

Catastrophic Medical Expenditures

Reflections on Three Issues

Adam Wagstaff



WORLD BANK GROUP

Development Economics

Development Research Group

November 2018

Abstract

There is a large empirical literature analyzing ‘catastrophic’ medical expenses. The ‘basic’ approach defines expenses as catastrophic if they exceed a prespecified percentage of total consumption or income; the approach, in effect, tells us whether expenses cause a large percentage reduction in living standards. An alternative ‘ability-to-pay approach’ defines expenses as catastrophic if they exceed a prespecified percentage of consumption or income less actual expenses on nonmedical necessities or an allowance for them. This paper argues that this approach does not, contrary to what is sometimes claimed, tell us whether expenses are large enough to undermine a household’s ability to purchase nonmedical necessities. The paper also compares the income- and consumption-based variants of the ‘basic’ approach, showing that if the individual is a borrower after a health shock, the income-based ratio will exceed the consumption-based ratio, while the opposite is true when the individual continues to be a saver after

a health shock. The paper also explores the relationship between the income- and consumption-based ratios and the more theoretically correct ratio proposed by Flores et al. It shows that if the individual is a borrower after a health shock, both the consumption- and income-based ratios will exceed the Flores et al. ratio, with the income-based ratio overestimating it by more, while in the case when the individual continues to be a saver after a health shock, the income-based ratio will overestimate Flores et al.’s ratio by less and may not overestimate it at all. Last, the paper proposes a lifetime money metric utility approach that defines medical expenses as catastrophic in terms of their lifetime (rather than single-period) consequences. It shows that under certain assumptions, the money-metric utility approach and the Flores et al. approaches are identical, and both are operationalizable without data on how households finance their medical expenses.

This paper is a product of the Development Research Group, Development Economics. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at <http://www.worldbank.org/research>. The author may be contacted at awagstaff@worldbank.org.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

Catastrophic Medical Expenditures: Reflections on Three Issues

By

Adam Wagstaff

Development Research Group, World Bank, Washington DC, USA

Key words: catastrophic medical expenditures; impoverishing medical spending; impoverishment; out-of-pocket medical expenditures; financial protection; financial hardship

Corresponding author: Adam Wagstaff, Development Research Group, World Bank, Washington DC 20433, USA, awagstaff@worldbank.org

JEL classification: I10, I14, I15

Acknowledgements

My thanks to Eddy van Doorslaer and Patrick Eozenou for helpful comments on a draft on which this paper is based. The findings, interpretations and conclusions expressed in this paper are entirely mine, and do not necessarily represent the views of the World Bank, its Executive Directors, or the governments of the countries they represent.

Introduction

An extensive applied literature exists, dating back to the 1980s (see e.g. Berki 1986; Wyszewianski 1986), measuring ‘catastrophic’ out-of-pocket medical expenses — expenses that exceed a specific fraction of a household’s total consumption or income. Indeed, the measure has become so popular that it has become one of the UN Sustainable Development Goals (indicator 3.8.2).¹ Despite its popularity and intuitiveness, there are several unresolved issues surrounding the measure. This paper explores three of them: whether to relate out-of-pocket expenses to income (or consumption) or to a transformation of income (or consumption) that might better proxy a household’s ‘ability-to-pay’; whether to relate out-of-pocket expenses to income or consumption; and how to capture the lifetime consequences of the various ways of financing (or ‘coping’ with) out-of-pocket expenses.

In this literature, we implicitly think of out-of-pocket medical expenses as involuntary and being incurred in response to a health shock, allowing the individual to return to their previous utility level but not conferring any utility per se; indeed, the receipt of medical care per se likely confers disutility. The individual’s welfare or utility is therefore captured by their consumption net of any out-of-pocket expenses. This is consistent with most recommendations on how to treat medical expenditures in studies of poverty and income distribution (Blinder 1985; Citro and Michael 1995; Deaton and Zaidi 2002).² While it is accepted among advocates of this view that there are some medical expenditures (or at least some elements thereof) that may be discretionary (plastic surgery for aesthetic reasons is the commonest example), and while this view contrasts with the branch of economic theory that treats medical expenditures as fully under the individual’s control (cf. e.g. Grossman 1972), the consensus seems to be that, on balance, it makes sense when measuring poverty and income inequality to treat medical expenses as involuntary, netting them out from income, and not including them in the consumption aggregate.³

¹ The catastrophic expense indicator is used to capture the financial protection side of universal health coverage. See indicator 3.8.2 https://unstats.un.org/sdgs/indicators/Global%20Indicator%20Framework%20after%20refinement_Eng.pdf.

² As Blinder (1985) put it: “...dollars [spent on medical care] do not buy happiness; they just maintain human capital”. Or as the US National Academy of Sciences (NAS) Panel on Poverty and Family Assistance put it: “A sick person with high medical care expenditures is not made better off than a healthy person with no or relatively low expenditures; at best, the added expenditures serve only to restore the sick person to a healthy state” (Citro and Michael 1995, p236). Or as Deaton and Zaidi (2002) put it when discussing whether a household’s consumption aggregate should include medical expenditures: “By including health expenditures for someone who has fallen sick, we register an increase in welfare when, in fact, the opposite has occurred.” (p32).

³ This advice is not always heeded in practice. The view that out-of-pocket medical expenses are nondiscretionary underpins the approach of the U.S. Census Bureau which subtracts them (and health insurance premiums) from income in its supplemental poverty measure (SPM)(Fox 2017). Interestingly, despite the recommendations of Deaton and Zaidi (2002), in practice 95% of World Bank poverty assessments either do not take medical spending out of the picture or do not appear to do so. A search in September 2017 of poverty assessments in the World Bank’s documents and reports (<http://documents.worldbank.org/curated/en/home>) yielded 178 documents of which 76 were actual poverty assessments. Of these, only 4 explicitly excluded all medical spending from the consumption aggregate. Of the rest, 5 used income (and did not subtract medical expenses), 24 explicitly included medical expenses in the consumption aggregate, 43 did not explicitly exclude them.

Should ‘ability-to-pay’ be the benchmark?

