

How Large Is the Government Spending Multiplier?

Evidence from World Bank Lending

Aart Kraay

The World Bank
Development Research Group
Macroeconomics and Growth Team
December 2010



Abstract

This paper proposes a novel method of isolating fluctuations in public spending that are likely to be uncorrelated with contemporaneous macroeconomic shocks and can be used to estimate government spending multipliers. The approach relies on two features unique to many low-income countries: (1) borrowing from the World Bank finances a substantial fraction of public spending, and (2) actual spending on World Bank-financed projects is typically spread out over several years following the original approval of the project. These two features imply that fluctuations in spending on World

Bank projects in a given year are in large part determined by fluctuations in project approval decisions made in previous years, and so are unlikely to be correlated with shocks to output in the current year. World Bank project-level disbursement data are used to isolate the component of public spending associated with project approvals from previous years, which in turn can be used to estimate government spending multipliers, in a sample of 29 aid-dependent low-income countries. The estimated multipliers are small, reasonably precisely estimated, and rarely significantly different from zero.

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

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

How Large Is the Government Spending Multiplier? Evidence from World Bank Lending

Aart Kraay
The World Bank

1818 H Street NW, Washington, DC 20433, USA, akraay@worldbank.org. I would like to thank Robert Barro, Michael Clemens, Joseph Formosa, Constantino Hevia, Ethan Ilzetzi, Justin Lin, Claudio Raddatz, Martin Ravallion, Luis Servén, and Susan Yang for helpful discussions. Javier Morales Sarriera provided exceptional research assistance. Financial support from the Knowledge for Change Program of the World Bank is gratefully acknowledged. The views expressed here are the author's, and do not reflect those of the World Bank, its Executive Directors, or the countries they represent.

1. Introduction

Empirically identifying fiscal multipliers requires a strategy to isolate changes in public spending and/or taxes that are plausibly uncorrelated with contemporaneous economic shocks. In this paper I propose a novel method of identifying such fluctuations in public spending that relies on two features unique to many low-income countries: (1) borrowing from the World Bank finances a substantial fraction of public spending, and (2) actual spending on World Bank-financed projects is typically spread out over several years following the original approval of the loan. The first fact means that fluctuations in spending on World Bank-financed projects are a significant source of fluctuations in overall public spending in these countries. The second fact means that fluctuations in World Bank-financed spending in a given year are largely determined by fluctuations in project approval decisions made in previous years, and thus are unlikely to be correlated with shocks to output in the current year. I use World Bank project-level disbursement data to isolate this component of public spending associated with past project approval decisions, and use it to estimate government spending multipliers in a sample of 29 mostly low-income countries where this source of fluctuations in public spending is large relative to the size of the economy.

The recent financial crisis has renewed interest in the long-standing question of the size of fiscal multipliers. Knowledge of the size of the multiplier is crucial to informing policy discussions about the appropriate scale and duration of fiscal stimulus packages in response to macroeconomic crises. Years of intensive and creative research have, however, yielded a bewildering array of estimates of the multiplier, ranging from zero and even negative to well above one. Nearly all of this evidence comes from a handful of developed economies, and is based on one of three primary identification strategies. The first are VAR-based identification schemes, of which Blanchard and Perotti (2002) is a leading example. These studies rely on the availability of quarterly data, together with the assumption that discretionary changes in fiscal policy take sufficiently long to implement that they cannot react to contemporaneous economic activity within a quarter.¹

¹ This identification strategy is infeasible in the majority of developing countries, and especially in the poorest low-income countries that are the focus of this paper, as most do not report fiscal or macro data on a quarterly basis. For a sample of 27 middle-income countries, Ilzetki and Végh (2008) and Ilzetki, Mendoza and Végh (2010) are able to assemble quarterly data in order to analyze the cyclical effects of fiscal policy in these countries using standard VAR-based identification strategies that have been applied to industrial countries. However, there is no overlap between their sample of emerging market economies with available quarterly data and my sample of low-income aid-dependent countries.

A second strategy consists of finding an external instrument that generates fluctuations in spending that are unlikely to be correlated with contemporaneous macroeconomic events. For example, Cohen, Covall and Malloy (2010) use changes in Congressional committee chairmanships to identify changes in federal spending at the state level in the United States that are driven by national-level electoral outcomes. They find evidence that these spending changes are negatively correlated with private investment and employment at the state level. Fishback and Kachanovskaya (2010) also study the state-level effects of federal spending, but focus on the New Deal era. They use a measure of swing voting behaviour as an instrument for public spending and find output multipliers ranging from 0.9 to 1.7 depending on the type of spending, although no appreciable impact on employment.

The third identification strategy consists of isolating a subcomponent of spending or taxes that arguably is uncorrelated with contemporaneous economic shocks. For example Barro (1981), Ramey and Shapiro (1998), Ramey (2009), Hall (2009), Fisher and Peters (2010) and Barro and Redlick (2010) all argue that changes in US military expenditures during major wars can be thought of in this way. This argument relies on the fact that these conflicts occurred outside the US (so that there was no direct effect of conflict on the US economy), and that their timing was determined by geopolitical factors unrelated to US macroeconomic fluctuations. In the same spirit, but on the tax side, Romer and Romer (2010) develop a narrative description of the rationale for individual tax policy changes in the US, and use this to distinguish between those changes that were taken for countercyclical purposes and those that were motivated by other considerations, such as claimed benefits for long-run growth, or for ideological reasons. They then argue that the latter subset of tax policy changes are unlikely to be correlated with contemporaneous macroeconomic shocks and thus can be used to estimate tax multipliers.

My approach falls in this third category of identification strategies, as it also consists of identifying a subcomponent of public spending whose fluctuations are plausibly uncorrelated with contemporaneous macroeconomic events. In my case, this consists of changes in government spending in a given year that are attributable to changes in World Bank project approval decisions made in previous years. My basic identifying assumption is that project approval decisions made in previous years do not react to shocks to growth in the current year, and so current-year disbursements on projects approved in previous years are unlikely to be correlated with current-year shocks to growth. Using these fluctuations in spending associated with project approval decisions from previous years, I consistently find point estimates of the government spending multiplier that are small and sometimes

even negative, depending on the specification and methodology used. Moreover, the estimates of the multiplier are reasonably precise, with standard errors that are comparable to those obtained in studies of US military expenditures, as well as VAR-based estimates. Across the various specifications and methodologies I consider, I consistently do not reject the null hypothesis that the multiplier is zero, and I can usually also reject the null hypothesis that the multiplier is equal to one.

