Does government spending affect income poverty?
A meta-regression analysis

Abstract

This paper presents the results of a meta-regression analysis of the relationship between government spending and income poverty, with a focus on low and middle income countries. Through a comprehensive search and screening process, we identify a total of 19 cross-country econometric studies containing 169 estimates of this relationship. We find that the size and direction of the estimated relationship are affected by a range of factors, most notably the composition of the sample used for estimation, the control variables included in the regression model, and the type of government spending. Overall, we find no clear evidence that higher government spending has played a significant role in reducing income poverty in low and middle income countries. This is consistent with the view that fiscal policy plays a much more limited redistributive role in developing countries, in comparison with OECD countries. In addition, we find that the relationship between government spending and poverty is on average less negative for countries in Sub-Saharan Africa, and more negative for countries in Eastern Europe and Central Asia, compared to other regions. We also find that the relationship is less negative for government consumption spending, in comparison with other sectors. Finally, we find some evidence indicating the possibility of publication bias.

Keywords

Income poverty; government spending; meta-regression analysis; pro-poor growth; fiscal policy.

1. Introduction

In September 2015, the United Nations announced a new target to eradicate extreme poverty by 2030, as measured by the number of people living on less than $1.25-a-day. Recent research indicates that this target is unlikely to be met by economic growth alone. If there is no change in the distribution of income within countries, the global $1.25-a-day headcount is projected to remain at between 5 and 7 percent in 2030, even under fairly optimistic assumptions regarding rates of economic growth (Lakner et al., 2014; Yoshida et al., 2014; World Bank, 2015a,b). Meeting the new global poverty target requires not just growth, but growth that is combined with distributional changes that by themselves reduce poverty.

There has been much debate in the literature about the factors that affect the amount of poverty reduction associated with economic growth (e.g. Ravallion 2001, Verschoor and Kalwij 2006, Son and Kakwani 2008). Within this debate, the level and allocation of government spending is often argued to be one key influence. However, the empirical evidence in support of this view is not always forthcoming. Several cross-country econometric studies have investigated the relationship between government spending and income poverty, and show an interesting diversity of results. For example, Mosley et al. (2004) find that 'pro-poor' government spending has a negative and statistically significant effect on the $1-a-day poverty headcount, and Kwon and Kim (2014) find that health spending has a negative and statistically significant effect. However, Kraay (2006) finds that the effect of government consumption spending on the 'redistribution' component of $1-a-day poverty reduction is not statistically significant, while the effect on the 'growth' component is in fact positive.[1] Wagle (2012) finds that the size and significance of the effect of government consumption spending on income poverty varies substantially, according to the sample and specification used.

The aim of this paper is to explain some of the heterogeneity in the estimated relationship between government spending and income poverty found in the literature, through the use

of meta-regression analysis (MRA). The aim of MRA is to provide a reliable and objective way of summarising research findings (Stanley & Doucouliagos, 2012). It uses multiple regression analysis to uncover the reasons why estimates vary, and to establish whether there are any consistent and generalisable results which apply across contexts and methods. Although evidence on the impacts of government spending on poverty comes from a range of different sources, we restrict our attention to cross-country econometric studies, on the grounds that a large number of such studies can be found in the literature, which do generate a diverse range of findings.

MRA has been used very widely in the literature on to date, on issues such as the effects of aid on economic growth (e.g. Doucouliagos & Paldam, 2008, 2009, 2015), the effects of distance on international trade (e.g. Disdier & Head, 2008), and the impact of foreign direct investment on domestic firms (e.g. Irsova & Havranek, 2013). However, we are not aware of any previous studies using MRA to study the determinants of income poverty at the national level, despite the relatively large empirical literature on this topic. One recent study uses MRA to look at the effects of government spending on income inequality, and finds some evidence of a moderate negative relationship, which is strongest when using the Gini coefficient or the top income share as the inequality measure (Anderson et al., 2016a). However, the change in inequality is not always an accurate guide to how a change in income distribution affects poverty (Datt & Ravallion, 1992). It is possible for government spending to affect income inequality without affecting income poverty – for example, if it leads to a redistribution of income from the richest households to households in the middle of the income distribution. It is also possible for government spending to affect income poverty without affecting inequality – if for example it leads to higher incomes across the entire distribution. As a result, we cannot necessarily infer the relationship between government spending and poverty from its relationship with inequality.

The remainder of the paper is organised as follows. Section 2 begins by briefly discussing the linkages between government spending and income poverty in theory, and the reasons why estimates of this relationship may vary. Section 3 then describes the inclusion criteria used to identify the studies included in our analysis, and the search process used. In total, we identify 19 cross-country econometric studies, containing 169 estimates of the relationship between a measure of government spending and a measure of income poverty. Section 4 then presents the results of the MRA, asking whether there is any consistent evidence of a relationship between government spending and income poverty across the 19 studies, and what explains the apparent heterogeneity in the estimated size and direction of this relationship. Overall, we find no clear evidence that higher government spending has played a significant role in reducing income poverty in low and middle income countries. However, we do find that the relationship between government spending and poverty is on average less negative for countries in Sub-Saharan Africa, and more negative for countries in Eastern Europe and Central Asia, compared to other regions. We also find that the relationship is less negative for government consumption spending, and health and education spending, in comparison with other sectors. Finally, we find some evidence indicating the possibility of publication bias, in that positive estimates of the relationship between government spending and poverty appear to be somewhat under-represented in the literature. Section 5 summarises these main findings, and discusses the implications for policy and for future research.

2. Theoretical background

The relationship between government spending and income poverty is complex and may vary for a number of reasons. First, and most obviously, it is likely to depend on the type of spending being considered. Government spending on transfers and indirect subsidies can reduce poverty directly, by raising the real disposable ('post-fiscal') income of poor households. It can also do so indirectly, by leading to better nutrition, health and education

among poor households, which in turn leads to higher market ('pre-fiscal') income. Government spending on basic health and education services and certain types of infrastructure (e.g. rural roads, water and sanitation, housing) is also widely considered to reduce poverty, by increasing the productivity and earnings potential of poor households (McKay, 2004; Mosley et al., 2004; Paternostro et al., 2007). These types of government spending are, at least in theory, most likely to reduce income poverty, and are often labelled as 'pro-poor' for that reason.[2]

Nevertheless, a large part of government spending on transfers and subsidies in developing countries does not reach poor households, due to imperfect targeting. In Indonesia for example, it is estimated that over 80% of the benefits of subsidised gasoline go to households in the top half of the income distribution (Rhee et al., 2014). Similarly, much of the benefits of government health and education spending are received by the middle classes, particularly in urban areas (e.g. Castro-Leal et al., 1999; Davoodi et al., 2010). Thus the actual impact of spending on transfers and other 'pro-poor' sectors will depend critically on how well the spending is targeted towards poor households, which may vary across countries. Transfers and subsidies can also have side-effects, e.g. on household labour supply, or on receipts of private transfers, which can offset their effect on income poverty (e.g. Cox & Jiminez, 1995; Sahn & Alderman, 1996; Cox et al., 2004). Thus even when well targeted, the overall impact of transfers and subsidies on income poverty is ambiguous.

