



The Effect of Government Health Expenditure on the Income  
Distribution: A Comparison of Valuation Methods in Ghana

*Jeremy Barofsky and Stephen D. Younger*

COMMITMENT TO EQUITY



**CEQ INSTITUTE**  
COMMITMENT TO EQUITY

**Tulane University**

Working Paper 66  
June 2019

## The CEQ Working Paper Series

The CEQ Institute at Tulane University works to reduce inequality and poverty through rigorous tax and benefit incidence analysis and active engagement with the policy community. The studies published in the CEQ Working Paper series are pre-publication versions of peer-reviewed or scholarly articles, book chapters, and reports produced by the Institute. The papers mainly include empirical studies based on the CEQ methodology and theoretical analysis of the impact of fiscal policy on poverty and inequality. The content of the papers published in this series is entirely the responsibility of the author or authors. Although all the results of empirical studies are reviewed according to the protocol of quality control established by the CEQ Institute, the papers are not subject to a formal arbitration process. Moreover, national and international agencies often update their data series, the information included here may be subject to change. For updates, the reader is referred to the CEQ Standard Indicators available online in the CEQ Institute's website [www.commitmenttoequity.org/datacenter](http://www.commitmenttoequity.org/datacenter). The CEQ Working Paper series is possible thanks to the generous support of the Bill & Melinda Gates Foundation. For more information, visit [www.commitmenttoequity.org](http://www.commitmenttoequity.org).

The CEQ logo is a stylized graphical representation of a Lorenz curve for a fairly unequal distribution of income (the bottom part of the C, below the diagonal) and a concentration curve for a very progressive transfer (the top part of the C).



## ABSTRACT

Government spending on services affects the level and distribution of welfare, but measuring its value is a challenge. To assess how publicly funded in-kind health care affects the income distribution, we must estimate its monetary value to beneficiaries. We describe and compare three approaches to measuring the distributional consequences of government health spending: average cost of provision, willingness-to-pay, and health outcomes. In addition, we estimate the value of financial risk protection from insurance, which is a benefit of health spending that can be added to each of the aforementioned approaches. Average cost is the standard method used in benefit-incidence studies (Lustig, 2018). This method values utilization of each unit of care at the government's average cost of provision, calculated with national accounts data and administrative records. Willingness to pay uses revealed preference to estimate compensating variations for health care subsidies. The health outcomes method estimates the effect of government health spending on mortality and values those mortality reductions in monetary terms.

We provide example applications for each of these methods using a national cross-section from Ghana for 2012/13. We estimate a willingness to pay model for outpatient services and find that, on average, users value those services at less than what the government pays for them. The estimated marginal effect of health spending for outpatient care on inequality are modest and somewhat smaller than those for the average cost approach. In contrast, the health outcomes method finds that the marginal effects of health spending for three causes of death and five health interventions are very large. Health interventions to reduce malaria mortality such as indoor residual spraying and distribution of insecticide-treated bed nets are strongly progressive and the averted mortality from providing anti-malarial medication dwarfs the distributional effects of any other public expenditure or tax in Ghana. Adopting the health outcomes approach dramatically changes our assessment of how public spending in Ghana affects the welfare distribution. The benefit of financial risk protection from Ghana's National Health Insurance Scheme equals 0.25% to 0.5% of income for the three poorest quartiles and between 0.5% and 1% of income for the wealthiest, yet insurance is still distributed somewhat more equally than income itself.

**JEL Codes:** I14, I15, I32, I13, H51, H40

**Keywords:** Health, Economic Inequality, Poverty, Mortality, Ghana, Full Income

---

\* Thanks to Ian Heffernan for providing valuable research assistance for this research project. We also thank Henry Aaron, Gary Burtless, Sean Higgins, Carol Levin, Nora Lustig, Eric Osei-Assibey, John Quattrochi, and Marcia Weaver for helpful comments and suggestions. This research project was conducted for the Commitment to Equity (CEQ) Institute at Tulane University and was made possible by the generous support of the Bill & Melinda Gates Foundation. For more information on CEQ, visit [www.commitmenttoequity.org](http://www.commitmenttoequity.org).

# The Effect of Government Health Expenditure on the Income Distribution: A Comparison of Valuation Methods in Ghana

Jeremy Barofsky and Stephen D. Younger<sup>1</sup>  
Commitment to Equity Institute  
Version: June 2019

## Abstract

Government spending affects the level and distribution of welfare, but measuring its value is challenging when used to buy in-kind services. To assess how publicly funded in-kind health care affects the income distribution, we must estimate its monetary value to beneficiaries. We describe and compare three approaches to measuring the distributional consequences of government health spending: average cost of provision, willingness-to-pay, and health outcomes. In addition, we estimate the value of financial risk protection from insurance, which is a benefit of health spending that can be added to each of the aforementioned approaches. Average cost is the standard method used in benefit-incidence studies (Lustig, 2018). This method values utilization of each unit of care at the government's average cost of provision, calculated with national accounts data and administrative records. Willingness to pay uses revealed preference to estimate compensating variations for health care subsidies. The health outcomes method estimates the effect of government health spending on mortality and values those mortality reductions in monetary terms.

We provide example applications for each of these methods using a national cross-section from Ghana for 2012/13. We estimate a willingness to pay model for outpatient services and find that, on average, users value those services at less than what the government pays for them. The estimated marginal effect of health spending for outpatient care on inequality are modest and somewhat smaller than those for the average cost approach. In contrast, the health outcomes method finds that the marginal effects of health spending for three causes of death and five health interventions are very large. Health interventions to reduce malaria mortality such as indoor residual spraying and distribution of insecticide-treated bed nets are strongly progressive and the averted mortality from providing anti-malarial medication dwarfs the distributional effects of any other public expenditure or tax in Ghana. Adopting the health outcomes approach changes dramatically our assessment of the distributional consequences of public spending in Ghana. The benefit of financial risk protection from Ghana's National Health Insurance Scheme equals 0.25% to 0.5% of income for the three poorest quartiles and between 0.5% and 1% of income for the wealthiest, yet insurance is still distributed somewhat more equally than income itself. In addition to the applications, we discuss the conceptual and practical advantages and disadvantages of each method and provide guidance on when each method should be used.

---

<sup>1</sup> Thanks to Ian Heffernan for providing valuable research assistance for this research project. We also thank Henry Aaron, Gary Burtless, Sean Higgins, Carol Levin, Nora Lustig, Eric Osei-Assibey, John Quattrochi, and Marcia Weaver for helpful comments and suggestions. This research project was conducted for the Commitment to Equity (CEQ) Institute at Tulane University and was made possible by the generous support of the Bill & Melinda Gates Foundation. For more information on CEQ, visit [www.commitmenttoequity.org](http://www.commitmenttoequity.org).

## 1. Introduction

Health spending in developing nations has expanded rapidly since 2000. Health expenditure now constitutes 37 percent and 23 percent of total government expenditure in low and low-middle income nations, respectively.<sup>2</sup> Many developing nations have expanded health insurance coverage, while the goal of universal health coverage even for the poorest countries has received increasing support among multilateral institutions and researchers (World Health Organization 2010; Jamison et al. 2013).<sup>3</sup> Clearly, any comprehensive attempt to understand the distributional consequences of government taxation and spending must come to grips with the benefit of publicly-funded health care. This is more difficult than most other parts of the budget because the benefits are in-kind, not cash. We need a way to value those benefits in monetary terms.

A standard CEQ Assessment does this with an estimate of the government's cost of providing the health services people receive in-kind.<sup>4</sup> The motivation for this chapter is that the standard approach may not be very accurate. We discuss criticisms of the standard approach in section 2 and then turn to two alternative methods for estimating the monetary value of in-kind health services to their beneficiaries. Throughout, we focus on methods that seem likely to be applicable in many countries. In particular, we take as a constraint that we should be able to implement the method with standard multi-purpose cross-sectional survey data similar to what a CEQ Assessment uses, perhaps along with other data that are readily available in most countries.

In addition to the standard average cost approach, we consider two alternative methods to estimate the benefits that patients receive from publicly-funded health services and a fourth method that is an add-on to the others rather than a substitute. The first uses actual demand for health services to estimate the benefits consumers receive from utilization, a revealed preference approach usually referred to as "willingness to pay." However, because consumers exhibit positive income elasticity for healthcare, if the poor and the rich exhibit the same level of health need, health care demand will be higher among the rich than the poor. Observed health care demand therefore reflects both an individual's willingness *and ability* to pay for care. Consequently, instead of referring to demand calculated using revealed preference as "willingness to pay," throughout the text we emphasize this distinction by using the term willingness and ability to pay (WATP). The demand estimates use survey respondents' choice of health care provider along the lines of the seminal papers by Gertler, Locay, and Sanderson (1987) and Gertler and van der Gaag (1990).

Our second alternative estimates the monetary value of the improved health that publicly-funded health care services generate. In particular, we estimate the reduction in mortality produced by government spending across five health interventions and three causes of death and value this averted mortality in monetary terms using the approach of Jamison et al. (2013). We calculate mortality averted through government action by comparing the mortality rate that obtains with Ghana's current level of health intervention coverage against an assumed counterfactual mortality rate that would have occurred had health intervention coverage been at the lowest level observed

---

<sup>2</sup> Authors' calculations using World Bank national accounts data.

<sup>3</sup> Recent health coverage expansions include Mexico, Ghana, Thailand, China, and India. See Cotlear et al. (2015) for a summary.

<sup>4</sup> This is usually called the "average cost" approach, but we prefer "cost of provision" to include the possibility of the "insurance value" approach discussed in Chapter 1 (Lustig, 2018).

among other West African nations. While it is also possible to generate such estimates for the mortality effects of entire health systems, our example estimates the value of mortality reduction from several specific health interventions, including the two largest causes of premature death in Ghana – malaria and HIV.

While these alternative methods each address some of the limits of the average cost approach, they have shortcomings of their own, both conceptual and practical. We discuss these in turn. It is not clear that one of the three methods is superior to the others, so we discuss situations in which researchers might want to use each of them. Our goal is to provide a menu of options for valuing publicly-funded health care benefits and some guidance on how to choose among the methods.

A common theme is that both new methods estimate the benefit of health care services to recipients rather than the cost to the government, thus addressing an important criticism of the standard approach (section 2). A more subtle point is that by divorcing the estimated benefit from the government's actual expenditure, we allow the expenditure to have positive (or negative) rates of return: our estimate of the monetary value of in-kind benefits can be more or less than the amount spent to provide it. This is not the case for most of the budget. When government makes a cash transfer, its monetary value to the recipient is the amount of the transfer, no more, no less. When government collects a tax, its monetary value is the amount paid, to a first order approximation.<sup>5</sup> But as we will see, government spending on in-kind health services can have very large rates of return. Jamison et al. (2013) for example, estimate the benefit-cost ratio of a basic package of health services in low income countries at between nine and twenty. Such high returns provide considerable leverage for public health spending's effect on the ex post welfare distribution.

Apart from the benefit of specific health services to their beneficiaries, the mere existence of publicly-funded health services provides insurance against the financial risk of catastrophic health care expenditures.<sup>6</sup> Indeed, policy makers' rationale for Britain's National Health Service and Medicare in the United States was not so much to improve people's health, but to protect them from large financial losses (Finkelstein and McKnight 2008). Note that this is different from the "insurance value" approach to measuring the value of health spending, which averages total health expenditures over all eligible beneficiaries rather than over actual users of publicly funded health services. The insurance here is financial and valuable for risk averse individuals because it lowers the probability of catastrophically large health expenditures. We estimate the change in household health spending risk across the income distribution and then calculate the monetary value of this risk change (usually a reduction) from health insurance using an expected utility framework where value depends on an assumed degree of risk aversion. The benefit of this financial risk reduction is additional to any ex post benefits associated with actual health care services received, so it can be added to any one of those benefit estimates to generate an overall benefit of public spending on health care.<sup>7</sup>

---

<sup>5</sup> [Insert reference to Younger box on how monetary value of tax to first approximation is equal to amount paid in tax.]

<sup>6</sup> See and Alam and Mahal (2014) and Acharya (2013) et al. for reviews.

<sup>7</sup> In principle, many forms of government spending provide financial risk protection: disability insurance, unemployment insurance, retirement pensions, and targeted transfers for those with low income. The theoretical literature goes back as far as the 1960s (Arrow, 1963) and there are many empirical studies. See Feldstein (1973) on health insurance, Bernheim (1987) for retirement pensions, and Bound et al. (2004) for disability insurance. Even

We apply each of these methods, along with the standard average cost of provision approach, to data for Ghana, the same data used in a CEQ Assessment for that country (Younger et al. 2016). We find that the revealed preference method produces benefit estimates that are somewhat smaller and somewhat less progressive than the standard average cost approach. The methods based on the monetary value of health improvements produce a wide range of estimates, some of which are so large that they alter dramatically the ex post income distribution, yielding enormous reductions in inequality.

This paper continues in the following way: section 2 describes the conceptual weaknesses of the current average cost method. These weaknesses constitute the basis for developing new methods to value in-kind government health spending. Section 3 describes the willingness and ability to pay approach for valuing health care using consumer observed behavior. The section applies the method to Ghana and summarizes results compared to the average cost method. Section 4 defines the health outcomes method, applies the method to Ghana, and summarizes results. Section 5 provides guidance on the circumstances in which each of the three methods should be used. Section 6 develops a method to estimate the value of insurance from reduced out-of-pocket health spending risk. It then applies the method to calculate the value of Ghana's National Health Insurance Scheme across the income distribution. Section 7 concludes.

## **2. What's Wrong with the Cost of Provision?**

Perhaps, nothing. Asking "How much does government spend to provide health service X to its beneficiaries and how is that spending distributed across the income distribution?" is a reasonable question. Call this "expenditure incidence." As long as the variation in expenditure per patient for service X is small,<sup>8</sup> using the average cost of provision for that service gives an estimate with which to calculate the expenditure incidence and allows governments to compare spending amounts across policy options.

In most instances, though, we are interested in the value of the service to the recipient, the "benefit incidence." This can be very different from the expenditure incidence because government can spend money inefficiently or corruptly (Gauthier and Wane 2008). Government may pay public sector health care staff more than what comparable private sector staff are paid (Lindauer and Nunberg 1994; Lakin 2010) or those staff may not perform the duties for which they are paid (Das and Hammer 2005; Das et al. 2012). In such cases, part of the benefit from public spending accrues to the staff, not patients. In addition, government may pad contracts with suppliers in exchange for kickbacks to functionaries, producing a similar effect. Or government may simply bungle the job for reasons of inattention or political economy, offering services of little value which patients nevertheless accept because they are free or because they assume that the health care staff know better. In all these cases, government expenditure is greater than the benefit to patients and, if it is inefficient, greater than the benefit to anyone.

---

though the topic is not limited to health spending, we include it here because it has not yet been treated in a CEQ methodological publication.

<sup>8</sup> Chapter XX in this Handbook discusses ways to keep the variation small.

On the other hand, there are at least two reasons to believe that the benefits of publicly-funded health care can be greater than what government spends to provide them. First, the marginal benefits for many health services are a step function. A first consultation to diagnose a sinus infection is valuable; a second is worth much less or nothing at all. The same is true for vaccinations, many surgeries, and infectious disease treatments. As a result, the demand for these services is discrete: high demand for one unit and usually zero demand for any further units. So even if the cost-of-provision is very low, patients will not demand more than one unit and the marginal value to them of that one unit may be much greater than its cost. It may make sense, then, to estimate the service's value and use that estimate in a benefit incidence analysis. In essence, this approach captures the consumer surplus associated with that one unit demanded. In doing so, it diverges from standard practice in national income accounting and most incidence analyses, where we value all units consumed at the marginal cost (usually the market price), thus ignoring consumer surplus. But doing so in the case of many health care services seems more egregious than for most other goods and services because, unlike most goods, health provides both intrinsic and instrumental value (Sen 1988). Better health is valuable in itself as the basis for future utility, while also permitting individuals to achieve other goals they value such as education and greater labor productivity. In addition, there exists an objective and widespread way to measure this benefit since we can credibly assume that the marginal utility of health is positive and constant without resorting to revealed preference as is necessary with other goods.

