and that of the UN or national metropolitan estimates could impact the degree to which an EMR under or over performed in relation to its national contexts because EMRs which largely exceed what is generally considered as the metropolitan region will include more rural or sparsely inhabited land, which usually perform less well compared to the national average than more urbanised territory. I thus made sure I included only EMRs which remained within reasonable bounds of the UN or national metropolitan estimate.

The percentage of the national population living in an EMR could also have an impact on the relation between EMR and national context, through the weight the EMR values would have on national averages. Indeed, because of data availability constraints, it was not possible to take the EMR value out of national averages: in many national contexts, no full dataset of all administrative units exists for all indicators, which means that the administrative units making up the EMRs had to be manually extracted. An EMR that makes up a very small percentage of the national population could, all others things held equal, see a much larger divergence from national conditions than an EMR which makes up a large percentage of the national population. There were four countries with large populations (China, India, the United States and Brazil) in which EMRs represented a percentage of the national population that was much smaller than in the other countries within my sample. This meant that the EMRs in these countries under or over performed their national context to a larger degree than other EMRs. To avoid this, I decided to use in these four countries. In China, provinces were used as national context and sub-provincial cities to build the EMRs, in India states and districts were used, in the U.S. the 9 regional divisions of the Census Bureau stood as national contexts and congressional districts were used for the EMRs, and in Brazil the 5 regions of the Brazilian Institute of Geography and Statistics (IBGE) and municipalities were used. Shifting down one level was found to be necessary to align the percentage of national population an EMR represented to the level of other EMRs.

With this method, a geographically representative sample of 126 EMRs from all world regions was assembled: 21 in Central and South America, 15 in East Asia, 8 in Eastern Europe, 11 in the Middle East and North Africa, 14 in North America and Australia, 21 in South Asia, 12 in South East Asia, 15 in Sub-Saharan Africa and 9 in Western Europe. To make data collection more manageable, only EMRs that exceeded the 1.5 million inhabitants mark were included. City states such Hong Kong and Singapore and were not included given that the focus here is on the discrepancies that exist between metropolitan regions and their national contexts. National population figures were used as a guideline to decide how many EMRs each country should contribute to the dataset, but the final list of EMRs was determined by the countries where the necessary data was available. The list of EMRs in this database was arrived at by balancing the number of EMRs in the countries available in each main region with the overall number of EMRs per million inhabitants of the geographical region that country is located in. Here is a breakdown of the number of EMRs per million inhabitants in each of these regions: 1 EMR for 26 million inhabitants for the 11 national contexts in Central and South America, 1 EMR per 64 million for the 13 in East Asia, 1 EMR per 34 million for the 6 in Eastern Europe, 1 per 28 million for the 6 in the Middle East and North Africa, 1 per 29 million for the 8 in North America and Australia, 1 per 69 million for the 14 in South Asia, 1 per 57 million for the 8 in South East Asia, 1 per 29 million for the 11 in Sub-Saharan Africa, and 1 per 40 million for the 9 in Western Europe.

The full list of these EMRs and the administrative sub-divisions that constitute them can be found in Annex 1.While there are limitations inherent in producing such estimates, they have the advantage of taking the spatial extent of cities on the ground into consideration. They are comparable because they represent a consistent relationship between the form of the urban fabric on the ground and the administrative boundaries that abstractly divide up a territory that fabric is located in.

