



UNITED NATIONS
UNIVERSITY

UNU-WIDER

World Institute for Development
Economics Research

Discussion Paper No. 2008/04

Measurement and Explanation of Inequality in Health and Health Care in Low-Income Settings

Eddy van Doorslaer¹
and Owen O'Donnell²

April 2008

Abstract

This paper describes approaches to the measurement and explanation of income-related inequality and inequity in health care financing, health care utilization and health and considers the applicability and the feasibility of these methods in low-income countries. Results from a comparative study of 14 Asian countries are used to illustrate the main issues. The structure of health finance in low-income countries, in particular the heavy reliance on out-of-pocket payments, means that the equity issues in finance are quite different from those of concern in high-income countries. Primary concern is not with the distribution of contributions to pre-payment mechanisms but with the deterrent effect of payments on utilization and the distribution of uninsured payment risks. Measurement of inequity in utilization of health care in low-income countries is constrained by the lack of reliable measures of health that can be used to standardize .../

Keywords: health inequality, equity

JEL classification: H22, H42, H51

Copyright © UNU-WIDER 2008

¹Erasmus University; ²University of Macedonia

This study is a revised version of the paper presented at the 29-30 September 2006 UNU-WIDER conference on Advancing Health Equity.

UNU-WIDER gratefully acknowledges the financial contributions to the research programme by the governments of Denmark (Royal Ministry of Foreign Affairs), Finland (Ministry for Foreign Affairs), Norway (Royal Ministry of Foreign Affairs), Sweden (Swedish International Development Cooperation Agency—Sida) and the United Kingdom (Department for International Development).

for need. Nonetheless, sufficient is known of the distribution of need in many circumstances in order to make inferences about equity from inequality in health care use. The empirical analyses demonstrate that, in low-income countries, the better-off tend to pay more for health care, both absolutely and in relative terms. But they also consume more health care. Health care is financed is largely according to the benefit principle. Assessing the distributional performance of health systems in low-income settings therefore requires examination of finance and utilization simultaneously.

Acknowledgements

The authors are grateful to the European Commission, INCO-DEV programme for funding the EQUITAP project (ICA4-CT-2001-10015) from which this paper derives and the EQUITAP partners for permission to use illustration material from the joint working papers.

Corresponding author: Eddy van Doorslaer, Erasmus School of Economics, Erasmus University, Rotterdam, The Netherlands; E-mail: vandoorslaer@few.eur.nl

The World Institute for Development Economics Research (WIDER) was established by the United Nations University (UNU) as its first research and training centre and started work in Helsinki, Finland in 1985. The Institute undertakes applied research and policy analysis on structural changes affecting the developing and transitional economies, provides a forum for the advocacy of policies leading to robust, equitable and environmentally sustainable growth, and promotes capacity strengthening and training in the field of economic and social policy making. Work is carried out by staff researchers and visiting scholars in Helsinki and through networks of collaborating scholars and institutions around the world.

www.wider.unu.edu

publications@wider.unu.edu

UNU World Institute for Development Economics Research (UNU-WIDER)
Katajanokanlaituri 6 B, 00160 Helsinki, Finland

Typescript prepared by Lorraine Telfer-Taivainen at UNU-WIDER
Printed at UNU-WIDER, Helsinki

The views expressed in this publication are those of the author(s). Publication does not imply endorsement by the Institute or the United Nations University, nor by the programme/project sponsors, of any of the views expressed.

ISSN 1609-5774

ISBN 978-92-9230-083-8 (printed publication)

ISBN 978-92-9230-084-5 (internet publication)

1 Introduction

It has been argued that in matters of health and health care, the public attaches greater importance to the achievement of equity than to efficiency (MacLachlan and Maynard 1982). Whether this is true or not can be debated but even if it is not given primary importance, health equity is certainly a goal that attracts strong support in many countries. This is beginning to be reflected in academic research. Recent decades have witnessed a dramatic expansion of the literature on health equity. Not only has the number of articles with the word ‘equity’ in the abstract grown rapidly, but their share of all articles published in Medline, for instance, has grown by 260 per cent in the last 25 years (O’Donnell et al. 2007b). Various factors have contributed to this development. An increased interest and awareness among international organizations, governments and NGOs worldwide is certainly one factor. But the increased availability of micro datasets and the development of new analytic methods also must have played an important role.

In the health economics literature, the work on the measurement and explanation of inequalities in health and health care has drawn a lot on analogies with the literature on the measurement of inequalities in, and the redistribution of, income. Wagstaff and van Doorslaer (2000a) reviewed this literature until 2000 with respect to three key topics: equity in health care finance, equity in health care delivery, and inequalities in health. The focus was exclusively on income-related inequality and inequity in finance, utilization and health using measures that derive from egalitarian-based concepts of equity which were at the heart of the development of most health care systems in OECD countries. As a result, all inequity analyses proposed and used draw heavily on rank-based measures of inequality, such as concentration curves and indices. Measurement methods for tax progressivity and income redistribution from the public finance literature could be used without much adaptation to analyze health care finance. Concentration curves and indices also proved very useful for the analysis of health inequality and, with some elaboration to take account of variation in need, of inequity in health care utilization. These methods were widely used for comparisons both within and across countries of the extent to which health care systems were achieving their egalitarian goals with respect to health finance (e.g. Wagstaff et al. 1999), health care delivery (e.g. van Doorslaer et al. 2000) and population health (e.g. van Doorslaer et al. 1997). While these analyses undoubtedly provided useful comparative descriptive information to the policy debates in OECD countries, invariably they all referred to high-income countries, with health care systems that typically had (nearly) universal coverage of their (relatively homogeneous) populations for a fairly comprehensive package of services.

In recent years, attempts have been made to apply similar methods to the analysis of health equity in low-income countries, which typically lack breadth of coverage with respect to both population and services. Some of the methods that have proved fruitful for the examination of equity in the finance and delivery of OECD health systems appear to be almost directly transferable to low-income contexts. But in other cases conceptual and/or practical considerations make the methods less suited to the analysis of the primary health equity issues of low-income countries. In this paper we aim to

review what we have learned from a recently completed large-scale cross-country comparative research project about health equity and the challenges of analyzing it in low-income countries. The EQUITAP Project on equity in the finance and delivery of health care in the Asian-Pacific Region (www.equitap.org), was funded by the European Union as a research collaboration between two European and fourteen Asian research teams with the explicit aim to see to what extent the methods developed for the analysis of equity in European health care systems could provide similarly useful information for the equity comparison of Asian-Pacific health systems.

For each of three distributions relevant to equity analysis—health payments, health care utilization and health status—we will review briefly (a) the standard measurement approach, (b) the required adaptation or qualification for the analysis of equity in low-income settings and (c) the findings derived to date in the Asian comparisons. The concern about these three distributions derives mostly from the widely perceived social aversion to inequality in the distributions of both health and income. A social welfare function in which welfare is rising with the means of health and income and falling with their variances might be appealing to many. Consistent with this position, some have argued for analysis of all health inequality, just as for income inequality (Gakidou et al. 2000). A less demanding concept of social justice recognizes that much of the variation in health is a matter of luck but takes exception to health disparities that derive from social, economic or ethnic status. Income-related inequality in health, for example, may be interpreted as inconsistent with the principle that all have an equal right to a long and healthy life. From this perspective, equity concerns about the distribution of health care by income derive from the more fundamental concern about the distribution of health by income, which they are expected to influence. Equity concerns about the distribution of health care payments by income are partly derived from concerns about how this affects the distribution of health care, and consequently health, but also, and importantly, because they affect the distribution of income. So, the distribution of income is not always only of instrumental interest to health sector equity. Sometimes it is the distribution of fundamental interest and the equity concern is about how the health sector impact on it. In what follows, we aim to make this clearer using some examples taken from the Asian comparative study. We start with a discussion of the distributional consequences of health care payments in Section 2. Section 3 discusses equity analyses of the distribution of health care utilization while Section 4 addresses the measurement and explanation of health inequality. Section 5 provides conclusions and a discussion.

2 The distribution of health care payments

The analysis of equity in the finance of health care has traditionally focused on the measurement of progressivity (Wagstaff et al. 1992; Wagstaff et al. 1999) and the income redistributive effect (van Doorslaer et al. 1999) of health sources in high-income countries. The main motivation was to determine the extent to which alternative financing mechanisms complied with the *ability to pay principle*. Progressivity analysis was used to measure *vertical equity*, i.e. the extent to which those with unequal incomes make unequal contributions to the financing of health care, by assessing the deviation of health payment concentration curve from the Lorenz curve describing the income distribution. The degree of to which payments departed from proportionality to incomes


was summarized using the progressivity index proposed by (Kakwani 1977). *Horizontal inequity*, i.e. the degree to which equals contribute unequally, was measured as the additional redistributive effect, over and above that due to departures from proportionality, as a result of the differential payments of those on equal incomes using an approach introduced by (Aronson et al. 1994).