A second argument applies to public services that are non-rivalrous, where one person's use of the service does not prevent another's use. These include natural monopolies like water and sanitation systems, and public goods like information on good health care practices and vector control. Here, too, the marginal benefit of many such services is a step function, so Samuelson (1969)'s standard efficiency condition – provide the public good until its marginal cost is equal to the sum of everyone's marginal benefits -- may not apply. Spraying once for mosquitos may have joint benefits much greater than the cost, but once the mosquitos are dead, further spraying has no value, so again, the benefits are greater than the marginal cost.

If there are good reasons to believe that the benefit incidence of publicly-funded health services differs significantly from their expenditure incidence, then we need methods to estimate that benefit incidence.

### **3. Using Health Care Consumers' Choices to Estimate the Compensating Variation for Public Health Care Expenditures**

The compensating variation is the amount of money one would need to receive to keep utility constant in the face of a change in prices and/or quality. Public spending on health care usually reduces the price that patients pay for health services, so the compensating variation of that price change is an exact measure of what the public spending is worth to them in monetary terms. For a good or service with continuous demand, the compensating variation for a price change is the area under the demand curve between the old and new prices, so if we can estimate the demand function, we can calculate the compensating variation.

The demand for most health care services is discrete, not continuous. This complicates both the demand estimation and the calculation of the compensating variation, but both are still possible.

Economists working on public transport first developed models to estimate the demand for any one of a few choices for commuting to work (walk, drive, take the bus, for example) using only the fact that if a consumer chooses one option over the others, her utility from that option must be greater than the utility she would derive from any of the others (McFadden 1981; Cameron and Trivedi 2005 provide a concise textbook exposition). Suppose that consumers can choose between  $J$  health care providers and define the utility derived from option  $j$  as:

$$U_j = V_j + \varepsilon_j$$

where  $V_j$  is a deterministic component of utility and is a function of variables we observe, but  $\varepsilon_j$  is random and unobservable. McFadden calls this a “random utility model” (RUM). It is possible to estimate the demand for each option  $j$  as the probability that the consumer chooses that option, which is the probability that the  $U_j > U_k$  for all  $k \neq j$ . Let  $y$  be the option chosen, then:

$$\Pr[y = j] = \Pr[U_k - U_j < 0 \forall k \neq j] = \Pr[\varepsilon_k - \varepsilon_j < V_j - V_k \forall k \neq j]$$

If we specify the deterministic component of utility as a function of observed variables, say,  $V_j = X_j\beta_j$ , and we specify a joint distribution for the  $\varepsilon$ 's, we can estimate these demand functions with maximum likelihood. If the  $\varepsilon$ 's are distributed multivariate normal, this model is a multinomial probit.<sup>9</sup> If they are distributed type I extreme value, this is either a multinomial logit (if the  $X$ 's are constant across the options, i.e., not option-specific), a conditional logit (if the  $X$ 's are option-specific but the  $\beta$ 's are constant across options), or a mixture of the two.<sup>10</sup> And if the  $\varepsilon$ 's have the generalized extreme value distribution, this is a nested logit.<sup>11</sup>

The multinomial probit model is more general than the logit models because it allows for correlations among the random components, the  $\varepsilon$ 's. But the logit models are usually easier to estimate and so are more often used in practical applications. Both multinomial and conditional logit models assume no correlation among the  $\varepsilon$ 's. The nested logit model generalizes these models to allow some limited correlation among options' random components, assuming that options nested together are more similar to one another than are options outside their “nest”.

The logit models have the additional practical advantage that calculating compensating variations is easier compared to the multinomial probit. The compensating variation for a price change for one option,  $i$ , is defined implicitly as:

$$\max_{\{j=1,\dots,J\}} U(y - p_j^0, X_j, \varepsilon_j) = \max_{\{j=1,\dots,J\}} U(y - CV - p_j^1, X_j, \varepsilon_j)$$

where  $y$  is income and  $y - p_j$  is net income after paying for option  $j$ , i.e., income available for consuming other goods and services that provide utility;  $X_j$  are other determinants of the utility associated with choosing option  $j$ , which includes the quality of option  $j$  but also characteristics of the consumer; and  $CV$  is the compensating variation. In this simple example with one price change, the only difference between  $p_j^0$  and  $p_j^1$  is when  $j=i$ , but the formula can be applied as well to

<sup>9</sup> Stata estimates this with *asmprobit*.

<sup>10</sup> Stata calls this “alternative-specific conditional logit” and estimates it with the command *asclogit*.

<sup>11</sup> Stata estimates this with *nlogit*.

multiple price (or quality) changes. The compensating variation depends on variables we can observe, but also on the unobservable errors. To get around this, we calculate the expected value of the compensating variation. In general, this requires integrating over the joint distribution of the  $\varepsilon$ 's, something that is computationally intensive. But Herriges and Kling (1999) show how to approximate this expected value for these models, and Dagsvik and Karlstrom (2005) give another approach that reduces the integration to one dimension. Even more remarkably, if the marginal utility of income is constant, Small and Rosen (1981) derive a closed form solution for the expected value of the compensating variation.<sup>12</sup>

This procedure yields the expected value of the compensating variation for each observation (person or household) which we use as the value of the subsidy implicit in publicly-funded health services to that person or household. An important criticism of the WATP method is that since there is a positive income elasticity for health care, the compensating variation will be lower for the poor than the rich because demand reflects willingness *and* ability to pay. To overcome this issue, we first estimate individual WATP for outpatient care (described in greater detail below), but value that care to each beneficiary using the average WATP across the population in our estimated demand model. It is this average WATP value rather than the cost of provision that we use to estimate the benefit incidence of public spending on that service.<sup>13</sup>

### 3.1 Example Application – Demand for Health Care Consultations in Ghana

We model the choice of seeking health care conditional on having been ill or injured in the previous two weeks.<sup>14</sup> We assume that utility for each care option is a separable function of consumption of non-health goods and services, the quality of health care received, and a random component:

$$U_j = V_C(C_j)_j + V_Q(Q_j)_j + \varepsilon_j = \ln(Y - P_j) \cdot \alpha + X \cdot \beta_j + \varepsilon_j$$

The second equality assumes that utility is logarithmic in consumption to ensure that the marginal utility of income is declining in income, but always positive.<sup>15</sup> It also imposes a budget constraint that consumption must equal income less the cost of the chosen health care option. We assume the latter includes user fees plus the time costs of getting to a health care facility and any time spent waiting for attention there. Specifically, net consumption  $C_j$  is defined in the following way where the term in parentheses defines the cost of care choice  $j$ .

$$C_j = Y - (OOP_j + wT_j + wH_j)$$

---

<sup>12</sup> The solution is complicated; we give the specific formula for the case of two options in section 3.1 below where we apply these methods to demand for health services in Ghana.

<sup>13</sup> For the particular case we estimate in Ghana, results using each observation's WATP are not very different from the results we present here using the average WATP across individuals.

<sup>14</sup> This is simpler than the usual approach that assumes people choose between no care, public care, and private care as in Gertler, Loay, and Sanderson (1987). We do this because there is not a clear distinction between publicly-funded and private health care in Ghana. Some private providers accept payment from government for National Health Insurance Scheme participants, and some public providers charge fees to those who are not NHIS participants, so we limit our model to the no care vs. care decision.

<sup>15</sup> Estimates using a quadratic yielded a negative marginal utility of income at the highest incomes.

The price for the no care option is zero.  $OOP_j$  is the out-of-pocket (OOP) price for provider  $j$ ,  $wT_j$  is the opportunity cost of time for travel to and from the health facility, and  $wH_j$  is the opportunity cost of wait time at the health facility. We use a combination of health system structure and self-reports to calculate the OOP price of outpatient care. For those with membership in the NHIS, no coinsurance, copayment, or deductible is required at the point of service (Nguyen et al. 2011). Therefore, OOP prices for NHIS members (66% and 61% of the self-reported sick or ill sample in urban and rural areas, respectively) are set to zero. This is consistent with what is observed empirically as approximately three-quarters of individuals accessing outpatient care report paying no OOP fees at the point of service. The data set used, the Ghana Living Standards Survey (GLSS) 2012/2013, asks questions about outpatient price paid for care overall and the price by stage of care (registration, consultation, diagnosis, drugs and treatment). Because the price paid for the stages of care does not always sum to the reported total, we use the maximum of these two measures. Once defined, we estimate prices for outpatient care using the median self-reported costs for those without NHIS coverage per cluster. We set OOP outpatient fees paid for those without NHIS coverage at the district median.

We calculate the opportunity cost of time as annual household income divided by the number of work hours in a year to obtain an hourly wage per household. For households that report zero total wages earned, we replace their opportunity cost of time with the minimum reported hourly wage in their district. For the other components of cost, the GLSS 2012-13 included a community survey for rural areas that asks about travel time to the nearest health care facility and wait time at the nearest facility and cluster level median values are used from this survey. For urban areas without the community survey, we use median self-reported travel and wait times by survey cluster and, if missing, by region.

We assume that quality depends on household characteristics  $X$  that are the same across options.<sup>16</sup> Because a logit model can only identify the  $\beta$ 's relative to one option, those for the no care option are assumed to be zero, so quality is normalized such that it is zero for the no care option. Note that the  $\alpha$  coefficient is not subscripted. This constrains the marginal utility of income to be the same across options.<sup>17</sup>

Table 1 gives the results of the conditional logit demand model for outpatient care. Most of the coefficients have the expected sign, but few are statistically significant. Demand increases with income, for children under five years old, and as the number of days sick in the last two weeks increases. It is unusual that the coefficients on urban residence and the education variables are not significantly different from zero. This reflects Ghana's considerable efforts to expand health care coverage in recent years, including through the National Health Insurance System (NHIS).<sup>18</sup>

---

<sup>16</sup> In theory,  $X$  could also include option-specific characteristics such as quality measures, but we have no such data in Ghana nor is this type of data included in most income and expenditure surveys.

<sup>17</sup> Most statistical packages, including Stata's `mlogit` and `clogit` commands, normalize all the coefficients in the base option to zero which implies a different marginal utility of income for the base option. Stata's `asclogit` command allows us to constrain the marginal utility of income to be the same across all options.

<sup>18</sup> Appendix figure A.1 shows the approximately equal access to NHIS coverage across income deciles in 2012/13.

Table 1 – Conditional Logit Estimates of the Demand for Outpatient Health Consultations

Variable	Coefficient	Std Error
ln(consumption)	14.1	0.9
Urban	0.0361	0.0531
Age	-0.0043	0.0050
Age <sup>2</sup>	4.15E-05	6.14E-05
Male	-0.0832	0.0518
Child < 5 years old	0.0674	0.0298
Older than 70 years	0.0653	0.0679
Days sick	0.0797	0.0067
Primary School	-0.0729	0.0908
Junior High	-0.0159	0.1090
Senior High	0.0114	0.1194
Above High School	0.2023	0.1739
Missing Educ.	0.3228	0.0868
Use Mother's Educ.	-0.2285	0.1121

Source: GLSS-6 and authors' calculations

We use these demand estimates to calculate the compensating variation for a price change in the cost of a health care consultation equal to the average cost of provision used in the CEQ assessment for Ghana, 33.6 cedis (Younger, Osei-Assibey, and Oppong, 2017). Given that this amount is small relative to income, we assume that the marginal utility of income is constant, which makes the calculation of the compensating variation straightforward, as given by Small and Rosen (1981):

$$CV = \left(\frac{1}{\lambda}\right) \left( \ln(\exp(V_{no\ care}^{ex\ ante}) + \exp(V_{care}^{ex\ ante})) - \ln(\exp(V_{no\ care}^{ex\ post}) + \exp(V_{care}^{ex\ post})) \right)$$

where  $\lambda$  is the marginal utility of income and the V's are the estimated utility functions evaluated at the care and no care options and before and after the price change. Table 2 gives distributional statistics for the compensating variation for a price change for outpatient consultations from zero to 33.6 cedis, the average cost of provision at public health centers in 2012. It also gives the estimated benefits using the average cost of provision, the standard approach in most CEQ assessments. On average, the compensating variations are about 10 percent lower than the average cost estimates. This could reflect inefficiency or corruption in the provision of public health care services that affects demand for care, but it is also to be expected insofar as beneficiaries cannot value an in-kind benefit more than the cash required to provide that benefit.

The compensating variations are also distributed across all survey respondents in the regression sample (those reporting being sick or injured in the last two weeks) because the probability of using outpatient services among this population is almost always positive. The average cost method only gives benefits to actual users and so is more concentrated – note that the benefit for the average cost method at the 25<sup>th</sup> percentile and median are zero. Actual users, though, are more

concentrated among the poor as witnessed by the more negative concentration coefficient for the average cost method. So, in addition to lowering the overall estimate of total benefits, use of the compensating variation spreads the benefits away from the poorer people who actually used outpatient services in this sample to all those who reported themselves sick or injured.<sup>19</sup> The marginal effects<sup>20</sup> for both methods are relatively small, producing reductions of 0.8 and 0.7 percentage points for the Gini coefficient.

Table 2 – Distributional Statistics for Estimated Compensating Variation and Average Cost of Provision for Outpatient Consultations

Method	Mean	25th	Median	75th	min	max	c.c.	Gini – Marginal Effect
Average Cost of Provision	311	0	0	830	0	966	-0.187	0.008
Compensating Variation	264	157	365	365	0	365	-0.099	0.007

Source: GLSS-6 and authors’ calculations; Younger, Osei-Assibey, and Oppong, 2017.

Notes: Values in 2012 Ghanaian cedis, annualized.

The quartiles in the column headers are for estimated benefits, not income.

All statistics are on the sample of those reporting an illness or injury in the past two weeks except the concentration coefficients which are on the entire GLSS sample.

The compensating variation is for a price decrease from 33.6 cedis, the average cost of provision, to zero.

### 3.2 Discussion of the Revealed Preference Approach to Valuing Publicly-Funded Health Care

An important advantage of the compensating variation over the average cost of provision is that it anchors the estimate of the value of care in consumers’ observed behavior. In addition, because the demand estimates can be conditional on consumers’ characteristics, the value we estimate can vary across the population according to those characteristics, including need for health services. Perhaps most importantly, large discrepancies between what government pays for services and consumers willingness and ability to pay for them may reflect inefficient or corrupt government spending. Using WATP avoids erroneously attributing that expenditure to beneficiaries.

But as with all the approaches, WATP has limitations. Conceptually, this approach is applicable only to services that are private goods because we must observe demand. Practically, using survey data to estimate the demand for health care services is an order of magnitude more effort than that of the average cost approaches. It is certainly possible to estimate the demand for health care

<sup>19</sup> If the concentration coefficient for health care services using the average cost method were positive, we would also expect that the concentration coefficient for the compensating variation to be closer to zero.

<sup>20</sup> The “marginal effect” is the amount that this spending changes the Gini coefficient. [INSERT REFERENCE TO APPROPRIATE HANDBOOK CHAPTER].

services with the single cross-section of data used for a *CEQ Assessment* as we have done here. But a skeptical econometrician could easily cast doubt on whether this approach can successfully identify the demand function needed to calculate the compensating variation for a price change. In addition, as in the average cost approach, while it is theoretically possible to estimate willingness and ability to pay for many different publicly funded health services, in practice we are forced to aggregate those services into a few groups, which we assume have the same value. In our example, we aggregate all outpatient consultations into a single group.