When measuring catastrophic medical expenses, it has been suggested (Wagstaff and van Doorslaer 2003) that one might relate out-of-pocket medical expenses to income (or consumption) or to a transformation of income (or consumption) that might be thought to better proxy a household’s ‘ability-to-pay’ (ATP). The ATP approach, it was suggested, might be operationalized either by deducting actual expenses on nonmedical necessities, or by deducting an allowance for such expenses. In practice, as reported in Wagstaff et al. (2018a), articles published in health and medical journals have tended to follow the ATP approach, often citing Xu et al. (2003) who trailblazed the approach, while papers in economics journals have typically employed the ‘basic’ approach, i.e. relating out-of-pocket expenses to consumption or income.⁴

The ‘basic’ approach and large percentage reductions in living standards

Assume a simple one-period model, where there are no savings and borrowing, and all income, Y , is consumed. Let C_{PRE} and C_{POST} denote consumption before and after (or without and with) the health shock necessitating out-of-pocket medical expenses, M . (Or we can think of C_{PRE} as total consumption and C_{POST} as nonmedical consumption.) Then we have the ‘basic’ catastrophic payment ratio:

$$CATA_{BASIC} = \frac{M}{Y} = \frac{M}{C_{PRE}} = \frac{C_{PRE} - C_{POST}}{C_{POST} + M}.$$

The denominator is what, under the assumptions made here, nonmedical consumption would have been in the absence of the health shock necessitating out-of-pocket medical expenses. The numerator is equal to out-of-pocket medical expenses, and therefore tells us the drop in nonmedical consumption caused by the health shock. Thus, the ratio $CATA_{BASIC}$ can be interpreted as the (negative of the) proportionate reduction in (nonmedical) consumption associated with the out-of-pocket medical spending necessitated by the health shock. So, if medical spending is 10% of pre-shock consumption, medical spending causes a 10% reduction in nonmedical consumption. If half of households have medical spending exceeding 10% of their consumption, we can infer that out-of-pocket medical spending caused a 10% or larger reduction in nonmedical consumption among half of households.

If nonmedical consumption is the accepted welfare indicator, and if policy makers are concerned about sharp drops in household welfare, we have a clear rationale for being concerned about catastrophic out-of-pocket medical expenses. This is in line with Berki’s (1986) rationale for his study: “An expenditure for medical care becomes financially catastrophic when it endangers the family’s ability to maintain its customary standard of living”.

The ‘ability-to-pay’ approach and absolute hardship

There is nothing in the above that leads us to expect that catastrophic out-of-pocket expenses necessarily involve financial hardship in an absolute sense. An individual might see her nonmedical consumption drop by, say, 25% following a health shock, but might still be comfortably off afterwards. The obvious question is whether the alternative ATP-based approach might serve as a

⁴ Some papers in economics journals also use a variant of the ability-to-pay approach, almost always deducting actual food consumption and expressing out-of-pocket expenses as a share of nonfood consumption.

useful complement to the ‘basic’ approach by shedding light on this issue, i.e. whether medical spending might be so large relative to an individual’s resources as to undermine her ability to purchase nonmedical necessities such as food. This is, after all, the rationale of the ATP approach.

Let N be the amount of spending required to cover nonmedical necessities. Then, the ATP approach redefines the catastrophic spending ratio as:

$$CATA_{ATP} = \frac{M}{Y - N} = \frac{M}{C_{PRE} - N} = \frac{C_{PRE} - C_{POST}}{C_{PRE} - N}.$$

As suggested by Wagstaff and van Doorslaer (2003), N could be actual food expenditure or a flat-rate allowance for food spending. Xu et al. (2003) adopt the latter approach, setting N equal to an estimate of the amount of spending required to reach 2000 calories a day. Of course, one might argue, as pointed out by Wagstaff and van Doorslaer (2003), that N should probably capture more than food expenses – there are other nonmedical expenses, such as housing, heating, clothing, etc. that are also likely to be of concern to a policy maker. This leads Wagstaff and Eozenou (2014) to suggest setting N equal to the poverty line on the grounds that this is society’s definition of overall nonmedical needs. Thomson et al. (2016) build on this suggestion empirically.

The ATP-based ratio allows us to classify people into three groups (cf. Wagstaff and Eozenou 2014):

- a) individuals who have insufficient resources even before medical spending to meet their nonmedical needs: $M/(C_{PRE} - N) < 0$;
- b) individuals who, despite spending on medical care, can still finance their nonmedical needs: $0 < M/(C_{PRE} - N) < 1$; and
- c) individuals who would have been able to meet nonmedical needs in the absence of medical spending but cannot ‘after’ medical spending: $M/(C_{PRE} - N) > 1$.

If N is set equal to the poverty line per Wagstaff and Eozenou’s suggestion, individuals in group (a) are pushed further into poverty by out-of-pocket expenses, while those in group (c) are pushed below the poverty line having been above it ‘before’ the health shock that necessitated the out-of-pocket expenses (i.e. they are ‘impoverished’ by out-of-pocket expenses).

All of those in group (c) and some in group (b) incur catastrophic expenses in the ATP approach. Within these two groups, however, the size of the catastrophic spending ratio does not shed light on the extent of, or risk of absolute hardship. This is illustrated by the examples in Table 1, where the expenditure required to meet nonmedical needs (or the poverty line) is assumed to be \$1.90-a-day and the catastrophic payment threshold is assumed to be set at 10%.

Table 1:
Four hypothetical cases illustrating the catastrophic payment ratio defined with respect to
consumption less nonmedical needs

Case	Indiv.	C_{PRE}	M	C_{POST}	C_{PRE-N}	C_{POST-N}	$M/(C_{PRE-N})$
a	1	\$1.80	\$0.15	\$1.65	-\$0.10	-\$0.25	-150%
a	2	\$1.77	\$0.20	\$1.57	-\$0.13	-\$0.33	-150%
b	3	\$2.14	\$0.02	\$2.12	\$0.24	\$0.22	8%
b	4	\$3.10	\$0.10	\$3.00	\$1.20	\$1.10	8%
b	5	\$4.30	\$0.80	\$3.50	\$2.40	\$1.60	33%
b	6	\$6.00	\$1.35	\$4.65	\$4.10	\$2.75	33%
c	7	\$3.10	\$1.30	\$1.80	\$1.20	-\$0.10	108%
c	8	\$4.25	\$2.55	\$1.70	\$2.35	-\$0.20	108%

Individuals 1 and 2 illustrate case (a). In such cases, $CATA_{ATP}$ is negative because even before the health shock necessitating the out-of-pocket medical expenses, the individual's consumption is less than the poverty line. So, in this case, neither individual is recorded as having catastrophic expenses. Yet in both cases, out-of-pocket expenses push the individual still further into poverty. Both individuals have the same value of $CATA_{ATP}$, yet individual 2 ends up further below the poverty line.