To a large extent, the distribution of the economic burden of health care is determined by the structure of financing: the split between direct payment and pre-payment and the relative contributions of taxation, social insurance and private insurance to the latter. The balance between direct payment and pre-payment determines the extent to which actual or potential users pay for health care. Direct payments are charged in relation to actual costs, while private insurance premiums are related to (pooled) expected costs. Taxation and social insurance break the link between use of health care, realised or expected, and financial liability. Instead, liability can be made a function of ability to pay. Taxation addresses the bill for health care to the taxpayer and, indirectly, it is mostly workers and consumers that pick up this bill. Social insurance places the main burden on workers. Private insurance accumulates funds from those that choose to insure against the cost of future illness. Analyses of the distribution of health payments in OECD countries was instructive in showing how alternative mixes of health care financing sources (taxes, social insurance, private insurance and direct payments) may lead to very different consequences for progressivity and redistribution. We just give two examples here. It was, for instance, very revealing to find that mainly tax-financed systems (like the Scandinavian ones) are not necessarily more progressive than mainly social-insurance financed health systems (like France or Italy). Similarly, it was interesting to learn that the income redistributive effect (from rich to poor) of tax payments going into health care was larger in the US (a largely privately funded system) than in the UK (a largely publicly funded system), not because US direct taxes are more progressive, but simply because the size of the fraction of income going to health funding and raised through taxes is higher in the US than in the UK. To a large extent, the distribution of the economic burden of health care is determined by the structure of financing: the split between direct payment and pre-payment and the relative contributions of taxation, social insurance and private insurance to the latter. The balance between direct payment and pre-payment determines the extent to which actual or potential users pay for health care. Direct payments are charged in relation to actual costs, while private insurance premiums are related to (pooled) expected costs. Taxation and social insurance break the link between use of health care, realised or expected, and financial liability. Instead, liability can be made a function of ability to pay. Taxation addresses the bill for health care to the taxpayer and, indirectly, it is mostly workers and consumers that pick up this bill. Social insurance places the main burden on workers. Private insurance accumulates funds from those that choose to insure against the cost of future illness.

2.1 Distribution of health financing in Asia

O'Donnell et al. (2008) have analysed the structure and the distribution of health care financing in 13 territories that account for 55 per cent of the Asian population. Private insurance plays a relatively minor role in most of the health systems considered. The main distinguishing factor in these health financing systems is the balance between public pre-payment and private out-of-pocket (OOP) payment. The latter accounts for at least 30 per cent of total expenditure on health (TEH) in all territories, except Japan.

Figure 1: Out-of-pocket (OOP) and general government taxes (TAX) as share of total expenditure on health




Source: O'Donnell et al. (2008).

As is clear from the health financing triangle in Figure 1, these Asian territories conform to the stylised fact that reliance on OOP payments declines with the level of development. Poorer territories such as Nepal, Bangladesh, Kyrgyz, Punjab and Sri Lanka are all very close to the 45° line, indicating that health care is financed almost exclusively from OOP and general government revenues (GGR). Nepal and Bangladesh rely more heavily on OOP, while the burden is close to being evenly split in Kyrgyz, Punjab and Sri Lanka. The distance from any point to the 45° line gives the share contributed by insurance, social and/or private. The high/middle-income territories—Japan, Taiwan and South Korea—lying furthest below the 45° line have significant social insurance systems. Hong Kong, the only other high-income territory, also relies predominantly on publicly financed pre-payment, but does this through taxation rather than social insurance. Japan and Taiwan collect more than half of health system funding from social insurance, while high co-payments in Korea mean that it still collects half of all financing from direct payments. The remainder of financing in China comes from social insurance (16.5 per cent) and community financing arrangements (8.2 per cent).

Using survey data on household payments, incidence assumptions and health accounts data on aggregate expenditures by source, O'Donnell et al. (2008) have estimated the distributions of each financing source and of total health financing by ability to pay. A selection of results is presented in Figure 2, showing concentration and Kakwani indices for total payments, and obtained as a weighted sum of the various sources (i.e. taxes, social and private insurance premiums and direct payments), weighted by their (macro) revenue shares. The positive concentration indices indicate that in all these territories, high-income households contribute more than low-income households to the financing of health care. But only in low- and lower middle-income countries and in Hong Kong

are the Kakwani indices positive, indicating that the better off contribute more as a proportion of their ability to pay. The disproportionality is in the opposite direction in three high/middle income territories operating universal social insurance. Disaggregated results presented in O'Donnell et al. (2008) show that direct taxation is the most progressive source of finance and is most progressive in poorer economies with a narrow tax base. The distribution of out-of-pocket (OOP) payments also depends on the level of development. In high-income economies with widespread insurance coverage, OOP payments absorb a larger fraction of the resources of low-income households. In poor economies, it is apparently the better off that spend relatively more OOP. At first sight, this appears to contradict much of the evidence on regressiveness of direct payments in high-income economies (Wagstaff et al. 1999), but it may merely illustrate that the poor simply cannot afford to pay for health care in low-income economies. Among the high-income territories, Hong Kong is the one example of progressive financing arising from reliance on taxation, as opposed to social insurance, and an ability to shield those on low-incomes from OOP payments. Thailand has a similar financing structure and achieves a similar distributional outcome.

Figure 2: Concentration and Kakwani indices for total health financing



Source: O'Donnell et al. (2008).

So the short answer to the question ‘who pays for health care in Asia?’ is that the better-off pay more. Does this mean that health care financing in these countries is very redistributive? No, the picture is somewhat misleading because with the exception of Hong Kong, Korea, Taiwan and Japan, and more recently Thailand, none of these countries have ensured close to universal coverage of a fairly comprehensive package of health services. Only these higher income countries have divorced the link between payment for care (on the basis of the *ability to pay principle*) and receipt of health care (largely on the basis of the *need principle*, see below) and a redistribution interpretation can be placed on the progressivity of health payments as in other OECD countries. For all others, the lack of coverage for a substantial share, if not the majority, of the population and the reliance on out-of-pocket direct payments for funding at least 50 per

cent—and often much more—of their health care, means that it is the *benefit principle* that predominantly governs the distribution of payments for health care. The distribution of OOP payments largely reflects the distribution of benefits (in the absence of fee waivers) and must be interpreted as such. While in high-income economies the distribution of health financing is of interest largely because of its implications for the distribution of income, in low-income economies it is the consequences of health care financing for the distribution of health care, and subsequently health, that is of primarily of interest. Nonetheless, health payments can have a substantial effect on the economic welfare of households in low coverage situations. But it is the income risk arising from these payments, rather than income redistribution, which is of greatest interest. With restricted health insurance cover, large, unforeseen expenditures on health care can have catastrophic consequences for living standards and, in the extreme, may push households into, or further into, poverty (Wagstaff and van Doorslaer 2003; Xu et al. 2003). In the next two sections we discuss how the threat that out-of-pocket health payments pose to living standards can be analysed and illustrate the methods with findings for Asian countries.


2.2 Catastrophic health payments in Asia

Health care can be expensive. In the absence of insurance cover, households with severe and immediate medical needs can be forced to spend a large fraction of the household budget on health care. Such spending must be accommodated by cutting back on consumption of other goods and services, by accumulating debt, by running down savings or by selling assets. Whichever the financing strategy adopted, the household suffers a cost that may be labelled ‘catastrophic’. The concept of catastrophic payments has been put into operation by defining them as occurring once OOP payments cross some threshold share of household total expenditure (Berki 1986; Wyszewianski 1986; Pradhan and Prescott 2002; Wagstaff and van Doorslaer 2003; Xu et al. 2003). While it is acknowledged that the choice of threshold is arbitrary, 10 per cent of total expenditure has been a common choice (Pradhan and Prescott 2002; Ranson 2002; Wagstaff and van Doorslaer 2003); with the rationale that this represents an approximate threshold at which the household is forced to sacrifice other basic needs, sell productive assets, incur debt, or be impoverished (Russell 2004).

Prevalence of catastrophic payments can be measured by a headcount—the percentage of households spending more on health care than some threshold fraction of resources. The threshold may be defined as a fraction of total or, given that food spending is close to subsistence level and is less discretionary in very poor households, of non-food expenditure (Wagstaff and van Doorslaer 2003). Van Doorslaer et al. (2007) have estimated the prevalence of catastrophic health payments in fourteen countries and territories in Asia, together accounting for 81 per cent of the Asian population. The findings can be summarized as follows. First, heavy reliance on OOP financing has important consequences for household living standards. Figure 3 presents prevalence estimates at four different threshold values. The headcount necessarily falls as the threshold is raised but changing the threshold does not substantially affect the ranking of countries with the highest/lowest incidence of catastrophic payments. OOP payments for health care absorb more than 10 per cent of household resources in at least one-tenth

of all households in Bangladesh, China, India, Korea and Vietnam.¹ Alternative measures of catastrophic payments that, analogous to the poverty gap, reflect their intensity as well as prevalence can be calculated. Since the majority of the population does not incur catastrophic payments, the prevalence tends to dominate such statistics and the general pattern across countries is similar to that for the headcount (van Doorslaer et al. 2007).

Figure 3: Percentage of households incurring catastrophic payments: various thresholds for OOP as % of total expenditure



Source: Van Doorslaer et al. (2007).


Second, the lower income countries, with usually higher OOP finance shares, have a higher prevalence of catastrophic payments (see Figure 3), although there is still substantial variation at similar OOP financing contributions. China, for instance, relies on OOP financing only slightly more than Indonesia, but the prevalence of catastrophic payments is much higher in China than Indonesia. This does not appear to be simply a reflection of the fact that, on average, Chinese are better-off than Indonesians since the difference exists even when catastrophic payments are defined with respect to non-food expenditure. Clearly, the propensity to spend on medicine is higher in China than Indonesia and there is less protection against very high medical bills that exhaust a substantial share of household resources. There is some evidence that government intervention in Indonesia is effective in reducing exposure to catastrophic health payment risks (Pradhan and Prescott 2002).