There are numerous possible objections to this basic identification strategy. One immediate concern is that past project approval decisions are in fact correlated with contemporaneous shocks, because these shocks are persistent or otherwise predictable in some way. I address this concern by controlling for lagged growth, and by allowing for longer lags between project approval and actual disbursements. Another concern is that the timing of individual disbursements on World Bank projects approved in previous years is driven by contemporaneous shocks, even if the project approval decisions are not. I address this concern by constructing an alternative artificial measure of disbursements based on typical disbursement rates for similar projects in other countries that, by construction, do not reflect domestic shocks. Another potential issue with my identification strategy is that World Bank-financed spending might crowd in or crowd out other forms of public spending, and so bias estimates of the overall government spending multiplier. To deal with this issue, in my core results I use changes in World Bank-financed spending as an instrument for changes in total government spending in order to estimate spending multipliers.

Several caveats about the evidence in this paper are worth acknowledging at the outset. First, this paper shares with much of the literature the difficulty that empirically-estimated government spending multipliers are not “deep” structural parameters but rather reflect the confluence of a wide variety of factors including preferences, technology, the nature of spending, and the eventual taxes that finance them. For this reason the term “multiplier” is perhaps best understood simply as short-hand for the empirical correlation between plausibly exogenous changes in government spending and changes in output. Second, my results by design are based on data for a particular set of poor and/or small economies where World Bank lending is an important source of financing for public expenditure, and the effects of government spending on output in other countries outside my sample might very well be different. Finally, I emphasize that the empirical work here is designed only to assess the the short-run impact on output of changes in public spending associated with changes in disbursements on World Bank loans. This question is of course related to, but distinct from, that of the long-run growth impacts of foreign assistance more generally, which has been debated endlessly in the vast empirical aid-growth

literature. With respect to this literature, I claim nothing more than that the short-run multipliers estimated here are potentially consistent with a wide variety of long-run estimated impacts of World Bank lending in particular, or aid in general, on growth.

Section 2 of the paper presents the empirical framework I use to estimate the government spending multiplier. Section 3 describes the project-level disbursement data I use to construct alternative measures of fluctuations in World Bank-financed spending that arguably are uncorrelated with contemporaneous macroeconomic shocks. Section 4 contains my core estimates of the multiplier and subjects them to a variety of robustness checks. Section 5 explores several hypotheses as to why the estimated government spending multiplier is so small, and Section 6 concludes.

2. Empirical Framework and Identification Strategy

I consider variants on the following minimal empirical framework that can be used to quantify the short-run cyclical effects of government spending on output:

$$(1) \quad \frac{y_t - y_{t-1}}{y_{t-1}} = \alpha + \beta \frac{g_t - g_{t-1}}{y_{t-1}} + \varepsilon_t$$

Here, y_t denotes real GDP (measured in constant local currency units); g_t denotes total government spending; and ε_t denotes all other sources of GDP fluctuations, such as other fiscal or monetary policy changes, terms of trade shocks, changes in productivity, natural disasters, and many other shocks. Out of necessity, all data are measured at annual frequency, given the unavailability of quarterly data in the sample of low-income countries I will be working with. Although later I will be combining information from multiple countries, for notational convenience I suppress country subscripts. The key parameter of interest is β which captures the government spending multiplier, i.e. the contemporaneous change in output due to a change in government spending. If an additional dollar of government spending does lead to reductions in any of the other expenditure components of GDP, the multiplier would be one. As noted in the introduction, this is not a deep structural parameter, but rather simply an empirical summary of the contemporaneous relationship between government spending changes and output fluctuations.

For notational convenience, let Δx_t denote the deviation of $\frac{x_t - x_{t-1}}{y_{t-1}}$ from its mean for any variable x , so that Equation (1) can be re-written as:

$$(2) \quad \Delta y_t = \beta \Delta g_t + \varepsilon_t$$

The standard difficulty in identifying the government spending multiplier is that changes in government spending are likely to be correlated with other contemporaneous shocks to output, i.e. $E[\Delta g_t \varepsilon_t] \neq 0$, and so OLS estimation of Equation (2) will lead to biased estimates of the multiplier. In developed countries, where automatic stabilizers are important and governments are able to borrow to finance countercyclical increases in spending, it is plausible to think that $E[\Delta g_t \varepsilon_t] < 0$, so that OLS estimates of the multiplier would be biased downwards by virtue of the fact that government spending increases endogenously during downturns. In contrast, in many developing countries with limited automatic stabilizers, and where governments have limited access to finance, the more likely concern is that government spending is procyclical, i.e. $E[\Delta g_t \varepsilon_t] > 0$, and so OLS estimates of the multiplier would be biased upwards.²

The high-frequency VAR-based approach to identification hinges on the assumption that $E[\Delta g_t \varepsilon_t] = 0$ when the data is observed at quarterly frequency (conditional on the lags that are included in the VAR, and after netting out in some way the effects of automatic stabilizers). The rationale for this assumption is that discretionary fiscal policy changes take sufficiently long to implement that they cannot react to economic activity within a quarter. The instrumental variables approach to identification involves finding some external source of variation in public spending that arguably is uncorrelated with the error term in Equation (2), often due to political factors.

The approach to identification taken here involves isolating a subcomponent of public spending for which it is reasonable to believe that its fluctuations are uncorrelated with the error term. In particular, let $g_t = g_{1t} + g_{2t}$ be a decomposition of total public spending for which it is likely that $E[\Delta g_{1t} \varepsilon_t] = 0$. For lack of a better term, I will refer to changes in this subcomponent of total spending as exogenous to contemporaneous macroeconomic shocks. The US-based military expenditure approach to identification labels g_{1t} as military expenditures. In that case, the identifying assumption that $E[\Delta g_{1t} \varepsilon_t] = 0$ can be supported by the observations that (1) the major conflicts occurred outside

² See Ilzetzki and Végh (2008) for extensive empirical evidence suggesting that fiscal policy is indeed procyclical in a sample of 27 mostly middle-income countries for which they were able to assemble quarterly macro and fiscal data. Fatás and Mihov (2002) also present empirical evidence for procyclical government spending in a large cross-section of countries using annual data.

the US, so that there was no direct adverse effect of wartime destruction on the US economy; and (2) their timing was driven by geopolitical considerations orthogonal to fluctuations in the US economy. Since changes in military and non-military spending in the US are more or less uncorrelated, a simple OLS regression of Δy_t on Δg_{1t} will then deliver a consistent estimate of the multiplier.