Individuals 3, 4 and 5 illustrate case (b). Individuals 3 and 4 are not classified as having incurred catastrophic spending, both having a value of $CATA_{ATP}$ of 8%. Yet individual 3 comes closer to being pushed into poverty by out-of-pocket medical expenses than individual 4. Individual 5 also has sufficient income before and after the health shock to keep out of poverty, but her out-of-pocket expenses are sufficiently high relative to her income for her expenses to be classified as catastrophic. Yet, even after paying her medical expenses, individual 5 – whose expenses are are classified as catastrophic – is further from the poverty line than individuals 3 and 4 whose expenses are not classified as catastrophic.

Individuals 7 and 8 illustrate case (c). Both suffer from health shocks that results in them going from being above the poverty line 'before' medical expenditures to being below the poverty line 'after' them. Both are thus impoverished by out-of-pocket expenses. Both also incur catastrophic spending at the 10% threshold using the ATP definition. Moreover, both have the same value for the $CATA_{ATP}$ ratio. Yet individual 8 ends up \$0.20 below the poverty line while individual 7 ends up only \$0.10 below the poverty line.

Conclusions regarding the 'basic' vs. ATP approaches

Thus the 'basic' approach to catastrophic medical expenditures can be understood as an attempt to see how far out-of-pocket expenses on health cause relative hardship in the sense they cause a large percentage reduction in a household's living standards compared to what they would have been in the absence of the health shock necessitating the out-of-pocket expenses. While at first glance, the ATP approach might seem a useful complement to this basic approach by potentially shedding light on absolute hardship in the sense that out-of-pocket expenses might be so large relative to a household's resources as to undermine its ability to purchase nonmedical necessities, such as food, in practice, as shown in the examples above, this turns out not to be the case. The ATP approach tells us whether impoverishment occurs (case (c)), whether an already-poor individual is pushed still further into poverty (case (a)), or whether, despite incurring medical expenses, the individual

manages to stay above the poverty line (case (b)). But within each of these cases, the ATP-based ratio does not tell us how far out-of-pocket spending eats into resources required for necessities, nor how close this is to happening.

Consumption vs. income

The second issue explored in this paper is whether when computing the incidence of catastrophic spending one should relate out-of-pocket expenses to income or consumption. Wagstaff and van Doorslaer (2003) developed their arguments using income. But then in the empirical illustration, because they used data from a developing-country setting, where income has traditionally been argued to be harder to measure than consumption, they replaced income by consumption, proxying pre-payment income by consumption gross of out-of-pocket expenditures, and post-payment income by consumption net of out-of-pocket expenditures. The questions arise: Which measure produces the higher incidence of catastrophic out-of-pocket medical expenses? And which is to be preferred? The answer to both questions, it turns out, depends on whether, after the health shock necessitating the out-of-pocket expenses, the individual in question is a borrower or a saver.

The case of zero borrowing (or saving)

Assume the individual derives utility from nonmedical consumption in just two periods (the current period (period 0) and the future (period 1)), but needs to be in good health to be able to consume. For simplicity, assume initially that the individual has the same income in the two periods (we relax this assumption in a moment), and chooses to consume the same amount in both periods, so point a in Figure 1 is her income endowment and her consumption choice in the absence of the health shock. (The results that follow can be generalized to the case where point a is off the 45°-line – what matters is the individual either cannot borrow or can do but borrows only for medical spending.) The individual starts on indifference curve I_1 and spends nothing on medical care. She then experiences a health shock requiring a fixed amount of medical care spending M^0 in period 0.

Assume for the moment the individual is unable to borrow. After the shock, she moves to point b on indifference curve I_2 where her nonmedical consumption is lower by the full amount of M^0 , and equal to the distance OB . If we observe income (OA) and medical spending (BA) in our data, we can compute the income-based catastrophic payment ratio:

$$CATA_{INC} = \frac{M^0}{Y^0},$$

which in this case is equal to:

$$CATA_{INC} = \frac{BA}{OA}.$$

If we do not observe income, but do observe nonmedical consumption (OC), we can compute the consumption-based catastrophic payment ratio. We add medical spending (BA) to nonmedical consumption (OB) to get ‘consumption gross of medical expenditures’ (OA), which is our proxy for pre-payment income, and the denominator for our consumption-based catastrophic payment ratio. Thus, our consumption-based catastrophic payment ratio is equal to:

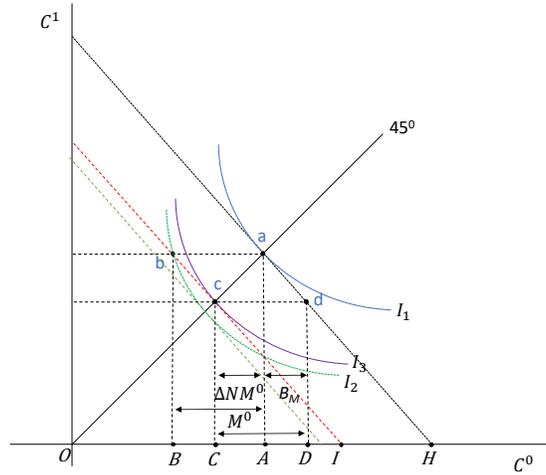
$$CATA_{CONS} = \frac{M^0}{NM^0 + M^0},$$

which in this case is also equal to:

$$CATA_{CONS} = \frac{BA}{OA}.$$

So, in the case where borrowing is not possible, we have $CATA_{INC} = CATA_{CONS}$, and there is no choice to make.