Third, there are also important differences across countries in the distribution of catastrophic payments. Figure 4 presents concentration indices for the incidence of catastrophic spending. It shows that in high-income territories, catastrophic payments tend to be evenly distributed, or even slightly concentrated on the less well-off. In most low-income countries, however, it is households with higher total expenditure that are

¹ When the catastrophic threshold is defined as a proportion of non-food expenditures, poor economies like Nepal and Kyrgyztan move to a higher—and richer ones like Korea, Taiwan and Hong Kong to a lower—rank in the distribution.

more likely to spend a large fraction of those resources on health care. This reflects the inability of the poorest of the poor to divert resources from basic needs. But in some countries, it also seems to reflect the protection of the poor from user charges. In China, Kyrgyz and Vietnam, where there are no exemptions of the poor from charges, the poor are as likely, or even more likely, to incur catastrophic payments.

Figure 4: Concentration indices for catastrophic headcount (OOP>10% total exp.)



Source: Van Doorslaer et al. (2007).

In sum, there is cross-country variation in the prevalence and distribution of catastrophic payments that seems to be attributable to differences in national income, financing structure and user charging policy. Economic development is certainly an important determinant of the degree to which household welfare is put at risk by health payments, but there is no iron law that condemns the households of low-income countries to suffer financial hardship because of these payments. Some countries, in particular Malaysia, the Philippines, Sri Lanka and Thailand have managed to contain the OOP health financing share below the average level at their national incomes. In contrast, Bangladesh, China, India, and Vietnam, stand out in relying heavily on OOP financing and having a high incidence of catastrophic payments. While the second group of countries is, in general, poorer, there is, for instance, little difference between the average incomes of China and Sri Lanka.

Descriptive studies of catastrophic payments such as described above have obvious limitations. First, they do not examine the impact of OOP payments on the quantity and quality of care consumed and need to be complemented by studies of health care utilization. Second, they do not capture all potentially catastrophic effects of illness or disability, such as lost earnings. Third, they do not investigate whether health shocks are absorbed by using borrowing or dis-saving to smooth consumption (Gertler and Gruber 2002). If households can cope in the short run by using borrowing, savings or assets to finance health care and so protect other consumption, then the impact in the current period may be overstated. Flores et al. (2008) take into account self-reported coping behaviour with high outlays for inpatient care in India and find that coping-adjusted catastrophic expenditure ratios are substantially lower.

2.3 Impoverishment

Paying for health care may push households into, or further into, poverty. Such impoverishment is not captured by the standard measures of poverty that compare total household resources, including those exhausted by health care, with a poverty line that reflects needs for food and possibly those for other basic necessities but cannot take full account of health care needs. The variability and unpredictability of health care costs means that they cannot be reflected in a given poverty line. If expenditures on health care were completely non-discretionary, constituting resources that are not available to meet other basic needs, then it would be appropriate to assess poverty on the basis of household resources net of payments for health care.² Of course, not all expenditures on health care are made without discretion. There is ample evidence that such expenditures are responsive to incomes and prices. Nonetheless, it is likely that households make great sacrifices in order to meet needs for vital health care. It seems inaccurate to categorise a household as non-poor simply because high medical expenses raise its total spending above the poverty line, while spending on food, clothing and shelter is below subsistence levels.

The difference between poverty estimates derived from household expenditures gross and net of OOP payments for health care provides a rough approximation to the poverty impact of such payments (Wagstaff and van Doorslaer 2003; Gustaffson and Li 2004). Van Doorslaer et al. (2006b) have estimated the change in the poverty headcount ratios for 11 low-to-middle-income countries in Asia by comparing household consumption/expenditure both gross and net of OOP health payments relative to two poverty lines. Their baseline poverty estimates, shown in Table 1, are quite consistent with those of the World Bank (Chen and Ravallion 2004). At the \$1.08 poverty line, subtracting OOP payments from total resources results in a 3.8 percentage point increase in the headcount in Bangladesh, equivalent to almost 5 million people, a 3.7 percentage point increase in India (over 37 million) and a 2.6 percentage point increase in China (32.4 million). The total estimated increase in the poverty headcount is 78.25 million people, or 2.7 per cent of the population of these eleven low/middle-income Asian countries. This does not, of course, provide an estimate of how poverty would change if some form of pre-payment replaced OOP financing of health care. Identification of such an effect would require tracing the impact of such a reform on households' utilization of health care, work effort, consumption and savings. Nonetheless, the figure is informative of the magnitude of the impoverishing effect of payments for health care that is not currently reflected in poverty estimates. It tells us how many individuals are not counted as poor despite the fact that the value of their consumption of all goods and services other than health care is less than the extreme poverty line of \$1.08 per day.

In absolute percentage point terms, the largest increases in poverty at the lower poverty line are in Bangladesh, India, China and Nepal. Of course, the number of individuals pushed into poverty by OOP payments is greatest in India and China. The relative increase in poverty is greatest, by far, in Vietnam, where the poverty rate rises by a third. It rises by 18.9 per cent in China, 16.8 per cent in Bangladesh and 11.9 per cent in India. As we saw in previous sections, these are the countries with the highest OOP

² A National Academy of Sciences Panel made this recommendation as the appropriate approach to measuring poverty in the USA (Cirto and Michael 1995). Alternative estimates of USA poverty based on the approach are available (Short et al. 1999; Short and Garner 2002).

budget shares and prevalence of catastrophic payments. It would appear to be both the high levels of OOP payments and their even distribution throughout the income distribution that is responsible for the very high poverty impacts in Vietnam and China. But there are still large poverty impacts in Bangladesh and India, where OOP payments are more heavily concentrated on the better-off.

Regression analysis confirms that the percentage point change in the poverty headcount is positively correlated with the OOP financing share and, as would be expected, with the initial headcount (van Doorslaer et al. 2006b). Deviations from the positive correlation between the initial poverty headcount and the absolute poverty impact are more interesting than the relationship itself. For example, the initial headcount is higher in the Philippines than it is in China but the poverty impact is more than four times greater in China. And initial headcounts are similar in Sri Lanka and Vietnam but the poverty impact in Vietnam is four times that in Sri Lanka. These differences point to the consequences of high reliance on OOP financing in China and Vietnam. The proportion of the population at risk of falling below the \$1 threshold, defined as those initially between the \$1 and \$2 thresholds, is not significantly correlated with the poverty adjustment, suggesting that there is substantial cross-country variation in the extent to which vulnerable individuals are protected from the impoverishing effects of health payments. This can be seen directly in Figure 5. Roughly one-half of the population live on between \$1 and \$2 per day in Bangladesh, India and Indonesia. However, while 3.7 per cent of the population slip below the \$1 threshold in both Bangladesh and India after subtracting payments for health care, only 0.7 per cent of Indonesians are impoverished. In the five countries in which 30-40 per cent of the population lie between the two poverty thresholds, there are substantial differences in the poverty impacts. Over 2 per cent are impoverished in China and Nepal, 1.2 per cent in Vietnam and much less than 1 per cent in the Philippines and Sri Lanka. Again, these differences reflect different degrees of reliance on OOP financing. But this does not explain all the differences. Vietnam is more heavily reliant on OOP payments than China but is apparently more successful in limiting their impoverishing effect.

The impoverishing effect of health care payments is further illustrated graphically for the example of China in Figure 6 using a Pen Parade graph as suggested by (Wagstaff and van Doorslaer 2003). The chart plots household pre-payment consumption per capita against the cumulative percentage of individuals ranked by pre-payment consumption. At the intersection of this curve with the poverty line, the x-coordinate is the poverty headcount: 13.7 per cent in China at the \$1.08 poverty line. The area below the poverty line but above the pre-payment curve gives the poverty gap. From the pre-payment curve, the vertical lines show what happens to consumption after payments for health care are subtracted, and which individuals are pulled below the poverty line by such payments. It can be seen that mainly individuals in the bottom third of the gross expenditure distribution are pulled below the \$1-per-day threshold by medical expenses, while those in the middle third of the pre-payment distribution are most likely to be left below the \$2-per-day threshold after netting out health payments. There are also many individuals high up in the pre-payment distribution that incur large (or catastrophic) medical expenses, without these necessarily driving the household into absolute poverty.


Table 1: Poverty headcounts: effect of taking account of OOP payments for health care

Poverty line	\$1.08 per day					\$2.15 per day				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Pre- payment headcou nt	Post- payment headcount	Percentage point change	Number of individuals	Percentage change	Pre- payment headcoun t	Post- payment headcount	Percentage point change	Number of individuals	Percentage change
Bangladesh	22.5%	26.3%	3.8%	4940585	16.8	73.0%	76.5%	3.6%	4653875	4.9
China	13.7%	16.2%	2.6%	32431209	18.8	44.6%	46.4%	1.8%	23198460	4.1
India	31.1%	34.8%	3.7%	37358760	11.9	80.3%	82.4%	2.1%	20638361	2.6
Indonesia	7.9%	8.6%	0.7%	1440395	8.7	58.2%	59.9%	1.7%	3493767	2.9
Kyrgyz Rep.	2.2%	2.7%	0.1%	5989	4.7	32.2%	34.1%	1.9%	94793	6.0
Malaysia	1.0%	1.1%	0.1%	10562	4.4	11.8%	12.1%	0.3%	58626	2.1
Nepal	39.3%	41.6%	2.2%	515933	5.7	80.4%	81.7%	1.3%	290280	1.6
Philippines	15.8%	16.4%	0.6%	445680	3.7	50.2%	51.2%	1.1%	790333	2.1
Sri Lanka	3.8%	4.1%	0.3%	60116	8.3	39.1%	40.8%	1.7%	325783	4.3
Thailand	2.1%	2.3%	0.2%	100201	7.9	24.2%	24.9%	0.7%	417626	2.8
Vietnam	3.6%	4.7%	1.1%	848870	30.1	36.9%	41.4%	4.5%	3492321	12.1
TOTAL	19.3%	22.0%	2.7%	78158299	14.0	58.8%	60.8%	2.0%	57454225	3.4

Notes: Column (3)=(2)-(1). Column (4)=(3)*population. Column (5)=(3)/(1). Bold figures in columns (3) and (8) denote statistically significantly different from zero at the 5% significance level.