The impact of government spending on poverty also depends on the way it is financed (McKay, 2004). Direct income taxes are considered to have little direct impact on poverty, either because households living below the poverty line are exempt, or because they are outside the direct tax system altogether. In many countries however, a significant share of tax revenue comes from indirect taxes (e.g. value-added tax, excise duties). In Latin America for example, around 60 per cent of tax revenue comes from VAT, in comparison with 40 per cent in OECD

countries (Goni et al., 2011). Such taxes can increase poverty, by raising the prices of goods and services consumed by poor households. Monetary financing of government spending could also have an adverse effect on poverty, by causing higher inflation (Easterly & Fischer 2001).

There are good reasons therefore to expect the impact of government spending on poverty to vary, according to the sector of spending, how well it is targeted, and the way in which it is financed. The effect may also differ according to the time-period of analysis, since some types of spending have direct, immediate impacts on poverty (e.g. transfers and indirect subsidies), while others only have more indirect, medium-term effects (e.g. health, education and infrastructure spending). The measure of poverty could also make a difference; if some types of spending are more effective in reaching the very poorest households, we would expect to see a stronger relationship with the squared poverty gap or the Watts index, rather than the poverty headcount (see for example Bourguignon and Fields 1997). We seek to investigate these and other potential sources of variation in the estimated relationship between government spending and income poverty in our MRA.

3. Search strategy and inclusion criteria

3.1 Inclusion criteria

Study design. We restrict our attention to cross-country econometric studies which regress a measure of income poverty on a measure of government spending, plus other explanatory variables. This type of regression can be written as follows:

$$h_{it} = \beta_0 + \beta_1 G_{it} + \sum_k \beta_k Z_{itk} + \varepsilon_{it} \quad (1)$$

where h is a measure of income poverty, G is a measure of government spending, Z_k is a set of other explanatory variables considered to affect income poverty (e.g. GDP per capita, trade

openness, inflation, good governance), and ε is the error term, with subscripts *i* and *t* indicating country and year respectively. We include any study that estimates a version of equation (1), in terms of levels or differences over time.[3]

Outcome measure. Although there are a number of domains in which poverty can be assessed, we focus on income poverty in order to avoid excessive heterogeneity in the results. We require that income poverty is measured using an absolute poverty line which is fixed in real terms over time; any studies focusing solely on relative poverty are omitted. We also require that income poverty is measured using a comprehensive measure of income, including market and non-market income, and that data are drawn from a representative national household survey.[4] We do however include studies which use total expenditure rather than total income to measure poverty, on the grounds that expenditure is a more reliable indicator when data on income are difficult to collect. We include studies using any measure of income poverty, including the poverty headcount, the poverty gap, the squared poverty gap, and the Watts index.[5]

Spending measure. We include studies focusing on any type of government spending, either total spending or a specific category of spending (e.g. health, education, social welfare). Each may be measured in nominal terms, as a share of total GDP, or (for disaggregated measures) as a share of total government spending. We do however require that the spending measures exclude spending by private households; for example, health spending must exclude private 'out of pocket' health expenses. We also require that data on government spending be derived from a source other than the household survey used to measure poverty. This means that we exclude studies which focus on household receipts of government income transfers, as recorded in household surveys, on the grounds that this is not a measure of government spending per se.

Population. The analysis is restricted to studies focusing on low and/or middle income countries, as defined by the World Bank. Studies including some former low or middle income countries which are now classified as high-income are also included, on the grounds that the vast majority of countries included in such cases are low or middle income. However, we exclude any studies which focus predominantly on high income or OECD countries.

Other criteria. We include refereed and non-refereed journal articles, working papers, conference proceedings, book chapters, government reports, NGO reports and other technical reports. However, we restrict the review to studies published since 1990; this is mainly on the grounds that reliable, cross-country data on income poverty have only been available since the early 1990s, so that any studies before this date would not meet basic requirements in terms of data quality. Our review is also restricted to studies published in English, Portuguese, or Spanish.

3.2 Search strategy

Studies meeting the above inclusion criteria were identified as part of a wider systematic review carried out by the authors (see Anderson et al., 2015, 2016b), which looked at the relationship between income poverty and government policy more generally. The search process was carried out between June and November 2014. The search strategy used for this review is set out in detail in Anderson et al (2015, 2016b); here we provide a brief summary.

We searched a number of online databases for relevant articles, including Scopus, IBSS, Web of Knowledge, Econlit, IDEAS, SSRN and Google Scholar, as well as additional databases and websites maintained by organisations such as the World Bank, the OECD, the IMF and the ILO.[6] Each database was searched using a combination of broad search terms, such as 'poverty', 'the poor', 'pro-poor' and 'government'. When using foreign language databases, each term was translated into the appropriate language, i.e., Portuguese or Spanish. In order

to decrease the risk of overlooking relevant studies, we did not include the more precise seach term 'income poverty'. This is because our initial searches revealed many relevant studies which do not, at least in the abstract and keywords, explicitly mention income poverty but refer only to 'poverty'.[7]

Our electronic searches identified 11,986 records in total. Apart from a small minority which could not be located (197 records), we screened all the remaining records first by title and abstract, and then (if necessary) by full text, to determine which of them met our inclusion criteria.[8] In addition to the electronic searches, we also identified a number of additional studies relevant to the review via hand-searching. In particular, we checked the reference lists of all the academic journal articles identified via our electronic searches which met our inclusion criteria, to see if there were any other additional studies relevant to our review which we had missed.

3.3 Search results

In total, we identified 19 studies that meet our inclusion criteria, containing 169 estimates of the relationship between a measure of government spending and a measure of income poverty. A full list of the studies is contained in Appendix 1. Of the 19 studies, nine have been published in peer-reviewed academic journals, and all have been published since the year 2000.[9] Table 1 provides an overall description of the sample, as well as the distribution of results. Of the 169 estimates, 124 show a negative relationship – i.e., higher (lower) government spending is associated with lower (higher) poverty – of which 31 are statistically significant at the 10% level or below. The remaining 45 show a positive relationship, of which 21 are statistically significant.

[Table 1 about here]

The results in Table 1 to a large extent confirm the view noted in the Introduction, namely that the cross-country econometric literature provides a diverse set of findings regarding the relationship between government spending and income poverty. One might be tempted to conclude that the balance of evidence points toward a negative relationship, on the grounds that more estimates show a negative and statistically significant relationship than a positive one. Such a conclusion would clearly be dangerous however, since just under half of the estimates show no statistically significant relationship at all. Moreover, the results from 'vote counting' exercises such as Table 1 need to be treated with caution (Higgins & Green 2011; Stanley & Doucouliagos, 2012). On the one hand, they give only limited information about the distribution of estimates; they say nothing about the size and strength of the relationship for example. On the other hand, they fail to take into account potential biases in the way that estimates are reported (e.g. publication bias), or differences in the way that estimates are derived – for example, between estimates based on different types of government spending, or different measures of poverty. Meta regression analysis provides a more informative and robust way of synthesising evidence of this sort.