Figure 1: The cases where borrowing is not possible and borrowing occurs only for medical care



The case where borrowing can occur, but only to finance medical expenses

Things look different if the individual starting at a in Figure 1 can, in fact, borrow to finance medical expenses. Instead of consuming at point b , she now consumes at point c , reaching a higher indifference curve, I_3 . She finances her medical expenditure in part by a cut in her nonmedical consumption of an amount ΔNM^0 , and in part by borrowing an amount equal to B_M . Her total consumption (including medical expenditures) has increased, as a result of the health shock, from OC to OD . If we observe medical spending (BA) and income (OA), we can compute the income-based catastrophic payment ratio, which is equal, as before, to:

$$CATA_{INC} = \frac{BA}{OA}.$$

If we observe nonmedical consumption (OC) and medical spending ($BA=CD$), we can compute the consumption-based catastrophic payment ratio, we can compute consumption gross of medical expenditures which is now equal to OD . Thus, we now have:

$$CATA_{CONS} = \frac{BA}{OD} = \frac{CD}{OD}.$$

So, in the case where borrowing occurs only to finance health spending, we have

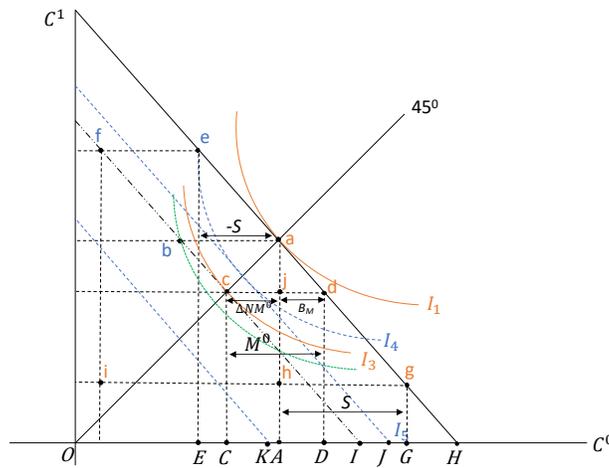
$$CATA_{CONS} \left(= \frac{BA}{OD} \right) < CATA_{INC} \left(= \frac{BA}{OA} \right).$$

The reason is that medical expenditures inflate the individual's total consumption in the current period which we use as the denominator in the consumption-based approach.

The case where borrowing can occur for any reason

What if the individual borrows or saves in general, and not just for medical care? Consider the case of a borrower who borrows even more after a health shock. Suppose such an individual starts with an income endowment shown by point e in Figure 2, and borrows an amount $-S$ to reach the consumption bundle at point a . The health shock causes her to move to point c which she reaches by borrowing an additional amount equal to B_M and cutting back nonmedical consumption by an amount equal to ΔNM^0 . Or consider the case of a saver who reduces savings to finance health spending necessitated by a health shock. This individual starts with the income endowment shown by point g in Figure 2, and saves an amount S to reach the consumption bundle at point a . The health shock causes her to move to point c which she reaches by reducing her savings by an amount equal to B_M and cutting back nonmedical consumption by an amount equal to ΔNM^0 .

Figure 2: The cases of a borrower who borrows even more after a health shock, and a saver who saves less after a health shock



How do the two catastrophic expenditure shares compare in this more general case? We can see the relationship between them by noting that $CATA_{CONS}$ can be written:

$$CATA_{CONS} = \frac{M^0}{NM^0 + M^0} = \frac{M^0}{Y^0 - S + B_M}.$$

Thus, we have:

$$CATA_{CONS} < CATA_{INC} \Rightarrow \frac{M^0}{Y^0 - S + B_M} < \frac{M^0}{Y^0} \Rightarrow S < B_M$$

So, we have $CATA_{CONS} < CATA_{INC}$ for sure when $S < 0$ (e.g. the income endowment is point e in Figure 2) or when $S = 0$ (i.e. the income endowment is point a). We also have $CATA_{CONS} < CATA_{INC}$ when $0 < S < B_M$ (i.e. the income endowment lies between points a and d), and we have $CATA_{CONS} = CATA_{INC}$ when $S = B_M > 0$ (i.e. the income endowment is point d). We only have $CATA_{CONS} > CATA_{INC}$ when $S > B_M > 0$ (i.e. the income endowment lies to the southeast of point d , e.g. point g). Thus, when the individual is a borrower after a health shock, the income-based ratio

will exceed the consumption-based ratio, while the opposite is true when the individual continues to be a saver even after a health shock.

Comparing with the ideal measure

So far, we have simply compared the two measures. The analysis begs the question: Which of the two measures is closer to the truth? Ultimately, when we compute the catastrophic payment ratio we are trying to learn whether an individual's (current) (nonmedical) consumption is substantially reduced by the out-of-pocket medical expenses necessitated by a health shock. We are therefore interested in computing the percentage reduction in (current) (nonmedical) consumption due to the out-of-pocket medical expenses necessitated by the health shock (cf. Berki 1986).

As pointed out in Wagstaff and van Doorslaer (Wagstaff and van Doorslaer 2003), the consumption-based catastrophic expenditure ratio overestimates the counterfactual current consumption: in Figures 1 and 2, the counterfactual consumption bundle is shown by point *a*, not by point *d* as the consumption-based approach assumes. Moreover, both the income- and consumption-based approaches assume that the reduction in current consumption caused by the health shock is equal to the amount spent on medical care; in fact, the fall in consumption caused by the out-of-pocket expenditures is equal to *CA* not *BA* or *CD* as is assumed.

In the light of this, Flores, Krishnakumar, O'Donnell and van Doorslaer (2008) propose a coping-adjusted catastrophic payment ratio:

$$CATA_{FKOV} = \frac{M^0 - B_M}{NM^0 + M^0 - B_M},$$

which in Figures 1 and 2 is equal to:

$$CATA_{FKOV} = \frac{CA}{OA}.$$

This measure is hard to compute because household surveys rarely contain the necessary data on the amount borrowed or dissaved to finance out-of-pocket medical expenses. In fact, in their empirical illustration, Flores et al. had to go back to an Indian survey from 1996; the later Indian survey (conducted in 2004) did not contain the necessary breakdown. So, it is of interest to know how the theoretically incorrect (but readily computed) measures – the income- and consumption-based measures – relate to the (more) theoretically correct (but more data-demanding) coping-adjusted measure.