Even the apparently attractive feature of relying on consumers' choices has been challenged in the literature because this approach assumes that consumers are rational in their health care decisions. But given the limits to consumer sovereignty in health care generally (Akerlof 1995), and particularly in low-income settings, many of the rational model's assumptions do not hold. That is, revealed preference is limited because it relies on the assumption that individuals are able to accurately estimate the expected health benefit of obtaining care. Psychological biases such as underweighting low probability events and tunneling to the present moment and information asymmetry between patient and provider mean that individual decisions often differ systematically from utility optimization. Difficulties in accurately estimating the benefit of health care choices means that the observed price elasticity of demand is often high and the implied value of care low, even when health technologies are highly effective and disease burdens are substantial. Moreover, for certain health choices, identifying benefits is more challenging than for others, making observed choice valuations lower. For highly effective health technologies for which effects are relatively difficult to observe, such as better water quality in protected springs (Kremer et al. 2011) and insecticide-treated bednets (ITN) in malaria endemic areas (Cohen and Dupas 2010), very high levels of price elasticity of demand are observed.<sup>21</sup> These results indicate that, particularly for health care services whose effects are difficult to observe (eg: most preventive care), WATP will systematically underestimate expected health benefits. Noting this discrepancy, Greenstone and Jack (2015) observe that, because individuals in high disease burden areas do not exhibit high willingness-to-pay to avoid that burden, there is "hardly a more important topic for future study than developing revealed preference measures...that capture the aesthetic, health, and/or income gains from environmental quality [such as clean water]."

In addition, low-income households experience liquidity constraints that impede decision-making<sup>22</sup> about health care, and they lack information, or the education to process information, on the returns to health care. Even when individuals intend to use health care, they may have trouble fulfilling those intentions (Laibson 1997). The limited studies in the developing world that measure WATP do indeed find values for reduced mortality risks lower by several orders of magnitude than estimates in high-income countries.<sup>23</sup> This contradiction between high health burdens (and therefore high returns to health care) and low WATP on the part of consumers challenges the

---

<sup>21</sup> Cohen and Dupas (2010) offer pregnant women in rural Kenya an ITN at prenatal clinics and find that net acquisition declines from 99% to 39% when price increases from 0 to \$US0.60.

<sup>22</sup> Mani et al. (2013). There is a response to Mani and others' argument, though. If the reason that consumers do not demand health care despite its high value is that they are liquidity constrained, then the problem is the liquidity constraint, not lack of demand for health care. If government were to relieve the liquidity constraint with a cash transfer, say, it is possible that beneficiaries would spend the money on something that provided even greater value than health care. To be consistent with the "consumer surplus" approach we take here, a CEQ assessment would then need to value the cash transfer at greater than its monetary value, too.

<sup>23</sup> Examples include Kremer et al. (2011) and León and Miguel (2017) and are discussed further below.

rational model. Health care may be worth more than low-income consumers’ observed behavior implies. That is the motivation for the health outcomes approach we address next.

#### 4. The Health Outcomes Approach

Publicly-funded health services should improve people’s health, reducing both mortality and morbidity. This has instrumental value: healthier workers are more productive; healthier children learn better. But it also has intrinsic value: health is the basis upon which all other utility is enjoyed since extending life allows individual to purchase additional utility (Hall and Jones 2007).<sup>24</sup> Given that the primary goal of a health system is to improve health status, the health outcomes approach estimates a monetary value for those improvements derived from public health care spending.

The approach begins with an estimate of the effect of government health care spending on mortality.<sup>25</sup> We estimate this by comparing health outcomes in the nation under study against counterfactual health outcomes, which represent what would have occurred without government health care spending. We assume that this counterfactual is the mortality level that would have obtained if a country experienced the minimum level of health intervention coverage observed in peer countries over a similar time period. We use readily available epidemiological models to estimate how health intervention coverage rates affect mortality. Because the mortality changes are usually small, Jamison et al. (2013) use “standardized mortality units” (SMU), a change in the probability of death of 0.0001 ( $10^{-4}$ ). Our estimate of the change in mortality for people of age  $a$  caused by a particular publicly-funded health care intervention  $h$  is:

$$\Delta SMU_{c,h,a} = M_{c,h,a} - M_{c_{cf},h,a}$$

The change in mortality,  $\Delta SMU_{c,h,a}$ , is the difference in the mortality rate among age group  $a$  in country  $c$  for the nation’s current coverage level of health intervention  $h$  minus the same mortality units in a comparator counterfactual country (“cf”) that has the lowest level of health intervention  $h$  coverage among all comparator countries. We assume that this is the mortality that would occur if the government in the country of interest spent nothing for health intervention  $h$ .<sup>26</sup>

Data for the mortality estimates must come from a source other than the household survey used for *CEQ Assessments* because income/expenditure or living standards surveys do not usually ask about mortality, and even when they do, they do not ask about health care the deceased may or may not have received. The alternate data source is usually an epidemiological model or clinical trial whose main purpose is to identify the effect of the health service of interest. These include medical trials of very specific health care interventions – one drug or care practice, for example –

---

<sup>24</sup> In addition to this reference, the health outcomes approach is derived from an earlier literature that measures the social value of mortality and morbidity from 1970-1990 in the U.S. (Cutler and Richardson 1998) and mortality only across the 20<sup>th</sup> century in the U.S. (Murphy and Topel 2006).

<sup>25</sup> In principle, the approach could be expanded to consider the value of reducing morbidity as well. We are unaware of any applications that do that. Given that its benefits are larger and easier to measure across nations, we limit our discussion to mortality.

<sup>26</sup> An alternative approach would assume that there would be no health care of type  $h$  at all in the absences of public spending on  $h$ , but this is unrealistic because there will always be some private provision of the health intervention unless it is a pure public good.

but could also include estimates for entire health systems. Results for specific health care interventions may be few for any one country, but for many aspects of health care, it is reasonable to borrow results from studies done in other countries, particularly if the two countries share similar socioeconomic, environmental, and disease transmission characteristics.

One particularly useful example of an epidemiological modelling tool for developing countries is the Spectrum system of policy models that allows researchers and policymakers to estimate the impact on mortality (but not morbidity) of health interventions for HIV, malaria, and a series of maternal, childhood, and non-communicable diseases.<sup>27</sup> These models are based on demographic data and projection models combined with epidemiological disease transmission models across multiple conditions, adjusted for the specific demographic and health data from a country of interest.

Once we have estimates for the impact of publicly-funded health care interventions on mortality, we must put a monetary value on those changes. We do this using a large literature that examines the behavior of people who systematically and voluntarily increase their mortality risk by, say, pursuing an occupation like policing or coal mining, and the additional income they earn for accepting that risk. That additional pay divided by the increased mortality risk gives an estimate of the value of small changes in mortality risk, which can be understood as the sum of what a cohort would pay for risk reductions.<sup>28</sup>

While most survey data used for a *CEQ Assessment* are sufficient to estimate simple wage equations with variables to indicate the premium for risky professions, they do not have sufficient data to estimate the mortality probabilities associated with those professions, so here too, the health outcomes approach uses secondary sources. There are many such studies with a wide range of results for the value of an SMU, but most are for developed countries. One important and uncomfortable result in those studies is that people's willingness to accept higher mortality risk varies substantially with income. This is true within countries and also between them. Hammitt and Robinson (2011) review the literature and conclude that a reasonable value for a mortality risk reduction of one SMU (VSMU) at age thirty-five is 1.8 percent of annual GDP per capita.<sup>29</sup> To adjust this value for age, Jamison et al. (2013) suggest multiplying the value of an SMU for 35-year-olds by the ratio of life-expectancy at one's current age to life expectancy at 35 years old. They also halve this value for children under five years old. We follow both these conventions in our examples, but as in Jamison et al. (2013), hold the value of an SMU constant within any one country.

For a given health intervention  $h$ , the formula to value mortality change is the following:

---

<sup>27</sup> Avenir Health (2018). Examples of Spectrum's use include Korenromp et al. (2015), Korenromp et al. (2016) and Stover et al. (2012). See section Appendix Section A.1 for further details on mortality calculations.

<sup>28</sup> This literature often aims to calculate the "value of a statistical life", and it is still known by that name (Viscusi and Aldy, 2003). But for our calculations, we do not need to value an entire (statistical) life, but only small changes in mortality probabilities measured in SMUs. In this, we follow Jamison et al. (2013) and Hammitt and Robinson (2011).

<sup>29</sup> More commonly, the value placed on changes in micro-risk of mortality are expressed as the value of a statistical life (VSL). A VSL is calculated as the aggregated value of  $10,000 \times \text{VSMU}$ . For reference, Ghana's GDP per capita in 2013 is US\$1730 (World Bank, Atlas Method) meaning the VSL used implicitly in this analysis is US\$311,400.

$$V(e'(a), e(a), y) = 0.018y \cdot \int_0^{\infty} n(a) \cdot \frac{e(a)}{e(35)} \cdot \Delta SMU_{c,h,a}(e'(a), e(a)) da$$

where:

- $e(a)$  is life expectancy at age  $a$  before the health spending of interest;
- $e'(a)$  is life expectancy at age  $a$  after the health spending;
- $V(e'(a), e(a), y)$  is the monetary value of changing from the ex ante to the ex post life expectancies;
- $y$  is GDP per capita
- $n(a)$  is the population density at age  $a$ ; and
- $\Delta SMU_{c,h,a}(e'(a), e(a))$  is the change in standardized mortality units – a change in mortality risk of  $10^{-4}$  – at each age  $a$  for country  $c$  and health intervention  $h$  that results from government health expenditure.<sup>30</sup>

Note that the integral is across the age distribution at a point in time, not across a person's life. This is the value for one year of health spending that reduces mortality probability by 0.01 percent. In words, we estimate the monetary value of a health intervention by calculating for each age group the change in mortality probability it induces, adjust for life expectancy at each age, sum those changes, weight by the age-group population share, and multiply that sum by 0.018 times GDP per capita. Again, although the value of an SMU is assumed to be proportional to income per capita across nations, within a nation, we hold it constant. The implicit assumption is that from the policy maker's perspective, eliminating the death of one citizen is equally valuable irrespective of a citizen's income.

#### 4.1 Example Application

We use the Spectrum system of policy models to estimate the mortality reduction due to five specific health interventions and three causes of death: indoor residual spraying (IRS) for mosquito control; distribution of insecticide-treated bed nets (ITNs); distribution of antimalarial drugs (mainly artemisinin) – all three to reduce malaria mortality; the distribution of antiretroviral therapy (ART) to treat HIV/AIDS; and diabetes control. In each instance, we use the Spectrum software to predict the mortality reduction across the age distribution caused by government intervention defined as the difference in health intervention coverage rates in Ghana versus the minimum coverage counterfactual rate – the lowest health intervention coverage rate among peer nations between 2011 and 2015.<sup>31</sup> We then apply Jamison et al.'s (2013) calculation of the monetary value of that increased mortality risk.

Table 3 shows mortality averted by age group due to each of the five health interventions. Note in particular the large reductions due to the distribution of anti-malarial drugs. For children up to four

<sup>30</sup> Jamison et al. (2013). Although VSL is commonly used in the literature, we focus our discussion on the value of SMUs because the maximum change in mortality risk we investigate is 46 SMUs or 0.0046 of a VSL.

<sup>31</sup> Ghana's peer nations are based on geography and data availability. They include Benin, Togo, Ivory Coast, Guinea, Liberia, Sierra Leone, and Nigeria. Data on health interventions coverage levels for IRS, ITNs, anti-malarial medication use, and diabetes prevalence come from Measure DHS and UNAIDS for ART.

years old, this is almost a one-half percentage point per year reduction in mortality. While much smaller, the other malaria interventions also have substantial effects on mortality. For antiretroviral therapy, the effects are smaller overall than ITNs or antimalarial medication, but comparable to ITNs for prime age adults. The effects of diabetes control on mortality are minimal.

Table 3 - Change in age-specific mortality rates from reducing health intervention coverage rates from current levels to the minimum rates in peer nations by age group and health intervention per 10,000 people

	<b>Indoor Residual Spraying</b>	<b>Insecticide Treated Bednets</b>	<b>Anti-Malarial Medication</b>	<b>Adult HIV Treatment - Male</b>	<b>Adult HIV Treatment -Female</b>	<b>Diabetes Control</b>
<b>&lt; 1 year</b>	0.37	4.98	46.44	-0.07	-0.07	0.00
<b>1 to 4</b>	0.37	4.98	46.44	-0.07	-0.07	0.00
<b>5 to 9</b>	0.29	3.06	35.94	0.00	0.00	0.00
<b>10 to 14</b>	0.29	3.06	35.94	0.00	0.00	0.00
<b>15 to 19</b>	0.22	1.32	13.33	0.05	0.08	0.00
<b>20 to 24</b>	0.22	1.32	13.33	0.13	0.13	0.00
<b>25 to 29</b>	0.22	1.32	13.33	0.58	0.65	0.00
<b>30 to 34</b>	0.22	1.32	13.33	1.11	0.97	0.05
<b>35 to 39</b>	0.22	1.32	13.33	2.75	2.18	0.05
<b>40 to 44</b>	0.22	1.32	13.33	4.27	2.48	0.05
<b>45 to 49</b>	0.22	1.32	13.33	4.24	1.94	0.05
<b>50 to 54</b>	0.22	1.32	13.33	3.64	1.38	0.10
<b>55 to 59</b>	0.22	1.32	13.33	2.44	0.92	0.10
<b>60 to 64</b>	0.22	1.32	13.33	1.69	0.64	0.25
<b>65 to 69</b>	0.22	1.32	13.33	1.23	0.46	0.25
<b>70 to 74</b>	0.22	1.32	13.33	0.99	0.38	0.30
<b>75 to 79</b>	0.22	1.32	13.33	0.85	0.32	0.30
<b>80 to 84</b>	0.22	1.32	13.33	0.61	0.14	0.20
<b>85+</b>	0.22	1.32	13.33	0.61	0.14	0.20
<b>Total</b>	0.26	2.23	23.21	0.92	0.59	0.10

Age-specific mortality rate (per 10,000 population). Changes calculated using Spectrum system of policy models (Avenir Health 2018).

Use of each of the modeled health interventions is not necessarily equal across the income distribution, so the next step in the analysis is to distribute the calculated reductions in mortality probability across the income distribution according to actual use. GLSS 2012/13 does not include sufficient information on use of these specific health interventions across the income distribution, so we use the 2014 Ghana Demographic and Health Survey (DHS), which does. Table 4 shows how the coverage of each health intervention or disease prevalence varies across the distribution of wealth. We take the distribution of wealth based on an index of asset ownership from the DHS as a proxy for the distribution of income. Individuals are much more likely to report receiving indoor residual spraying in the last 12 months in the poorest quintile compared to others, a

consequence of the fact that this is mostly deployed in poor, rural areas with high malaria prevalence. Distribution of insecticide treated bed nets is also more common in the poorer quintiles. These two expenditures will thus have a strongly progressive impact on the income distribution. Coverage of antimalarial drugs is lower in the poorest two quintiles so spending on them will be less progressive. It is also worth noting that the use of malaria drugs is approximately constant across the income distribution compared to measured malaria prevalence, which is highly skewed by income level.

For antiretroviral drugs and diabetes control, the Ghana 2014 DHS does not include information on actual use, but does have information on disease prevalence. We also assume that use of the corresponding health care services is proportional to prevalence, though it seems likely that this will be biased towards being too equalizing since in practice we would expect that richer people with HIV or diabetes would be more likely to get treatment in the absence of government intervention. HIV prevalence is concentrated in the middle three quintiles, while diabetes is much more common at the two highest quintiles.