A final point worth noting is that the number of studies and estimates included in the MRA by Anderson et al (2016a) on the relationship between government spending and income inequality – 84 and 952 respectively – is significantly larger than in this paper. This is an interesting finding, highlighting the much larger number of cross-country econometric studies looking at the relationship between government spending and income inequality than between government spending and income poverty. However, there is relatively little overlap between the two sets of studies: only seven of the 19 studies included in this paper also feature in the MRA by Anderson et al (2016a). This to a large extent justifies the separate and independent systematic review of the literature carried out in this paper.

4. Meta regression analysis

In this section we present our meta-regression analysis designed to establish whether there is any evidence of a relationship between government spending and income poverty, and to explore the reasons for heterogeneity in both the size and direction of this relationship. We follow the MAER-NET guidelines to report the findings (see Stanley et al., 2013). We first discuss our effect size measure, our initial tests for publication bias, and our overall modelling approach (Section 4.1). We then present the results of the analysis (Section 4.2).

4.1 MRA approach

Effect sizes

Despite our relatively strict inclusion criteria, the regression coefficients contained in each of our 19 studies are not directly comparable. This is due to differences across studies in the measure of spending and the poverty metrics used. For example, some studies use spending as a percentage of GDP, while others use spending in US$; some studies express poverty in logarithmic terms, while others do not. All regression-based estimates were therefore converted into a comparable measure, the partial correlation coefficient. In our context, the partial correlation measures the strength of association between income poverty and government spending, holding all other factors constant. This was the best choice given our particular context, given that its value is not affected by the units in which either the spending variable or the poverty variable is measured in. It is calculated as follows:

$$r = \frac{t}{\sqrt{t^2 + df}} \tag{2}$$

where t is the t-statistic of the regression coefficient and df reports the degrees of freedom from the t-statistic (Stanley & Doucouliagos, 2012). If the t-statistic was not reported we calculated it by dividing the regression coefficient by its standard error.

We were able to calculate the partial correlation coefficient corresponding to each of the 169 regression coefficients contained in the 19 studies. The values range from -0.698 to 0.322; the simple (unweighted) average is -0.150 (s.e. 0.056), while the weighted average is -0.089 (s.e. 0.054).[10] It has been suggested that a partial correlation coefficient of less than 0.07 in absolute terms can be considered small, with 0.17 or above considered to be moderate, and 0.33 or above large (Abdullah et al., 2015). This is in line with what Cohen (1988) suggests, who argues that for partial correlation coefficients, the effects are considered to be small when r = 0.1, medium when r = 0.3 and large when r = 0.5.

Publication bias

Publication bias is a serious issue in the context of systematic reviews as it can introduce serious biases in meta-analytical results. It is argued that studies reporting statistically significant findings are more likely to be published in peer-reviewed journals than studies reporting statistically non-significant findings. This bias in the literature will then also be reflected in the meta-analysis, as published studies are more likely to be included (Borenstein et al., 2009).

The funnel plot is one of the most common methods to detect for the presence of publication bias (see for example Egger et al., 1997). Figure 1 illustrates a funnel plot which plots the effect size on the x-axis (here the partial correlation coefficient between a measure of government spending and a measure of income poverty), and precision on the y-axis (here the inverse of the standard error of the partial correlation coefficient). At the bottom of the graph we find the estimates with less precision, i.e. larger standard errors, while the estimates with more precision, i.e. smaller standard errors, are found towards the top of the funnel plot. If there is no publication bias, the estimates will be distributed in the shape of an inverted funnel: the dispersion of estimates will fall as the level of precision rises, but the dispersion will be symmetrical either side of the true effect size. However, if there is publication bias – for

example, if some studies remain unpublished or unreported because they do not report any statistically significant effects – then this leads to asymmetry in the funnel plot. This in turn implies that the average effect size calculated in the meta-analysis will tend to overestimate the true effect size. The amount of publication bias is more substantial, the more severe the degree of asymmetry in the funnel plot.

[Figure 1 about here]

In our case a visual inspection of the funnel plot suggests a lack of symmetry, suggesting the presence of publication bias. Note however that visual inspection of funnel plots can be subjective (Borenstein et al., 2009; Abdullah et al., 2015) and thus Stanley (2005, 2008) suggests the FAT-PET (Funnel-Asymmetry Precision-Effect) regression as an empirical test to check more reliably for any publication bias. We carry out this test as part of our meta-regression analysis.

Modeling heterogeneity

The results in Figure 1 suggest that a certain degree of heterogeneity remains in the studies in the meta-regression analysis, in that the reported estimates are rather spread out. To better understand what drives this heterogeneity we adopt the following meta-regression model to explore heterogeneity in the reported estimates:

$$r_{ij} = \beta_1 + \sum_m \beta_m Z_{ijm} + \beta_0 SE_{ij} + \nu_{ij} \quad (3)$$

where r is the partial correlation coefficient between government spending and income poverty of the ith estimate from the jth study (19 studies and 169 estimates), Z_m is a set of variables that capture differences in the relationship between government spending and income poverty, SE is the standard error of the partial correlation coefficient and ν is the error term.[11] The following variables are included in the set of variables Z_m:

Measures of the dependent variable: The vast majority of estimates (86% of the total) use the headcount measure of poverty; we include a dummy variable equal to 1 for estimates using the poverty gap measure (9% of the total). The remaining 5% of estimates use either the squared poverty gap or the Watts index.[12]

Measures of government spending: We include three dummy variables which are equal to 1 for each of the following measures of spending: social welfare spending, consumption spending, and health and/or education spending. These accounted for 17%, 38% and 14% of estimates respectively. The remaining estimates refer to total government spending, or other categories of expenditure (e.g. defence, housing, agriculture) for which we had only a few observations in each case. Social welfare spending includes components such as pensions, social security, social protection, and other welfare spending. Government consumption spending refers to government final consumption expenditure. Health and/or education spending refers to health spending, education spending, or the combination of the two. We tried adding another dummy variable for the other categories of spending in our dataset (e.g. housing, agriculture), but this made little difference to the results.

Country composition: More than half of the estimates (60% of the total) are based on samples including countries from all developing country regions. We include a dummy variable equal to 1 for samples including countries from Sub-Saharan Africa (79% of the total), and another for samples including countries from Eastern Europe and Central Asia (60% of the total).

Estimation method: We include a dummy variable equal to 1 if estimates are derived using ordinary least squares (OLS; 58% of the total). The remaining estimates use a method other than OLS, e.g. dynamic panel estimators such as generalised method of moments (GMM), panel data analysis using random or fixed effects, and other econometric approaches such as instrumental variables and 2-stage or 3-stage least squares.

Sample period: The average year of the data was included to account for the different time periods used by different studies, since the relationship between government spending and income poverty may vary over time. This variable was transformed by subtracting the average year of data across all estimates (1992) from the average year of data for each estimate.