Source: Van Doorslaer et al. (2006b).

Figure 5: Headcount increase and population at risk



Source: Van Doorslaer et al. (2006b).

Figure 6: Pen's Parade and out-of-pocket health payments, China 2000



Source: Van Doorslaer et al. (2006b).

Under conditions, the difference between poverty estimates derived from household resources gross and net of OOP payments for health care may be interpreted as a rough approximation to the impoverishing effect of such payments. Only if OOP payments for health care were completely non-discretionary and total household resources fixed, the difference between the two estimates would correspond to poverty due to health payments. Neither of the two conditions holds perfectly. A household that chooses to spend excessively on health care is not pushed into poverty by OOP payments. A household may borrow to cover health care expenses. Then, household expenditure gross of OOP payments does not correspond to the resources that would be available in the absence of those payments. For such reasons, our comparison of poverty estimates cannot be interpreted as the change in poverty that would arise from some policy reform that eliminated OOP payments for health care. Nonetheless, the comparison is indicative of the scale of the impoverishing effect of health payments and it does show the extent to which poverty is currently underestimated (or hidden) by ignoring the amount of household resources that are exhausted by payments for health care.

Impoverishment will be observed only if a household spends on health care. But many households may be so poor that they are unable to spare any resources to spend on health care and must go without. Considering the links from forgone health care to health and subsequently earnings, the long-term poverty impact of charges may be substantially greater than the short-term effect on current consumption. It is important therefore to consider the impact of payments on the distribution of health care and not only on the distribution of economic resources.

3 Inequality and inequity in health care utilization

The distribution of health care in relation to income is of interest for many reasons. Foremost are consequences for health inequality. If the poor are relatively deprived of effective health care interventions, income-related inequalities in health will be exacerbated. Not only is this of immediate concern, it will strengthen the health-poverty trap that can retard economic growth (Sala-i-Martin 2005). The distribution of health care may also be examined to assess whether there is equity in the allocation of health care resources. This is the motivation that has been predominant in OECD countries focused research on the distribution of health care (van Doorslaer et al. 1992; van Doorslaer 2000; van Doorslaer et al. 2006a). The aim has been to establish whether the distribution of health care obeys the principle of horizontal equity, defined as equal treatment for equal need (ETEN). After standardizing for differences in need, any remaining income-related inequality in the utilization of health care is interpreted as horizontal inequity (Wagstaff and van Doorslaer 2000). The evidence suggests that many OECD countries are close to achieving their horizontal equity objectives, although the results often differ markedly by type of utilization. This approach has also proved feasible in three of the higher-income countries in the Asian study with near universal coverage (Lu et al. 2007).

Application in low-coverage settings runs into two problems. The first is conceptual. The horizontal equity approach uses the observed average relationship between use and need characteristics (while appropriately controlling for other factors) as the implicit

norm to identify needed or need-expected use. In other words it assumes that, ‘on average, the system gets it right’ (van Doorslaer et al. 2000b). This assumption may be appropriate in systems which have a long experience with near universal and comprehensive coverage, but is clearly more likely to be violated in systems with very partial coverage of population and services. The second problem concerns data requirements, which for this analysis include measures of income, health care use and need for the same individuals. In the OECD focused research, need has been proxied by demographics and self-reported measures of health. Its application to low-income countries is currently constrained by the availability and reliability of self-reported health measures, which often fail to show the income gradients that are observable in more objective, but less general, measures of health, such as infant mortality (Murray 1996; Wagstaff 2002). In fact, self-reported measures often fail to show any income gradient in health, or even a pro-poor gradient (Baker and Van der Gaag 1993), suggesting substantial income-related differences in the conception and reporting of health problems. This reporting bias issue will be examined in the next section. First, let us consider how equity in the utilization of health care can be assessed in low-income countries without reliable self-reported measures of health and without an estimate of the relationship between use and need that can be used as the norm in standardizing for need differences.

There is no problem if there is little or no variation in need. Then, inequality in utilization represents inequity. Variation in need will often be limited provided that the population group and treatment of interest are defined to be sufficiently homogeneous. For example, all children within a certain age band are in need of immunization against measles, tuberculosis, etc. While it could legitimately be argued that the benefit from immunization varies with prevalence of the disease in the child’s locality, such variation in need is limited in comparison with that for adults’ visits to a doctor. *Demographic and Health Survey* data show clear pro-rich disparities in child immunization rates, use of antenatal care and medically attended births (Gwatkin et al. 2003), for example, that can be interpreted as inequity and not merely inequality.


Conclusions about equity can sometimes be inferred from the distribution of utilization if there is prior knowledge of the distribution of need. For example, it is well established that rates of mortality and malnutrition are higher among poorer children in many low-income countries (Gwatkin et al. 2003). Poor children, and quite probably poor adults, are in greater need of medical interventions. If the distribution of health care is not skewed toward the poor, then this is sufficient to reject the proposition of equity. Pro-rich inequality in utilization then provides a lower bound to the degree of pro-rich inequity. In high-income countries this approach is not helpful since the poor tend to use health care more than the rich. There the question is whether the greater utilization by the poor matches their greater need such that equity is achieved. In low-income countries, however, inequality in utilization often does not favour the poor and it is possible to make a statement about inequity without simultaneously measuring need. It is evidently not possible, however, to make precise estimates of the degree of inequity, which limits the ability to make cross-country or other comparisons.

Besides concern for health sector equity, redistribution of economic welfare provides a further motivation for public provision of health care in low-income countries. With severe informational and administrative constraints on redistribution through tax and cash transfers, in-kind transfers, such as health care, can be used to effect redistribution of final incomes provided the poor make greater use of such transfers (Besley and Coate

1991). Establishing whether such redistribution takes place requires identification of the incidence of public spending on health care in relation to the distribution of income. If public health care is more concentrated on the poor than is income, then it raises the economic welfare of the poor relative to that of the better-off and inequality in final incomes is reduced. Note that this does not require that the poor receive absolutely more health care than the rich. Public spending on health care could be pro-rich but still inequality-reducing.

The EQUITAP study has examined whether public spending on health care was (a) pro-poor, and (b) inequality-reducing in eight Asian countries and three Chinese provinces/regions (O'Donnell et al. 2007a). These questions were addressed by testing dominance of the concentration curve of the public health subsidy against the 45° line, representing an equal distribution, and the Lorenz curve of household consumption respectively. Formal statistical tests of dominance were used for inference (Bishop et al. 1992). By way of example, subsidy concentration curves and Lorenz curves for Malaysia and India are presented in Figure 7. The concentration curve in Malaysia lies (significantly) above the diagonal—it is pro-poor. In India, by contrast, the curve lies below the diagonal—the poor get less than a proportionate share of public health care. While the subsidy concentration curve in India appears to be slightly inside the Lorenz curve, the difference is not statistically significant. The hypothesis that public spending on health care in India has no impact on inequality in living standards cannot be rejected.

Figure 7: Concentration curve of public health subsidy and Lorenz curve of household consumption per equivalent adult



Source: Derived from data and analysis reported in O'Donnell et al. (2007a).

Dominance results for all countries included in the study are summarized in Table 2. Public spending on health care is pro-poor only in the three highest income countries. Sri Lanka is the only low-income country that distributes public health care equally. It is pro-rich in all the other low-income countries, but in most cases is also inequality-reducing. The two exceptions are India and Nepal, where the share of public spending received by the poor does not even surpass their share of total consumption.

The concentration indices presented in Figure 8 measure the degree of income-related inequality in the distribution of public health spending. The pro-poor inequality in Hong Kong, indicated by a negative concentration index, is much greater than that in Malaysia and Thailand. As anticipated from the result of the dominance test, the index is zero in Sri Lanka, indicating equality. It is only very slightly positive in Vietnam suggesting that the pro-rich inequality detected by the dominance test is marginal. The pro-rich bias is more substantial in the other low-income countries and is greatest in the poorest country—Nepal. The Kakwani index—equal to the concentration index less the Gini coefficient—is negative in all countries but for Nepal, indicating inequality-reduction, as could be anticipated from the dominance tests. The index is greatest in absolute value in Hong Kong, indicating that is where the distribution of public health spending diverges most from proportionality to income. Since the level of public health spending relative to national income is also greatest in Hong Kong, it can be inferred that the reduction in inequality is greatest there. In fact, with the exception of the Chinese provinces, aggregate government spending on health as a proportion of GDP is greatest in the countries where this spending is distributed least proportionately to income and so inequality-reduction will be greatest in these countries.


Table 2: Distribution of public health subsidy in Asia

	Inequality-reducing	Inequality neutral
Pro-poor	Hong Kong SAR Malaysia Thailand	
Equal	Sri Lanka	
Pro-rich	Vietnam Bangladesh Indonesia Gansu (China) Heilongjiang (China)	India Nepal

Source: Derived from results reported in Table 1 of O'Donnell et al. (2007a).