2.2 Ensuring Data Comparability

Once the spatial comparability of the EMRs was established, a second methodological step was to guarantee the comparability of the data collected for the administrative divisions making up the EMRs. In order to make sure that these estimates were comparable across EMRs, I developed a two-tiered estimation technique. In the first instance, a standardised ratio between the performance of an EMR and that of its national context in a dimension was computed based on the systematic use of indicators available in that specific national context. This ratio, which tracks the degree to which a particular EMR over or under performs its national context in the dimensions of health and wealth, was then applied to internationally comparable national level health and wealth indicators. The EMR level health and wealth indices obtained can thus be seen as qualifications of the national indices that depend on how the EMR performs relative to its national context in that particular dimension. This estimation technique was found to be the best available at responding to the existing state of data at the sub-national level: there is no fixed set of indicators that exists for all sub-national entities needed, and much less for the same time period. The strength of this estimation technique is that it allows for different indicators within different national contexts to give us a picture of how EMRs perform relative to their nations. This allows for greater flexibility when faced with different indicators in different contexts, and also more flexibility with different time periods. Indeed, while the values of a particular indicator can change quite significantly over a period of time, the ratio between the EMR and national value of that indicator is likely to change much less. This has allowed me to look for data for the 2000 to 2010 period. I will now present the methods used to select the indicators at the EMR and national levels and the standardisation techniques used to make them comparable, and then move in the following section to the selection of the internationally comparable national indicators.

2.2.1 Health and Wealth Indicators at the Metropolitan Level

A first step was to select, within each national context, a series of health and wealth indicators that were available for both the nation and the administrative sub-divisions used to build the EMRs in that national context. This was made necessary by the absence of single indicators in either of the dimensions of health or wealth available internationally for sub-national administrative units. This can be explained by the wide variety of circumstances existing globally, which are translated into different national statistical capacities as well as priorities in data collection. Not only do Bangladesh and the United Kingdom have different resources and capacities that will affect the type of data they can collect and their level of spatial disaggregation, but they are also not necessarily interested in collecting the same kinds of indicators. For the purposes of this study, it was thus necessary to use the indicators available at a sub-national level in each national context, even if they differed across countries, and to develop a mechanism to guarantee their comparability. This meant identifying the indicators in each dimension from which it was possible to systematically implement the following procedure:

1. From within the smallest possible set of all available and adequate indicators that cover the set of administrative units making up the EMRs, identify two priority indicators and a set of other second-order indicators, and;

2. In each national context, use only the two priority indicators if they are available and assign them a 50% weight each, or, if only one of the two priority indicators is available, use the one that is available, assign it a 50% weight and give the geometric average of all other second-order indicators present the other 50% (or 75% and 25% respectively if there is only one second-order indicator available), or, if none of the two priority indicators are available, use the geometric mean of all second order indicators present.

For health, the two priority indicators were the life expectancy and the infant mortality rate. Second-order indicators that could be used were immunisation rates, skilled assistance at delivery, doctors and hospital beds per capita, and the percentage of mothers protected against tetanus. For the wealth dimension, the two priority indicators were GPD per capita and income per capita. The second-order indicators that could be used were the poverty rate, household characteristics (access to safe water, sanitation or electricity) and indicators measuring malnourishment. This procedure allowed the standardised ratio between EMR and nation to be based on as limited a set of indicators as possible.

To standardise the ratio between the value obtained by the EMR on these indicators and that of its national context was problematic because of the wide range of distributions different indicators can take. For this estimation technique to produce results that allowed for valid comparisons to be made across different national contexts, it was crucial to make sure that the ratios calculated between indicators at the EMR and national levels were comparable across indicators with different numerical distributions. To do this, I grouped indicators according to the numerical distributions they tend to take and used different standardisation procedures to calculate the EMR to national ratio for each one of those groups. A first group is made up of all variations on the life expectancy. I have chosen to take the simple ratio between EMR life expectancy and national life expectancy in this case because there are very few instances of a drastic difference between life expectancies at these two scales. A second group is made up of all indicators that are expressed as a percentage. This is the group with the most indicators, as they are usually derived from censuses and surveys. The ratio standardisation procedure I used for this group is the ratio of the square roots of the EMR and national values. Given the wide range of ratios that can be obtained from this group of indicators, I chose to use the square roots in order to reduce the overall size these ratios can take. A third group is made up of all indicators that lend themselves to a ratio standardisation procedure based on the logarithmic function, and thus indicators that are susceptible to decreasing marginal returns. These include measures of wealth (GDP per capita, household income per capita, etc.) and measures of health (infant mortality rate, doctors per 10,000). The ratio standardisation procedure chosen here takes the ratio of the logarithms of the EMR and national values (both raised or decreased by as many orders of magnitude as it is necessary to get the national value to a magnitude of 10^2).