In this paper I adopt the same general approach of isolating a plausibly exogenous subcomponent of public spending, although in a very different setting. In particular, I argue that fluctuations in spending in a given year that are associated with World Bank projects approved in previous years are unlikely to be correlated with contemporaneous macroeconomic shocks. This is because these disbursements in a given year primarily reflect project approval decisions made in previous years, before current-year macroeconomic shocks are known. To make this case, some institutional background is useful. The lending activities of the World Bank are organized by project. A project typically consists of an agreement between the World Bank and a developing country to engage in some kind of public spending, financed by loans provided by the World Bank. For example, a project might consist of an agreement to build a particular infrastructure project, or to fund a teacher training project, or to support a particular health intervention, or a myriad of other potential development-oriented government actions that the World Bank finances. In some cases, the project simply provides general budget support, and the associated spending priorities are then chosen by the recipient government.

Projects are identified through a consultative process between World Bank staff and the government of the country in which the project is to be implemented. Crucially for my purposes, these projects typically are designed to be carried out over several years. A document describing the project is prepared by World Bank staff, and includes a proposed amount of World Bank funding for the entire project, together with a timeline of planned expenditures over the life of the project. The project is then approved by the Board of Executive Directors of the World Bank. Once the project is approved, it is implemented over time, with spending on the project financed by disbursements on World Bank loans. Thus, in any given year, total spending financed by the World Bank reflects project approval decisions and spending plans made in many previous years as well as in the current year. My approach to identification consists of isolating the part of spending in each year that reflects project approval decisions from previous years. In particular, for each country and year, I measure g_{1t} as disbursements on World Bank projects approved in previous years, but not in the current year. I refer to this as *disbursements on previously-approved projects*, as opposed to *total disbursements* which include

disbursements on projects approved in the same year, and hence are potentially correlated with contemporaneous events.

There are (at least?) three immediate potential objections to this basic identifying assumption. The first is that while project approval decisions from previous years were made prior to the realization of current macroeconomic shocks, the latter may have been predicted by World Bank decision-makers at the time of project approval. For example, World Bank decision-makers may have good information on which to base forecasts of growth in future years, and tailor project approval decisions to anticipated future growth shocks. Or more simply, if shocks to growth are serially correlated, then a project approved in response to a contemporaneous macroeconomic shock will also be correlated with future macroeconomic shocks. I address this first concern in two ways. First, I construct an alternative measure of disbursements on previously-approved projects that excludes disbursements on projects approved not only in the current year, but also in the previous year.³ Second, as a robustness check I control for lagged growth in my regressions. This is a natural and direct way of controlling for a large variety of macroeconomic shocks that might be persistent over time, and to the extent that they contemporaneously influence World Bank project approval decisions, would undermine my identification strategy.

The second potential objection to my identification strategy is based on the observation that actual disbursements on World Bank projects do not always unfold as originally planned at the time of project approval. Deviations from initially-planned disbursements reflect a wide range of factors, including unforeseen technical problems in the implementation of the project, procurement delays, unexpected delays to investigate possible financial irregularities in the project, and many other considerations. As long as these factors are uncorrelated with contemporaneous macroeconomic shocks, deviations from planned disbursements do not undermine my identification strategy.

However it is also possible that deviations from planned disbursements are in fact correlated with contemporaneous shocks, although the direction of this correlation is ambiguous. It could be that disbursements on projects approved in previous years are accelerated in response to an adverse shock in the current year, as a way for the World Bank to deliver resources quickly to countries affected by negative shocks. On the other hand, adverse macroeconomic shocks could make project implementation more difficult, and thus lead to a reduction in disbursements relative to original plans.

³ Although not reported for reasons of space, results are also similar if I exclude disbursements on projects approved in the current and two previous years.

Yet another possibility is that subsequent disbursements on a project are triggered by the recipient government meeting various conditions for policy improvements. If the latter lead to better aggregate growth performance, these subsequent disbursements would be spuriously positively correlated with growth through this channel of policy conditionality. Any of these possibilities would undermine my key identifying assumption that disbursements are uncorrelated with contemporaneous shocks simply because they are associated with projects approved in previous years.

I address this second concern with the identification strategy in two ways as well. First, I construct an artificial set of disbursements for each project based on typical rather than actual disbursement profiles. Specifically, for each project, I construct artificial predicted disbursements for each year of the life of the project as the total size of the project, multiplied by the average disbursement rate in the same year for all projects in the same sector, region, and approval year. I then aggregate all these predicted project-level disbursements up to the country-year level, again excluding disbursements on projects approved in the same year. This measure of disbursements by construction now reflects only project approval decisions at the country level, but not country-level deviations from planned disbursements.⁴ Second, to the extent that disbursements are triggered by successful policy reforms that in turn cause higher growth, a straightforward solution is to control directly for empirical proxies for such reforms, as I do below.

A third concern with my identification strategy is that changes in World Bank-financed spending, Δg_{1t} , may be correlated with changes in non-World Bank financed spending, Δg_{2t} . A perennial concern among aid donors is the extent to which aid-financed expenditures are additional to, or alternatively supplant, other forms of public spending. In some cases, World Bank-financed projects involve cofinancing by other donors or by the recipient government, suggesting a positive correlation between World Bank-financed spending and non-World Bank-financed spending. On the other hand, given that money is fungible, it is also possible that increases in World Bank-financed spending allow recipient governments to cut back on spending in other areas, or induce other aid donors to cut back on their support. Yet another possibility is that disbursements on World Bank loans are simply used to pay

⁴ The region-sector-year averages are based on (a) the World Bank's standard regional groupings (East Asia, South Asia, Middle East and North Africa, Sub-Saharan Africa, and Latin America and the Caribbean), and (b) the following major sector classification of projects (Agriculture and Rural Development, Energy and Mining, Transport, Education, and Other), and (c) each approval year beginning in 1985. An alternative would be to use information on the schedule of disbursements for the project that is projected at the time of project approval to construct these synthetic disbursements. This information exists in project-level documents, but unfortunately is electronically retrievable only for a subset of projects starting in the mid-1990s.

off existing loans from the World Bank or from other creditors, and so lead to no net increases in government spending. If these latter effects dominate the first, the two types of spending will be negatively correlated. In this case, a simple regression of growth on changes in World Bank-financed spending alone will lead to downward-biased estimates of the multiplier, since it will not control for any output effects of the accompanying reductions in other non-World Bank-financed spending.