We can show that we always have:

$$CATA_{CONS} \left(= \frac{BA}{OD} \right) > CATA_{FKOV} \left(= \frac{CA}{OA} \right).$$

To see this, note that $CATA_{FKOV}$ can be written:

$$CATA_{FKOV} = \frac{M^0 - B_M}{NM^0 + M^0 - B_M} = \frac{M^0 - B_M}{Y^0 - S}.$$

We must have $CATA_{CONS} > CATA_{FKOV}$, since the opposite can only be true if $B_M < 0$:

$$CATA_{CONS} < CATA_{FKOV} \Rightarrow \frac{M^0}{NM^0 + M^0} < \frac{M^0 - B_M}{NM^0 + M^0 - B_M} \Rightarrow B_M < 0.$$

By contrast, the ordering of $CATA_{FKOV}$ and $CATA_{INC}$ depends on the values of B_M , S , M^0 and Y^0 . Specifically, we have:

$$CATA_{INC} < CATA_{FKOV} \Rightarrow \frac{M^0}{Y^0} < \frac{M^0 - B_M}{Y^0 - S} \Rightarrow \frac{B_M}{S} < \frac{M^0}{Y^0} \Rightarrow \frac{B_M}{M^0} < \frac{S}{Y^0}.$$

The above inequality cannot hold when $S < 0$, so in this case we have $CATA_{FKOV} < CATA_{INC}$. This is labeled zone 1 in Table 2, which collects the various results for convenience. The inequality is also reversed when $S = 0$, which is labeled zone 1/2. In the case where $0 < S < B_M$ (zone 2), we have $B_M/S > 1$, and hence in this case too we have $CATA_{FKOV} < CATA_{INC}$. Likewise, in the case where $S = B_M > 0$, we can also conclude that $CATA_{FKOV} < CATA_{INC}$, since otherwise we would have $M^0 > Y^0$ (zone 2/3). When $S > B_M > 0$, there are two possible outcomes: (a) the case where $B_M/S > M^0/Y^0$ or equivalently $B_M/M^0 > S/Y^0$, and where we still have $CATA_{FKOV} < CATA_{INC}$ (zone 3); and (b) the case where $B_M/S < M^0/Y^0$ or equivalently $B_M/M^0 < S/Y^0$, and where we have $CATA_{FKOV} > CATA_{INC}$ (zone 4). This latter case is, in fact, impossible if, as is assumed in Figure 2, the individual chooses to equalize consumption in the two periods before and after the health shock. The ratio B_M/M^0 in Figure 2 is equal to the ratio jd/cd . This is the same as the ratio hg/ig . The numerator of the latter ratio is equal to S , but the denominator is less than Y^0 ; hence in Figure 2 we have $B_M/M^0 = jd/cd = hg/ig > S/Y^0$.

Table 2: Summary of results

Zone	Figure 2	Description	Ordering of catastrophic ratios
1 $S < 0$	NW of point a , e.g. e	Individual borrows before the health shock, and borrows even more afterwards	$CATA_{FKOV} < CATA_{CONS} < CATA_{INC}$
1/2 $S = 0$	Point a	Individual is neither a saver nor a borrower before the health shock	$CATA_{FKOV} < CATA_{CONS} < CATA_{INC}$
2 $0 < S < B_M$; hence $B_M/S > 1 > M^0/Y^0$	Point between a and d	Individual saves before the health shock but is a borrower afterwards	$CATA_{FKOV} < CATA_{CONS} < CATA_{INC}$
2/3 $S = B_M > 0$; hence $B_M/S = 1$	Point d	Individual saves before the health shock but is neither a saver nor a borrower afterwards	$CATA_{FKOV} < CATA_{CONS} = CATA_{INC}$
3 $S > B_M > 0$; and $1 > B_M/S > M^0/Y^0$	Point to southeast of d , e.g. g	Individual saves before and after the health shock, but saves less afterwards	$CATA_{FKOV} < CATA_{INC} < CATA_{CONS}$
4 $S > B_M > 0$; and $B_M/S < M^0/Y^0$	Not possible in Figure 2 if consumption choices are on 45°-line	Individual saves before and after the health shock, but saves less afterwards	$CATA_{INC} < CATA_{FKOV} < CATA_{CONS}$

Conclusions regarding income vs. consumption

When the individual is a borrower after a health shock, the income-based ratio will exceed the consumption-based ratio, while the opposite is true when the individual continues to be a saver

even after a health shock. This does not answer the question of which is to be preferred. If the aim of the exercise is to calculate the percentage reduction in (current) consumption due to the out-of-pocket medical expenses necessitated by the health shock, the (more) theoretically correct catastrophic expenditure measure is that proposed by Flores et al. (2008), equal to the reduction in consumption due to the health shock necessitating the out-of-pocket expenses, expressed as a ratio of counterfactual consumption. The above shows that when the individual is a borrower after a health shock, the consumption- and income-based catastrophic payment ratios will both exceed Flores et al.'s more theoretically correct ratio, with the income-based ratio overestimating it by more. However, when the individual is still a saver even after a health shock, the income-based ratio will overestimate the Flores et al. ratio by less, and may not actually overestimate it at all.

Capturing the lifetime consequences of coping

The coping-adjusted catastrophic payment measure proposed by Flores et al. (2008) is superior to the traditional measure that assumes, incorrectly, that current consumption falls by the full amount of the health expense and hence overestimates the counterfactual consumption. However, Flores et al.'s approach, like the traditional approach, focuses on the implications of out-of-pocket medical expenses (and coping strategies) for current consumption; they ignore the implications for future consumption. These measures do not therefore get at the lifetime welfare reduction caused by the medical expenditures necessitated by the health shock (cf. Wagstaff 2010). This last section asks how one might modify the approach to capture the lifetime consequences of out-of-pocket medical expenses and the coping strategies used to finance them.

A lifetime perspective

A natural ratio to get at the lifetime welfare reduction (in percentage terms) of out-of-pocket expenses is IH/OH in Figures 1 and 2, or equivalently CD/IH . Assuming the individual can borrow, IH is the amount of money she would need to receive after the health shock to restore her to her utility before it, i.e. the equivalent variation (EV).⁵ IH in this case is equal to the health expense \mathcal{M} . If the individual is unable to borrow, she moves to indifference curve I_2 rather than I_3 , and her EV will exceed \mathcal{M} . Thus, when borrowing is allowed, the amount spent on health care gives us the money-metric utility reduction; but when borrowing is not allowed, health spending likely understates the money-metric utility reduction.

It turns out that under certain assumptions, including those made in Figures 1 and 2, the lifetime money metric utility (LMMU) reduction measure IH/OH and the Flores et al. measure CA/OA are identical. The assumptions are that the individual's consumption in one period is a constant fraction or multiple of the previous period's consumption. In the two-period case, the assumption means that consumption choices, before and after a health shock, lie on a linear consumption expansion path through the origin, which includes as a special case the case illustrated in Figures 1 and 2.