Other explanatory variables: We include a set of dummy variables to capture the different explanatory variables which are included when estimating equation (1), in particular: GDP per capita, trade policy, inflation, governance, education, foreign aid and income inequality. Each dummy variable was coded as 1 if the relevant explanatory variable was included in the model specification, and as 0 if otherwise. The variables were chosen after careful review of all included studies. The trade category incorporates all variables that were considered valid measures of trade policy, e.g. average import tariffs or export duties, or trade policy indices. The governance category reflects various aspects of good governance and institutional quality, including voice and accountability, control of corruption, rule of law, and so on.[13] The education category includes years of education and other schooling-related variables such as enrolment rates. Foreign aid, income inequality and inflation appeared frequently in the specifications and it was decided to include them as well.

Publication: The standard error of the partial correlation coefficient is included to account for publication bias. We also include a dummy variable equal to 1 for estimates derived from unpublished studies, e.g. working papers (62% of the total).[14]

A full list of all explanatory variables included in the MRA, and their mean values and standard deviations, is reported in Appendix 2.[15] The estimations are carried out using a regression procedure with a weighted least squares (WLS) routine advocated by Stanley and Doucouliagos (2017). They demonstrate how an unrestricted WLS-MRA is likely to be as good as and often better than both random-effects and fixed-effects meta-regression analysis in practical applications. The majority of the studies we included reported more than one

regression coefficient that could be used to calculate the partial correlation coefficient, none of the studies specified a preferred result, and thus we were faced with multiple estimates per study. This needs to be dealt with appropriately to avoid bias due to data dependency (Lipsey & Wilson, 2001). The literature suggests a number of approaches to dealing with multiples estimates per study (see for example Lipsey & Wilson, 2001; Borenstein et al., 2009) and there is no consensus on the preferred approach. Following Abdullah et al (2015), our preferred approach to accounting for multiple estimates per study is to use precision squared (i.e., the inverse variance) as weights with study level clustered standard errors. However, we also explore a range of different approaches to dealing with multiple estimates per study as robustness checks, which are outlined in more depth in the next section.

4.2 Results

Our main results are reported in Table 2. Regression 1 reports the results of the FAT-PET model, where the standard error of the partial correlation coefficient is regressed on the partial correlation coefficient. The FAT-PET regression is an empirical check to explore publication bias. The results of regression 1 indicate that there is publication bias, as the coefficient for the standard error is statistically significant. This supports the visual inspection of the funnel plot (Figure 1), which also suggests publication bias is present. The coefficient on the standard error term is negative, indicating that the estimated partial correlation coefficients are skewed towards negative values; positive estimates of the relationship appear to be under-reported in the literature.

The constant in regression 1 indicates the average relationship between government spending and income poverty, after correcting for publication bias. This takes the value of 0.057 (s.e. 0.083), which is quite different from the average relationship before correcting for publication bias (-0.089, s.e. 0.054). Controlling for publication bias therefore makes a clear difference, turning a small to moderate negative relationship into a small positive relationship. However,

in neither case is the relationship statistically significant. Overall therefore, the FAT-PET model provides evidence of publication bias, but no evidence of a relationship between government spending and income poverty, after correcting for publication bias.

Regression 2 is our main model as it includes all of the moderator variables described above. To help interpret the coefficients, it helps to start by considering a base case in which all moderator variables are equal to zero. In this case, the average relationship between government spending and poverty, after correcting for publication bias, is given by the constant term. This takes the value of -0.274 (s.e. 0.141), indicating a moderate negative relationship, although this is statistically significant at the 10 per cent level only. However, the more interesting point is whether and if so how this average figure is affected by our moderator variables. The results indicate that three sets of variables are important.

The first relates to the composition of the sample used for estimation. The coefficient on the SSA variable is positive, indicating that the relationship between government spending and poverty is less negative in this region, compared to other regions. By contrast, the coefficient on the EECA variable is negative, indicating that the relationship is more negative in this region, compared to others. These effects are both statistically significant at the 1% level, and large in size. For example, the average relationship across samples including countries from all regions is -0.180 (s.e. 0.082). However, for samples excluding countries from Sub-Saharan Africa, the average relationship is -0.774 (s.e. 0.271); for samples excluding countries from Eastern Europe and Central Asia, it is 0.321 (s.e. 0.174).

The second set of variables corresponds to the control variables used during estimation. The coefficient on the education variable is positive and statistically significant, indicating that the relationship between government spending and poverty is on average less negative when controlling for education. Studies which fail to control for education may therefore be overestimating the effect of government spending on poverty. By contrast, the coefficients on

the inequality and foreign aid variables are both negative and statistically significant, indicating that the average relationship is more negative when controlling for either of these variables.

The third set of variables refers to the measure of spending used. The relationship tends to be more negative for social welfare spending, but this effect is not statistically significant. By contrast, the relationship tends to be less negative for health and education spending; this effect is statistically significant at the 1% level. We also see that the relationship between spending and poverty tends to be less negative when focusing on government consumption spending, which includes health and education spending, as well as spending on other public goods and services.

[Table 2 about here]

The results for the other moderator variables in regression 2 are more inconclusive. We find no evidence that the measure of poverty used, the time period of the sample, the estimation method, or inclusion of any of the other control variables (GDP per capita, trade policy, governance and inflation) makes a difference to the results; the coefficients on these moderator variables are not statistically significant. We also see that the coefficient for the standard error is no longer statistically significant in regression 2, suggesting that the evidence for publication bias in regression 1 is not entirely robust. Nevertheless, in overall terms, our model is able to account for nearly three quarters (72 percent) of the total variation in the estimated relationship between government spending and income poverty found in the literature. This is a large amount, given the significant amount of heterogeneity across the different studies.

In regression 3 we follow Leonard et al. (2014) and employ a general-to-specific modelling strategy, removing the variable that has the largest p-value until all p-values are less than 0.05.

The rationale for employing a general-to-specific approach can be found in Stanley and Doucouliagos (2012), who prefer a more specific model as it makes the underlying associations clearer. In this model, we observe that the two variables for sample composition (SSA, EECA) are retained, together with three of the control variables (inequality, education and foreign aid), and two of the spending measures (consumption spending and health and education spending) – all of which are also statistically significant in our main model. This more restricted model is still able to account for nearly two thirds (65 per cent) of the total variation.

In Table 3 we report the results of our robustness tests. We first re-estimate our main model using four different approaches to dealing with multiple estimates per study. As discussed in Section 4.1, a number of approaches have been used in the literature to control for this problem, but there is no consensus on the optimal approach. In column 1 we use the sum of precision squared for each study as weights. The regression in column 2 further refines this approach by using the sum of precision squared for each study as weights and the standard error of the mean of the partial correlation coefficient instead of the standard error of the partial correlation coefficient. Column 3 shows the results when using the inverse number of estimates per study as weights. The approach used in column 4 is similar to that in column 2, except that the weighted mean of the partial correlation coefficient is used as the dependent variable, rather than the partial correlation coefficient. Finally, we re-estimate the model using robust regression techniques, designed to control for outliers and influential observations; this is shown in column 5.