The natural solution to this problem is to use changes in World Bank-financed spending, Δg_{1t} , as an instrument for changes in total government spending, Δg_t . This implies the following first-stage regression:

$$(3) \quad \Delta g_t = \gamma \Delta g_{1t} + u_t$$

The slope coefficient in the first-stage regression captures how non-World Bank-financed spending responds to World Bank-financed expenditures. In particular, since $\Delta g_t = \Delta g_{1t} + \Delta g_{2t}$, the probability limit of the OLS estimate of γ is $1 + COV[\Delta g_{1t}, \Delta g_{2t}] / V[\Delta g_{1t}]$. This is recognizable as one plus the slope coefficient of a regression of the non-World Bank-financed spending on World Bank-financed spending. If the first-stage slope coefficient is less than (greater than) one, World Bank-financed spending is associated with decreases (increases) in non-World Bank-financed spending.

The reduced-form regression of changes in output on changes in World Bank-financed spending is also of interest. Substituting Equation (3) into Equation (2) gives:

$$(4) \quad \Delta y_t = \beta \gamma \Delta g_{1t} + \varepsilon_t + \beta u_t$$

This expression clarifies that a simple OLS regression of Δy_t on Δg_{1t} will deliver a consistent estimate of the overall government spending multiplier only if $\gamma = 1$, i.e. only if there are no accompanying changes in non-World Bank-financed spending. Note also that by the principle of indirect least squares, the two-stage least squares estimate of the multiplier is simply the ratio of the reduced-form slope, $\beta \gamma$, to the first-stage slope, γ . This shows how the IV estimator corrects for the problem of concurrent changes in the non-World Bank-financed component of government spending.

3. Data

I rely on disbursement data for individual World Bank projects over the period 1985-2009. Over this period I have information on actual quarterly disbursements by project over the life of each project, for the universe of all projects financed by the two main lending arms of the World Bank: non-concessional lending to middle-income countries by the International Bank for Reconstruction and Development (IBRD), and concessional lending to low-income countries through the International Development Association (IDA).⁵ For each project, I sum the quarterly data within each calendar year to arrive at annual disbursement flows for each project. Then for each country and year, I sum across all active projects to obtain *total disbursements* on World Bank loans. I then subtract disbursements on projects approved in the same year to arrive at my measure of *disbursements on previously-approved projects* on World Bank loans. Disbursements on previously-approved projects by construction reflect project approval decisions made in previous years, and my basic identifying assumption is that they are unlikely to be correlated with contemporaneous macroeconomic events. As discussed above, I also construct an alternative measure of disbursements on previously-approved projects that excludes disbursements on projects approved in the current *and* previous year.

Figure 1 shows the fluctuations over time in total disbursements and disbursements on previously-approved projects as a fraction of GDP for one country in my sample, Zambia. The height of the bars shows total disbursements on World Bank loans. These are large in Zambia, averaging 3.4 percent of GDP, and moreover are also very volatile, ranging from close to zero in the late 1980s and late 2000s, to well over five percent of GDP in many other years. The dark-shaded lower portion of each bar isolates my measure of disbursements on previously-approved projects, while the remainder of the bar shows disbursements on projects approved in the same year. In most years, the bulk of disbursements on World Bank loans are associated with projects approved in previous years, and my core identifying assumption is that these can be thought of as plausibly exogenous to contemporaneous macroeconomic shocks. However, in a few years such as 1991 and 1999, there are large disbursements on projects approved in the same year. In 1991, for example, a large (for Zambia) project worth \$210 million was approved, and approximately three-quarters of it was disbursed in the same year. This project was an “Economic Recovery Credit” and was intended to “support economic reforms aimed at

⁵ My dataset in principle covers all 10475 projects approved between the first World Bank project in 1948 and early 2010 when my data stop. However, electronic records on quarterly project-level disbursement flows are not available before 1985, and the cost of manually entering this data from archived paper records is prohibitive. I therefore rely only on data from 6529 projects approved since 1985.

macroeconomic stabilization” following a period of zero and negative GDP growth. These types of disbursements are clearly responding to current macroeconomic events and so cannot be thought of as plausibly exogenous to contemporaneous shocks. This is why I exclude disbursements in a given year on projects approved in the same year from my measure of disbursements on previously-approved projects. However, it is worth noting that this approach is conservative because it also excludes approval-year disbursements on many projects that may have been undertaken for non-cyclical reasons.

A key ingredient in my identification strategy is that there are substantial lags between the approval of a project and the eventual disbursement of all of the funds approved for it. I document this in Figure 2, which reports the average across all projects of the fraction of total spending that is disbursed in year t of the project ($t=0, \dots, 10$ with $t=0$ indicating the year in which the project was approved). For the average World Bank project, just over 12 percent of the original approved amount is disbursed in the year in which the project is approved, and the remaining 88 percent of the total is disbursed over subsequent years. Another way of seeing the importance of disbursement lags is to consider the fraction of total disbursements in a given country-year that is associated with project approval decisions made in previous years. In my core regression sample (described in more detail below), the median across country-years of this ratio is 99 percent. This means that for a typical country-year observation in my sample, 99 percent of total disbursements on World Bank loans are associated with project approval decisions made in previous years. Because of this, even changes in total disbursements are unlikely to be very strongly correlated with contemporaneous events.

In order to obtain meaningfully-precise estimates of the multiplier, it is important that disbursements on World Bank loans are also large relative to borrowing-country GDP. To ensure that this is the case, I focus on a set of aid-dependent countries where (a) annual disbursements on World Bank loans are available for at least 20 of the 25 years between 1985 and 2009, and (b) total disbursements on World Bank loans as a share of GDP averaged over the sample period 1985-2009 exceed one percent. This results in a set of 41 mostly low-income countries where World Bank lending has been an important source of financing for public spending over the past 25 years.

My country sample is further limited by the availability of data on total public spending required to construct Δg_t . My primary source for this data is the IMF’s World Economic Outlook database, which provides information on total general government expenditures, typically beginning in the late 1980s. I supplement this with information taken from the World Bank’s African Development Indicators which also reports data on total general government expenditures, going back to the early 1980s for many

countries. I then drop 12 countries for which there are remaining gaps in the government spending data. This results in a final sample of 29 countries reported in Table 1. Given that I am focusing on the most aid-dependent countries in the world, it is not very surprising that the majority of countries in my sample are located in Sub-Saharan Africa. The only four exceptions are Bolivia, Morocco, Tunisia, and Jordan. For these 29 countries, World Bank spending is large not only as a share of GDP, but also as a share of total government spending. The time-averaged share of World Bank-financed spending in total spending ranges from a low of 3.3 percent in Jordan to a high of 18.5 percent in Uganda, and averages 9.3 percent.