The health and wealth factors, or the standardised ratios between EMR and national context, were arrived at by applying the standardisation techniques just presented to indicators selected systematically in each national context. These factors are then applied to the national health and wealth indices computed at the international level. It is to these that we now turn.

2.2.2 Internationally Comparable National Health and Wealth Performance Indices

Now that the systematic procedure used to select the indicators at the EMR level and the standardisation techniques used to make them comparable have been presented, we can move on to the national level indices to which the EMR health and wealth factors were then applied. To ensure that these national level indices represented an accurate picture of the situation in the country under consideration, I chose to follow the spirit of the UNDP's recently revised set of indicators and standardisation techniques (UNDP 2010). For the wealth index, I used the UNDP's indicator, gross national income per capita (PPP 2008 $). To measure health, the UNDP uses a single indicator, the life expectancy of a country's population. In order to get a more complete picture of health performance at the national level, I chose to supplement the life expectancy with the infant mortality rate from the United Nation's World Population Prospects database. Both of these indicators were weighted equally. I chose to use both indicators to measure national health performance because of the wide differences in national performance that exist for two those indicators. Using only one or the other would have led to very different assessments of the health performance of the selected countries. Some, like Russia, perform much better internationally with respect to infant mortality (rank 44 out of 143) than life expectancy (rank 92), while others, like Albania rank higher in life expectancy (31st) than in infant mortality (60th). Only 7 countries rank equally on both measures, and all nations experience an average absolute rank difference of 11.2 between the ranks obtained on each individual indicator. Using both the life expectancy (the indicator used by the UNDP and also one of the indicators most commonly used to assess health levels), and the infant mortality rate (which is more health systems based and commonly used to measure progress in development), provides a more comprehensive assessment of national health performance than either one in isolation. Sources for the comparable figures of life expectancy, infant mortality rate and GNI per capita for Chinese Provinces, Indian States, U.S. States and Brazilian Regions used here are in the bibliography.

In contrast to the UNDP which has developed a different standardisation procedure for each of the indicators it uses, I decided to use a single standardisation technique for all 3 indicators, based on the minimum and maximum values achieved on each indicator by the 143 countries in my sample (z = x – xmin / xmax – xmin). In order to make sure the standardised values obtained on each indicator either followed a normal distribution or were relatively well spread out over the 0 to 1 range, the frequency distributions of each of the 3 indicators were studied and transformations were used where necessary. For the life expectancy, the frequency distribution showed that the values were relatively well distributed, and I followed the UNDP in not applying any transformation to the data. The standardisation procedure based on the minimum and maximum values of the sample of 143 countries was applied to the life expectancy figures directly.

The frequency distribution for the infant mortality rate, in contrast, showed a large range of values (a minimum of 1.9 per 1,000 live births in Singapore and a maximum of 135.9 per 1,000 live births in Afghanistan) with a large concentration of values at the lower end and low frequencies in the middle and higher end of the distribution. The UNDP does not use the infant mortality rate to assess national health performance, but it uses a logarithmic transformation to normalise GNI per capita figures at the national level. While the logarithmic transformation does not go so far as to normalise the infant mortality rate distribution, it is warranted in this case by the more even spread it gives the data, allowing to better account for differences in infant mortality rates at the national level.

The frequency distributions for the GNI per capita with and without a logarithmic transformation are shown below. As mentioned above, the UNDP uses the logarithmic transformation to normalise the GNI per capita figures and I have adopted this procedure here. This procedure allows for a very large range of values (a minimum value of USD 176 for Zimbabwe and a maximum of USD 79,426 for Qatar) to be more evenly distributed across the 0 to 1 range.

In order to avoid values of 0 and 1, the minimum and maximum values of the data on which the standardisation procedure was applied were respectively decreased and increased by a small percentage. The values of these percentages were decided by looking at the average standardised value they would yield across 143 countries in the sample. I wanted to ensure that the average values of the health and wealth indices were similar in order to minimise any biases when comparing the values obtained by countries on these two indices.