[Table 3 about here]

In all the robustness tests shown in Table 3, the moderator variables for SSA, EECA, inequality, foreign aid, consumption spending and health and education spending remain statistically significant at the 1% level or below, and retain the same sign and order of magnitude as in

Table 2. One relatively small difference is that the education variable is no longer statistically significant in columns 3-5. However, the main difference between Tables 2 and 3 is that more of the moderator variables are statistically significant in Table 3. First, the dummy variable for unpublished studies is now positive and statistically significant, indicating that estimates of the relationship between government spending and poverty from unpublished studies are on average less negative than estimates from published studies.[16] The year of data variable is also positive and statistically significant, indicating that the relationship between spending and poverty is becoming less negative over time. By contrast, the OLS variable is negative and statistically significant, indicating that estimates of the relationship between spending and poverty derived using OLS are on average more negative, in comparison with estimates derived from other approaches, e.g. two-stage least squares (2SLS), fixed or random effects. The coefficients for GDP per capita, governance, inflation and the standard error term are also now statistically significant, in at least some cases.

4.3 Discussion

To summarise, the results of our MRA model show that the estimated relationship between government spending and poverty is affected by a range of factors. The most robust evidence points to three sets of factors: the regional composition of the sample used for estimation, the control variables included in the regression model, and the type of government spending. In addition, there is some evidence that the estimation method and the time period of the sample affect the results, although in these cases the evidence is slightly less robust. There is also some evidence of publication bias, in terms of a negative relationship between the estimated partial correlation coefficient and its standard error, and studies not published in academic journals showing on average a less negative relationship between spending and poverty. Again however, this evidence is not robust across all our specifications.

Many of these results are in line with prior expectations, and mirror other findings in the literature. The absence of any clear evidence of a negative relationship between government spending and income poverty is consistent with evidence showing that fiscal policy plays a more limited redistributive role in developing countries, in comparison with OECD countries (e.g. Goni et al., 2011; Lustig et al., 2014; Lustig 2015). This is partly due to lower levels of spending and taxation, and less progressive taxation, but it is also due to less progressive government spending. For example, several studies have shown, using benefit incidence analysis, that public spending on education and health in many developing countries has not been very effective in reaching the poorest households (e.g. Castro-Leal 1999; Demery, 2000; Davoodi et al., 2010). Our findings are similar, in that we find no evidence that government spending on health and education in developing countries has been particularly 'pro-poor', in comparison with other sectors, in terms of its effect on income poverty.[17] We also see only limited evidence that government spending on social welfare is more strongly associated with poverty reduction than spending in other sectors.

The evidence that the relationship between government spending and poverty is on average less negative in Sub-Saharan Africa, and more negative in Eastern Europe and Central Asia, is also consistent with other evidence. For example, several studies using benefit incidence analysis show that fiscal policy has played a significant redistributive role in Eastern Europe, via well-targeted cash transfers and high and progressive levels of spending (e.g. Lelkes & Sutherland, 2009; Cok et al., 2013). By contrast, there is evidence that anti-poverty programmes have tended to be less well-targeted in countries with lower levels of governance, and lower levels of per capita GDP (Coady et al., 2006). In addition, Gupta and Verhoeven (2001) find that government spending on health and education is, at least on average, less effective in terms of raising health and education outcomes in Sub-Saharan Africa, compared to other regions. Baldacci et al. (2008) and Rajkumar and Swaroop (2008) further show that the effectiveness of health and education spending is lower in countries

with lower levels of governance. Although these findings relate to the effects of spending on health and education outcomes, our results suggest that a similar tendency may apply to income poverty.

The results for the control variables are more surprising. If government spending reduces inequality as well as poverty, and if poverty and inequality are positively correlated, we would expect the estimated relationship between spending and poverty to be less negative when controlling for inequality. In fact however, we see the opposite. The results for foreign aid are also surprising. If government spending and foreign aid both reduce poverty, and if government spending and foreign aid are positively correlated, we would again expect the relationship between spending and poverty to be less negative when controlling for aid. The finding that the relationship is in fact more negative could perhaps reflect a positive relationship between aid and poverty, due to aid selectivity. It is of course difficult to predict how the omission of a relevant explanatory variable is likely to affect an estimated regression coefficient, except under somewhat restrictive assumptions (Wooldridge, 2010). More broadly however, it is clear that the choice of control variables included in the regression model can make a large difference to the results, and most studies use quite different sets of control variables. We discuss the implications of this in the conclusion.

Finally, our results are also consistent with those of Anderson et al (2016a), who use MRA to investigate the relationship between government spending and income inequality. This study found some evidence of a moderate negative relationship between government spending and income inequality, which was strongest for social welfare and other social spending. However, the relationship was much weaker when the share of the poorest quintile is used as the measure of inequality, suggesting (consistent with this paper) that the redistributive impact of government spending is often concentrated in the upper half of the income distribution. The study also found that the relationship between government spending and income

inequality was more negative in Eastern Europe and Central Asia, and when using OLS as the estimation method, which is again consistent with the results in this paper. Nevertheless, there are also some interesting differences between the two sets of results. For example, while both studies find some evidence of publication bias, the results in Anderson et al (2016a) suggest that negative estimates (of the relationship between government spending and income inequality) are under-represented in the literature. By contrast, the evidence in this paper suggests that positive estimates (of the relationship between spending and poverty) are under-represented. It is not possible to say precisely what might be driving this finding, but one possible explanation is that positive estimates of the relationship between government spending and poverty conflict more strongly with researchers' prior expectations than they do for inequality.

5. Conclusion

In recent years, debates around income poverty have again risen to prominence, following the new UN target to eradicate extreme poverty by 2030. The aim of this paper is to better understand the factors influencing income poverty at the national level, through the use of meta-regression analysis (MRA). We focus on the role of government spending, on the grounds that decisions with respect to its level and allocation are widely considered to play a central role in shaping income distribution, and generating a pattern of growth that is effective in reducing poverty – even though it is clearly not the only factor. We also focus on evidence from cross-country econometric studies, on the grounds that a large number of studies can be found using this approach, which have generated a diverse set of findings – even though it is clearly not the only type of evidence.

Our main findings may be summarised as follows. Overall, we find no clear evidence from the cross-country econometric literature that government spending has played a significant role in reducing income poverty in developing countries. Many of the 19 studies included in our

review do of course find a negative and statistically significant relationship between government spending and income poverty, and the average relationship across all 169 estimates contained is negative, and not negligible in size. We have shown however that this overall average is not statistically significant, and its size may be exaggerated by publication bias – in this case, a possible under-reporting of positive estimates of the relationship. More significantly however, we show that estimates of the relationship between government spending and poverty are affected by a range of factors, most notably the regional composition of the sample used for estimation, the control variables included in the regression model, and the type of government spending.

In terms of the implications for policy, the key issue is how the contribution of government spending to the reduction of poverty can be strengthened. Current proposals in this area include a reduction of spending on universal price subsidies, particularly for fuel and energy, which are often quite regressive in their impact, as well as greater and better targeted spending on social welfare, health and education (see for example Bastagli et al. 2012, IMF 2014). The expansion of cash transfer programmes in many developing countries since the early 1990s is one promising development, given evidence of their progressivity (Goni et al., 2011). However, such programmes typically remain small as a share of overall spending. Moreover, our evidence suggests that if anything, the relationship between government spending and income poverty in developing countries has become less negative over time. This has parallels with trends seen in OECD countries, where the redistributive impact of public spending has tended to fall in recent decades (OECD 2011). The challenges in strengthening the poverty reducing and redistributive impact of public spending are clearly not to be under-estimated.