Merely restricting attention to countries that are large recipients of World Bank loans on average is however not enough to ensure that fluctuations in disbursements on World Bank loans are large – it could for example be the case that countries receive a large but steady flow of project approvals, leading to a large but stable flow of disbursements. Fortunately for my purposes, this is not the case, as disbursements on World Bank loans fluctuate significantly over time. Table 2 documents summary statistics on the magnitude of fluctuations in the four measures of disbursements on World Bank projects, Δg_{1t} , as well as total government spending, Δg_t and output growth, Δy_t , pooling all country-year observations in my sample. Fluctuations in total disbursements are quite substantial, with a standard deviation of 1.4 percent of GDP, and of 1.1 percent of GDP for disbursements on previously-approved projects. By way of comparison, the standard deviation of GDP growth rates is 3.9 percent in this sample, and the standard deviation of changes in total government spending is 2.9 percent.

4. Estimates of the Government Spending Multiplier

4.1 Basic Results

My benchmark estimates of the government spending multiplier are reported in Table 3. The top panel reports reduced-form regressions of growth on changes in the four measures of disbursements on World Bank loans: (1) total disbursements, (2) disbursements on previously-approved projects (excluding disbursements on projects approved in the same year), (3) disbursements on previously-approved projects (excluding disbursements on projects approved in the same year *and* the previous year), and (4) predicted disbursements. The estimated slope coefficients in the reduced-form regressions are small in absolute value, ranging from -0.24 to 0.14, depending on the measure of

disbursements used. In all four specifications, I cannot reject the null hypothesis that the simple correlation between growth and changes in World Bank-financed spending are zero.⁶

Although I have argued that changes in World Bank-financed spending are plausibly uncorrelated with contemporaneous shocks, the estimated slope coefficients in the reduced-form regressions will not deliver a consistent estimate of the multiplier because they do not account for concurrent changes in other forms of public spending that may occur when World Bank-financed spending increases. To address this issue, I turn to the first-stage and second-stage regressions described above. The first-stage regressions of changes in total government spending on changes in World Bank-financed spending are reported in the middle panel of Table 3. Crucially for identification purposes, the first-stage regressions are all quite precisely estimated, with first-stage F-statistics ranging from 10.5 to 13.1. In all four cases, these exceed the Staiger and Stock (1997) rule of thumb of 10, indicating that weak instrument pathologies are unlikely to be a concern in the IV regressions that follow in the bottom panel. It is also noteworthy that the estimated slope coefficients are all less than one, and significantly so in all cases except for the final predicted disbursements measure, indicating that total government spending increases less than one-for-one when World Bank-financed spending increases. The magnitude of this effect is non-trivial: for example, using disbursements on previously-approved projects, the estimated first-stage slope indicates that when World Bank-financed spending increases by one dollar, total government spending increases by only 40 cents.

The bottom panel of Table 3 reports the 2SLS estimates of the government spending multiplier. In the first column, I report the OLS estimate of the multiplier to provide a useful benchmark for comparison with the IV results that follow. The OLS estimator delivers a slope of 0.32 that is very strongly significantly different from zero. The remaining columns report the IV estimates of the multiplier. For the two measures of disbursements on previously-approved projects, and the predicted disbursements measure as well, the IV estimates of the multiplier are smaller than the OLS estimates,

⁶ All of the specifications in this and subsequent tables are estimated pooling country-year observations and imposing common slopes and intercepts across countries. An obvious relaxation of this parameter homogeneity assumption is to allow for country fixed effects. Results obtained in this way are nearly identical to those reported here. This reflects the fact that in my sample of countries, most of the variation in annual fluctuations in growth, government spending, and World Bank disbursements, occurs within rather than between countries. Unfortunately, however, the available time series for each country are so short (between 20 and 25 annual observations) that allowing both the intercepts and slopes to vary across countries results in extremely imprecise estimates of country-specific multipliers.

ranging from -0.31 to 0.27. This is consistent with the idea that the OLS estimates are biased up due to procyclicality in overall government spending.

The multipliers in the bottom panel of Table 3 are also reasonably precisely estimated, with standard errors ranging from 0.30 to 0.39. By way of comparison, the standard errors are not too much larger than those reported in Barro and Redlick (2010), who estimate similar specifications using data over the past century for the United States and obtain standard errors for the estimated coefficient on defense spending ranging from 0.06 to 0.27 (their Table 2, first row). They are also similar to those in Blanchard and Perotti (2002) – their Figure 5, for example, reports an impact multiplier of 0.84 with confidence bands that imply a standard error of 0.35. In all four columns I cannot reject the null hypothesis that the multiplier is zero, and in the last two columns, I can also reject the null hypothesis that the multiplier is equal to one at the 95 percent confidence level.⁷

When reading the results in Table 3, it is useful to keep in mind that the four measures of changes in disbursements on World Bank loans become more credibly exogenous as we move from left to right across the table. As discussed above, changes in total disbursements are potentially endogenous to the extent that they include disbursements on projects approved in the same year, possibly motivated by macroeconomic events during the year. Such disbursements are excluded from the remaining three measures of disbursements on previously-approved projects, making them more plausibly exogenous than total disbursements. And finally, the predicted disbursements measure also cleans out potentially endogenous responses of disbursements on previously-approved projects to current macroeconomic shocks by using predicted rather than actual disbursement profiles. This observation aids in the interpretation of the fact that the IV estimates of the multiplier in the bottom panel become smaller as we move to more plausibly exogenous measures of disbursements. For example, the IV estimate of the multiplier is actually slightly larger than the OLS estimate when changes in total disbursements are used as an instrument. A possible explanation for this difference is that changes in total disbursements are *positively* correlated with the macroeconomic shocks, and this failure

⁷ While in the last column of Table 3 the predicted disbursements measure is a generated instrument (consisting of actual project approvals multiplied by estimated average disbursement rates), this does not matter for the asymptotic distribution of the 2SLS estimator as long as actual project approvals in year t are not correlated with macroeconomic shocks in year $t+1$ and higher, as per my core identifying assumption. See Wooldridge (2002) Chapter 6.1.2.

of the exclusion restriction leads to an upwards bias in the IV estimator.⁸ Moving to more credibly exogenous measures of disbursements reduces this upwards bias and leads to successively smaller IV estimates.