Take the case where there are just two periods, and the individual's consumption choices are as in Figure 1, i.e. consumption is equalized across periods, before and after (or with and without) the

⁵ It is also the amount of money she would be willing to pay before the health shock to avoid being pushed down to the level of utility after the health shock, i.e. the compensating variation (CV).

health shock. The LMMU measure is IH/OH or equivalently M^0/OH . If r is the interest rate, and we assume $C_{PRE}^0 = C_{PRE}^1 = C_{PRE}$, we have:

$$OH = C_{PRE}^0 + \frac{C_{PRE}^1}{1+r} = C_{PRE} \frac{2+r}{1+r}$$

So, then we have:

$$LMMU = \frac{IH}{OH} = \frac{M^0}{C_{PRE} \frac{2+r}{1+r}} = \frac{M^0 \frac{1+r}{2+r}}{C_{PRE}}$$

The measure of Flores et al. (hereafter FKOV) is equal to CA/OA . We know the denominator is C^0 . From the above, if we assume consumption is equalized across periods, we can write:

$$OI = C_{POST} \frac{2+r}{1+r}$$

But we know $OH - OI = M^0$. Hence:

$$M^0 = C_{PRE} \frac{2+r}{1+r} - C_{POST} \frac{2+r}{1+r} = \frac{2+r}{1+r} (C_{PRE} - C_{POST}),$$

and therefore

$$CA = (C_{PRE} - C_{POST}) = M^0 \frac{1+r}{2+r}$$

So, the Flores et al. measure is:

$$FKOV = \frac{CA}{OA} = \frac{(C_{PRE} - C_{POST})}{C_{PRE}} = \frac{M^0 \frac{1+r}{2+r}}{C_{PRE}}$$

which is the same as the LMMU measure. The two measures put medical expenditures and consumption on a common footing, on the assumption that medical expenses are a one-off expense while consumption keeps occurring period after period: one can think of the measures as relating only a fraction of medical expenses to current consumption (the first formulation in the expression for LMMU above) or as relating the entire amount of medical expenses to the present value of consumption (the second formulation).

We can generalize this result to multiple periods. Where we have T periods, we have:

$$OH = C_{PRE}^0 + \frac{C_{PRE}^1}{1+r} + \frac{C_{PRE}^2}{(1+r)^2} + \dots + \frac{C_{PRE}^{T-1}}{(1+r)^{T-1}}$$

which if consumption is constant across periods becomes

$$OH = C_{PRE} \left[1 + \frac{1}{(1+r)} + \frac{1}{(1+r)^2} + \dots + \frac{1}{(1+r)^{T-1}} \right] = \tau(r, T) C_{PRE},$$

where

$$\tau(r, T) = \frac{1 - \left(\frac{1}{1+r}\right)^T}{1 - \frac{1}{1+r}}.$$

So, then we have:

$$LMMU = \frac{IH}{OH} = \frac{M^0 / \tau(r, T)}{C_{PRE}}.$$

We can see how this compares to the Flores et al. measure CA/OA. From the above, if we assume consumption is equalized across periods, we can write:

$$OI = \tau(r, T) C_{POST}.$$

But we know $OH - OI = M^0$ and therefore

$$M^0 = \tau(r, T) C_{PRE} - \tau(r, T) C_{POST} = \tau(r, T) (C_{PRE} - C_{POST}).$$

Thus, the Flores et al. measure is

$$FKOV = \frac{CA}{OA} = \frac{(C_{PRE} - C_{POST})}{C_{PRE}} = \frac{M^0 / \tau(r, T)}{C_{PRE}},$$

which is again the same as the LMMU measure.

We can still get the equivalence result under a weaker assumption. Suppose that rather than consumption being equalized over time, consumption in period $t+1$ is a constant fraction or multiple of consumption in period t . Thus, we assume $C_{PRE}^1 = \alpha C_{PRE}^0$, $C_{PRE}^2 = \alpha^2 C_{PRE}^0$. In this case, we have (see Annex):

$$LMMU = \frac{IH}{OH} = \frac{M^0 / \varphi(r, T)}{C_{POST}^0} = FKO V = \frac{CA}{OA} = \frac{(C_{PRE}^0 - C_{POST}^0)}{C_{PRE}^0} = \frac{M^0 / \varphi(r, T)}{C_{POST}^0},$$

where

$$\varphi(r, T) = \frac{1 - \left(\frac{\alpha}{1+r}\right)^T}{1 - \frac{\alpha}{1+r}}.$$

Operationalizing the FKOV and LMMU approaches

So, under certain assumptions, the Flores et al. approach gives the same answer as the (lifetime-consistent) lifetime money metric method. The challenge with the Flores et al. method (and the LMMU method) is that with the typical cross-section household survey, we do not have the necessary information to operationalize the method. It turns out we can, in fact, operationalize the Flores et al. and LMMU methods if we are willing to make the same assumptions that ensure the equivalence of the two methods.

Our problem is that we do not observe C_{PRE} . But we can substitute in for the variables we do observe using the fact that:

$$C_{PRE} = \frac{M^0}{\tau(r, T)} + C_{POST}.$$

So, we have

$$LMMU = FKOV = \frac{M^0/\tau(r, T)}{C_{PRE}} = \frac{M^0/\tau(r, T)}{M^0/\tau(r, T) + C_{POST}}.$$

So, for example, if $r=10\%$, $T-1=25$, $C_{POST}=90$ and $M^0=10$, then the regular consumption-based catastrophic spending ratio is 10%, while the LMMU and FKOV measures (under the assumptions we have made) are only 1.1%. In empirical work, one would of course have to make an assumption about the lifecycle length T , which could vary from household to household depending on their age structure. And, depending on the availability of data in the survey on interest rates charged on loans (perhaps specifically for medical expenses), one may also have to make an assumption about the interest rate r . One could examine the sensitivity of one's estimates to the assumptions made about r and T .

More general assumptions

The assumptions we have made thus far are quite strong. The assumption that people borrow or dissave to finance medical expenses, but do not finance them fully out of borrowing and dissaving is consistent with studies that find that people borrow to finance medical expenses (Demirguc-Kunt and Klapper 2012) and that nonhealth consumption falls after a health shock (see e.g. Gertler and Gruber 2002; Jack and Suri 2014). The evidence is consistent, in other words, with one implication of our assumptions, namely that people partially smooth consumption after a health shock.