In terms of the implications for research, one might be tempted to conclude that there is little one can learn from cross-country econometric studies, and that other sorts of evidence – e.g.

fiscal incidence analysis, computable general equilibrium analysis, single-country econometric studies, and qualitative case studies – should be the focus. We argue instead that our results provide suggestions as to how cross-country econometric research in this area might usefully proceed in future. Three directions are most apparent. First, future work should look more closely at whether and how the effect of government spending on poverty varies by region. None of the 19 papers in our review explored this possibility, but our results, combined with evidence from other sources discussed above, strongly indicate that effects do vary. Second, our results suggest that more careful consideration is needed with regard to the choice of control variables. It is surprising that the 19 studies in our review contain such a wide range of different control variables, often without clear justification. Although there is understandably much discussion in most cross-country econometric studies regarding the estimation methods used (e.g. OLS vs. 2SLS or panel data methods), our results suggest that the control variables included in the model make a greater difference to the results.

Third, future work should consider whether and if so how the impact of government spending on poverty varies according to country characteristics, through the use of interaction terms. During our searches we did in fact find six studies using interaction terms. Three of these include interaction terms between government spending and average income (Verschoor & Kalwij, 2006; Carmignani, 2011; Balakrishnan et al., 2013); two include interaction terms between government spending and a measure of financial crisis (Baldacci et al., 2002; Assimaidou et al., 2013); the other includes interaction term between government spending and a globalisation index (Bergh & Nilsson, 2011). The results from these studies are difficult to compare directly with those from studies not including interaction terms, and were therefore excluded from our MRA.[18] Nevertheless, more studies of this sort are arguably warranted. For example, we found no studies testing whether the effect of spending on poverty reduction is greater in countries with better governance, similar to the studies by

Baldacci et al. (2008) and Rajkumar and Swaroop (2008) which focused on health and education outcomes.

Finally, future work should consider how sometimes diverse results from different econometric studies are best explained and synthesised. While there have been many excellent reviews of the literature on the determinants of pro-poor growth, and on the relationship between fiscal policy and income poverty (e.g. Klasen, 2003; McKay, 2004; Paternostro et al., 2007), this paper is to our knowledge the first attempt to carry out a systematic review of this literature. The 19 studies included in our analysis were identified using a clear and transparent set of inclusion criteria, and a comprehensive search strategy covering both published and unpublished studies. The analysis was carried out using a robust empirical approach (MRA) which has been widely applied to the study of economic growth, but as yet applied much less widely to the study of distributional issues. The overall approach has the potential to be extended into a number of other related areas, including other possible determinants of income poverty (e.g. governance, trade reform, tax policy), as well as other socio-economic outcomes affected by government spending decisions.

References

Abdullah, A., Doucouliagos, H., & Manning, E. (2015). Does education reduce income inequality? A meta-regression analysis. *Journal of Economic Surveys*, *29,* 301-316.

Anderson, E., D'Orey, M. A. J., Duvendack, M., & Esposito, L. (2015). The impact of government policies on income inequality and the translation of growth into income poverty reduction: protocol for two systematic reviews. *Journal of Development Effectiveness, 7*(4), 484-498.

Anderson, E., D'Orey, M. A. J., Duvendack, M., & Esposito, L. (2016a). Does government spending affect income inequality? A meta-regression analysis. *Journal of Economic Surveys,* forthcoming. DOI: 10.1111/joes.12173.

Anderson, E., D'Orey, M. A. J., Duvendack, M., & Esposito, L. (2016b). What policies and interventions have been strongly associated with the translation of growth into reductions in income poverty? A systematic review. EPPI-Centre, Social Science Research Unit, Institute of Education, University of London (2016), 122p.

Assimaidou, K., Kiendrebeogo, Y. and Tall, A. (2013). Social Protection for Poverty Reduction in Times of Crisis. *Serie Etudes et Documents du CERDI.* Clermont Ferrand: Centre D'Etudes et de Recherches sur le Developpement International (CERDI).

Balakrishnan, R., Steinberg, C., & Syed, M. (2013). The elusive quest for inclusive growth: growth, poverty and inequality in Asia. *IMF Working Paper, 13/152.* Washington D.C.: International Monetary Fund.

Baldacci, E., Mello, L. D., & Inchauste, G. (2002). Financial crises, poverty, and income distribution. *IMF Working Paper.* Washington D.C.: International Monetary Fund.

Baldacci, E., Clements, B., Gupta, S., & Cui, Q. (2008). Social spending, human capital, and growth in developing countries. *World Development, 36*(8), 1317-1341.

Bastagli, F., Coady, D., & Gupta, S. (2012). Income inequality and fiscal policy. *IMF Staff Discussion Note 12/08.* Washington D.C.: International Monetary Fund.

Bergh, A., & Nilsson, T. (2011). Globalisation and absolute poverty: a panel data study. *IFN Working Paper, 862*. Stockholm: Research Institute for Industrial Economics.

Borenstein, M., Hedges, L. V., Higgins J. P. T., & Rothstein, H.R. (2009). *Introduction to Meta-Analysis.* Chichester: Wiley.

Bourguignon, F., & Fields, G. (1997). Discontinuous losses from poverty, generalized P α measures, and optimal transfers to the poor. *Journal of Public Economics*, *63*(2), 155-175.

Carmignani, F. (2011). The Making of Pro-Poor Growth. *Scottish Journal of Political Economy, 58***,** 656-684.

Castro-Leal, F., Dayton, J., Demery, L., & Mehra, K. (1999). Public spending on health care in Africa: do the poor benefit? *World Bank Research Observer, 14*(1), 49-72.

Chen, S., & Ravallion, M. (2010). The developing world is poorer than we thought, but not less successful in the fight against poverty. *The Quarterly Journal of Economics*, *125*(4), 1557-1625.

Coady, D., Grosh, M., & Hoddinott, J. (2006). *Targeting of Transfers in Developing Countries*. Washington D.C.: World Bank.

Cohen, J. (1988). *Statistical Power Analysis for the Behavioural Sciences. 2nd ed.* Hillsdale, NJ: Erlbaum Associates.

Čok, M., Urban, I., & Verbič, M. (2013). Income redistribution through taxes and social benefits: The case of Slovenia and Croatia. *Panoeconomicus, 60***,** 667-686.

Cox, D., & Jiminez, E. (1995). Private transfers and the effectiveness of public income redistribution in the Philippines. In D. van de Walle & K. Nead (Eds.), *Public Spending and the Poor: Theory and Evidence*. Baltimore: John Hopkins University Press.

Cox, D., Hansen, D., & Jimenez, E. (2004). How responsive are private transfers to income? Evidence from a laissez-faire economy. *Journal of Public Economics, 88*(9-10), 2193-2219.