The picture that emerges from these first basic results is clear: the government spending multipliers in Table 3 are reasonably-precisely estimated, quite small, and even negative in the case of the most plausibly-exogenous measure of predicted disbursements. In all four specifications, 95 percent confidence intervals around the IV point estimates include zero, and in two cases also exclude one, suggesting a very limited impact effect of government spending on output in the short run. The rest of this section presents several robustness checks on this basic finding, and the next section of the paper explores a variety of hypotheses as to why the estimated spending multipliers are so small.

4.2 Robustness Of Basic Results

Figure 3 provides a visual summary of the benchmark results, using the predicted disbursements measure as an instrument. The three panels of Figure 3 in turn report the reduced-form relationship between growth and changes in disbursements on World Bank loans; the first-stage relationship between changes in total government spending and changes in the World Bank-financed component of public spending; and the structural relationship between growth and changes in total spending. A striking feature of the data is that there are very large fluctuations in all three variables in my sample of low-income countries. In light of this, a natural concern is that the results in Table 3 could be sensitive to a small number of influential observations. I explore this possibility in Table 4, by means of three robustness checks. First, I re-estimate the regressions in Table 3 29 times, dropping one country at a time from the sample. In the top panel of Table 4 I report the minimum and maximum across these 29 samples of the IV slope, IV standard error, and first-stage F-statistic. In the bottom two panels, I consider the robustness of my results to dropping potentially influential individual data points rather than entire countries. I do this by using two standard rules of thumb to identify potentially influential observations in the first-stage and reduced-form OLS regressions, and then re-estimate the reduced-form, first-stage, and second-stage regressions eliminating this set of possibly influential data points. In the middle panel, I use the covariance ratio statistic, which measures changes in the precision of the OLS estimates as individual observations are dropped from the sample, while in the bottom panel, I use the

⁸ This might be the case if there are difficulties in project implementation when countries experience negative growth shocks. An extreme case might be a country that falls into civil conflict, triggering both lower growth and a suspension of World Bank activity.

DFITS measure which captures changes in the OLS slopes as individual observations are dropped from the sample (see Belsely, Kuh and Welsch (1980) for details).

Looking at the top panel of Table 4, it is apparent that dropping individual countries has relatively little impact on my results. In nearly all cases, my first-stage regressions are reasonably strong, with first-stage F-statistics greater than 10. There are some fluctuations in the estimates of the multiplier, but these are moderate, and in no case do I find an estimated multiplier that is significantly different from zero. Removing influential observations, as is done in the bottom two panels of Table 4, affects the results somewhat more. In nearly all cases, the estimated multiplier is smaller when influential observations are dropped from the reduced-form and first-stage regressions. However this difference is in most cases small. Interestingly, in several cases, dropping influential observations strengthens identification relative to the benchmark results: in five out of eight specifications, the first-stage F-statistics are greater than those reported in the default specifications in Table 3. Overall, this first set of robustness checks based on eliminating influential observations is broadly consistent with the benchmark estimates: the multiplier is small, in most cases reasonably precisely estimated, and generally the estimates are smaller as successively more credibly exogenous measures of disbursements are used as instruments.

Table 5 presents three further sets of robustness checks, that are designed to address potential concerns about the exclusion restriction. In the top panel of Table 5, I consider the possibility that shocks to growth are persistent over time. If World Bank project approvals also are correlated with contemporaneous shocks to growth, then subsequent disbursements on these projects would also be correlated with shocks to subsequent growth, in violation of my exclusion restriction. As noted above, the most straightforward way to address this concern is simply to control for lagged GDP growth. In all four specifications, I find that lagged growth is significantly correlated with contemporaneous growth. However, its inclusion has almost no effect at all on my estimates of the multiplier, which are virtually unchanged from those in Table 3. Controlling for lagged growth also does not appreciably weaken identification: the first stage F-statistics for the excluded instrument are above 10 in all four specifications.

Another possible problem is that changes in spending might only affect output growth with a lag. If these lagged changes in total spending are correlated with contemporaneous changes in disbursements, perhaps because disbursements themselves are persistent over time, this too would lead to violations of the exclusion restriction. In the middle panel of Table 5, I also include lagged

government spending to address this possibility. In these specifications, I use current and lagged changes in my measures of disbursements on World Bank loans as instruments. A first caveat is that these richer dynamics are much more weakly identified. In the bottom row of this panel, I report Cragg-Donald statistics which provide a summary of instrument strength in the case of multiple instruments and endogenous variables. Comparing these with the Stock-Yogo critical values reported in the table suggests that there are likely non-trivial size distortions in hypothesis tests based on the usual asymptotic approximations. With this caveat in mind, I do find estimated coefficients on lagged government spending that are positive in all four specifications, suggesting that longer-run multipliers could be larger than the impact multiplier I have been estimating thus far. However, these are quite imprecisely estimated, and I cannot reject the null hypothesis that the estimated coefficient on lagged government spending changes is zero.

A third possible violation of the exclusion restriction might occur if disbursements on existing World Bank projects are triggered by policy reforms that are required as part of the conditionality associated with the project. If policy reforms lead to faster growth, this would induce a spurious correlation between contemporaneous changes in disbursements and growth that is driven by omitted policy reforms. To investigate this possibility, I include a measure of changes in policy as an additional control variable in the bottom panel of Table 5. The specific measure I use is the World Bank's Country Policy and Institutional Assessment (CPIA) ratings, that are produced annually by World Bank country economists for all client countries. These provide a rating on a six-point scale of the quality of policies and institutions, based on a checklist of various policy areas.⁹ Annual changes in the CPIA are significantly positively correlated with changes in output. Consistent with the idea that disbursements respond positively to policy reforms, I find that after controlling for policy changes, the estimated multipliers are slightly smaller than those reported in the benchmark specifications in Table 3. Notably the estimated multipliers are still reasonably strongly identified, with first-stage F-statistics greater than 10 in all four specifications. However, in all four cases the estimated multipliers remain insignificantly different from zero.

⁹ The checklist used for the CPIA ratings has evolved over time. A description of the current format can be found at www.worldbank.org/ida. The CPIA rating process is taken quite seriously, as countries' eligibility for concessional World Bank loans significantly depends on these ratings. This is reflected in an elaborate set of benchmarking and review procedures that are applied throughout the CPIA rating process. The other main virtue of the CPIA data in this context is that it has full coverage of all country-year observations in my dataset.