Notes: Pro-poor, concentration curve statistically dominates (lies above) the 45° line of equality. Pro-rich, concentration curve is statistically dominated by the 45° line. Equality, concentration curve is statistically indistinguishable from the 45° line. Inequality-reducing, concentration curve statistically dominates the Lorenz curve. Inequality neutral, concentration curve is statistically indistinguishable from (India) or crosses (Nepal) the Lorenz curve.

Figure 8: Concentration and Kakwani indices for public health subsidy



Source: O'Donnell et al. (2007a: Table S.3).

Cross-country differences in the distribution of public health care are assessed by tests of dominance between concentration curves. Results are reported in Table 3. The distribution in Hong Kong dominates, i.e. is more pro-poor, than that in all other countries. The distributions in Malaysia, Thailand and Sri Lanka are indistinguishable. Malaysia and Thailand dominate all other distributions and Sri Lanka dominates all others but for Vietnam and Bangladesh. Vietnam dominates the remainder of countries. So, countries/provinces can be broadly grouped as follows in relation to the distribution of public health care: (a) Hong Kong (very pro-poor); (b) Malaysia, Thailand and Sri Lanka (mildly pro-poor to neutral); (c) Vietnam (mildly pro-rich); (d) Bangladesh, Indonesia, India, Gansu, Heilongjiang and Nepal (very pro-rich).

The interpretation of these results requires consideration of two questions. First, in the many low-income countries in which the poor do not get their share of public health care, does this necessarily represent a failure of public policy? Second, why is it that public health care is more pro-poor in Malaysia, Thailand, Sri Lanka and, to a lesser extent, Vietnam than in other developing countries of Asia?

The answer to the first of these questions largely depends on the objectives of the public health spending. If the aim is to ensure that the poor get more public health services than the better-off, then the objective is clearly not being achieved in most cases. Alternatively, subsidising health care may be part of a wider policy to reduce relative differences in living standards between rich and poor. With the exceptions of India and Nepal, the subsidy achieves this objective. It is inequality reducing. But those concerned about inequalities within the health sector may not be content with reduction in general economic inequality. From this perspective, is the fact that the poor get less of the subsidy necessarily a failure? Despite its name, benefit incidence analysis informs of the incidence of public health expenditures, rather than the benefits from these expenditures. Even though the poor get a lower than proportionate share of the subsidy, the impact on their health can be greater if the marginal product is declining with the

initial level of health (given a positive relationship between income and health) (Filmer et al. 2002). Further, the analysis describes the distribution of the subsidy and does not identify the impact of the subsidy on the distribution of health care, nor health. The poor would be less able to afford health care in the absence of public provision, while the crowding-out effect will be greater for the better-off. Consequently, the poor may get less of the subsidy but experience a larger net effect on total health care utilization (Filmer et al. 2002). While the distribution of the subsidy is not pro-poor, the subsidy can still shift the distribution of health care in a pro-poor direction. These hypotheses are consistent with evidence showing that public spending has no significant effect on health of the non-poor but a positive marginal impact on the health of the poor (Bidani and Ravallion 1997; Gupta et al. 2003; Wagstaff 2003).

Table 3: Cross-country dominance of public health subsidy concentration curves


	Malaysia	Thailand	Sri Lanka	Vietnam	Bangladesh	Indonesia	India	Gansu	Heilongjiang	Nepal
Hong Kong SAR	D*	D*	D*	D	D*	D*	D*	D*	D*	D*
Malaysia		ns	ns	D	D	D*	D*	D*	D*	D*
Thailand			ns	D	D	D*	D*	D	D*	D*
Sri Lanka				ns	ns	D	D	D	D*	D*
Vietnam					D	D*	D	D*	D	D*
Bangladesh						ns	ns	ns	ns	ns
Indonesia							ns	D	ns	D
India								D	ns	D
Gansu (China)									ns	ns
Heilongjiang (China)										ns

Notes: D indicates rejection of the null that curves are indistinguishable in favour of dominance (more pro-poor) of the row country over than column country using the multiple comparison test and 5% significance. ns indicates failure to reject the null that the curves are indistinguishable. There are no cases of curves crossing. * indicates that the intersection-union test rejects the null of nondominance against the alternative of *strict* dominance at 5%. If no * appears, this test does not reject its null.

Source: O'Donnell et al. (2007a: Table 4).

The evidence shows that, on average, the better-off typically receive most of the subsidy. But this is informative of the distributional implications of a policy change only if marginal changes in the subsidy were delivered in strict proportion to current utilization (Younger 2003). Of course, many policy reforms will deliver marginal gains that differ from average gains (Lanjouw and Ravallion 1999). For example, the political power of higher-income classes may allow them to capture most of the initial public spending on health but as their lower health needs are satisfied, additional programmes may disproportionately benefit the poor. Then additional public spending will shift the distribution in a pro-poor direction. There is some evidence, admittedly not particularly strong, that this holds for primary care in Indonesia (Lanjouw et al. 2002). So, caution should be exercised in interpreting evidence of pro-rich bias in the distribution of the public health subsidy. It does not necessarily mean that public policy is not shifting health care resources toward the poor.

Figure 9: Public health spending: magnitude and incidence



Source: Derived from results presented in O'Donnell et al. (2007a).

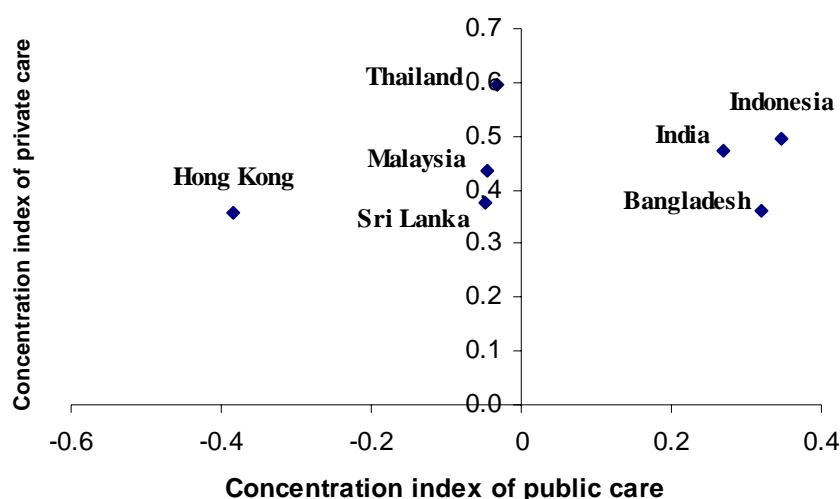
Characteristics of the health care systems of Malaysia, Thailand and Sri Lanka that can be expected to contribute to a more pro-poor distribution of resources can be identified. In particular, all three emphasise universality and seek to minimize user charges or have effective systems of exempting the poor from such charges. The scale of public funding to health care in these countries is another possible contributory factor. Figure 9 shows a clear negative relationship between the scale of public health spending and its concentration index—public spending is more pro-poor where it is higher.³ Of course, this is only an association and many common factors, including political institutions and preferences, may drive both the level and the distribution of public health spending. Nonetheless, there are mechanisms through which the scale of funding could affect its incidence. More resources allow a wider geographic distribution of resources bringing health services closer to the rural poor. Malaysia, Sri Lanka and Vietnam all have public health facilities in close proximity to the rural population. For the trickle-down mechanism described in the previous paragraph to take effect, spending levels have to reach sufficiently high levels such that the health needs of the better-off begin to be satisfied.

While relatively high levels of public spending can ensure adequate technical quality in the public sector, universality of access leads to long waiting times and minimal amenities, creating incentives for the better-off to opt-out of the public sector. These incentives become stronger as the economy grows, as it has done most impressively in Malaysia and Thailand, and there is an expansion of middle and higher income groups

³ The two Chinese provinces are exemptions to this relationship. This is partly due to the fact that the public spending figures include social health insurance, which is a much larger share of the total in China than in the other study countries.

with not only the desire but also the means to purchase higher quality care in the private sector (Hammer et al. 1995). Figure 10 shows clear discrepancies between the distributions of public and private inpatient care in Hong Kong, Malaysia, Sri Lanka and Thailand. But in Bangladesh, India and Indonesia, the distribution of private sector care is only slightly more pro-rich than is care in the public sector. In these countries there does not yet exist a large middle-class with the effective demand for the greater convenience of private sector care. In addition, quality differentials between the sectors can be limited. In Bangladesh and India, the poor make extensive use of unqualified private sector providers that may be cheaper and more accessible than public sector alternatives. In Hong Kong, Malaysia, Sri Lanka and Thailand, private sector care is pro-rich while the public sector is pro-poor or neutral. The combination of (near) universal public provision, a private sector offering an attractive alternative, and incomes that make demand for this alternative effective leads to redistribution through public provision in the way that theory predicts (Besley and Coate 1991).

Figure 10: Distribution of public and private hospital inpatient care



Source: Somanathan et al. (2005).

In summary, the analysis shows that the pervasive outcome of a pro-rich distribution of public health care in most Asian and other low-income countries is not unavoidable but that effective targeting is easier to realise at higher national incomes. Hong Kong, Malaysia, Sri Lanka and Thailand have demonstrated that the allocation of sufficient public resources coupled with a policy of universal access can ensure far greater benefits to the poor than may have hitherto been assumed. Higher incomes not only make such policies more feasible, it also makes them more effective, with respect to the target efficiency of spending, by availing the private sector opt-out.

Descriptive analysis such as that presented above is useful in identifying a problem. It shows that public spending on health care does not predominantly reach the poor in most low-income countries. Hypotheses can be offered to explain this and solutions can be proposed. But from a descriptive analysis it is not possible to identify the likely effectiveness of any policy reforms in shifting the distribution of health care toward the poor. For this, an evaluative approach is required.