Table 4 -- Access to health interventions or disease prevalence by wealth quintile, % of population

Wealth Quintile	IRS <sup>1</sup>	ITN <sup>2</sup>	Antimalarial Medication <sup>3</sup>	Malaria Prevalence <sup>4</sup>	HIV Prevalence (ages 15-49)			Diabetes Prevalence <sup>5</sup>		
					Women	Men	Total	Women	Men	Total
<b>Lowest</b>	29.2	52.2	41.4	42.1	1.2	0.5	0.9	3.66	0.61	2.24
<b>Second</b>	8.7	53.6	46.4	39.5	3.1	1.8	2.5	1.7	2.27	1.97
<b>Middle</b>	8.1	43.3	54.6	24.6	3.2	1.7	2.5	3.11	1.88	2.49
<b>Fourth</b>	4.7	32.9	51.4	13.9	4	1	2.5	7.03	3.73	5.24
<b>Highest</b>	5.5	29.5	52.8	7.5	2.5	0.8	1.7	7.82	6.81	7.2
<b>Total</b>	9.7	43	48.5	26.7	2.8	1.1	2	4.664	3.06	3.828

Source: Ghana DHS 2014 Final Report and Gatimu et al. (2016)

Notes: [1] Percentage of households that received indoor residual spraying in last 12 months

[2] Percentage of children < 5 who slept under an LLITN last night

[3] % who took any antimalarial medication among children under 5 with fever in the last two weeks

[4] Malaria prevalence using microscopy measurement.

[5] Diabetes prevalence for adults > 50 by wealth quintile from Table 2, Gatimu et al. (2016).

The last step in the analysis is to monetize the mortality reductions we have distributed across the income quintiles. We do this using the estimated probabilities of treatment from Table 4 and the monetization formula from Jamison et al. (2013) defined above. Table 5 summarizes the result of applying this valuation method to the estimates of mortality averted through government health expenditure. To put the values in perspective, average expenditure per capita in the GLSS data was 2261 cedis per year in 2012 and the official poverty line is 1314 cedis per adult equivalent. Clearly,

the estimated value of mortality reduction due to malaria drugs distribution is substantial. But the other malaria interventions and antiretroviral drugs also have high rates of return.<sup>32</sup>

Table 5 – Distributional Statistics for the Value of Mortality Reduction Due to Specific Health Interventions in Ghana

Intervention	Mean	25th	Median	75th	min	max	c.c.	Gini – Marginal Effect
Malaria Drugs	1890	845	1217	3131	0	4430	-0.020	0.117
Indoor Residual Spraying	23	9	15	27	0	97	-0.372	0.008
Treated Bed Nets	175	71	125	286	0	417	-0.175	0.039
Antiretroviral Drugs	41	0	7	58	-6	393	0.106	0.005
Diabetes	1	0	0	2	0	16	0.359	0.000

Source: GLSS-6 and authors' calculations

Notes: Values in 2012 Ghanaian cedis per year.

Statistics in the first six columns are for the distribution of benefits, not income.

The concentration coefficients in Table 5 show distributions of benefits consistent with the assumptions based on Table 4. The benefits of malaria drugs are spread evenly across the population.<sup>33</sup> Indoor spraying and bed nets are both highly progressive. To put these in perspective, the most progressive expenditures in the original CEQ assessment for Ghana are a targeted school feeding program (concentration coefficient of -0.40) and the Livelihood Empowerment Against Poverty (LEAP) conditional cash transfer scheme (concentration coefficient of -0.29). The benefits of antiretroviral drugs and diabetes control go more to richer people than the poor, but neither is as concentrated as income itself.

The marginal contributions of the malaria interventions to inequality reduction is also huge. Those for indoor spraying, are comparable to the largest marginal effects for any of the budget items considered in the original CEQ assessment. Those for bed nets and malaria drugs dwarf other line items. In fact, the marginal effect for the entirety of government taxes and social spending in the CEQ assessment is 0.025 for the Gini (Younger, Osei-Assibey, and Oppong, 2017). The marginal effects for antiretroviral drugs look smaller, but they are in fact similar to many of the marginal effects for social expenditures in the original CEQ assessment. Only diabetes control has no perceptible effect on inequality. It is important to remember that what drives these results is the extraordinary rates of return from averted mortality through the malaria interventions, the validity of which we consider in the next section.

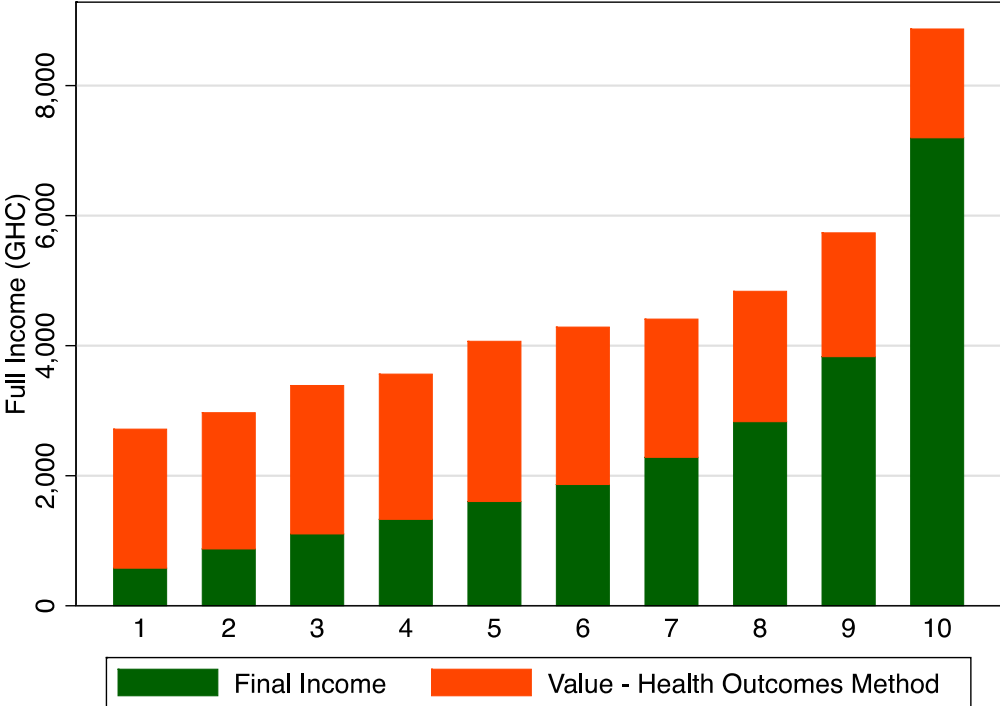
We do not explicitly estimate the change in headcount poverty rates when using the health outcomes approach to value in-kind health spending. We refrain from making this estimate because the monetized value of reduced mortality risk is sufficiently large to represent a substantial

<sup>32</sup> We do not have information on the costs of these programs, but given that public health care expenditure was less than two percent of GDP in 2012, the costs of these programs must be far less than the benefits calculated here.

<sup>33</sup> This might seem at odds with Table 4 which shows somewhat higher use in richer quintiles. But the largest mortality gains are for children, and children tend to be concentrated in the poorer quintiles, equalizing the effect.

proportion of overall income. To make an accurate estimate of changes to poverty headcount rates would require re-defining the poverty threshold inclusive of health value. Given uncertainty about what health need should be, we instead show the combined monetary income and the monetized value of health benefits by decile (figure 1). We observe that relative value of health benefits is more than three times income at the lowest decile and about two times income at the third decile. Monetary income and the monetized value of health become equal only at the seventh income decile. This relative distribution of monetized value from health versus income drives the significant inequality reductions we observe in table 5 when using the health outcomes approach.

Figure 1: Final Income and Value of Health using the Health Outcomes Method by Income Decile



### 4.2 Discussion of the Health Outcomes Approach

Even though the estimated benefits in Table 5 are enormous relative to other tax and spending programs,<sup>34</sup> we must start our discussion by noting that these are for just the few health interventions for which we can calculate effects on mortality probabilities. Although a comprehensive assessment of health interventions using the health outcomes approach would yield larger estimates of their value, this estimate includes both the first and second largest causes of

<sup>34</sup> In principle, the mortality benefits of non-health policies should also be added to a benefit-incidence analysis using the health outcomes approach. However, given the scope of this Chapter is focused on health and the method is applied to a high-mortality, weak social safety net nation such as Ghana, we expect most health benefits to flow from health sector spending.

premature death in Ghana in 2012 (malaria and HIV respectively, IHME 2017).<sup>35</sup> In addition, Ghana has experienced rapid progress in reducing the burden of both diseases as malaria- and HIV-related mortality have declined by 28% and 43.5%, respectively between 2005-2016. This means that the current analysis already includes a substantial proportion of Ghana's government-provided mortality gains. Nevertheless, given the overwhelming size of the estimated monetary effects, health spending measured with the health outcomes approach has a much larger effect on the income distribution than any other public expenditure or tax in Ghana. Malaria drugs alone reduce the Gini coefficient by 11 percentage points. So, while Jamison et al. (2013)'s argument in favor of greater health spending is based solely on the high benefit/cost ratios, it is also clear that there is a strong distributional argument in favor of increased spending on certain types of health, particularly those that benefit rural areas and children. Indeed, it seems likely that the most equalizing thing a country could do is to put in place Jamison et al. (2013)'s basic health care package.<sup>36</sup>

Anyone's first reaction to these results, including our own, is to question whether the estimates are too high. It is unlikely that the change in mortality estimates are far off as there is an abundance of epidemiological modelling and demographic and medical data with which to estimate the effects of specific health interventions, and those models and data underpin the Spectrum models we use to estimate mortality reductions. One limitation is that the definition of the counterfactual minimum level of intervention coverage, if shifted, would also affect mortality change estimates. However, adjustments are unlikely to produce substantively different mortality impacts. And indeed, the way in which we have estimated changes in mortality risk may well be biased downward as even in the comparator country with the worst coverage for any specific intervention, there is probably still *some* public spending on that intervention while the appropriate counterfactual for a CEQ assessment should be zero public spending.

The monetary value of mortality reduction, though, is a normative question and so open to more debate. Almost all studies to estimate the value of SMUs are done in developed countries. While estimates of the value of reduced mortality vary widely, they do not vary by more than an order of magnitude. A summary report from the OECD (2011) that provides practical guidelines for valuing reduced mortality from health, environmental, and transport policies in high income nations notes that the U.S. Office of Management and Budget (OMB) recommends a range of US\$100 to US\$1000 per reduced SMU for all government agencies.<sup>37</sup> The U.S. Environmental Protection Agency (EPA) has produced the most cost-benefit analyses of policies that affect mortality and uses a central estimate of 2007 US\$750 per SMU. Other agencies in the U.S. government use central estimates that vary between US\$500-US\$680.<sup>38</sup> Given lower incomes in the European Union (EU) and the OECD overall, the OECD report also recommends a central SMU value of

---

<sup>35</sup> Premature death refers to the total number of years of life lost (YLLs) per cause of death where YLLs are defined as the average number of years lost for a given death compared to what that person would have lived given Ghana's average life expectancy.

<sup>36</sup> An explicit assessment of the equalizing effect of introducing this basic health care package is, however, beyond the scope of the current paper.

<sup>37</sup> Throughout this paper, we refer to the value of reduced mortality using the value of a micro-mortality risk of SMU instead of the more commonly used VSL. We translate guidance from the OECD (2011) in VSL terms to SMUs. Additional background on valuing change in mortality risk can be found in appendix section A.2.

<sup>38</sup> These agencies include the Department of Transportation, Food and Drug Administration, and the Department of Homeland Security.

2005 US\$350 for the EU and 2005 US\$290 for an analysis applied across the OECD. To compare these values with our analysis, valuing an SMU at 1.8% of GDP means that each SMU is equal to 2007 US\$865.<sup>39</sup> The value we employ therefore is larger than the EPA's central estimate, but well within the U.S. government's most frequently used value range for regulatory decision-making.

While ethically uncomfortable, it seems necessary for public policy decisions to adjust the value of reduced mortality risk by national income because of government budget constraints. Poor countries simply cannot spend money as if mortality reduction were as valuable as the rich countries estimates given above. The challenge is how to make the adjustment across countries. Several authors have used cross-country regressions of SMU estimates on GDP per capita to estimate an elasticity that can then be used to predict the value of an SMU in countries for which there are no estimates. We have used the estimate of Hammitt and Robinson (2011) of 1.8 percent of annual GDP per capita for a mortality risk reduction of 0.01 percentage points at age thirty-five.<sup>40</sup> Projecting that constant ratio from the income range of the countries for which we have studies that value reduced mortality risk to poor countries like Ghana can be seen as problematic. But it is consistent with the recommendations and practice in both Jamison et al. (2013) and the OECD (2011) study, both of which apply a unit income elasticity for the value of mortality reduction to adjust across countries.

The other option would be to use estimates of SMU values from studies done in the developing world, but there are very few of those and they produce estimates that vary by orders of magnitude from developed country estimates. One study looks at options for transport to the airport in Freetown, Sierra Leone and finds values of reduced mortality only slightly lower than those from developed countries. But international airport users are on the wealthier end of the income distribution. Another study of willingness and ability to pay for improved water quality among the rural poor in Kenya finds an SMU value of \$US0.077 from reduced micro-risk of child death from diarrheal disease. SMU estimates are calculated by comparing the health benefit of cleaner water against the opportunity cost of walking longer distances to obtain it. This estimate is four orders of magnitude smaller than developed country estimates, and the authors provide several reasons why those estimates are likely to be unreasonably low (Kremer et al. 2011).

Given this range of mortality risk values, we investigate how much we would need to reduce an SMU's value from 1.8% of GDP per capita to make aggregate benefits from the health outcomes approach equaled those from the average cost method. This calculation takes as given the estimated change in mortality risk from government health spending in Ghana and explores what SMU value would equalize total benefits using the health outcomes and average cost methods. Performing this analysis, we find that an SMU's value would need to be divided by 40 for the aggregate values of each method to be equal. In other words, the VSMU would need to be reduced to 0.045% of GDP per capita instead of 1.8% (2007 US\$21.6) for the value of health spending to be equal under the health outcomes and average cost methods. A reduction in the value of averted mortality that large – almost 1.5 orders of magnitude – is outside of the plausible range, given valuations in the developed world. Put another way, an SMU of less than US\$22 is small enough to be rejected on normative grounds. Thus, even though we acknowledge that uncertainty exists in what exact SMU

---

<sup>39</sup> Given U.S. GDP per capita of \$48,061 in 2007 current dollars.

<sup>40</sup> Other estimates are lower – Miller (2000) for example estimated a relation of 1.2 – but again, not two or three orders of magnitude lower.

should be used, our qualitative results remain the same, unless we reduce the SMU to a level too low to be credible.

Based on this robustness check and the magnitude of our results, it should be clear that even if we used the lowest value of reduced mortality risk in the OECD report and adjusted for Ghana's income, our qualitative results would not substantively change. The minimum SMU value mentioned in the OECD (2011) summary across all high-income nations is US\$100, the low end of OMB guidance. Although, to our knowledge, no government agency has used an SMU value that low for policy analysis, if we had used that value instead of US\$865, our estimates of the value of public spending for health would be reduced by almost 90%. However, even with that value, the three causes of death we analyze would still generate larger reductions in inequality than any other area of government expenditure. In addition, the provision of malaria medication alone would still represent the largest single contribution to reduced inequality among the budget items analyzed in CEQ's previous analysis in Ghana (Younger et al. 2016).

One weakness of the health outcomes approach as applied here is that it addresses only individual health interventions, those for which we have readily available estimates of mortality effects. So even though the estimated benefits are sometimes quite large and represent the two largest causes of premature mortality in Ghana, they are partial. In a country where we do not have estimates for the link between publicly-funded health spending and the main causes of mortality, this will be a more important limitation.

There is one further conceptual question to consider in this discussion. If government spends, say, a dollar to distribute an artemisinin tablet to a malaria patient and that tablet reduces mortality by an amount we estimate to be worth \$1000, should we count the benefit that government has transferred to the patient as \$1 or \$1000? We have already mentioned that standard national income accounting would value this at its \$1 cost, which is also the standard average cost approach of most incidence analyses. To justify the \$1000, we must argue that we should treat health benefits differently because more health is both utility-enhancing on its own and, crucially, allows the purchase of additional utility in the future.