5. Why Is the Estimated Government Spending Multiplier So Small?

Although the government spending multipliers I have estimated in the previous section are close to -- and insignificantly different from -- zero, simple Keynesian models imply (and many policy discussions assume) that the spending multiplier is at least one. In this section, I investigate several potential explanations for why my estimates of the multiplier are so small relative to this benchmark. I first discuss the scope for attributing the small size of the multiplier to attenuation bias due to measurement error. I then consider certain special features of World Bank lending that may account for a small estimated multiplier, as well as special features of the recipient countries. I also consider in more detail the implications for the interpretation of the estimated multipliers of the fact that the fluctuations in spending financed by World Bank loans, while plausibly exogenous to contemporaneous shocks, are also anticipated by private agents. Finally, I investigate whether the multiplier for World Bank-financed spending is larger than the multiplier for overall government spending.

5.1 Measurement Error

A prosaic -- but potentially important -- reason why the estimated multiplier is so small is simply that there is attenuation bias due to measurement error in government spending. This possibility is particularly relevant in my sample of very poor countries where statistical capacity is weak and national accounts and fiscal data are often of poor quality. There are however two reasons why it is difficult to attribute the small size of the estimated multiplier to the effects of measurement error. The first is simply that measurement error would need to be very large in order to explain the small size of the multiplier. To illustrate this point, it is instructive to ask how much measurement error would be required to overturn my conclusion that the multiplier is significantly less than one. Consider for example the basic OLS regression in the first column of Table 3, which delivered an estimated multiplier of 0.32 with a standard error of 0.06. Holding fixed the estimated standard error, the estimated multiplier would need to be at least 0.88 in order to barely not reject the null hypothesis that it is equal to one at the five percent significance level (i.e. $0.88 + 1.96 * 0.06 = 1$). Thus, if classical measurement error were the culprit for finding a multiplier significantly less than one, it must be sufficiently severe as to create an attenuation bias of $0.32/0.88$. Simple textbook calculations tell us that the ratio of the variance of measurement error to the variance of true government spending would need to be at least 1.75 in order to generate this much of attenuation bias. While this calculation is merely illustrative, it does suggest that appealing to measurement error to explain the small estimated multipliers would

require it to be very severe indeed in order to overturn the conclusion that the multiplier is significantly less than one.

Second, as long as the measurement error in government spending is uncorrelated with changes in disbursements on World Bank loans, the IV estimates of the multiplier are still consistent even in the presence of measurement error in total spending. Moreover, if measurement error were important, then other things equal, we should expect to find IV estimates that are larger than the OLS estimates. Yet in nearly all the specifications we have seen so far, the opposite is true. Taken together these arguments suggest that while measurement error in government spending surely is nontrivial for these countries, it is unlikely that my estimates of the multiplier suffer from major attenuation bias as a result.

5.2 Special Features of World Bank-Financed Spending

The identification strategy in this paper hinges on the correlation between growth and changes in public spending that are associated with changes in disbursements on World Bank loans. This opens the possibility that there are unique features of World Bank-financed spending that might reduce its short-run stimulative effect on output relative to other forms of spending. And since the composition of total spending shifts towards World Bank-financed spending when the latter increases, this may account for why the estimated overall government spending multiplier is so small. One such possibility is that the distribution of World Bank-financed government spending between domestic production and imports is different from that of non-World Bank-financed spending. To take an extreme scenario, if a World Bank-financed project consists exclusively of purchases of machinery or consultancy services imported from abroad, then one might expect the stimulative effects of this spending in the recipient country operating through expansions in domestic aggregate demand to be limited.

To investigate this issue further, I match up the project-level data on disbursements with contract-level data on individual procurement contracts awarded on World Bank projects. The World Bank's procurement records include information on the country of origin of the supplier to which each procurement contract was awarded, as well as, of course, the total value of the contract. I aggregate these individual contracts up to the project level, and construct the share of total procurement contract value associated with a project which is tendered to suppliers situated in the country where the project is located. This is only an imperfect indicator of where disbursements on World Bank loans are spent, for at least three reasons. First, this information is available only as an aggregate for the entire project, and not for the individual year-over-year disbursements. Absent better information, I make the

assumption that the domestic procurement share is the same for all the annual disbursement flows over the life of the project. Second, this information is only available electronically for projects approved since 1990. To address this limitation, I take country-sector averages of the domestic procurement share and apply them to disbursements occurring during the earlier period 1985-1990. Third, knowing the location of the vendor does not necessarily indicate where the disbursements are spent, since I do not have systematic information on what the vendors do. In some cases, foreign vendors might supply specific imported machinery or equipment to a project, while in other cases the vendor might be a firm based abroad who then hires locally to perform the services specified in the contract. Similarly, a domestically-located vendor might very well purchase goods and services both at home and abroad over the course of fulfilling its contract.

With these limitations in mind, I aggregate these project-level estimates of domestically-spent annual disbursement flows up to the country-year level to arrive at the same four measures of disbursements as before, but now excluding the portion awarded to foreign vendors. The domestic share of World Bank spending is substantial: pooling all country-year observations, the median share of total disbursements that is contracted domestically is 48 percent, and the 5th and 95th percentiles of this ratio are 22 percent and 72 percent, respectively. Changes in domestic disbursements are also highly correlated with changes in total disbursements, with a pooled correlation of 0.88. However, changes in domestic disbursements are substantially less volatile than total disbursements, with a standard deviation roughly half as large, and this will result in weaker identification of the multiplier.

I then re-estimate Equation (1), using changes in these alternative measures of domestic disbursements as an instrument for changes in total government spending. The results are reported in the top panel of Table 6. The point estimates of the multiplier are not very different from those in the benchmark specifications in Table 3. However, they are less precisely estimated, as can be seen from the larger second-stage standard errors and smaller first-stage F-statistics. As a result, it is not possible to conclusively identify any significant differences between these multipliers based only on changes in domestic spending and those in the benchmark specifications. Qualitatively, however, this suggests that this compositional effect is likely not important in accounting for the small size of the estimated multiplier.

Another unique feature of spending associated with World Bank projects is that it is financed by loans that are highly concessional. This feature has implications for the standard neoclassical mechanism for a positive government spending multiplier. According to this mechanism, when the

government borrows to finance spending, private agents' wealth is diminished by the present value of future taxes required to repay the debt. This effect is accentuated by any distortionary effects of the eventual taxation required to pay down the debt. Private agents respond by increasing labour supply to restore their desired wealth levels, and output rises. When the governments of the low-income countries in my sample borrow from the World Bank to finance public spending, they primarily borrow from the International Development Association (IDA), the concessional lending arm of the World Bank. Standard IDA credits are zero-interest loans, with 40 year maturities and an initial 10 year grace period. These very favourable terms imply that the present value of future taxes required to pay back an IDA credit are much lower than if the government had borrowed on market terms. For example, taking a discount rate of 5 percent per year, the present value of future repayments associated with \$1 of borrowing from IDA is only 28 cents. In contrast, loans offered by the non-concessional arm of the World Bank, the International Bank for Reconstruction and Development (IBRD), are much closer to being on market terms. If the main neoclassical mechanism for a positive multiplier is operative, this concessionality effect might explain why the estimated multiplier is so small.