Datt, G., & Ravallion, M. (1992). Growth and redistribution components of changes in poverty measures: a decomposition with applications to Brazil and India in the 1990s. *Journal of Development Economics, 38*, 275-295.

Davoodi, H., Tiongson, E., & Asawanuchit, S. (2010). Benefit Incidence of Public Education and Health Spending Worldwide: Evidence from a New Database. *Poverty and Public Policy, 2*(2), 5-52.

Demery, L. (2000). Analysing the incidence of public spending. In F. Bourguignon & L. Pereira (Eds.), *The Impact of Economic Policies on Poverty and Income Distribution*. New York: World Bank & Oxford University Press.

Disdier, A. C., & Head, K. (2008). The puzzling persistence of the distance effect on bilateral trade. *Review of Economics and Statistics, 90*(1), 37-48.

Doucouliagos, C., & Paldam, M. (2008). Aid effectiveness on growth: a meta-study. *European Journal of Political Economy, 24*, 1-24.

Doucouliagos, C., & Paldam, M. (2009). The aid effectiveness literature: the sad results of 40 years of research. *Journal of Economic Surveys*, *23*(3), 433-461.

Doucouliagos, C., & Paldam, M. (2015). Finally a breakthrough? The recent rise in the size of the estimates of aid effectiveness. In Arvin, B. M. and Lew, B. (eds), *Handbook on the Economics of Foreign Aid*. Edward Elgar: Cheltenham.

Easterly, W., & Fischer, S. (2001). Inflation and the poor. *Journal of Money, Credit and Banking, 33*(2), 160-178.

Egger, M. D., Smith, G., Schneider, M., & Minder, C. (1997). Bias in meta-analysis detected by a simple, graphical test. *BMJ*, *315*, 29–634.

Goni, E., Lopez, J. H., & Serven, L. (2011). Fiscal Redistribution and Income Inequality in Latin America. *World Development*, *39*(9), 1558–1569.

Gupta, S., & Verhoeven, M. (2001). The efficiency of government expenditure: experiences from Africa. *Journal of Policy Modeling, 23*, 433-467.

Gupta, S., Verhoeven, M., & Tiongson, E. (2002). The effectiveness of government spending on education and health care in developing and transition economies. *European Journal of Political Economy, 18*, 717-737.

Gupta, S., Verhoeven, M., & Tiongson, E. (2003). Public spending on health care and the poor. *Health Economics, 12*, 685-696.

Higgins J. P. T., & Green S. (Eds.). (2011). *Cochrane Handbook for Systematic Reviews of Interventions Version 5.1.0 [updated March 2011].* The Cochrane Collaboration. Available from www.handbook.cochrane.org.

Hosmer, D. W., & Lemeshow, S. (2000). *Applied logistic regression. 2nd ed.* New York: John Wiley & Sons, Inc.

International Monetary Fund (IMF). (2014). Fiscal policy and income inequality. *Policy Paper*. Washington D.C.: International Monetary Fund.

Iršová, Z., & Havránek, T. (2013). Determinants of Horizontal Spillovers from FDI: Evidence from a Large Meta-Analysis. *World Development*, *42*, 1-15.

Kalwij, A., & Verschoor, A. (2007). Not by growth alone: the role of the distribution of income in regional diversity in poverty reduction*, European Economic Review*, *51*, 805-829.

Kaufmann, D., A. Kraay and M. Mastruzzi (2004). "Governance Matters III: Governance Indicators for 1996, 1998, 2000, and 2002". *World Bank Economic Review* 18:253-287.

Klasen, S. (2003). In search of the Holy Grail: how to achieve pro-poor growth. In B. Tungodden & N. Stern (Eds.), *Towards Pro-Poor Policies. Proceedings from the Annual Bank Conference on Development Economics – Europe*. Washington D.C.: World Bank.

Kraay, A. (2006). When is growth pro-poor? Evidence from a panel of countries. *Journal of Development Economics*, *80*, 198-227.

Kwon, H. - J., & Kim, E. (2014). Poverty Reduction and Good Governance: Examining the Rationale of the Millennium Development Goals. *Development and Change, 45***,** 353-375.

Lakner, C., Negre, M., & Prydz, E. (2014). Twinning the goals: how can promoting shared prosperity help to reduce global poverty? *Policy Research Working Paper, 7106*. Washington D.C.: World Bank.

Lelkes, O., & Sutherland, H. (Eds.). (2009). *Tax and benefit policies in the enlarged Europe: assessing the impact with microsimulation models.* Ashgate: Farnham.

Leonard, de Linde M., Stanley, T. D., & Doucouliagos, H. (2014). Does the UK Minimum Wage Reduce Employment? A Meta-Regression Analysis. *British Journal of Industrial Relations*, *52*(3), 499-520.

Lipsey, M. P., & Wilson, D. (2001). *Practical Meta-Analysis*. London: Sage Publications.

Lustig, N., Pessino, C., & Scott, J. (2014). The impact of taxes and social spending on inequality and poverty in Argentina, Bolivia, Brazil, Mexico, Peru and Uruguay. *Public Finance Review* 42(3): 287-303.

Lustig, N. (2015). The redistributive impact of government spending on education and health: evidence from 13 developing countries in the Commitment to Equity Project. In Clements, B., Mooij, R., Gupta, S. and Keen, M. (eds) *Inequality and Fiscal Policy*. Washington D.C.: International Monetary Fund.

McKay, A. (2004). Assessing the impact of fiscal policy on poverty. In T. Addison & A. Roe (Eds.), *Fiscal Policy for Development: Poverty, Reconstruction and Growth*. Basingstoke: Palgrave Macmillan.

Mosley, P., Hudson, J., & Verschoor, A. (2004). Aid, poverty reduction and the 'new conditionality'. *Economic Journal, 113*, F217-F243.

Organisation for Economic Cooperation and Development (OECD) (2011). *Divided We Stand: Why Inequality Keeps Rising*. Paris: OECD.

Paternostro, S., Rajaram, A., & Tiongson, E. (2007). How does the composition of public expenditure matter? *Oxford Development Studies, 35*(1), 47-82.

Rajkumar, A.S., & Swaroop, V. (2008). Public spending and outcomes: does governance matter? *Journal of Development Economics, 86*, 96-111.

Ravallion, M. (2001). Growth, inequality and poverty: looking beyond averages. *World Development* 29 (11): 1803-1815.

Rhee, C., Zhuang, J., Kanbur, R., & Felipe, J. (2014). Confronting Asia's rising inequality: policy options. In R. Kanbur, C. Rhee & J. Zhuang (Eds.), *Inequality in Asia and the Pacific: Trends, Drivers and Policy Implica0tions*. London: Routledge.

Sahn, D. E., & Alderman, H. (1996). The effect of food subsidies on labor supply in Sri Lanka. *Economic Development and Cultural Change, 45*(1), 125-145.

Sala-i-Martin, X. (2006). The world distribution of income: falling poverty and … convergence, period. *Quarterly Journal of Economics* 121 (2): 351-397.