Another consideration is that in the absence of government provision, patients could presumably buy artemisinin on their own for about \$1. If government simply transferred cash of \$1 and the patient used that dollar to buy artemisinin, we would not give government credit for the spectacular rate of return to artemisinin in an incidence analysis; we would count it as a dollar transferred. The question, then, is if government provides the good or service in-kind rather than the cash one could use to buy it, should we "credit" the government with the rate of return in an incidence analysis that aims to understand the overall distributional impact of taxation and spending? We argue that the value provided by government spending is actually two-fold: 1) the monetary value of anti-malarial treatment itself and 2) the value of facilitating access to this technology when needed among patients without the information or education to do so otherwise. Finally, we note that this question would not pertain to genuine public goods provided by government because individual consumers could not buy those on their own in private markets. The case for using the health outcomes method is therefore stronger when incomplete provision of health-related public goods is a significant driver of national mortality.

In the end, the choice between the health outcomes approach and the other methods to valuing in-kind health services depends on a difficult normative question. We are persuaded that we should treat health differently, certainly for public goods, but probably also for most publicly funded health services. But we recognize that not everyone will be convinced. What our example in Ghana shows without doubt, though, is that where one comes out on this question has huge consequences for an incidence analysis. If we value health services at cost, they will have positive, but modest distributional effects. If we value them based on reduced mortality, they may overwhelm the other line items in the budget and add a strong distributional argument in favor of universal provision of basic health services.

## **5. Summary: Choosing Among the Options**

We now turn to providing guidance on how and when each of these methods can be applied beyond the Ghana context we investigate here. Each of the three options presented has strengths and weaknesses and will be the best options depending on a nation's health system and the questions a particular analysis intends to answer. To judge which option is most appropriate in a given country, we describe the positive and negative attributes of each method along five dimensions: 1) conceptual validity, 2) comprehensiveness in health budget coverage, 3) ability to address the health budget in detail, 4) data requirements, and 5) ease of use.

### **5.1 Conceptual validity**

The motivation for this chapter is that the conceptual validity of the average cost of provision is weak. There is no reason to believe that what government spends on health services is anywhere close to the value of those services to beneficiaries. A particular concern is that government may spend money corruptly or incompetently so that the average cost of provision overstates the benefits to recipients of any health care provided. But it is also true that for some health care spending, the value in terms of life and health far outweighs what government spends to provide it.

Nevertheless, the average cost of provision does have the advantage of precedence, which includes comparability with the many existing studies on health spending incidence. In addition, the average cost of provision is consistent with the way national income accounting treats all government spending: in practice, the value of anything purchased by government is measured by its cost.

Another subtle way in which the average cost of provision approach is consistent with both national income accounting and most incidence analyses is that it ignores consumer surplus when valuing consumption of health care services. The other two methods we present explicitly try to capture that surplus. In national income accounting, all units of a good and service purchased are valued at the market price, so the total "value" is the marginal benefit of the last unit purchased multiplied by the quantity. In most incidence analyses, if government provides a free or subsidized good, we value that benefit at the quantity times the market price, just as in national income accounting. Both WATP and the health outcomes approach, on the other hand, estimate something closer to compensating variations, the integral under the demand curve. Given the possibility for extraordinarily large compensating variations for some health care services, this may well be

appropriate. But using that estimate differs from existing accounting systems, which one might take as a conceptual advantage for the average cost of provision method.

Because the WATP approach relies on revealed preference it is considerably more attractive conceptually than the average cost of provision. Using people's own decisions in real circumstances to infer the value to them of the care they are buying is a natural approach for economists. But this approach assumes that people are rational consumers and in health care there is significant evidence that individual behavior departs from optimality, especially for preventive care because there is not always an obvious (to the consumer) cause-and-effect relation between a health care choice and a health outcome. In addition, preventive care is an intertemporal and probabilistic decision: cost today versus expected future benefit. Experimental economists have shown that many people do not make utility maximizing probabilistic or intertemporal choices. So, we might expect the WATP approach to significantly undervalue preventive care. Therefore, we recommend using this approach for curative care where its benefit is both immediate and obvious. Many surveys that form the basis of incidence studies provide such health care data – use of outpatient consultations and inpatient hospital stays. WATP is applicable in most practical circumstances to value these types of care despite the conceptual limitation of revealed preference.

Of course, WATP requires the ability to observe demand for publicly-funded health services. This is not possible for non-rivalrous goods such as public goods and services provided by natural monopolies. Important aspects of health-related spending (not necessarily in the health sector budget) have these characteristics: provision of information about health and vector control are public goods; water and sanitation systems are natural monopolies. A revealed preference approach will not work for such services.

The conceptual validity of the health outcomes approach depends on whether we can put a monetary value on reductions in mortality probabilities.<sup>41</sup> If we can, then this approach responds well to the limitations of the other two approaches. It can capture the fact that health care spending may provide benefits that are much larger than their cost of provision, but if government provides services incompetently, this approach will capture that in the consequent lack of mortality reduction. Further, this approach does not rely on rational choice in health care demand. And because it does not rely on revealed preference for health care services, it is applicable to non-rivalrous goods and services.

While putting a monetary value on mortality is uncomfortable for some, it is a fact that governments must continuously make resource allocation decisions about health policies that affect mortality risk. Some of these choices are based on an explicit and public value for reduced mortality. And indeed, governments across the developed world use cost-benefit analysis as one input into the broader policy decision-making process. However, deciding not to use an explicit mortality valuation only means that health resources are allocated in a potentially ad-hoc way. One policy may implicitly value mortality reduction more than another (by, for example, spending more to achieve a similar result). Since these decisions are ultimately political, avoiding an explicit valuation increases the likelihood that spending to reduce mortality among the better off and more powerful will be valued more than others.

---

<sup>41</sup> The same question applies to morbidity, if that were to be included.

There are three conceptual limitations to the health outcomes approach. The first objects to the use of consumer surplus from reduced mortality in an accounting exercise, as discussed above. The second challenges the possibility of finding a reasonable estimate of the monetary value of reduced mortality probability and/or the ethics of using such an estimate. And the third is the \$1/\$1000 question: if government provides a service that costs \$1 and reduces mortality probability by an amount worth \$1000, should we use \$1 or \$1000 in the accounting of government-provided benefits if the beneficiary could have purchased that service herself? In any other aspect of incidence analysis, we would use \$1. Each of these criticisms are addressed above in section 4. In low-income nations with a large rural population and incomplete coverage of health-related public goods the justification for using the health outcomes approach is particularly strong. This is recommended for two reasons: 1) neither of the other methods are able to include the benefits from public goods and 2) the justification for using consumer surplus from reduced mortality is strongest when income and informational constraints limit citizens' ability to obtain needed health care on their own.

## **5.2 Comprehensiveness**

CEQ assessments aim to include as much of the budget as possible, so a comprehensive treatment of health expenditures is important. Each of the three options has the potential to be comprehensive, though in practice, data availability limits each one. Average cost of provision can certainly be comprehensive if the survey questionnaire asks about use of all types of publicly-funded health care services. In practice, survey questions are mostly limited to consultations and hospitalizations, even though these represent the bulk of health care expenditures. The "insurance value" approach to average cost of provision avoids this problem by assigning benefits to all eligible beneficiaries whether they actually use publicly-funded health services or not, so it is completely comprehensive.

The WATP approach is limited to private goods and services which, again in practice, tend to be consultations and hospitalizations. The health care outcomes approach, as we have applied it, is limited to a few specific health care services we can easily link to mortality reductions with epidemiological models. But, in principle, this approach could also compare overall mortality across time and countries with similar disease vulnerabilities to estimate the reduced risk of death associated with all publicly-funded health care.

## **5.3 Detail**

In principle, all three approaches can be quite detailed. We could calculate the average cost of provision for very precisely defined medical services and assign the benefits to users of those services. But the demand for administrative data on the cost of a very large number and type of services would be daunting, something most ministries of health could not provide. And the demands on survey data would be similarly daunting. Sample sizes would need to be extremely large to have reliable samples of very specific health care services and in any event these data are not collected in standard income and expenditure or health surveys. The same problems will affect the WATP approach. While in principle we could estimate the demand for very precisely defined health care services, in practice, surveys will not collect sufficient information to do so reliably. We have seen that the health outcomes approach can be quite detailed. But in addition to the same

demand for information about who actually uses publicly-funded health care as the other two methods, the health outcomes approach also needs epidemiological models linking a specific service to mortality reductions. These models remain limited and focus, understandably, on infectious diseases.

#### **5.4 Data Requirements**

As mentioned above, to be both comprehensive and detailed, each method would require data that is not routinely available. All three methods require information on who uses publicly-funded health care services.<sup>42</sup> All three will generally suffer from lack of detail in this regard. The average cost of provision approach also requires information on the administrative cost of providing publicly-funded health care services and will suffer from lack of detail here, too. The health outcomes approach also needs information linking health care spending to mortality probabilities, something that is not available for all countries nor for all health care services.

In preparing this chapter we chose to limit ourselves to methods that can be applied with a single cross-sectional household survey since most CEQ studies accept that limitation. That survey clearly must ask about respondents' use of health care services or their affiliation with health insurance schemes (for the "insurance value" approach) in as much detail as possible. Beyond that, each method needs additional data. The average cost of provision approach needs administrative information on the cost of providing health care services, at a level of detail consistent with the survey information. The WATP approach requires that the survey also include adequate information on factors that influence demand for publicly-funded health care services and its substitutes, including measures of quality and the opportunity cost of using those services. This includes price, but also travel and waiting times. The health outcomes approach needs data that link specific health care services to reductions in mortality probabilities, which usually come from epidemiological models.

Which of these requirements is the least restrictive? Given the large number of existing studies of health care incidence, getting the administrative information for the average cost of provision seems to be generally feasible. Surveys certainly can ask about health care pricing and quality, but many do not, making the WATP approach less generally applicable. Estimates of the mortality reduction from health care services would seem to be the most difficult to find, but the Spectrum models are very helpful in this regard and there is a wealth of medical and epidemiological research that remains to be tapped for such information.

#### **5.5 Ease of Use**

Assuming that one has the data required, the average cost of provision approach is clearly the easiest to apply. Estimating WATP models is more difficult. While such models are quite common, our experience is that they require significant econometric expertise to both specify such that they converge to coefficient estimates consistent with their underlying behavioral models. The health outcomes approach also requires epidemiological knowledge about the drivers of disease burden in the nation under study. The Spectrum System of Policy Models was created to allow policy

---

<sup>42</sup> The "insurance value" approach is again an exception. It only requires information on who is covered by the insurance of interest.

makers to project how health system change would impact mortality without extensive training. For this reason, the health outcomes approach can be applied without significant specialized statistical knowledge as long as extensive health system background is also available.

## 5.6 Advice

We cannot make a blanket recommendation based on our experience to date. Clearly much depends on data availability, but even ignoring those limits, there are strong conceptual arguments for and against each method. Given this uncertainty, it probably makes sense to take the average cost of provision approach as a default if for no other reason than precedent. What is the strongest reason to override that default? For health care systems where government spends money corruptly or incompetently on mostly curative care services, there is a strong case to be made for switching to the WATP model. Because consumers will put little value on low-quality services, this approach will capture the lack of benefit from misspent funds.

The strongest case for the health outcomes approach occurs in nations where non-rivalrous health services are not provided universally and therefore generate a substantial impact on mortality. WATP cannot handle these, and the cost of provision may grossly underestimate the value of these services to beneficiaries. Given our stunning results in Ghana, we feel that any study that has access to reliable evidence on the impact of any health care expenditure on mortality probabilities should explore the health outcomes approach. Not doing so risks missing what may be by far the largest benefit government provides to its citizens.<sup>43</sup>

## 6. Insurance Value of Financial Risk Reduction

All public spending on health provides insurance to eligible beneficiaries. This is obvious in the case of social insurance schemes, but is equally true of generally provided health services as well. Health services paid through general revenue require taxes from all taxpayers and provide benefits to those who draw unfortunate outcomes by falling ill. Since most people are risk averse, this insurance has value to them over and above the cost of providing health services or the value of their health outcomes because it reduces the variance of their ex post income. As such, this approach identifies an additional value of public health spending that is ignored in benefit-incidence analysis and can be added to any of the previous methods.

We calculate the value of financial risk protection from health insurance in two steps. The first step calculates for each person a distribution of income after health care expenditures both with and without insurance. One way to do this is through matching methods that compare health expenditures for insured and uninsured people. In countries where only part of the population has access to publicly funded health care (as in a social insurance system limited to formal sector workers), we can estimate this difference by comparing the health spending of those inside and outside the system, usually using the same survey data from the *CEQ Assessment*. We show an example for Ghana in the next section. Another option is to examine changes in access to publicly funded health services over time to measure impacts on the distribution of health spending. This

---

<sup>43</sup> For practitioners interested in reproducing these estimates or applying them in a new context using this analysis as the starting point, we have posted all data and do-files used in the paper. Additional details can be found in appendix section A.5.

requires two separate surveys, one before and one after a significant change in health policy, but not necessarily a panel.

In either case, we first calculate the distribution of health expenditures for insured households compared to matched uninsured households using a quantile regression of out-of-pocket health expenditures on a dummy variable for insurance status at each percentile of the health expenditure distribution.<sup>44</sup> Second, we use a risk-averse utility function to evaluate household utility from reduced financial risk attributable to government health insurance. This approach has been used widely to estimate insurance value in the U.S. for Medicare (Finkelstein & McKnight 2008) and Medicare Part D (Engelhardt & Gruber 2010), Japan (Shigeoka, 2014), Thailand (Limwattananon et al., 2015), Ghana (Powell-Jackson et al. 2014), and Mexico (Barofsky 2015), although the distributional consequences have not been a focus of this research. We assume that households satisfy a per period budget constraint of  $c = y - m$  where  $y$  represents income,  $m$  household health spending,  $c$  non-health expenditure, and utility is determined under a constant relative risk aversion (CRRA) utility function:<sup>45</sup>

$$U(c) = \begin{cases} \left( \frac{c^{(1-\varepsilon)} - 1}{(1-\varepsilon)} \right) & \text{if } \varepsilon \neq 1 \\ \ln(c) & \text{otherwise} \end{cases}$$

Call the distributions of health spending calculated with the quantile regressions  $P_k(m)$  where  $k=[1,0]$  indexes those households with and without insurance, respectively. The difference between household income minus  $P_0(m)$  or  $P_1(m)$  determines the change in risk exposure from insurance. Household expected utility is:

$$EU(y, \gamma, P_k(m)) = \int_0^m u(\max[y - m, \gamma y]) P_k(m) dm$$

where  $\gamma$  represents an assumed minimum consumption value under which household expenditure does not fall irrespective of the cost of medical care. Previous studies set this limit between 20 and 40 percent of household expenditure in developed countries. But for a poor country that threshold seems far too low. It is unlikely that a poor person in Ghana could spend 60 percent of her household's income on health care. So, in addition to the standard assumption, we also impose a lower limit on expenditures equal to Ghana's extreme poverty line.

The risk premium represents the quantity of money a risk-averse household would be willing to pay to completely insure against a given financial risk distribution. The risk premium for a household is:

---

<sup>44</sup> Standard errors in the quantile regressions are clustered to adjust for correlation of outcomes within enumeration areas (Parente and Santos 2016).

<sup>45</sup> We vary the coefficient of relative risk aversion to check the sensitivity of the value of risk protection to this important parameter.

$$\pi_k = E_k(y - m) - CE_k] \\ = \left\{ \sum_{m=0}^m (\max[y - m, \gamma]) P_k(m) \right\} - \left\{ u^{-1} \left[ \sum_{m=0}^m u(\max[y - m, \gamma]) P_k(m) \right] \right\}$$

where  $E_k(y - m)$  represents the expected value of a household's non-health expenditure and  $CE_k$  is the household's certainty equivalent for the same distribution of health spending. The difference in risk premia between those with and without coverage,  $\pi_1 - \pi_0$ , represents the monetary value of financial risk protection provided by government health insurance.