4 Health inequality

Two issues of importance to the study of health inequality in low-income countries are considered in this section: (i) reporting heterogeneity in measures of health; (ii) decomposition of health inequality and of changes in inequality.

4.1 Reporting heterogeneity in health measures

The difficulty confronted in obtaining an accurate measure of health for a study of health inequality varies with the type of inequality one is seeking to examine. If the purpose is to measure the total variation in health across the population, then detailed measures of health available from administrative records or health surveys can be used. The task is more difficult if the aim is to measure socioeconomic-related inequality in health. Data that provide measures of health and of socioeconomic status for the same individuals are then required. There is usually a trade-off between the use of health surveys that provide more detailed measures of health but less detailed measures of economic status, such as income and consumption, and the use general household and expenditure surveys that measure income and/or consumption more accurately but have more crude instruments for the measurement of health. In recent years, the demographic and health surveys have weakened this trade-off. It has proved possible to construct an index of household wealth from data on assets and living conditions and to measure inequality in child mortality and nutritional status, for example, in relation to this indicator of economic status (Gwatkin et al. 2003). It has been demonstrated for some countries that the measurement of economic-related inequality in such health indicators is robust to the use of household consumption, rather than a wealth index, as the measure of economic status (Wagstaff and Watanabe 2003).⁴

The DHS do not provide a measure of general health in the adult population. In research on socioeconomic-related inequality in OECD countries, extensive use has been made of an indicator of self-assessed health that asks individuals to select a description of their health in general from four or five categories ranging, for example, from very good to very poor. This indicator has proven, perhaps surprisingly, to be highly informative of health status. Controlling for physiological measures of health, it is an independent predictor of mortality (Idler and Benyamini 1997). But self-reported measures of health have proven to be less useful in low-income countries. As noted in the previous section, they do not display the steep socioeconomic gradients that are apparent for more objective measures such as child mortality (Murray 1996; Wagstaff 2002) and they can even show perverse gradients in which the poor are apparently in better health than the rich (Baker and Van der Gaag 1993). This suggests that, for a given health condition, the poor report their health more favourably than the better-off. A possible explanation is that there are peer effects on the reporting of health. When an ailment is highly prevalent within a community or social group, it may not even be recognized as a health problem but accepted as the norm. Constrained access to health professionals will also

⁴ There is evidence from Mozambique that economic-related inequality in the utilization of health care is sensitive to the use of consumption or a wealth index as the indicator of economic status (Lindelow 2006).


limit the opportunities for the poor to learn of their health problems. The magnitude of the problem varies with the precise nature of the health questions. Gradients are least pronounced for a question, popular in Living Standards Measurement Surveys, which asks whether illness was experienced in the past four weeks. Questions about chronic conditions and self-assessed health are less common but show slightly more pronounced gradients (Wagstaff 2001).

Case vignettes have been proposed as an instrument to correct for systematic heterogeneity in the reporting of health (Tandon et al. 2003; King et al. 2004). The idea is to identify variation in reporting behaviour from individuals' ratings of defined cases. For example, sample respondents may be asked to rate, on a 5-point scale, the health of someone that cannot walk 100 metres without stopping to catch breath. Variation in the ratings with respondents' characteristics, such as gender, age, income, education and nationality, allows estimation of the effects of these characteristics on reporting behaviour. Assuming that respondents report on the vignettes in the same way that they report their own health, the identified reporting effect can be purged from the rating of own health. Inequality in this cleansed measure of health should then reflect variation only in true health conditions and not in the reporting of these conditions.

Bago d'Uva et al. (2008) use vignette data collected for Indonesia, Andhra Pradesh (India) and 3 Chinese provinces (Gansu, Henan and Shan-dong), as part of the WHO Multi-Country Survey Study of Health and Responsiveness (Sadana 2002), to examine the effect of removing systematic reporting heterogeneity from measured disparities in health by income, education, demographics and urban-rural location. This involves simultaneously estimating the effect of characteristics on rating of the vignettes and on rating of own-health under the restriction that the estimated threshold values that determine whether a vignette is placed in one health category or another also apply to the categorization of own-health. In effect, information from the vignette responses identifies how a given latent value of health is rated; leaving the information from the rating of own-health to identify how the latent health index varies with characteristics.


The reporting effects for income are summarized in Figure 11. Each bar shows the probability that an individual with given characteristics (i.e. male, 15-29 years, no primary-level education and living in a rural area) and with income at the threshold of the richest quintile will rate any given vignette as corresponding to very good health relative to the probability that an individual with the same characteristics but with income at the threshold of the poorest income quintile will rate the vignette in this way. So, the bars indicate rich-poor relative differences in the probability of reporting very good health. This is done for vignettes classified into five health domains. For the Indonesian sample, there appears to be little difference in the reporting of health by income. But in the Indian and Chinese samples the higher income groups are less likely to report any given condition as corresponding to very good health. In Andhra Pradesh, this is true for the domains for pain, self-care and functioning in usual activities. In the Chinese provinces, it is true for pain and affective behaviour. Homogeneity in reporting behaviour by income is rejected in all but one health domain for the Indian and Chinese samples and all but two domains for the Indonesian sample.

Figure 11: Probability of high-income individual to that of low-income individual reporting any vignette health condition as very good health



Source: Bago d'Uva et al. (2008: Figure 1).

Figure 12: Probability of high-income individual to that of low-income individual reporting own-health as very good health before and after purging reporting heterogeneity



Source: Bago d'Uva et al.(2008: Figure 2).

The effect of purging reporting heterogeneity from measured disparities in health by income level is illustrated in Figure 12. The bars show the probability that an individual with given characteristics, as described above, and with income in the highest quintile will report her health as being very good relative to the probability that someone in the poorest quintile with the same characteristics will report very good health. The lighter red bars are derived from probabilities estimated from a model that does not correct for reporting heterogeneity. These rich-poor differences may therefore reflect both true differences in health and differences in reporting behaviour. The blue bars show the relative probabilities of reporting own-health as very good after purging reporting heterogeneity. In the Indonesian sample, there is little evidence of health disparities by income before or after purging of reporting bias. The absence of inequality has little face validity and may in itself be taken as evidence of reporting differences. If true, it seems that the vignette instrument is not sufficiently effective in detecting and purging this heterogeneity. In Andhra Pradesh, there is evidence of some income-related inequality in health without taking account of reporting bias and when this is corrected the measured inequality increases by a relatively large degree in the domains of pain, self-care and usual activities but the adjustment is much smaller in the other domains. Health disparities are greater in China and the correction increases them further, particularly in the domains of pain and affective behaviour.

This study finds that the way in which individuals report their health varies significantly with their socioeconomic and demographic characteristics and this biases measures of income-related inequality in health downward. But in most cases, the magnitude of the effect is rather small. This may reflect the difficulty poorly educated individuals have in completing the vignette exercise, which requires a considerable degree of abstract thought. While it is certainly an approach that deserves further experimentation, the jury remains out on whether it does offer a sufficiently more accurate way of measuring health for the purpose of examining socioeconomic inequality in health in low-income countries.

4.2 Decomposition of health inequality

Measurement of health inequality is a first step toward understanding the socioeconomic determinants of health and of health sector inequities. A natural next step is the explanation of the sources of health inequality and the factors contributing to its change over time. Decomposition techniques can be used to provide such explanations. The method presented in this section decomposes income-related health inequality, as measured by the concentration index, into the contribution of various factors. The approach can also be used to explain change in the concentration index.

Wagstaff et al. (2003) demonstrate that the concentration index of health can be written as the sum of the contribution of factors, such as demographics, education, region, etc., to income-related health inequality, where each contribution is the product of the elasticity of health with respect to the factor and the concentration index of the factor. That is, the concentration index can be written as:

$$C = \sum_k (\beta_k \bar{x}_k / \mu) C_k + GC_\varepsilon / \mu \quad (1)$$

where μ is the mean of the health measure, \bar{x}_k is the mean of k th factor, β_k is its coefficient from least squares regression of health on all factors, C_k is the concentration index for the k th factor and GC_ε is the generalized concentration index for the error term of the regression. The term in brackets is the elasticity of health with respect to the k th factor, label this η_k .

Change in the concentration index can be explained by applying a Oaxaca-type decomposition (Oaxaca 1973) to (1) to give, for example,


$$\Delta C = \sum_k \eta_{kt} (C_{kt} - C_{kt-1}) + \sum_k C_{kt-1} (\eta_{kt} - \eta_{kt-1}) + \Delta(GC_{\varepsilon t} / \mu_t) \quad (2)$$

where t indicates time period and Δ denotes first differences. Change in income-related inequality in health is decomposed into changes in inequality in the determinants of health (the first term), on the one hand, and changes in the elasticities of health with respect to these determinants (the second term), on the other. This approach can also be used to examine differences in inequality across cross-sectional units, such as countries (see e.g. van Doorslaer and Koolman 2004).

Wagstaff et al. (2003) used these decomposition methods to explain the level and change in income-related inequality in child height-for-age z-scores (HAZ), a measure of long-term nutritional status, in Vietnam in 1993 and 1998. The results are summarized in Figure 13. The concentration index was negative in both years, indicating the nutritional status was lower among poorer children. The first two bars in the figure show the contribution of each factor to the concentration index. The largest contribution is the direct effect of household consumption, the measure of economic status against which inequality is assessed. There is also a very large contribution to consumption-related nutrition inequality from commune fixed effects. This suggests that nutrition and consumption both vary across communes and this commune level covariance makes an important contribution to the measured inequality in nutrition. There is a smaller contribution from access to sanitation and to safe drinking water in 1998. Age differences in nutrition and consumption shift the inequality in the other direction, toward lower nutrition among the better-off but this is outweighed by the other contributions. The concentration index increased in absolute value between 1993 and 1998 from -0.075 to -0.102 indicating an increase in consumption-related inequality in child height deficit. The third bar in the figure shows that most of this increase in inequality is explained by changes in the distribution and effect of household consumption and in the contribution of the commune level covariance between nutrition and consumption.