However, our assumption is stronger than partial smoothing, requiring that people equalize consumption over time, or at least plan their consumption such that consumption in period $t+1$ is a constant fraction or multiple of consumption in period t for all t . Even the latter assumption is strong. It is true that the evidence suggests that people choose a path for consumption over the lifecycle that is smoother than their income profile, by borrowing in periods where income is low,

and saving when it is high. But the evidence also suggests that people do not equalize consumption over time (see e.g. Browning and Crossley 2001). And even the assumption that consumption in period $t+1$ is a constant fraction or multiple of consumption in period t is unlikely to be consistent with evidence.

Without these assumptions, the Flores et al. approach will likely not give the same result as the LMMU method, which is the method one would want to use if one wants to factor in the full lifetime welfare impacts of health shocks. Moreover, the LMMU method will likely no longer be operationalizable using assumptions about T and r , and data on current-period medical expenses and nonmedical consumption. (The “will likely” is because the assumptions are sufficient to give the results above.)

To see the possible difference between the two approaches when the assumptions are relaxed, consider the following two-period example. The household’s income in both the present and the future is \$350. Initially it consumes \$350 in each period. After the health shock, the household has to spend \$100 on medical expenses, and to finance them reduces its current consumption by \$25, financing the rest through a reduction in future consumption. The ratio underlying the Flores et al. method is 7%. Suppose the interest rate is 10%. The present value of the initial consumption and income stream is \$668, making the ratio underlying the lifetime money metric utility method 15%, twice that underlying the Flores et al. method.

If we are not prepared to make the assumptions above, there are a couple of options. One is to see if there are weaker assumptions that also allow the LMMU method to be operationalized using only assumptions about T and r , and data on current-period medical expenses and nonmedical consumption.⁶ The alternative would be to find out how to operationalize the LMMU method without them. We would need to have data on current medical expenses and an estimate of the present value of the household’s consumption stream, which, since over the lifecycle the household must satisfy its budget constraint, is the same as the present value of its income stream, cf. point OH in Figure 1. If we have panel data on household income, we could quite possibly estimate the latter present value, and hence operationalize the lifetime money metric method.

Conclusions regarding capturing the lifetime consequences of coping

The LMMU approach proposed here, which gets at the lifetime consequences of coping with out-of-pocket medical expenses, has been shown under certain assumptions to give the same results as the Flores et al. approach. Moreover, both can be operationalized under these assumptions even in the context of a typical cross-section household survey where data are not available on how households finance their medical expenses. However, if, as seems likely, these assumptions are considered too strong, the two approaches will not be equivalent, and to operationalize the LMMU approach, one would need to use panel data containing data on income to compute the present value of the income stream.

Conclusions

With respect to the first issue considered (whether to relate out-of-pocket expenses to income (or consumption) or to ‘ability-to-pay’ (ATP)), the paper has argued that the ‘basic’ approach to

⁶ These weaker assumptions may or may not leave the lifetime money metric utility method and the Flores et al. method equivalent; that is of academic interest, but not of special interest for empirical work.

catastrophic medical expenditures can be understood as an attempt to see how far out-of-pocket expenses on health cause relative hardship in the sense they cause a large percentage reduction in a household's living standards compared to what they would have been in the absence of the health shock necessitating the out-of-pocket expenses. The paper argued that while, at first glance, the ATP approach might seem a useful complement to this basic approach by potentially shedding light on absolute hardship in the sense that out-of-pocket expenses might be so large relative to a household's resources as to undermine its ability to purchase nonmedical necessities, such as food, in actuality this turns out not to be the case. The ATP approach tells us whether impoverishment occurs, whether an already-poor individual is pushed still further into poverty, or whether, despite incurring medical expenses, the individual manages to stay above the poverty line. But within each of these cases, the ATP-based ratio does not tell us how far out-of-pocket spending eats into resources required for necessities, and therefore does not tell us how far such spending leads to absolute hardship nor how close this is to happening.

The implications with respect to the first issue is that if the goal is to get at relative hardship, one should express out-of-pocket expenses relative to consumption (or income) and not relative to ATP. By contrast, if the goal is to get at the absolute hardship associated with out-of-pocket expenses, one should eschew the ATP approach to catastrophic expenditures and instead focus explicitly on the concept of impoverishment (Wagstaff and van Doorslaer 2003; van Doorslaer et al. 2006; Wagstaff et al. 2018b).

With respect to the second issue (whether to relate out-of-pocket expenses to consumption or income), the paper has shown that the answer depends on whether, after the health shock necessitating the out-of-pocket expenses, the individual in question is a borrower or a saver. When the individual is a borrower after a health shock, the income-based ratio will exceed the consumption-based ratio, while the opposite is true when the individual continues to be a saver even after a health shock. This does not answer the question of which is to be preferred. If the aim of the exercise is to calculate the percentage reduction in (current) consumption due to the out-of-pocket medical expenses necessitated by the health shock, the question is which comes closer to the (more) theoretically correct catastrophic expenditure measure proposed by Flores et al. (2008), defined as the reduction in consumption due to the health shock necessitating the out-of-pocket expenses, expressed as a ratio of counterfactual consumption. The paper finds that when the individual is a borrower after a health shock, the consumption- and income-based catastrophic payment ratios will both exceed Flores et al.'s more theoretically correct ratio, with the income-based ratio overestimating it by more. However, when the individual is still a saver even after a health shock, the income-based ratio will overestimate the Flores et al. ratio by less, and may not actually overestimate it at all.

One implication for applied work with respect to the second issue is that the choice between income and consumption should not be based on what data happen to be available; rather the choice should be based on whether the households involved are net savers or borrowers after dealing with medical expenses. Where households have sizable savings, and are able to absorb the costs of medical care without becoming borrowers, the income-based measure gets closer to the Flores et al. gold standard. By contrast, where households are borrowers even before health shocks, or have insufficient savings to prevent them from becoming borrowers when health shocks occur, the consumption-based approach gets closer to the truth. Data from the World Bank's FINDEX project (Demirguc-Kunt et al. 2015) suggest that in the world as a whole 12% of households

borrow for health or medical purposes: rates are higher in low-income countries than in high-income countries.⁷ Of course, these numbers are percentages of the entire population not just medical care users. Still, they do suggest that the income-based measure may be more relevant to high-income countries and the consumption-based measure more relevant to low-income countries.