Son, H. and Kakwani, N. (2008). Global estimates of pro-poor growth. *World Development* 36 (6): 1048-1066.

Stanley, T. D. (2005). Beyond publication selection. *Journal of Economic Surveys, 19*, 309–345.

Stanley, T. D. (2008). Meta-regression methods for detecting and estimating empirical effects in the presence of publication bias. *Oxford Bulletin of Economics and Statistics, 70*, 103–127.

Stanley, T. D., & Doucouliagos, H. (2012). *Meta-regression Analysis in Economics and Business.* Abingdon: Routledge.

Stanley, T. D., Doucouliagos, H., Giles, M., Heckemeyer, J.H., Johnston, R.J., Laroche, P., Pugh, G., Nelson, J., Paldam, M., Poot, J., Rosenberger, R., & Rost, K. (2013). Meta-analysis of economics research reporting guidelines. *Journal of Economic Surveys, 27*(2), 390–394.

Stanley, T. D., & Doucouliagos, H. (2017). Neither Fixed nor Random: Weighted Least Squares Meta-Regression. *Research Synthesis Methods* 8(1): 19-42.

Verschoor, A., & Kalwij, A. (2006). Aid, social policies and pro-poor growth. *Journal of International Development, 18***,** 519-532.

Wagle, U. R. (2012). The economic footing of the global poor, 1980–2005: The roles of economic growth, openness and political institutions. *Journal of International Development, 197***,** 173-197.

Wooldridge, J. (2010). *Econometric Analysis of Cross Section and Panel Data. 2nd ed*. London: MIT Press.

World Bank (2015a). *Global Monitoring Report 2014/15: Ending Poverty and Sharing Prosperity*. Washington D.C.: World Bank.

World Bank (2015b). *A Measured Approach to Ending Poverty and Boosting Shared Prosperity: Concepts, Data and the Twin Goals*. Washington D.C.: World Bank.

Yoshida, N., Uematsu, H., & Sobrado, C. (2014). Is extreme poverty going to end? An analytical framework to evaluate progress in ending extreme poverty. *Policy Research Working Paper, 6740.* Washington D.C.: World Bank.

[1] The 'redistribution' component of poverty reduction over a given period of time is the amount of poverty reduction resulting from the change in income distribution over the period, holding the mean level of income constant. The 'growth' component is the amount of poverty reduction resulting from the change in the mean level of income, holding the distribution of income constant (see Kraay 2006; Datt and Ravallion 1992).

[2] Note that government spending on basic health and education services does not have direct, first-round impacts on income poverty, in the same way that taxes and transfers do. This is because most measures of poverty used in the literature are based on income measures which reflect a household's purchasing power over private goods and services, but not publicly provided goods and services such as health or education. The difficulties of including the latter in a broader measure of income are typically considered too great; see for example Chen and Ravallion (2010: 1591).

[3] One slight variation of the above approach is provided by Kraay (2006). This involves first decomposing changes in income poverty into their growth and redistribution components (see footnote 1 above). It then regresses the size of each component on a set of explanatory variables, including a measure of government spending. We include the results from this study, given the clear parallels with studies looking at overall changes in poverty – although we found no other studies using a similar approach.

[4] Note that income poverty is sometimes measured by combining information on the distribution of income from household surveys with data on average income from the National Accounts data; see for example Sala-i-Martin (2006). We include studies using poverty estimates derived in this way, although they make up a relatively small proportion of the total.

[5] The poverty headcount (or headcount ratio) is the proportion of the population that lives below the poverty line, while the poverty gap measures the extent to which individuals fall below the poverty line, as a proportion of the poverty line. The squared poverty gap is the average of the squared values of the poverty gap, and unlike the other two measures is affected by changes in inequality among the poor.

[6] A full list of databases searched, and the precise search terms used, is available in Anderson et al (2016b).

[7] The wider systematic review from which the studies analysed in this paper were taken focused on the effect of government policy on the translation of economic growth into poverty reduction, as opposed to poverty reduction per se. Our search terms therefore included the keyword 'growth' as well as 'poverty'. It is possible therefore that we would have identified more studies meeting the inclusion criteria for this particular analysis had we omitted the additional keyword 'growth', although our additional hand-searches was designed to pick up any relevant studies not captured by the electronic searches.

[8] Of the 11,789 records that were screened in total, 10,021 records were excluded following screening by title and abstract; the remaining 1,768 records were all screened by full text.

[9] Note that we found several working papers which were subsequently published as journal articles. We do not treat such papers as separate studies; instead, we treat them as 'linked papers', part of a single study. We also found one study that met our inclusion criteria but which did not report the estimated coefficient on the government spending measure, nor its standard error, t-statistic or p-value; the coefficient is instead simply listed as being negative and not statistically significant. This study (Bergh and Nilsson 2011) was therefore excluded from the analysis.

[10] Unless stated otherwise, all weighted averages in this paper use the inverse variance of the partial correlation coefficient as weights; all standard errors are cluster-robust standard errors (clustering by study).

[11] The standard error of the partial correlation coefficient is calculated by dividing the partial correlation coefficient by the t-statistic of the regression coefficient. Note that the standard error of the partial correlation coefficient differs from the standard error of the regression coefficient.

[12] We experimented with including an additional dummy variable in the model for this category, but it was not statistically significant and made little difference to the results. We therefore excluded it.

[13] The most common measures of institutional quality used are the Worldwide Governance Indicators (Kaufmann et al 2004): voice and accountability, control of corruption, government effectiveness, regulatory quality, rule of law, and political stability. However, only one study included all of these indicators in the regression analysis (Kwon and Kim 2014); the others used different combinations of

them, or an aggregate governance index reflecting all six. This meant that we were not able to test for the effect of any one specific governance indicator (e.g. rule of law, control of corruption) on the estimated relationship between government spending and poverty, due to a small number of observations in each case. For example, we had only 5 estimates (from one study) for regressions which control specifically for the control of corruption, and only 14 estimates (from four studies) for regressions which control specifically for voice and accountability.

[14] By unpublished we mean that a study was not published as a peer-reviewed journal article or book chapter. If a study was initially brought out as a working paper, before being published as a journal article, we include the estimates from the journal article. We only include the results from the working paper if additional or different to those in the journal article; in such cases, the working paper results are classified as unpublished, while the journal article results are classified as published. However, we found only one study where there was a difference in results between the working paper and the journal article (Kraay 2006), and the differences were small in size.

[15] The data used for the meta-regression analysis as well as the corresponding STATA do files are available from the authors on request.

[16] Note however that we find little evidence that published studies are more likely to contain statistically significant estimates than unpublished studies: the proportion of statistically significant coefficients is 52% for the published studies compared to 49% for the unpublished studies. This suggests that the difference in results between published and unpublished studies shown in Table 3 is due to factors other than selection by academic journals according to statistical significance.

[17] This does not of course deny that health and education spending plays an important role in other ways, for example in terms of raising overall health and educational attainments, for the population as a whole, and for poor households in particular (e.g. Gupta et al 2002, 2003).

[18] The main reason for not including studies including interaction terms in the MRA is that the partial correlation coefficient cannot be calculated in such cases.