### 6.1 Example Application – Financial Risk Reduction from Ghana's National Health Insurance Scheme

In Ghana, nearly all those with health insurance are covered through the NHIS. The NHIS began as separate district-based and mutual health insurance schemes and was rolled out nationally in 2004 (Duku et al. 2016). To increase access to care among the most vulnerable, statutorily the NHIS provides coverage without premiums to children under 18 years old, elderly aged 70 and above, pregnant women, and recipients of Ghana's conditional cash transfer program. In addition, formal sector workers pay into the system through payroll taxes, but are exempt from paying any premium. Given that over 90 percent of total health insurance coverage in 2012/13 GLSS is provided by the NHIS, this estimate is close to comprehensive for measuring financial risk protection from insurance coverage in Ghana. Figure A.1 shows health insurance coverage by income decile in Ghana in 2005 and 2012/13.

We match those without NHIS coverage to those covered using coarsened exact matching (Blackwell et al. 2009). This gives us a control group similar in a range of observable characteristics to those with coverage. Specifically, we match on the following variables: days sick in the last two weeks, days spent in the hospital in the last two weeks, the number of children under 5 and adults over the 70 in the household, and the household's disposable income.<sup>46</sup> Because we expect income to be highly predictive of health expenditures, we divide the sample into household income quartiles, matching households covered by NHIS to uninsured households separately for each one.<sup>47</sup> In all, then, we have  $99 \times 4 = 396$  regressions, one for each percentile of the health expenditure distribution and each quartile of the income distribution.

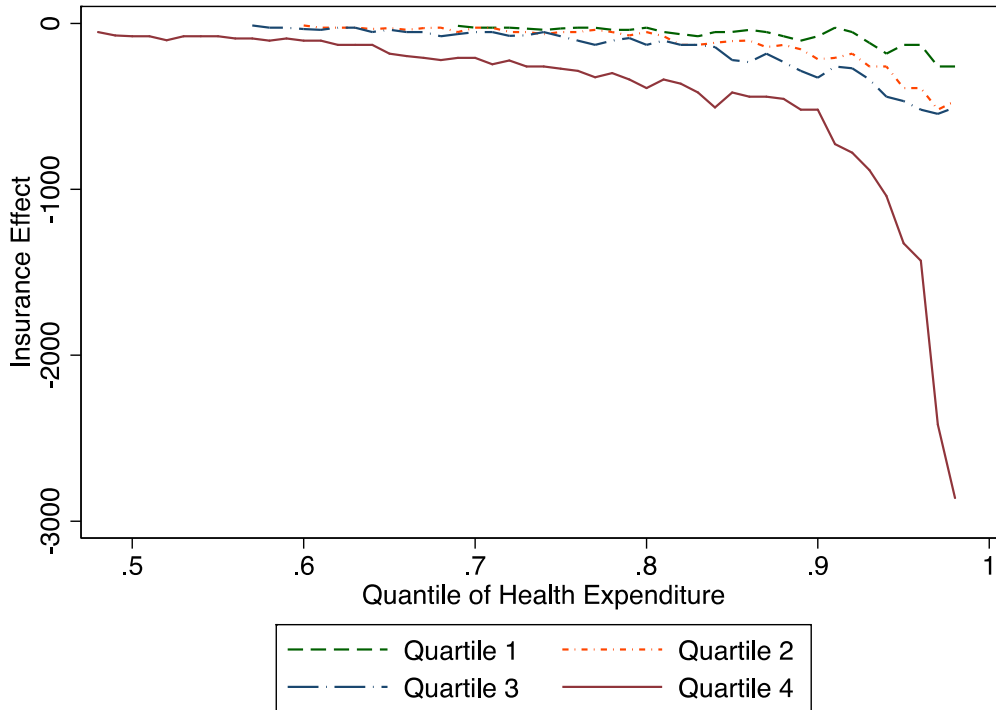
Figure summarizes the effect of NHIS insurance on out-of-pocket health expenditures for each income quartile using coarsened exact matching to create comparison groups without insurance. The horizontal axis is the quantile of health expenditure within each income quartile. Health spending is highly skewed within each quartile, so reductions in health spending from insurance coverage are small at lower centiles of health spending and rise rapidly above the 90<sup>th</sup> percentile, especially for the richest quartile. Even for the richest quartile, health spending in the 48<sup>th</sup> centile and below is zero, meaning that quantile treatment effects are mechanically zero below this level. For the highest spending groups, insurance provides increasingly important coverage against catastrophic expenditures due to health shocks. We should expect the benefits of NHIS financial risk protection to be skewed towards richer households even if the use of NHIS is not so skewed

<sup>46</sup> Disposable income for CEQ is household consumption, the standard welfare variable for GLSS.

<sup>47</sup> NHIS coverage is defined as a household in which more than half of residents have insurance.

because richer households would have been willing and able to spend significantly larger amounts on health care in the absence of NHIS insurance, making the risk reduction benefit of insurance correspondingly greater for those better-off households.

Figure 2 - Effect of NHIS insurance on medical expenditure by income quartile




**Note:** 'Insurance Effect' refers to the difference in annual out-of-pocket health spending for households with NHIS coverage compared to those without coverage (units are 2013 GHC).

Figure summarizes the value of financial risk protection from the NHIS by income quartile and by level of risk aversion. The consumption floor,  $\gamma$ , is set to Ghana's extreme (food) poverty line adjusted for household size. The figure shows that financial risk protection benefits are greatest for the richest quartile, ranging from 0.5% to 1% of quartile annual income across the most plausible estimates of risk aversion. For comparison purposes, the health benefits calculated using the standard CEQ average cost approach average 10%, 6%, 3.5%, and 1.8% of quartile mean income for the poorest to richest income quartiles, so the insurance values captured in the analysis are relatively small. The value of financial risk protection increases as risk aversion levels increase, as we would expect. But the largest difference is between the top income quartile and the rest. The value of risk protection as a percent of quartile income for the top income quartile is both more than double the lower three levels and also rises faster with increased risk aversion. So even though NHIS has targeting mechanisms meant to include poorer households, the largest benefits in terms of financial risk protection go to richer households because their health spending in the absence of insurance is greater. The insurance benefits estimated when using the extreme poverty line consumption floor are more concentrated among the top three income quartiles than income itself because the lowest quartile's insurance benefits are limited the most by this consumption floor assumption. Overall, however, insurance benefits with the extreme poverty consumption floor are

less concentrated than income, with a concentration coefficient of 0.30 compared to an income Gini coefficient of 0.41.

In addition to using the extreme poverty line, we also calculate financial risk protection benefits using a consumption floor proportional to 20% of household income. This assumes that all households, including the poor, will spend up to 80% of their income for health care, which is implausible in a setting such as Ghana. However, using this assumption allows us to calculate results that are comparable to other estimates of financial risk protection from insurance. Figure A.2 in the appendix shows the benefits by quartile and risk aversion as a percent of quartile income. The figure shows a similar pattern to those found with the extreme poverty line consumption floor, except benefits are larger for the poorest quartile. Consequently, the benefits are also relatively progressive compared to income with a concentration coefficient of 0.19.

Figure 3 - Insurance Value by Income Quartile and Risk Aversion



## 6.2 Discussion of the Financial Risk Reduction Methods

Because this approach does not need to be traded off against the others, the only consideration in using this method is whether the additional effort required to estimate households' counterfactual health spending is worthwhile. In a country like Ghana, where there is one primary health insurance system and plausible comparison groups can be constructed because insurance eligibility criteria include multiple groups across the income distribution, this approach is feasible. It would be more challenging for a country in which everyone is insured or substantial differences exist between insured and uninsured groups (eg: formal versus informal workers). Another

consideration is whether the size of the effect is relevant. In Ghana, it is. While not as large as the estimated health benefits from the standard average cost method, the financial risk reduction benefits are of the same order of magnitude and so merit consideration.

## 7. Conclusion

This paper compares three methods for estimating the benefits and incidence of health spending, while also defining how to measure the incidence of financial risk protection from government-provided health insurance. The average cost approach is the most common way to measure the benefits and incidence of government health spending. It is useful because of its ease of calculation and clarity. Most developing nations have sufficiently detailed nationally-representative surveys and national health spending accounts data such that the average cost of health care can be calculated widely and compared across nations and over time. In addition, these methods are accessible to researchers and government officials without extensive training in econometrics. But the government's cost of providing a service may have little to do with its actual value to beneficiaries. Governments may spend more than a service is worth if its provision is corrupt, inefficient, or misguided. At the same time, the service may be worth much more than what government spends if it has an appreciable effect on beneficiary health, something that has important and large intrinsic value. Each of these possibilities lead us to consider other options for valuing in-kind health services.

The first, the WATP approach, is conceptually attractive for an economist because it relies on revealed preference. Especially in cases where we suspect that government is spending far more than a service is worth, low willingness-and-ability-to-pay for those services will flag the problem and perhaps give a better estimate of the service's value.

The method does have limitations. Conceptually, we can only apply it if we can estimate a demand function for in-kind services, which rules out any services that are public goods. When using income and expenditure surveys as CEQ does, that also excludes most types of preventive care. Practically, the demand estimation requires considerably more work than the average cost method. In our particular example of Ghana, that extra work does not seem worthwhile as the results for the two methods are not too different. But that is just one example. We recommend the method's use for situations in which one suspects large differences between government costs and patient benefits due to corruption, inefficiency, or incompetence and especially if the analysis is most interested in valuing curative care only.

The second method we consider, health outcomes, takes a much different approach. It estimates the effect of government-funded health interventions on mortality probabilities and then monetizes those changes based on studies that look at how much money people will accept (pay) in exchange for small increases (reductions) in the probability of death.<sup>48</sup> This approach also estimates benefits rather than costs to government, but without relying on observed willingness-and-ability-to-pay for care. There are good reasons to be suspicious of WATP estimates for health care, especially preventive health care. People may have difficulty understanding the effect of health care services

---

<sup>48</sup> These valuations are, to be clear, revealed preference willingness and ability to pay for changes in mortality risk (instead of health care services) among adults in the developed world.

on their mortality and so under-value them considerably. It is also true that the severe liquidity constraints common in developing countries cause people to shun expenditures that offer future, uncertain benefits when their focus is survival for themselves and their children today.

Unlike the willingness and ability to pay for health care we calculate, our estimates for health benefits are sometimes radically different from cost of provision estimates. Indeed, the benefits to malaria control and treatment are so large that they swamp the distributional consequences of any other budget line or combination of budget lines, and this despite the fact that we are able to analyze only a few publicly funded health interventions (albeit the two largest causes of premature death). If this is the right approach, it is of substantial importance for distributional analysis. Unfortunately, whether this is the right approach depends on a normative argument about how to value changes in mortality probabilities in developing countries. However, for our estimates of the value of malaria medication to be similar to those for the average cost method, we would need to reduce the value we place on reduced mortality by about two orders of magnitude: an implausibly low valuation. For bednets provision on its own, we would need to reduce the value of mortality by about one order of magnitude. Clearly, our normative judgment (and most estimates of the value of mortality risk reduction) would need to be far off the mark for the health outcomes approach to be irrelevant.

There is a conceptual challenge to this approach: it is not consistent with how national income accounting treats publicly-funded services. To accept the health outcomes approach, we must argue that health is different because of its intrinsic importance and so should be treated differently in our accounting.

In practical terms, the health outcomes approach is limited to instances for which we have ready access to estimates of the effect of publicly-funded health interventions. The Spectrum models are very helpful in this regard, can be run for all developing nations, and include many but not all causes of death. Even though this approach requires more effort than the average cost approach, its results are so dramatic that the effort seems well worthwhile.

Lastly, we estimate the insurance value of reduced financial risk that comes from government health spending. This is an add-on to the other approaches discussed here because it is not valuing health services *per se*, but rather financial risk. The method does require an arbitrary assumption about risk aversion, but for a wide range of such assumptions, we find that the benefits are similar in value to average costs (and WATP). The effect of financial risk protection on the income distribution is somewhat less than average cost because its benefits tend to go more to richer households who spend more on health care in the absence of public funding.

Practically, our approach to financial risk requires us to match insured with uninsured households to compare their health expenditures. Ghana is a good case for such matching because the NHIS does not apply to everyone, but it is also not obviously correlated with income. If the matching is feasible, this method is also worthwhile in an incidence analysis.

## References

- Arrow KJ. Uncertainty and the welfare economics of medical care. *The American Economic Review* 1963;53:941-73.
- Aaron, H., & McGuire, M. (1970). Public goods and income distribution. *Econometrica: Journal of the Econometric Society*, 907-920.
- Acharya A, Masset Edoardo S, Taylor F, Massett E, Satija A, Burke M, Ebrahim S: The Impact of Health Insurance Schemes for the Informal Sector in Low-and Middle-Income Countries: A Systematic Review. Policy Research Working Paper. 2013, Washington, DC: World Bank
- Akerlof, G. (1995). The market for “lemons”: Quality uncertainty and the market mechanism. In *Essential Readings in Economics* (pp. 175-188). Macmillan Education UK.
- Alam, K. and Mahal A. (2014) “Economic impacts of health shocks on households in low and middle income countries: a review of the literature.” *Globalization and Health* (10)21.
- Ashenfelter, O., & Greenstone, M. (2004). Estimating the value of a statistical life: The importance of omitted variables and publication bias. *American Economic Review*, 94(2), 454-460.
- Avenir Health. 2018. Spectrum Manual, Spectrum System of Policy Models. In: Avenir Health, editor. *Spectrum Manual, Spectrum System of Policy Models*. Glastonbury, CT. <http://avenirhealth.org/Download/Spectrum/Manuals/SpectrumManualE.pdf>
- Baker, R., S. Chilton, M. Jones-Lee, and H. Metcalf. (2008), “Valuing Lives Equally: Defensible Premise or Unwarranted Compromise?” *Journal of Risk and Uncertainty*. 36, 125–138.
- Barofsky, J. (2015). Estimating the impact of health insurance in developing nations: Evidence from Mexico's Seguro Popular. Harvard School of Public Health, Cambridge, MA.
- Bastagli, F. (2015). Bringing taxation into social protection analysis and planning (Vol. 421). Overseas Development Institute Working Papers.
- Bernheim, Douglas B., (1987), “The economic effects of social security: toward a reconciliation of theory and measurement.” *Journal of Public Economics* 33, 273–304.
- Blackwell, M., Iacus, S. M., King, G., & Porro, G. (2009). cem: Coarsened exact matching in Stata. *The Stata Journal*, 9(4), 524-546.
- Bound, John, Cullen, Julie Berry, Nichols, Austin, Schmidt, Lucie (2004). “The welfare implications of increasing disability insurance benefit generosity.” *Journal of Public Economics* 88 (12), 2487–2514 (December).

Bourguignon, F., Da Silva, L. P., & Stern, N. (2004). Evaluating the poverty impact of economic policies. *Macroeconomic Policies and Poverty*, 36.

Braathen, N. A., H. Lindhjem and S. Navrud (2009), Valuing Lives Saved from Environmental, Transport and Health Policies: A Meta-analysis of Stated Preference Studies. OECD, Paris. See [www.oecd.org/env/policies/vsl](http://www.oecd.org/env/policies/vsl).

Cameron, A. C., & Trivedi, P. K. (2005). *Microeconometrics: methods and applications*. Cambridge university press.

Costa-i-Font, J., & Cowell, F. (2017). Does Insurance Expansion Alter Health Inequality and Mobility? Evidence from the Mexican Seguro Popular (No. 6788). CESifo Group Munich.

Cutler, D. M., & Richardson, E. (1998). The value of health: 1970-1990. *The American economic review*, 88(2), 97-100.

Dagsvik, J. K., & Karlström, A. (2005). Compensating variation and Hicksian choice probabilities in random utility models that are nonlinear in income. *The Review of Economic Studies*, 72(1), 57-76.