I have now explained the methods through which estimates of the health and wealth indices of 126 EMRs have been constructed by making use of available data.

Annex 2 lists the indicators and sources used in each national context to estimate the extent by which EMRs outperform their national contexts in health and wealth. And Annex 3 presents the dataset: EMR, region, EMR population, % of national population in the EMR, relation between the EMR's population and that of the corresponding UN WUP urban agglomeration (or national metropolitan estimate where necessary), health and wealth factors, national health and wealth indices and EMR health and wealth indices.

3. INVESTIGATING THE RELATIONSHIP BETWEEN HEALTH AND WEALTH AT THE EMR LEVEL

3.1 Is there a Health Advantage at the EMR Level?

Out of 126 EMRs, there are 15 with health outcomes worse than their national contexts and 12 with wealth levels lower than their national contexts. There is undeniably a health advantage at the EMR level but it seems to go hand in hand with a wealth advantage. Is the health advantage solely due to wealth?

A multiple regression model was used to ascertain the influence of the difference in wealth between national context and EMR and the respective difference in health. The dependent variable that will be investigated here is the degree to which EMRs outperform their national context in the dimension of health, what we called the health factor above (the standardised ratio between EMR and national performance on the indicators available in its national context). As mentioned above, calculating the health factor was crucial to achieving data comparability because it was then applied to internationally comparable national-level data to produce an estimate of EMR performance. However, this factor cannot be used here directly as the dependent variable because it shows a clear link to developmental levels: developing countries tend to show a much larger gap between EMR and national health performance than in developed countries. In order to standardize the extent to which EMRs outperform their national context across different developmental levels, the difference in the number of points between the index created out of the internationally comparable national-level data and that index once it has been multiplied by the EMR health factor will be used. This indicator thus measures the extent to which EMRs outperform their national context in health with respect to their level of development. For simplicity, we will call this variable the health performance of an EMR. The same technique is used to obtain the wealth performance of an EMR. We are thus investigating the relation between an EMR's health performance and its wealth performance.

Other available indicators were included in the model. The first are variables that were used to control that there were no systematic biases in the estimation of the EMR health and wealth performance: EMR population as a percentage of national population and the ratio between EMR population and that of its corresponding urban agglomeration from the UN WUP database (corrected with national estimates where necessary). A second set of variables were used to investigate what drives health performance at the EMR level: EMR population level, population growth of the UN WUP urban agglomeration that corresponds to the EMR for the 1950 to 2010 period, estimates of EMR net density (Paccoud 2011) and EMR wealth performance. As can be seen from Table 1, only EMR wealth performance had any significant effect on health performance at that same scale, explaining over 40% of the variation in health performance alone.

It is thus essential to take EMR wealth performance into consideration when discussing the issue of health advantage at the EMR level given how important the relation between EMR and national context in wealth outcomes is for health outcomes. We cannot discuss EMR health advantages in isolation: they need to be juxtaposed to the EMR wealth advantages.

3.2 Exploring the Decoupling and Coupling of Health and Wealth Performance

At closer look at the database reveals that the relationship between health and wealth performance can take two main forms: either health and wealth performance are relatively similar, or they are quite different. In the first case, we can say that health and wealth are coupled, in the second that they are decoupled. The 25 EMRs in which either health or wealth performance is negative while the other is positive are the easiest ones to categorise. It is clear that health and wealth performance are decoupled for these EMRs. The case of Berlin is a well known illustration of this situation: while it underperforms Germany in wealth by 0.722 to 0.745, its health index of 0.808 is higher than the German average of 0.803. Here, a health performance of 5 is contrasted to a wealth performance of -23.2. There are 15 EMRs in which health outcomes are lower than in the nation as a whole, but where wealth outcomes are better. Manila represents the inverse situation, with a health performance of -6.4 and a wealth performance of 43.6, thus underperforming the Philippines in health (with a health index of 0.56 at the EMR level and 0.57 at the national level), but exceeding it in wealth (0.53 against 0.49). There are 10 such EMRs.