Finally, with regard to the third issue (how to capture the lifetime consequences of alternative ways of coping with out-of-pocket expenses), the paper proposed a lifetime money metric utility (LMMU) approach, and showed that under certain assumptions, the LMMU approach gives the same results as the Flores et al. approach, and that both can be operationalized under these assumptions even in the context of a typical cross-section household survey where data are not available on the how households finance their medical expenses. The paper also showed, however, that if, as seems likely, these assumptions are considered too strong, the two approaches are not equivalent.

The implication for applied work with respect to the third issue is that to operationalize the more general LMMU approach to catastrophic medical expenses, one needs, in addition to data on current out-of-pocket expenses, panel data that would allow one to estimate the present value of the household's consumption stream in the absence of health shocks. Since over the lifecycle the household must satisfy its budget constraint, an obvious way to try to operationalize this approach is to use panel data on income to compute the present value of the household's income stream.

⁷ The highest rates are in Uganda (39%), Kenya (32%) and Malawi (31%). Borrowing rates in high-income countries average 5%, but Australia and the United States record rates higher than 7%.

Annex: Equivalence between the Flores et al. measure and the lifetime money metric utility measure under a more general assumption

Suppose we assume not that consumption is equalized over time, but rather that consumption in period $t+1$ is a constant fraction or multiple of consumption in period t . Thus, we assume $C_{PRE}^1 = \alpha C_{PRE}^0$, $C_{PRE}^2 = \alpha^2 C_{PRE}^0$. In this case, we have:

$$OH = C_{PRE}^0 \left[1 + \frac{\alpha}{(1+r)} + \frac{\alpha^2}{(1+r)^2} + \dots + \frac{\alpha^{T-1}}{(1+r)^{T-1}} \right] = \varphi(r, T) C_{PRE}^0,$$

where

$$\varphi(r, T) = \frac{1 - \left(\frac{\alpha}{1+r}\right)^T}{1 - \frac{\alpha}{1+r}}.$$

So, we have:

$$LMM = \frac{IH}{OH} = \frac{M^0 / \varphi(r, T)}{C_{POST}^0}.$$

To get the Flores et al. measure, we can write:

$$OI = \varphi(r, T) C_{POST}^0.$$

But we know $OH - OI = M^0$ and therefore:

$$M^0 = \varphi(r, T) C_{PRE}^0 - \varphi(r, T) C_{POST}^0 = \varphi(r, T) (C_{PRE}^0 - C_{POST}^0).$$

from which we get:

$$CA = C_{PRE}^0 - C_{POST}^0 = \frac{M^0}{\varphi(r, T)}.$$

Hence the Flores et al. measure is

$$FKOV = \frac{CA}{OA} = \frac{(C_{PRE}^0 - C_{POST}^0)}{C_{PRE}^0} = \frac{M^0 / \varphi(r, T)}{C_{POST}^0},$$

which is again the same as the lifetime money metric measure (LMM).

References

- Berki, S. (1986). "A look at catastrophic medical expenses and the poor." Health Affairs 5(4): 138-145.
- Blinder, A. S. (1985). Comment on "Measuring Income: What Kind Should Be In?" by David T. Ellwood and Lawrence H. Summers. Conference on Measurement of Noncash Benefits: Proceedings Vol. 1. Washington, DC, Bureau of the Census.
- Browning, M. and T. F. Crossley (2001). "The Life-Cycle Model of Consumption and Saving." Journal of Economic Perspectives 15(3): 3-22.
- Citro, C. F. and R. T. Michael, Eds. (1995). Measuring poverty: A new approach. Washington, D.C., National Academies Press.
- Deaton, A. and S. Zaidi (2002). Guidelines for constructing consumption aggregates for welfare analysis. Washington DC, World Bank.
- Demirguc-Kunt, A. and L. Klapper (2012). Measuring financial inclusion: the Global Findex Database. The World Bank, Policy Research Working Paper Series, 6025.
- Demirguc-Kunt, A., L. Klapper, D. Singer and P. Van Oudheusden (2015). The Global Findex Database 2014: measuring financial inclusion around the world. The World Bank, Policy Research Working Paper Series, 7255.
- 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 17 12: 1393-1412.
- Fox, L. (2017). The Supplemental Poverty Measure: 2016. Washington, DC, US Census Bureau.
- Gertler, P. and J. Gruber (2002). "Insuring consumption against illness." American Economic Review 92(1): 51-76.
- Grossman, M. (1972). "On the Concept of Health Capital and the Demand for Health." Journal of Political Economy 80 2: 223-255.
- Jack, W. and T. Suri (2014). "Risk Sharing and Transactions Costs: Evidence from Kenya's Mobile Money Revolution." American Economic Review 104(1): 183-223.
- Thomson, S., T. Evetovits, J. Cylus and M. Jakab (2016). "Monitoring financial protection to assess progress towards universal health coverage in Europe." Public Health Panorama 2(3): 357-366.
- 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 (2006). "Effect of payments for health care on poverty estimates in 11 countries in Asia: an analysis of household survey data." Lancet 368(9544): 1357-1364.
- Wagstaff, A. and E. van Doorslaer (2003). "Catastrophe and Impoverishment in Paying for Health Care: With Applications to Vietnam 1993-1998." Health Economics 12(11): 921-934.
- Wagstaff, A. (2010). Measuring financial protection in health. Performance Measurement for Health System Improvement: Experiences, Challenges and Prospects. P. C. Smith, E. Mossialos and S. Leatherman. Cambridge, Cambridge University Press: 114-137.
- Wagstaff, A. and P. Eozenou (2014). CATA meets IMPOV: a unified approach to measuring financial protection in health. Washington, DC, The World Bank.
- Wagstaff, A., G. Flores, J. Hsu, M. F. Smitz, K. Chepynoga, L. R. Buisman, K. van Wilgenburg and P. Eozenou (2018a). "Progress on catastrophic health spending in 133 countries: a retrospective observational study." Lancet Glob Health 6(2): e169-e179.
- Wagstaff, A., G. Flores, M. F. Smitz, J. Hsu, K. Chepynoga and P. Eozenou (2018b). "Progress on impoverishing health spending in 122 countries: a retrospective observational study." Lancet Glob Health 6(2): e180-e192.

- Wyszewianski, L. (1986). "Families with catastrophic health care expenditures." Health Services Research **21**(5): 617-634.
- Xu, K., D. B. Evans, K. Kawabata, R. Zeramdini, J. Klavus and C. J. Murray (2003). "Household catastrophic health expenditure: a multicountry analysis." Lancet **362**(9378): 111-117.