Das, J., & Hammer, J. (2005). Which doctor? Combining vignettes and item response to measure clinical competence. *Journal of Development Economics*, 78(2), 348-383.

Das, J., Holla, A., Das, V., Mohanan, M., Tabak, D., & Chan, B. (2012). In urban and rural India, a standardized patient study showed low levels of provider training and huge quality gaps. *Health affairs*, 31(12), 2774-2784.

De Ferranti, D. (1985). Paying for health services in developing countries: a call for realism. In *World Health Forum* (Vol. 6, No. 2, pp. 99-105).

Dmytraczenko, T., & Almeida, G. (Eds.). (2015). *Toward universal health coverage and equity in Latin America and the Caribbean: evidence from selected countries*. World Bank Publications.

Duku et al. Utilization of healthcare services and renewal of health insurance membership: evidence of adverse selection in Ghana. *Health Economics Review* (2016) 6:43.

Engelhardt, G. V., & Gruber, J. (2011). Medicare Part D and the financial protection of the elderly. *American Economic Journal: Economic Policy*, 3(4), 77-102.

Feldstein, Martin (1973). "The Welfare Loss of Excessive Health Insurance," *Journal of Political Economy* vol. 81, pp.251-280.

Finkelstein, A., & McKnight, R. (2008). What did Medicare do? The initial impact of Medicare on mortality and out of pocket medical spending. *Journal of public economics*, 92(7), 1644-1668.

Galárraga, O., Sosa-Rubí, S. G., Salinas-Rodríguez, A., & Sesma-Vázquez, S. (2010). Health insurance for the poor: impact on catastrophic and out-of-pocket health expenditures in Mexico. *The European Journal of Health Economics*, 11(5), 437-447.

Gauthier, B., & Wane, W. (2008). Leakage of public resources in the health sector: an empirical investigation of Chad. *Journal of African Economies*, 18(1), 52-83.

Gertler, P., & Gaag, J. V. D. (1990). *The willingness to pay for medical care: evidence from two developing countries*. Johns Hopkins University Press.

Gertler, P., Locay, L., & Sanderson, W. (1987). Are user fees regressive?: The welfare implications of health care financing proposals in Peru. *Journal of econometrics*, 36(1), 67-88.

Ghana Statistical Service (2018). *Ghana Population and Housing Census 2010*. <http://ghana.opendataforafrica.org/vbbbb/population-and-housing-census-of-ghana-2010>. Accessed 3-3-2018.

Global Burden of Disease Study 2015. *Global Burden of Disease Study 2015 (GBD 2015) Causes of Death and Nonfatal Causes Mapped to ICD Codes*. Seattle, United States: Institute for Health Metrics and Evaluation (IHME).

Greenstone, M., & Jack, B. K. (2015). Envirodevonomics: A research agenda for an emerging field. *Journal of Economic Literature*, 53(1), 5-42.

Hamilton, M., Mahiane, G., Werst, E., Sanders, R., Briët, O., Smith, T., ... & Gething, P. W. (2017). Spectrum-Malaria: a user-friendly projection tool for health impact assessment and strategic planning by malaria control programmes in sub-Saharan Africa. *Malaria journal*, 16(1), 68.

Hammit, J. K., & Robinson, L. A. (2011). The income elasticity of the value per statistical life: transferring estimates between high and low income populations. *Journal of Benefit-Cost Analysis*, 2(01), 1-29.

Herriges, J. A., & Kling, C. L. (1999). Nonlinear income effects in random utility models. *Review of Economics and Statistics*, 81(1), 62-72.

Higgins, S. (2015). *Can a Poverty-Reducing and Progressive Tax and Transfer System Hurt the Poor?* Working Paper 405 (No. 405). Center for Global Development.

Institute of Health Metrics and Evaluation (2017) Accessed 8/9/2017, <http://www.healthdata.org/ghana,%20IHME,%20>

Jamison, D. T., Summers, L. H., Alleyne, G., Arrow, K. J., Berkley, S., Binagwaho, A., ... & Ghosh, G. (2013). *Global health 2035: a world converging within a generation*. *The Lancet*, 382(9908), 1898-1955.

- Kremer, M., Leino, J., Miguel, E., & Zwane, A. P. (2011). Spring cleaning: Rural water impacts, valuation, and property rights institutions. *The Quarterly Journal of Economics*, 126(1), 145-205.
- Korenromp, E., Mahiané, G., Hamilton, M., Pretorius, C., Cibulskis, R., Lauer, J., ... & Briët, O. J. (2016). Malaria intervention scale-up in Africa: effectiveness predictions for health programme planning tools, based on dynamic transmission modelling. *Malaria journal*, 15(1), 417.
- Korenromp, E. L., Gobet, B., Fazito, E., Lara, J., Bollinger, L., & Stover, J. (2015). Impact and cost of the HIV/AIDS national strategic plan for Mozambique, 2015-2019—Projections with the Spectrum/Goals Model. *PloS one*, 10(11), e0142908.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2), 443-478.
- Lakin, J. M. (2010). The end of insurance? Mexico's Seguro Popular, 2001–2007. *Journal of Health Politics, Policy and Law*, 35(3), 313-352.
- Leon, G., & Miguel, E. (2017). Risky transportation choices and the value of a statistical life. *American Economic Journal: Applied Economics*, 9(1), 202-28.
- Limwattananon, S., Neelsen, S., O'Donnell, O., Prakongsai, P., Tangcharoensathien, V., Van Doorslaer, E., & Vongmongkol, V. (2015). Universal coverage with supply-side reform: The impact on medical expenditure risk and utilization in Thailand. *Journal of Public Economics*, 121, 79-94.
- Lindauer D. and Nunberg B., editors (1994) *Public Sector Wage and Employment Policies in Sub-Saharan Africa*, Washington, DC: The World Bank.
- Lustig, N. (Ed.). (2018). *Commitment to equity handbook: Estimating the impact of fiscal policy on inequality and poverty*. Brookings Institution Press.
- Lustig, N., & Higgins, S. (2013). *Commitment to Equity Assessment (CEQ): Estimating the Incidence of Social Spending, Subsidies, and Taxes-Handbook*. Subsidies, and Taxes-Handbook (September 1, 2013).
- Murphy, K. M., & Topel, R. H. (2006). The value of health and longevity. *Journal of political Economy*, 114(5), 871-904.
- Nguyen, H. T., Rajkotia, Y., & Wang, H. (2011). The financial protection effect of Ghana National Health Insurance Scheme: evidence from a study in two rural districts. *International Journal for Equity in Health*, 10(1), 4.
- O'Dea, C., & Preston, I. (2012). *The distributional impact of public spending in the UK*. IFS Working Paper 2012 (6), Institute for Fiscal Studies, London, UK.

O'Donnell, Owen, Eddy van Doorslaer, Adam Wagstaff and Magnus Lindelow. 2008. "Analyzing Health Equity Using Household Survey Data: A Guide to Techniques and Their Implementation." WBI Learning Resources Series. Washington, DC: The World Bank.

OECD (2011), "Valuing Mortality Risk Reductions in Regulatory Analysis of Environmental, Health and Transport Policies: Policy Implications", OECD, Paris, [www.oecd.org/env/policies/vsl](http://www.oecd.org/env/policies/vsl).

Parente, P.M.D.C. and Santos Silva, J.M.C. (2016), Quantile Regression with Clustered Data, *Journal of Econometric Methods*, 5, 1-15.

Powell-Jackson, T., Hanson, K., Whitty, C. J., & Ansah, E. K. (2014). Who benefits from free healthcare? Evidence from a randomized experiment in Ghana. *Journal of Development Economics*, 107, 305-319.

Spectrum-Malaria: a projection tool for health impact assessment and for malaria programs in sub-Saharan Africa strategic planning

Samuelson, P. (1969) "Pure Theory of Public Expenditure and Taxation," in *Public Economics; An Analysis of Public Production and Consumption and Their Relations to the Private Sectors*, J. Margolis and H. Guitton, eds., London and New York, 1969.

Shigeoka, H. (2014). The effect of patient cost sharing on utilization, health, and risk protection. *American Economic Review*, 104(7), 2152-84.

Small, K. A., & Rosen, H. S. (1981). Applied welfare economics with discrete choice models. *Econometrica: Journal of the Econometric Society*, 105-130.

Stover, J., Brown, T., & Marston, M. (2012). Updates to the Spectrum/Estimation and Projection Package (EPP) model to estimate HIV trends for adults and children. *Sex Transm Infect*, 88(Suppl 2), i11-i16.

Viscusi, Kip, and Joseph Aldy (2003). "The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World". *Journal of Risk and Uncertainty*, 27 (1): 5–76.

Wagner, Z., Barofsky, J., & Sood, N. (2015). PEPFAR funding associated with an increase in employment among males in ten sub-Saharan African countries. *Health Affairs*, 34(6), 946-953.

Wagstaff, A. (2009). Social health insurance vs. tax-financed health systems-evidence from the OECD. *World Bank Policy Research Working Paper*, (4821).

Wang, H., & Yu, Y. (2016). Increasing health inequality in China: An empirical study with ordinal data. *The Journal of Economic Inequality*, 14(1), 41-61.

Wodon, Q., & Moreno-Dodson, B. (Eds.). (2007). *Public finance for poverty reduction: concepts and case studies from Africa and Latin America*. World Bank Publications.

World Bank 2012. “Service Delivery Indicators Tanzania Technical Report.” <http://siteresources.worldbank.org/AFRICAEXT/Resources/SDI-Technical-Report-Tanzania.pdf>. Accessed 10/29/17.

World Health Organization. (2018). World Malaria Report 2017. Geneva: World Health Organization; 2017.

World Health Organization (2018). Global Health Observatory Data Repository. <http://apps.who.int/gho/data/?theme=main&vid=60630>. Accessed 3-6-18.

Zwane, A. P., & Kremer, M. (2007). What Works in Fighting Diarrheal Diseases in Developing Countries? A Critical Review. *The World Bank Research Observer*, 22 (1).

## **Appendix:**

### **Section A.1 Using the Spectrum Policy Models Software**

Spectrum is a system of policy models used to examine the impact of changes in health interventions for use by researchers and policymakers. Each projection starts with results from an underlying country-specific demographic model that projects population change using data on fertility, mortality, and migration rates. The demographic model comes prepopulated with country-specific data and population estimates from the United Nations Population Division. Built on these demographic projections, disease specific epidemiological models are used to estimate mortality from multiple causes of death including malaria, HIV, and a range of child, maternal health, and non-communicable diseases. In our application, the system of policy models is used to project how disease specific mortality rates will change across the age distribution when health intervention coverage rates change from their current level to their counterfactual minimum.

Each disease model combines our understanding of transition probabilities through different disease stages with measures of intervention efficacy using the scientific literature. Country specific data on underlying prevalence and disease type, environmental conditions that lead to transmission (for example, the entomological infection rate for malaria which measures exposure to infectious mosquitoes throughout the year) are also used. Intervention coverage data come from nationally representative surveys such as the Demographic and Health Surveys. Data sources are updated annually or as frequently as nationally representative surveys are conducted in a given country, while research literature is reviewed frequently as well to ensure efficacy parameters are up-to-date.

To implement projections for changes to HIV-, malaria-, and diabetes-related mortality, we estimated mortality rates first using all health intervention rates set at their Spectrum-provided current level. Modules used are DemProj, the base model that reflects population change, the AIDS Impact Model (AIM) which projects the consequences of the HIV epidemic including AIDS deaths by age and sex, and Malaria – a module which permits countries with endemic malaria to project over time malaria case load and malaria-attributable mortality, in 3 distinct age groups. The malaria transmission model underlying Spectrum’s estimates is based on OpenMalaria, which was developed by researchers at the Swiss Institute of Tropical Hygiene and Medicine and simulates the dynamics of malaria transmission and epidemiology in mosquito and human populations, and

the effects of malaria control. These statistical impact functions (described in Korenromp et al. 2016) are combined with a database of malaria endemicity and epidemiology at the subnational level to project future burden.

After mortality rates at current coverage are calculated, health intervention coverage levels are then changed one by one to reflect the counterfactual minimum coverage level and intervention-related age-specific mortality rates are again calculated. The difference in projected age-specific mortality rates when intervention coverage changes from current to counterfactual minimum levels is used as the change in mortality attributable to government spending. Since the Spectrum software only allows mortality rate projections in the future, intervention coverage levels are changed in year 2018 and the effect is taken from the first year after that change occurs – 2019. See Avenir (2018) for a detailed introduction on how to use the projection system and the contents of each module.

The counterfactual minimum level of health intervention coverage is estimated using DHS survey waves across other West African nations from 2011 to 2015 for malaria and diabetes interventions and UNAIDS for antiretroviral therapy coverage rates. The countries included are Benin, Togo, Ivory Coast, Guinea, Liberia, Sierra Leone, and Nigeria. For example, minimum reported use of insecticide treated bednets for children 5 years of age and below among Ghana's peer nations is 16.6% using the Nigeria 2013 DHS survey. This compares to a rate of 55% in Ghana. In addition, UNAIDS reports that the minimum ART coverage rate among these West African nations is Liberia in 2013 with a coverage rate of 13% compared to a rate of 65% in Ghana.

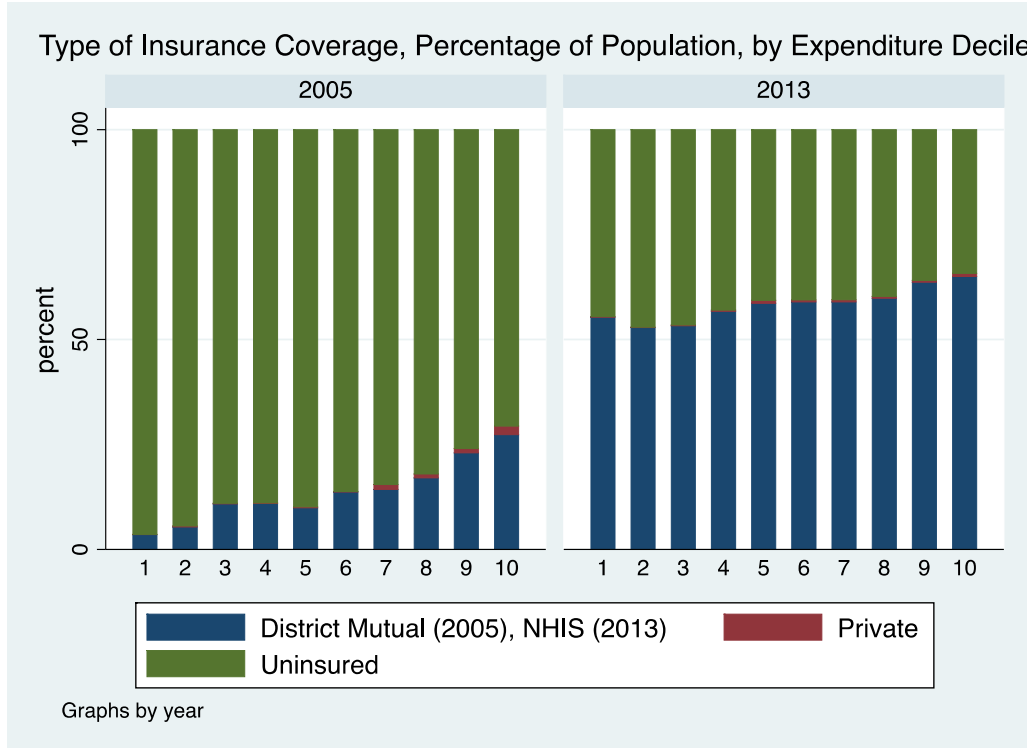
To finish the valuations of change in mortality, data on the distribution of the population in Ghana by age group and life expectancy by age is needed. The former is obtained from Ghana's Statistical Service (GSS) using the Ghana Population and Housing Census of 2010. To calculate total population in each age group in 2013, it is assumed that the distribution of population across age groups is constant from 2010 to 2013 and the GSS's total population projection for 2013 is used to estimate population in 2013 by age group (GSS 2018). Data on life expectancy in Ghana by five-year age group is obtained from the WHO's Global Health Observatory (WHO 2018).