We are now left with 101 EMRs in which both health and wealth performance are either positive or both negative. In order to decide in which category they should be placed, I looked at the ratio between health and wealth performance. After looking at the distribution of health and wealth performance, I decided that 52 EMRs with a health to wealth performance ratio larger than 0.5 and smaller than 2 would be labelled as coupled. An example of a coupled EMR is Xi'an, which has a health index of 0.599 and a wealth index of 0.592 while Shaanxi Province, of which it is the capital, has a health and wealth index of equal value 0.547. This translates into a health performance of 52.7 as compared to a wealth performance of 44.6, a ratio of 1.18. Hanoi is in a similar situation, with a health performance of 23 and a wealth performance of 17, yielding a health index of 0.628 for the EMR and of 0.605 for its national context, while in wealth the corresponding figures are 0.472 and 0.455. Here the ratio is 1.35.

The remaining 49 EMRs, in which health performance was either less than half or more than double wealth performance, were labelled as decoupled. For example, Medan, Indonesia's fourth EMR by population, has a health performance of 27 and a wealth performance of 9.9. In index terms, Medan outperforms its national average by 0.569 to 0.542 in health and 0.498 to 0.488 in wealth. Its ratio between health and wealth performance is 2.7. Lagos, another decoupled EMR, has the inverse relation between health and wealth, a health performance of 0.7 and a wealth performance of 40.2: it share a value of 0.20 on the health index with Nigeria, but significantly outperforms it in wealth with an index of 0.46 compared to Nigeria's 0.42. Here, the ratio between health and wealth performance is 0.02.

Within the set of 126 EMRs that are featured in this study, we can thus say that there are 52 for which health and wealth performance that are coupled and 74 for which they are decoupled. EMRs that are coupled show a much larger difference in health and wealth outcomes as compared to their national contexts, with both health and wealth performance significantly larger on average than for their decoupled peers (37.2 and 44.8 respectively for coupled EMRs, 6.7 and 28.6 for decoupled EMRs). This can be explained by the fact that EMRs that underperformed their national contexts in either health or wealth were included within the decoupled category. When taking these 25 EMRs out of the decoupled, the difference in wealth performance disappears. But there is still a significant difference as concerns health performance (10.8 for this subset of the decoupled EMRs and still 37.2 for coupled EMRs). This seems to show that EMRs in which health and wealth are coupled have an advantage in health performance over their decoupled peers.

A possible explanation for this situation can be found in the geographical distribution of EMRs in both these categories. Coupled EMRs are overwhelmingly made up of EMRs from South Asia (28.8%) and East Asian EMRs (15.4%). In contrast, those that are decoupled have a large number of Central and South American EMRs (20.3%) and Sub-Saharan African EMRs (13.5%). Seen through the percentage of EMRs from each region that are coupled or decoupled, this picture stays similar for the coupled EMRs, with 71.4% of South Asian, 53% of East Asian and 50% of South East Asian EMRs that are coupled. In the case of decoupled EMRs, we see that 89% of West European, 75% of East European, 73% of Middle-Eastern and North-African, 71% of Central and South American, 67% of Sub-Saharan African and 57% of North American and Australian EMRs are decoupled. Coupled EMRs are predominantly from regions that have the highest wealth performance: South Asian EMRs have a health performance of 42.2 on average, East Asian EMRs 29.8 on average and South East Asian EMRs 23.3 on average. In contrast, those regions that are most strongly represented in decoupled EMRs have much lower health performances: 3.5 for Western Europe, 20 for Eastern Europe, 1.5 for Middle Eastern and

North African EMRs, 12.2 for Central and South American EMRs, 15.8 for Sub-Saharan African EMRs and 8.5 for North American and Australian EMRs.