Measurement of horizontal inequity in the utilization of health care, as described in Section 3, involves measuring income-related inequality in utilization after standardising for differences in need. The decomposition approach provides a convenient way of doing this. One simply needs to deduct the contributions of the need standardizing variables from total inequality (van Doorslaer et al. 2004). An advantage is that the analyst can avoid imposing her judgments about what factors should be counted as need and therefore justifiably giving rise to variation in utilization. The full decomposition results can be presented and the user can choose on which factors to standardize.

Figure 13: Decomposition of level and change in concentration index for height-for-age z-scores of children less than 10 years old in Vietnam



Source: Wagstaff et al. (2003).

It should be emphasized that decomposition techniques, such as that described, are descriptive. They do not identify causal effects on the distribution of health or health care. While decomposition analysis is useful for explaining observed relationships, it does not allow inference about how a distribution would change in response to some policy intervention. The latter requires causal analysis and data containing some independent variation in the intervention variable. Another limitation of the above decomposition method is that it only applies to linear, additively separable models. In many cases, models to explain health and health care will be intrinsically non-linear and the simple linear decomposition will break down and will have to be replaced by other approaches that can handle intrinsically non-linear relationships (e.g. Wan 2004).

5 Conclusion

This review served to illustrate that, in moving attention from high-to low-income countries, some of the standard measurement methodology that has been proposed and used to measure income-related inequality in health payments, medical care use and health status, cannot straightforwardly be applied in low-income, low health insurance coverage settings. Some of the reasons for this are conceptual and relate to different equity priorities concerning the finance and provision of health care in such countries. Others derive from measurement problems.

With respect to finance, it is obvious that in systems with small revenue shares deriving from prepayments, questions of income redistribution or deviations from proportionality come secondary to questions of income protection. This is why measures of incidence and intensity of catastrophic out-of-pocket payments and their impoverishing effect are

required to assess the equity performance of health financing systems. Measures of progressivity and redistributive effect may even be misleading if used to examine the distribution of payments that are mainly driven by the benefit principle, rather than the ability to pay principle. In low-income countries, the rich clearly pay more, even proportionally more, but mainly for receipt of their own care. Unlike in universal coverage systems, which tend to require payments according to ability to pay and receipt according to need, in low coverage countries, payments and receipt of care are still very much linked and it can be misleading to analyze them separately. The measures of catastrophic payments and impoverishment that have been proposed and used to date are rather *ad hoc* and certainly do not have the sort of conceptual underpinnings that have been developed for progressivity measures. Further consideration should be given to the development of a conceptual basis for catastrophic health payments from which a measure can be derived.

With respect to equity in access and use of medical care, again, a straightforward application of approaches based on the need principle, which requires those in equal need to be treated equally, runs into two problems. One is, as above, that a lot of care is simply allocated on the basis of market principles—you get what you pay for—and therefore the standard approach of using the average need-use relationship as the norm breaks down. Secondly, self-reported health measures appear much less reliable in low-income/low-education settings and are therefore less appropriate for need standardization procedures. Given these problems, analysis is confined to examination of the distribution of health care with no standardization for need. Conclusions about equity can still possibly be drawn if there is prior knowledge of the distribution of need. Examination of the distribution of public health care is of additional interest since it reveals whether public spending dollars predominantly benefit the rich or the poor. This is integral to evaluation of the effectiveness of government anti-poverty policies. The Asian study shows that public spending on health care—despite being inequality-reducing—does not predominantly reach the poor in most low-income countries. It was found to be most pro-poor in Malaysia, Thailand and Sri Lanka, three countries which seem to have achieved better targeted public care subsidies as a result of a combination of (near) universal public provision with limited user charges on the poor, geographically dispersed health services and facilities, a private sector offering an attractive alternative, and incomes that make demand for this alternative effective.

With respect to inequalities in health itself, most studies hitherto have concentrated on objective measures like adult or child survival or nutrition because of presumed measurement biases in self-reported measures of health. This has generated a large and useful literature, ranging from descriptive over explanatory to evaluative; and such measures are clearly more appropriate the lower the level of development. But with rising levels of development, other aspects of health than nutrition and survival will gain importance and require measurement. Reports of weak, missing or even the ‘wrong’ gradients in self-reported measures of health and illness have generated skepticism concerning their usefulness and worries about strong and systematic reporting biases. The use of case vignettes has been suggested, both as a tool to testing for reporting tendencies, but also as a potential remedy to correct its biases. The evidence to date is still thin on the extent to which vignette-based modeling techniques will be able meet these promises.

While most of the evidence reviewed in this paper concerns inequality measurement and explanation, it is obvious that there is an even more urgent need to complement this

with evaluative evidence that is capable of demonstrating not only how inequality compares across time and place, or can be decomposed into (partial) associations with other inequalities, but also how it can be impacted upon by policy interventions. This calls for well-controlled, possibly even experimental designs, which may give us harder evidence on the causes and consequences of health inequalities, but may come at the price of lower generalizability. Advancing health equity will require accumulating evidence from a multitude of micro studies on what interventions in which settings are effective in changing the distributions of health payments, health care and health.

References

- Aronson, J. R., P. Johnson and P. J. Lambert (1994). 'Redistributive effect and unequal tax treatment', *Economic Journal* 104: 262-70.
- Bago d'Uva, T., E. van Doorslaer, M. Lindeboom, O. O'Donnell, et al. (2008). 'Does reporting heterogeneity bias the measurement of health disparities?', *Health Economics* 17(3): 351-75
- Baker, J. L. and J. Van der Gaag (1993). 'Equity in health care and health care financing: Evidence from five developing countries' in *Equity in the finance and delivery of health care* in E. Van Doorslaer, A. Wagstaff and F. Rutten. Oxford, Oxford University Press.
- Berki, S. E. (1986). 'A look at catastrophic medical expenses and the poor', *Health Affairs* 138-145.
- Besley, T. and S. Coate (1991). 'Public provision of private goods and the redistribution of income', *American Economic Review* 81(4): 979-84.
- Bidani, B. and M. Ravallion (1997). 'Decomposing social indicators using distributional data', *Journal of Econometrics* 77(1): 125-140.
- Bishop, J. A., J. P. Formby and P. D. Thistle (1992). 'Convergence of the South and Non-South income distributions, 1969-1979', *American Economic Review* 82(1): 262-72.
- Chen, S. and M. Ravallion (2004). 'How have the world's poorest fared since the early 1980s?' *World Bank Research Observer* 19(2): 141-69.
- Cirto, C. and R. Michael (1995). *Measuring poverty: A new approach*. Washington DC, National Academy Press.
- Filmer, D., J. Hammer and L. Pritchett (2002). 'Weak links in the chain II: A prescription for health policy in poor countries', *World Bank Research Observer* 17(1): 47-66.
- Flores, G, J. Krishnakumar, O. O'Donnell and E. van Doorslaer (2008). 'Coping with health care costs: Implications for the measurement of catastrophic expenditures and poverty', *Health Economics* (forthcoming).
- Gakidou, E., C. Murray and J. Frenk (2000). 'Defining and measuring health inequality', *Bulletin of the World Health Organization* 78(1): 42-54.
- Gertler, P. and J. Gruber (2002). 'Insuring consumption against illness', *American Economic Review* 92(1): 51-70.

- Gupta, S., M. Verhoeven and E. R. Tiongson (2003). 'Public spending on health care and the poor', *Health Economics* 12(8): 685-96.
- Gustaffson, B. and S. LI (2004). 'Expenditures on education and health care and poverty in rural China', *China Economic Review* 15: 292-301.
- Gwatkin, D. R., S. Rustein, K. Johnson, R. Pande, et al. (2003). *Initial country-level information about socio-economic differentials in health, nutrition and population*, Volumes I and II. Washington DC, World Bank, Health, Population and Nutrition.
- Hammer, J. S., I. Nabi and J. A. Cercone (1995). 'Distributional effects of social sector expenditures in Malaysia, 1974-89', in *Public spending and the poor*. D. van de Walle and K. Nead. Baltimore MD, Johns Hopkins University Press.
- Idler, E. and Y. Benyamini (1997). 'Self-rated health and mortality: a review of twenty-seven community studies', *Journal of Health and Social Behavior* 38(1): 21-37.
- Kakwani, N. C. (1977). 'Measurement of tax progressivity: An international comparison', *Economic Journal* 87(345): 71-80.
- King, G., C. J. L. Murray, J. Salomon and A. Tandon (2004). 'Enhancing the validity and cross-cultural comparability of measurement in survey research', *American Political Science Review* 98(1): 184-91.
- Lanjouw, P., M. Pradhan, F. Saadah, H. Sayed, et al. (2002). 'Poverty, education and health in Indonesia: Who benefits from public spending?' *Education and health expenditure and development: The cases of Indonesia and Peru*. C. Morrison. Paris, OECD.
- Lanjouw, P. and M. Ravallion (1999). 'Benefit incidence, public spending reforms and the timing of program capture', *World Bank Economic Review* 13(2).
- Lindelow, M. (2006). 'Sometimes more equal than others: how health inequalities depend upon the choice of welfare indicator', *Health Economics* 15(3): 263-80.
- Lu, R. J., G. M. Leung, S. Kwon, K. Y. K. Tin, E. van Doorslaer, and O. O'Donnell (2007). 'Horizontal equity in health care utilization – evidence from three high-income Asian economies', *Social Science & Medicine* 64: 199-212
- MacLachlan, G. and A. Maynard (1982). 'The public/private mix in health care: The emerging lessons', in *The public/private mix in health care: The relevance and effects of change*. G. MacLachlan and A. Maynard. London, Nuffield Provincial Hospital Trust.
- Murray, C. J. L. (1996). 'Epidemiology and morbidity transitions in India', in *Health, poverty and development in India*. M. Dasgupta, C. L.C. and T. N. Krishnan. Delhi, Oxford University Press: 122-147.
- Oaxaca, R. (1973). 'Male-female wage differentials in urban labor markets', *International Economic Review* 14: 693-709.
- O'Donnell, O., E. van Doorslaer, R. Rannan-Eliya A. Somanathan, S. R. Adhikari, D. Harbianto, C. C. Garg, P. Hanvoravongchai, M. N. Huq, A. Karan, G. M. Leung, C. W. Ng, B. R. Pande, K. Tin, K. Tisayaticom, L. Trisnantoro, Y. Zhang, and Y. Zhao (2007a). 'The incidence of public spending on health care: Comparative evidence from Asia', *World Bank Economic Review* 21(1): 93-123.