## **Section A.2 Financial Risk Protection with Consumption Floor Proportional to Income**

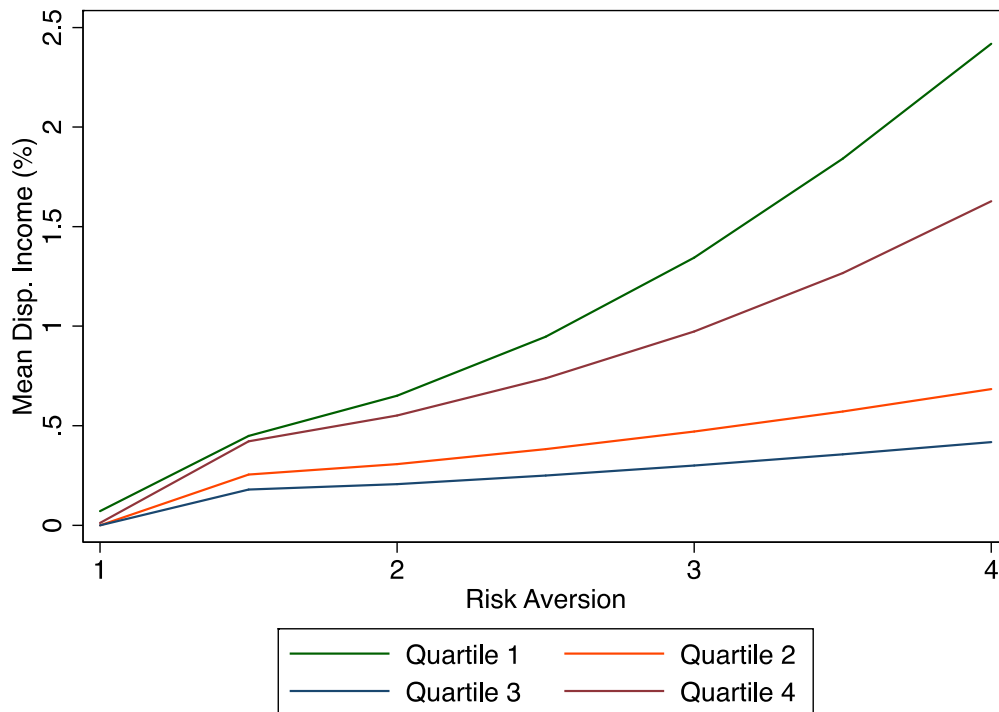
Figure A.1 displays the percentage of the population by income decile that is covered by the NHIS in both 2005 and 2013. Although in 2005 there was a clear income gradient for insurance access, by 2013 access was approximately equal across the income distribution.

Figure A.2 shows the values of insurance protection from health shock risk through the NHIS across income quartiles and levels of risk aversion, where the consumption floor is assumed proportional to household income. The results show, as expected, that the value of insurance increases with risk aversion. In contrast to using the extreme poverty line consumption floor, this figure shows that when a consumption floor proportional to income is used, the value of insurance is relatively progressive as benefits (as a percent of quartile mean disposable income) are largest for the lowest income group.

**Figure A.1: Insurance Coverage by Decile (Ghana) GLSS 2004/05 and 2012/13**



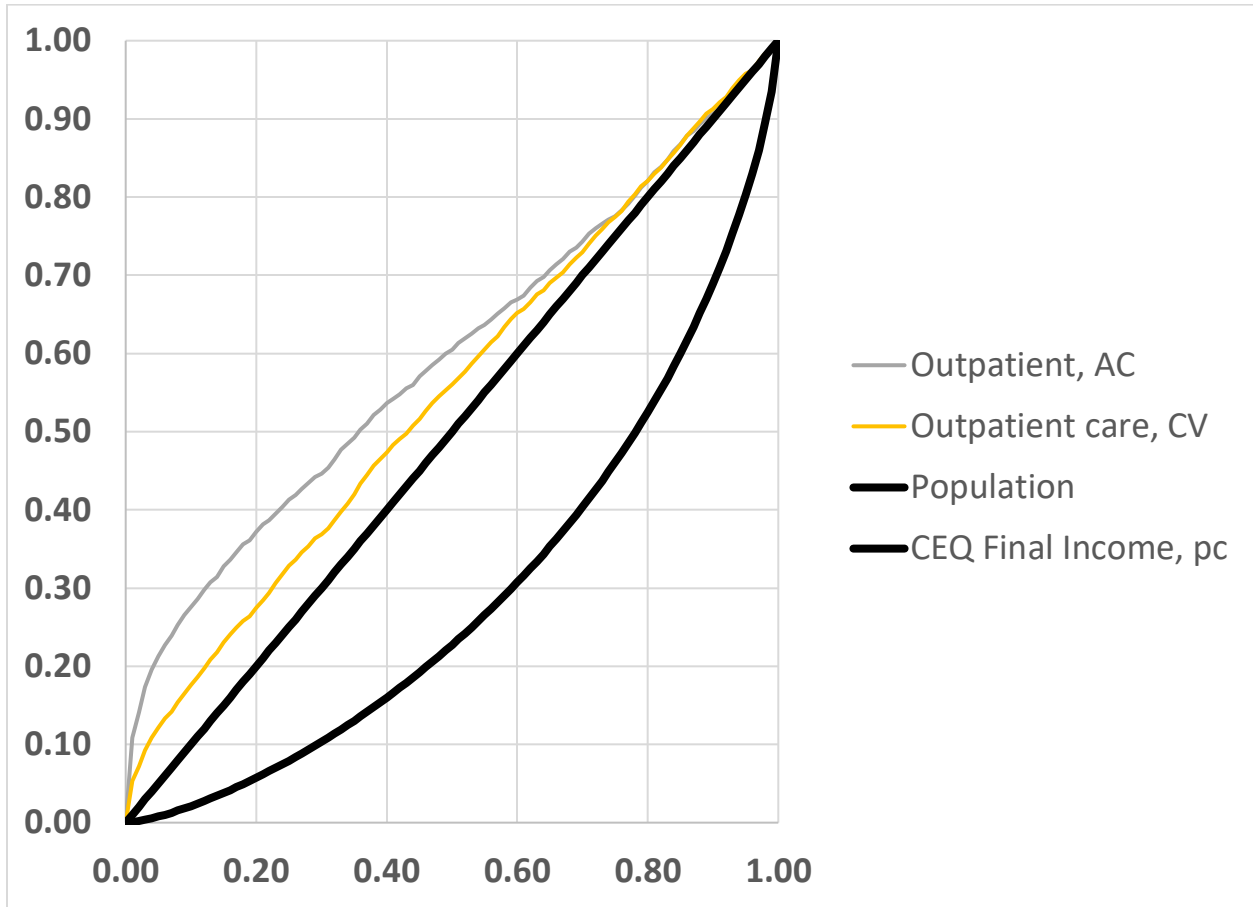
**Figure A.2 - Insurance Value by Income Quartile and Risk Aversion (consumption floor proportional to income).**



### Section A.3 – Concentration Curves By Valuation Method

This section of the appendix presents concentration curves for the benefits estimated with each health valuation method in the paper. Comparing concentration curves is more general than comparing concentration coefficients (the area between the curve and the 45-degree line), though in these examples, the information is qualitatively similar. Figure A.3 shows concentration curves for the average cost method. It shows that these methods produce a more progressive estimate of health care's value than the compensating variation WATP method because willingness-and-ability-to-pay increases with income, but both methods show a progressive distribution of benefits. Figure A.4 shows concentration curves for each health intervention analyzed in the health outcomes approach. It indicates that interior spraying for mosquitos is the most progressive public expenditure of this group, showing that, for example, almost 70% of the benefits from IRS goes to the bottom two income quintiles. The benefits from insecticide-treated bed nets also goes disproportionately to the poor, although less than IRS, because this program focuses on families with young children. The distribution of benefits from antimalarial drugs is spread evenly across the population, though as noted in the text, the need (as measured by malaria incidence) is greater among poorer households. The benefits of antiretroviral drugs and, especially, diabetes care, go more to richer households, though both are still distributed more equally than income itself. Figure A.5 shows the concentration coefficients for two estimates of financial risk protection using a consumption floor proportional to income. Specifically, one scenario uses a high level of risk aversion ( $\rho=4$ ) and low consumption floor of  $\gamma=10\%$ , which would maximize the estimate of insurance value. The other scenario uses a moderate level of risk aversion ( $\rho=2.5$ ) and higher consumption floor  $\gamma=20\%$ . Even with these changes we see that the distributional effects are similar, generating impacts that are more equal than income, but not absolutely progressive. In contrast, the concentration curve of insurance value with moderate risk aversion and a consumption floor  $\gamma$  equal to the extreme poverty line shows that benefits to the bottom 20% are lower than their proportion of income. However, the value is greater than their proportion of income for the middle 60% of the income distribution.

**Figure A.3: Concentration Curves for Outpatient Care Value**

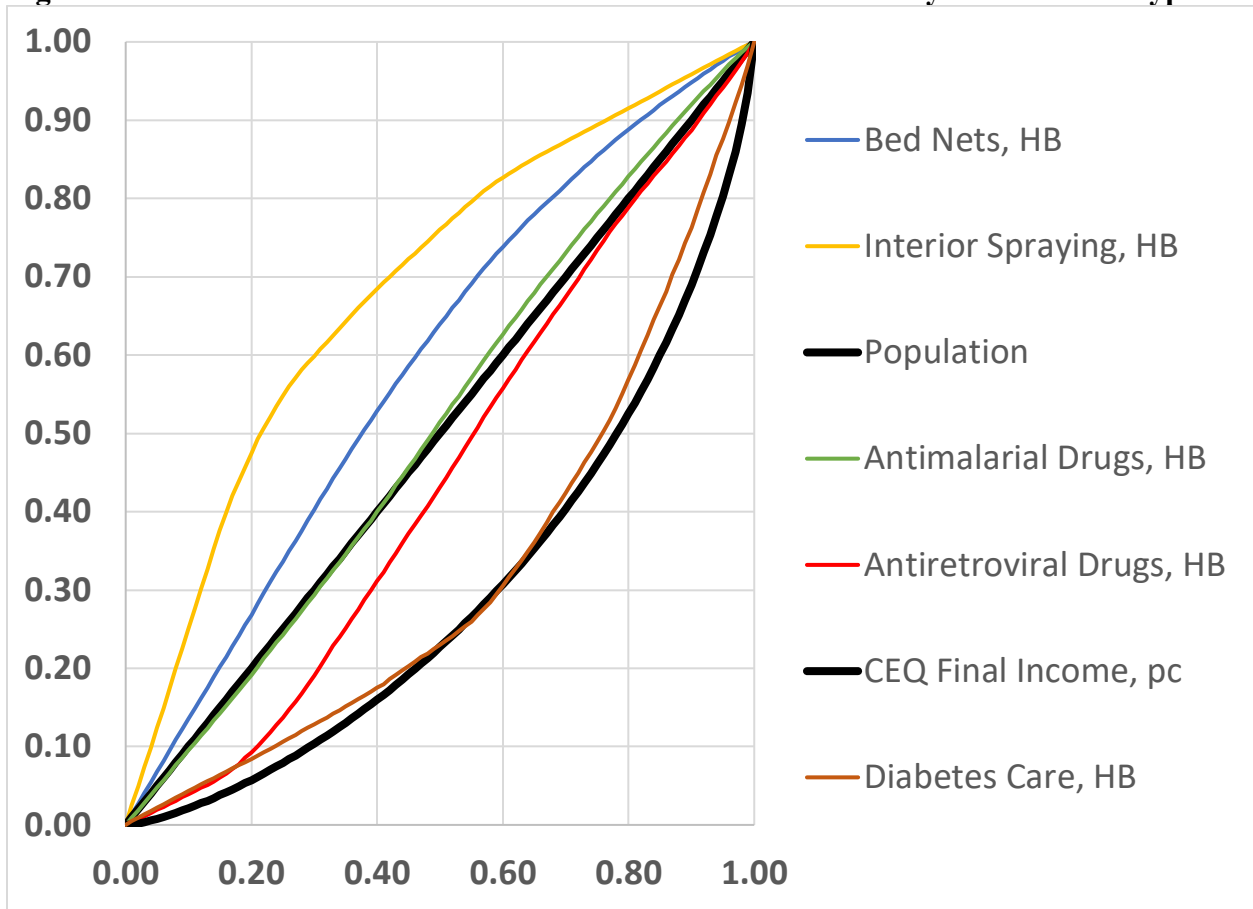


Source: GLSS-6 and authors' calculations

Notes: "AC" is for the average cost approach; "CV" is for the compensating variation.

Observations ranked by "Final Income" as defined by CEQ but *excluding* any benefits from public health care spending.


**Figure A.4 – Concentration Curves for Health Outcomes Benefits by Intervention Type**



Source: GLSS-6 and authors' calculations

Notes: Observations ranked by "Final Income" as defined by CEQ but *excluding* any benefits from public health care spending.

**Figure A.5 – Concentration Curves for Financial Risk Protection**



Source: GLSS-6 and authors' calculations

Notes:  $\rho$  is the coefficient of relative risk aversion.  $\gamma$  is the percent of household income used as a lower limit for non-health expenditures, except when  $\gamma=z$ , which indicates use of the extreme poverty line (792 cedis per adult equivalent per year) as the lower limit. Observations ranked by "Final Income" as defined by CEQ but *excluding* any benefits from public health care spending.

#### **Section A.4 Using Willingness and Ability to Pay by Matching Publicly-Funded Health Services to Private Health Services**

A straightforward approach to valuing health care services is to look at what people actually pay for health care services at private providers of comparable quality to the public services we want to value. This also relies on revealed preference: the service is obviously worth at least as much to users as they are willing to pay for it in a private market. Soares (2018) uses this approach for schooling in Brazil.

There are two challenges for this approach, one conceptual, the other practical. Conceptually, we need to control for different quality of services. We can do so by matching public to private services with one or more indicators of the quality of the service. Practically, we need data on the quality of services at public and private facilities and the fees charged at the private facilities.

Relatively few surveys used for a CEQ assessment include this information, which should be collected at the facilities level rather than at the household level. But the Demographic and Health Surveys' Service Provision Assessments, available in 15-20 countries, do provide the necessary facilities data. There may be other special purpose surveys in other countries done on an ad hoc basis.<sup>49</sup>

Once facilities are matched, we simply would use the matching private facility's price to estimate the value of the matching public facility's service. These values can then be assigned to users of public health facilities, perhaps disaggregated by geographic area and type of facility.

### **Example Application**

Although there has not been an example in the health sector of this method, it has been used in education to estimate the value of in-kind subsidies for schooling in Brazil. In the education context, there are clearer substitutes between private and public services than in health. However, there does exist a series of surveys called Service Provision Assessment collected through the Demographic and Health Surveys that may be used for this purpose. One exists for Ghana from 2002, outside of our timeframe of interest here.

### **Discussion of Public/Private Matching**

The advantage of the matching approach is that, like the discrete choice demand estimation, it relies on patients' actual choices. It is also much easier to implement. An important disadvantage is that the data on facility quality and prices may not be available in many countries. In addition, because this approach relies on demand for private services, it is not applicable to health spending on public goods or natural monopolies, services that have some of the highest returns for public spending.

### **Section A.5 Data and Do-files for Replication**

The do-files and data required to replicate this analysis in its entirety have been posted to [ADD LOCATION HERE]. These serve both to validate the results shown here as well as provide a starting point for similar analysis to be undertaken in another country using the five criteria we describe in section 5 to assess which method should be used and based on that nation's health system structure.

---

<sup>49</sup> Unfortunately, the World Bank's Service Delivery Indicators surveys do not collect information on fees.