Coupled EMRs are thus found in a limited set of geographical regions which all exhibit high health performances and share relatively similar developmental levels. In contrast, those that are decoupled are found across many different regions than span the spectrum of developmental levels: from Sub-Saharan Africa to Europe and North America. The coupled EMRs also share other characteristics that set them apart from their decoupled peers. These are derived from national level indicators given the lack of comparable data at the EMR level. This means that even though second level administrative entities were used to calculate the health and wealth performance of EMRs of China, India, U.S. and Brazil, we will have to use national level data to look at the differences between the coupled and the decoupled EMRs. We will thus have to assume here that the national level sets the overall frame which conditions possibilities at lower levels, and that the second level units that serve as national contexts in these four large countries do not differ too largely from the national context as a whole.

The first characteristic of these coupled EMRs is that their national contexts have, on average, a lower wealth index (0.54 against 0.58). Another contrast between these two groups concerns the national Gini index of income inequality (average for 2000 to 2006 from the UNU-WIDER database): coupled EMRs have a slightly lower Gini index on average than those that are decoupled (39.3 vs. 42.2). A final difference between these two groups comes from the percentage of total national health expenditure that came from private expenditure on health from the World Health Organisation's Global Health Expenditure database. Coupled EMRs had a significantly higher percentage of private expenditure on health (57.3%) than those EMRs where the performances are decoupled (48.4%). There is also a slight, but not significant, difference in national total expenditure on health as a percentage of GDP (6.2% for coupled EMRs and 7.2% for decoupled EMRs). Population growth for the UN urban agglomeration corresponding to the EMR over the 1950 to 2010 period does not differ significantly across groups (3.5% per year on average for coupled, 3.2 for decoupled).

The fact that 70% of Western European and Northern American and Australian EMRs are found within the decoupled EMR category can seem like an explanation of the differences in national wealth indices, net density, Gini index and private expenditure on health. However, even if we remove these EMRs (and the two Japanese EMRs from East Asia), these differences remain. The difference in national total expenditure on health as a percentage of GDP is now significant (4.9 for coupled and 5.8 for decoupled). Coupled EMRs are thus places that have a higher average health performance (41.3) than decoupled EMRs (7.5) while having a lower average national wealth index and a lower national percentage of GDP for health expenditure. This seems to be associated with a lower national Gini index and a higher percentage of total national health expenditure that came from private expenditure on health.

3.3 Health Systems, Inequality and Health Performance

It thus seems as though developing countries in which the health system is funded to a higher degree by private expenditure (mostly private health insurance and out of pocket payments) have a higher likelihood of seeing health and wealth performance coupled at the EMR level: could it be that wealth and health go hand in hand in these places because higher wealth directly translates into a higher capacity to pay for health and thus into a higher health performance? The way in which this could work is to focus on the effect that a high percentage of private expenditure in total health expenditure could mean for the territorial equality of health spending. A hypothesis here could be that high private expenditure means that there is little scope for public authorities to allocate resources equally among the administrative units that make up its territory. It is private expenditure that would drive the allocation of resources in particular places, depending on the financial resources of the population in each of those places. This in turn would have the effect of concentrating health outcomes where wealth outcomes are high. This is a way in which the relation between coupling of health and wealth performance at the EMR level and a high percentage of private expenditure in total health expenditure could be theorised. In this view, decoupled EMRs would have lower health performances than wealth performances on average because the extent to which different territories within the nation would show strong contrasts is limited by higher public spending on health, while this limit would not be imposed on wealth.

One way to test this hypothesis is to separate out from the decoupled EMRs from developing countries those with a low percentage of private expenditure in total health expenditure (15

EMRs where private expenditure represents 34.6% of total health expenditure). These 15 EMRs have a wealth performance equal to all other EMRs (40 as compared to 40.7) but a health performance that is only half of that achieved by the EMRs with a high percentage of private expenditure in total health expenditure (12 vs. 24.4). For the developing countries in my sample, the relation between health performance and the way in which the health system is financed seems to pass this simple test.