- O'Donnell, O, E. van Doorslaer, A. Wagstaff, M. Lindelöw (2007b). *Analyzing Health Equity using Household Survey Data: A Guide to Techniques and their Implementation*, World Development Institute, World Bank, Washington DC.
- O'Donnell, O, E. van Doorslaer, R. Rannan-Eliya, A. Somanathan, S. R. Adhikari, B. Akkazieva, D. Harbianto, C. C. Garg, P. Hanvoravongchai, A. N. Herrin, M. N. Huq, S. Ibragimova, A. Karan, S. M. Kwon, G. M. Leung, J. F. Lu, Y. Ohkusa, B. R. Pande, R. Racelis, K. Tin, K. Tisayaticom, L. Trisnantoro, Q. Wan, B. M. Yang, and Y. Zhao (2008). 'Who pays for health care in Asia?' *Journal of Health Economics* 27(2): 460-75
- Pradhan, M. and N. Prescott (2002). 'Social risk management options for medical care in Indonesia', *Health Economics* 11: 431-46.
- Ranson, M. K. (2002). 'Reduction of catastrophic health care expenditures by a community-based health insurance scheme in Gujarat, India: current experiences and challenges', *Bulletin of the World Health Organisation* 80(8): 613-21.
- Russell, S. (2004). 'The economic burden of illness for households in developing countries: A review of studies focusing on malaria, tuberculosis, and human immunodeficiency virus/acquired immunodeficiency syndrome', *American Journal of Tropical Medicine and Hygiene* 71(Supp. 2): 147-155.
- Sadana, R., A. Tandon, C. J. L. Murray, I. Serdobova, Y. Cao, W. J. Xie, S. Chatterji, and B. L. Ustun (2002). 'Describing population health in six domains: Comparable results from 66 household surveys'. Geneva, World Health Organization.
- Sala-i-Martin, X. (2005). 'On the health-poverty trap', in *Health and Economics Growth: Findings and Policy Implications*. G. Lopez-Casasnovas, B. Rivera and L. Currais. Cambridge, MA, MIT Press.
- Short, K. and T. Garner (2002). 'Experimental poverty measures: Accounting for medical expenditures', *Monthly Labor Review*: 3-13.
- Short, K., T. Garner, D. Johnson and P. Doyle (1999). 'Experimental Poverty Measures: 1990 to 1997', *U.S. Census Bureau, Current Population Reports, Consumer Income*, 60-205. Washington DC, US Government Printing Office.
- Somanathan, A., O. O'Donnell, E. Van Doorslaer, R. P. Rannan-Eliya, et al. (2005). 'Who gets health care in Asia?' EQUITAP Working Paper 4, Erasmus University, Rotterdam and IPS, Colombo.
- Tandon, A., C. J. L. Murray, J. A. Salomon and G. King (2003). 'Statistical models for enhancing cross-population comparability', in *Health systems performance assessment: debates, methods and empiricisms*. C. J. L. Murray and D. B. Evans. Geneva, World Health Organization: 727-46.
- Van Doorslaer, E. and X. Koolman (2004). 'Explaining the differences in income-related health inequalities across European countries', *Health Economics* 13(7): 609-28.
- Van Doorslaer, E., X. Koolman and A. M. Jones (2004). 'Explaining income-related inequalities in doctor utilization in Europe', *Health Economics* 13(7): 629-47.

- Van Doorslaer, E., C. Masseria, and X. Koolman for the OECD Health Equity Research Group (2006a). 'Inequalities in access to medical care by income in developed countries', *Canadian Medical Association Journal* 174: 177-183.
- Van Doorslaer E, O. O'Donnell, R. P. Rannan-Eliya, A. Somanathan, S. R. Adhikari, C. C. Garg, D. Harbianto, A. N. Herrin, M. N. Huq, S. Ibragimova, A. Karan, C. W. Ng, B. R. Pande, R. Racelis, S. Tao, K. Tin, K. Tisayaticom, L. Trisnantoro, C. Vasavid, and Y. Zhao (2006b). 'Effect of payments for health care on poverty estimates in 11 countries in Asia: an analysis of household survey data', *Lancet* 368: 1357-64.
- Van Doorslaer, E, O. O'Donnell, R. Rannan-Eliya, A. Somanathan, S. R. Adhikari, C. C. Garg, D. Harbianto, A. N. Herrin, M. N. Huq, S. Ibragimova, A. Karan, T. J. Lee, G. M. Leung, J. F. Lu, C. W. Ng, B. R. Pande, R. Racelis, S. Tao, K. Tin, K. Tisayaticom, L. Trisnantoro, C. Vasavid, and Y. Zhao (2007). 'Catastrophic payments for health care in Asia', *Health Economics* 16(11): 1159-84.
- Van Doorslaer, E., A. Wagstaff, H. Bleichrodt, S. Calonge, et al. (1997). 'Income-related inequalities in health: some international comparisons', *Journal of Health Economics* 16(1): 93-112.
- Van Doorslaer, E., A. Wagstaff, S. Calonge, T. Christiansen, et al. (1992). 'Equity in the delivery of health care: some international comparisons', *Journal of Health Economics* 11(4): 389-411.
- Van Doorslaer, E., A. Wagstaff, H. van der Burg, T. Christiansen, et al. (1999). 'The redistributive effect of health care finance in twelve OECD countries', *Journal of Health Economics* 18(3): 291-313.
- Van Doorslaer, E., A. Wagstaff, H. van der Burg, T. Christiansen, et al. (2000). 'Equity in the delivery of health care in Europe and the US', *Journal of Health Economics* 19(5): 553-84.
- Wagstaff, A. (2001). *Poverty and health*. WHO Commission on Macroeconomics and Health. Geneva.
- Wagstaff, A. (2002). 'Poverty and health sector inequalities', *Bulletin of the World Health Organization* 80(2): 97-105.
- Wagstaff, A. (2003). 'Child health on a dollar a day: Some tentative cross-country comparisons', *Social Science & Medicine* 57(9): 1529-38.
- Wagstaff, A. and E. van Doorslaer (2000a). 'Equity in health care finance and delivery', in *Handbook of Health Economics* A. J. Culyer and J. P. Newhouse. Amsterdam, Elsevier North-Holland. 1B: 1803-62.
- Wagstaff, A. and E. van Doorslaer (2000b). 'Measuring and testing for inequity in the delivery of health care', *Journal of Human Resources* 35(4): 716-33.
- Wagstaff, A. and E. van Doorslaer (2003). 'Catastrophe and impoverishment in paying for health care: With applications to Vietnam 1993-98', *Health Economics* 12: 921-34.
- Wagstaff, A., E. van Doorslaer, S. Calonge, T. Christiansen, et al. (1992). 'Equity in the finance of health care: Some international comparisons', *Journal of Health Economics* 11(4): 361-88.

- Wagstaff, A., E. van Doorslaer, H. van der Burg, S. Calonge, et al. (1999). 'Equity in the finance of health care: some further international comparisons', *Journal of Health Economics* 18(3): 263-90.
- Wagstaff, A., E. van Doorslaer and N. Watanabe (2003). 'On decomposing the causes of health sector inequalities, with an application to malnutrition inequalities in Vietnam', *Journal of Econometrics* 112(1): 219-27.
- Wagstaff, A. and N. Watanabe (2003). 'What difference does the choice of SES make in health inequality measurement?' *Health Economics* 12(10): 885-90.
- Wan, G. (2004). 'Accounting for income inequality in rural China: A regression based approach', *Journal of Comparative Economics* 32: 348-63.
- Wyszewianski, L. (1986). 'Financially catastrophic and high cost cases: Definitions, distinctions and their implications for policy formulation', *Inquiry* 23(4): 382-94.
- Xu, K., D. E. Evans, K. Kawabate, R. Zeramdini, et al. (2003). 'Household catastrophic health expenditure: A multicountry analysis', *Lancet* 362: 111-17.
- Younger, S. D. (2003). 'Benefits on the margin: Observations on marginal benefit incidence analysis', *World Bank Economic Review* 17(1): 89-106.