If we now turn to the remaining decoupled EMRs, these are by design similar to the coupled EMRs with respect to the percentage of private expenditure in total health expenditure. However, they have a much lower health performance, a higher health spending and a higher Gini index. There is no difference in terms of national wealth between these EMRs. If we follow the hypothesis laid out above, these decoupled EMRs can be seen as potential coupled EMRs that were not able to outperform their national average in health. Given the high proportion of private expenditure in total health expenditure at the national level, a sign that health outcomes will exhibit high territorial differentiation that follow wealth outcomes, these EMRs were expected to strongly outperform their national contexts in health. A possible cause for this failure of these decoupled EMRs to take advantage of local wealth outcomes could be the higher levels of inequality they experience. High inequality could affect access to health at the EMR level and thus block the translation of wealth performance into health performance. The fact that coupled EMRs have significantly lower inequality levels supports this hypothesis.

It thus seems as though we can distinguish three types of EMRs within those from developing regions: EMRs that have a low proportion of private expenditure in total health expenditure and low health performances (the decoupled EMRs), those with a high proportion of private spending in total health expenditure, a high Gini index and low health performances (the ‘failed coupled’ EMR) and those that have a high proportion of private spending in total health expenditure, a low Gini index and high health performances (the coupled EMRs).But what does this hypothesised relation between the financing of the health system, inequality and health performance mean for the developed EMRs in the sample? Those 25 EMRs show a strong split between North American EMRs on the one hand and the EMRs of Western Europe, Japan, Canada and Australia on the other. The United States has a much higher proportion of private spending in total health expenditure than the other developed countries in the sample (53.2% vs. 26.2%) and a higher Gini index (44.4 vs. 31.7). A higher percentage of private spending in health should push the health performance of U.S. EMRs up to the level of their wealth performance, while a high Gini index should push it down. The data shows that the U.S. EMRs have health performances that are on average double that of the other developed EMRs (8.8 vs. 4.4, but it is not a significant result), while wealth performances are very similar (12.3 vs. 15.1). The fact that the average health performance of U.S. EMRs is below the average of its wealth performance seems to show that the hypothesis has some validity for developed EMRs as well.

4. CONCLUSIONS

In this paper, I presented an estimation technique for the health and wealth performance of metropolitan regions that puts both spatial and data comparability on the same footing. On the basis of this dataset, I then argued that any discussion of health outcomes at the metropolitan level needs to take into account the wealth outcomes at that same scale. Two types of EMRs were presented, those which over or underperformed their national contexts to an equal degree, and those which did not. The characteristic that most differentiated these two types of EMRs was found to be the percentage of private expenditure in total health expenditure. An exploration of the dataset let to the positing of a relationship between health performance at the EMR level, the financing of the health system nationally and the national level of inequality: EMRs where the percentage of private expenditure in total health expenditure is high can have a health performance that matches its wealth performance only if the level of national inequality is low. This hypothesis was verified for developing and developed EMRs in the sample. The number of EMRs in the sample gives a degree of robustness to the conclusions arrived at through this dataset based on estimations but it is clear that further research is needed to better understand the dynamics at work here.

5. REFERENCES

Angel, S., Parent, J., Civco, D.L. and Blei, A.M. (2010), Atlas of Urban Expansion, Cambridge MA: Lincoln Institute of Land Policy, <http://www.lincolninst.edu/subcenters/atlas-urban-expansion/>

Brazilian Institute of Geography and Statistics (2012). Basic Data and Indicators database. < http://tabnet.datasus.gov.br/cgi/idb2010/a11tb.htm>

Brazilian Institute of Geography and Statistics (2012). Demographic Census 2010 database. < http://www.sidra.ibge.gov.br/cd/cd2010CGP.asp>

Brazilian Institute of Geography and Statistics (2012). Municipality GDP database 2005-2009. < http://www.ibge.gov.br/home/estatistica/economia/pibmunicipios/2005_2009/default.shtm>

Economist Intelligence Unit. Green City Index Reports. A research project conducted by the Economist Intelligence Unit, sponsored by Siemens. Siemens AG, Germany. <http://www.siemens.com/entry/cc/en/greencityindex.htm>

Indian Ministry of Health and Family Welfare (2012). Family Welfare Statistics in India. < http://nrhm-mis.nic.in/UI/FamilyWelfare2011/FamilyWelfareStatistics2011.pdf>

Indian Ministry of Finance (2012). Economic Survey 2009-2010 database.

< http://indiabudget.nic.in/es2009-10/esmain.htm>

Kearney, A. T. Global Cities Index, 2012. <http://www.atkearney.com/index.php/Publications/2012-global-cities-index-and-emerging-cities-outlook.html>

Measure of America (2011). American Human Development Database 2010-2011. < http://www.measureofamerica.org/maps/>

McKinsey Global Institute (2012). Urban World: Mapping the Economic Power of Cities. < http://www.mckinsey.com/Insights/MGI/Research/Urbanization/Urban_world>

Mercer (2011). Quality of Living City Rankings. <http://uk.mercer.com/articles/quality-of-living-survey-report-2011>

Paccoud, A. (2011), 'Cities, health and well-being: Methodology for an international analysis', Discussion paper, Urban Age Hong Kong Conference 18-10-11. < http://lsecities.net/files/2011/11/LSE-Cities-Metropolitan-HDI-and-Density-18-10-11.pdf>

Pricewaterhouse Coopers (2009). Which are the Largest City Economies in the World and how Might this Change by 2025? <http://www.ukmediacentre.pwc.com/imagelibrary/downloadMedia.ashx?MediaDetailsID=1562 >

United Nations Development Programme in China (2008). China Human Development Report 2007/2008. <http://hdr.undp.org/en/reports/national/asiathepacific/china/China_2008_en.pdf>

United Nations Development Programme. Human Development Report 2010 Technical Notes. <http://hdr.undp.org/en/media/HDR_2010_EN_TechNotes_reprint.pdf>

United Nations Development Programme. Human Development Report database. <http://hdr.undp.org/en/statistics/data/> (July 22, 2012)

United Nations Populations Division. (2012) World Urbanization Prospects, the 2011 Revision, 2012. <http://esa.un.org/unpd/wup/index.htm>

United Nations University and the World Institute for Development Economics Research (2008). World Income Inequality Database.
<http://www.wider.unu.edu/research/Database/en_GB/database/>

United Nations World Population Prospects database. < http://esa.un.org/unpd/wpp/Excel-Data/mortality.htm>

University of Michigan (2012). All China Data Centre's National, Provincial and Sub-Provincial Statistical Yearbooks 2011.
< http://chinadataonline.org/>

World Bank, World Development Report Database. <http://wdronline.worldbank.org/> (July 22, 2012

6. FIGURES

Figure 1: Summary output of regressing EMR health performance against population indicators, net density and wealth performance

SUMMARY OUTPUT

Regression Statistics	
Multiple R	0.6736
R Square	0.4538
Adjusted R Square	0.4264
Standard Error	19.7297
Observations	127

ANOVA

	df	*SS*	*MS*	*F*	*Significance F*
Regression	6	38801.9295	6466.9882	16.6135	0.0000
Residual	120	46711.1861	389.2599		
Total	126	85513.1156			

	Coefficients	*Standard Error*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*	*Lower 95.0%*	*Upper 95.0%*
Intercept	4.0777	9.4969	0.4294	0.6684	-14.7256	22.8810	-14.7256	22.8810
EMR to Urban Agglomeration population ratio (2010)	-5.6686	5.2798	-1.0736	0.2851	-16.1222	4.7851	-16.1222	4.7851
Urban Agglomeration population growth rate (1950-2010)	-0.0564	1.1828	-0.0477	0.9620	-2.3982	2.2853	-2.3982	2.2853
EMR Population (2010)	0.0000	0.0000	1.8603	0.0653	0.0000	0.0000	0.0000	0.0000
% of national population in EMR (2010)	-0.2346	0.2012	-1.1661	0.2459	-0.6330	0.1637	-0.6330	0.1637
Net density estimate (2011)	0.0005	0.0003	1.4310	0.1550	-0.0002	0.0011	-0.0002	0.0011
Wealth Performance (2000-2011)	0.5412	0.0599	9.0349	0.0000	0.4226	0.6597	0.4226	0.6597