A natural way to assess the role of concessionality in accounting for the small estimated multipliers is to re-estimate the multiplier separately using data on disbursements on IDA and IBRD loans. Unfortunately, however, this strategy is not feasible here. As noted above, I focus on a set of low-income countries where World Bank loans finance a substantial fraction of total expenditures. Nearly all of these countries borrow exclusively from IDA – in my sample, on average 88 percent of disbursements are on IDA credits, and so a sample corresponding to disbursements on non-concessional loans is impractically small to yield useful results. A different approach to assessing the role of concessionality is to exploit the fact that, starting in 1996, the World Bank, together with other multilateral lenders, for the first time began to forgive the debts owed to it by many low-income countries through the Heavily-Indebted Poor Countries (HIPC) Initiative. The anticipation of eventual debt relief might reduce even further the present value of future taxes expected by private agents, which in turn would imply that the estimated multiplier should be smaller during the period where there are reasonable expectations of eventual debt relief.

To implement this idea empirically, I return to my basic specification, and re-estimate it separately for the pre- and post-debt relief periods. The results are shown in the bottom two panels of Table 6, and are rather mixed. For three of the four disbursement measures, the IV estimates of the multiplier are actually smaller in the post-debt relief period, although the differences in estimated

multipliers in the two subperiods are small relative to their estimated standard errors and the differences are never statistically significant. A further important qualification is that identification is very weak in the pre-debt-relief period, with first-stage F-statistics well below 10 in all specifications. A final difficulty in interpreting these results is that the simple before- versus after-debt relief comparison conflates any potential impact of increased concessionality with other factors. One such notable factor is changes in world interest rates, which arguably might be used to discount the future tax obligations associated with current borrowing. During the pre-debt relief period 1985-1995, US long-term Treasury Bill rates averaged around 8 percent, while during the post-debt relief period they averaged 5 percent. This alone has sizeable effects on the present value of future taxes associated with IDA borrowing: as noted above, for a 5 percent discount rate, the present value of future taxes associated with one dollar of IDA borrowing is 28 cents, while at an 8 percent discount rate the present value of future taxes falls by nearly half, to 15 cents. This could offset the expected reduction in future taxes due to debt relief, and so obscure any effects of concessionality on the size of the estimated multiplier.

5.3 Special Characteristics of Recipient Countries

Another set of potential explanations for my small estimated multipliers has to do with the characteristics of the low-income countries. A leading concern might simply be that these countries have very weak institutional and governmental capacity, and this in turn weakens the potential for government spending to stimulate output. For example, an extreme case might be that corruption is rampant and World Bank-financed increases in government spending are illegally diverted abroad, with little domestic stimulative impact on output. To explore this possibility, I use the same CPIA measure of policy and institutional quality introduced earlier. Somewhat surprisingly, the countries in my sample are not so different in terms of this measure from other, mostly richer, developing countries outside of my sample. In my regression sample, the median CPIA score is 3.5, which is exactly the median score pooling all country-year observations for all developing countries over the same period since 1985.¹⁰ This basic fact casts doubt on the hypothesis that the estimated multiplier is low in my sample due to poor institutional or policy quality.

Nevertheless, it is plausible that the size of the estimated multiplier does vary with institutional quality. Fortunately, there is a lot of heterogeneity in measured policy and institutional capacity within my sample of 29 countries, which I can exploit to investigate further the hypothesis that the government

¹⁰ This is in part due to the fact that concessional IDA lending is allocated across countries using a formula which strongly rewards countries with better average policy performance as measured by the CPIA.

spending multiplier is larger in countries with better capacity. I do this by simply dividing my sample in half at the median CPIA score, and re-estimate the benchmark specifications in the strong policy and weak policy subsamples. The results are reported in Table 7. Unfortunately, identification is much weaker in the two subsamples, as reflected in much lower first-stage F-statistics and larger standard errors in most specifications. Qualitatively however, a surprising feature of the results is that the estimated multipliers, although very imprecisely estimated, are actually lower in the strong policy subsample in the top panel. While these differences are not statistically significant and should not be overinterpreted, these results nevertheless cast doubt on the hypothesis that the estimated multipliers are small in my sample of low-income countries simply because policy and institutional capacity is weak.

5.4 Role of Anticipation Effects

A fourth possible explanation for the small size of the estimated government spending multiplier has to do with anticipation effects. I identify the multiplier using fluctuations in World Bank-financed spending that, although plausibly uncorrelated with contemporaneous macroeconomic shocks, are also likely to be anticipated in advance by the private sector. This is because project approval decisions are public information, and so the spending plans set in motion by the project approvals are also known at the time of approval. To understand the implications of anticipated spending changes for my results, a useful starting point is a minimal neoclassical model with unproductive government spending financed by lump-sum taxes, of the sort considered by Ramey (2009). Absent anticipation effects, an increase in government spending lowers private wealth on impact by lowering the present value of future after-tax income. In response, consumers compensate for the loss of wealth by supplying more labour, consuming less, and investing more. The increase in labour supply means that output increases on impact. Subsequently, however, labour supply and output decline, and consumption rises, back to steady-state levels.

Using a calibrated model, Ramey (2009) shows that if spending increases are anticipated two quarters in advance of the actual increase in spending, then the standard neoclassical responses (i.e. the increase in labour supply, investment, and output, and the decline in consumption) all occur at the time that increased future government spending is anticipated. However the contemporaneous correlation of changes in these variables with the actual changes in spending once they eventually occur is very different. In particular, Ramey (2009) shows that the change in output contemporaneous with the actual change in spending is much smaller when the spending increase was anticipated in advance, than when it is unanticipated. Moreover, investment falls when the increase in spending occurs, and

consumption increases, which is just the opposite of the immediate impact effect when the spending is announced. The reason is simply that by the time the spending occurs, the initial increase in investment and decline in consumption have already occurred, and now investment (consumption) are declining (increasing) back to their steady-state levels.

At first glance, this provides a good candidate explanation for why my estimated multipliers are so small – it could simply be that most of the private sector response to the increase in spending has already occurred at the time that the original World Bank project was approved and announced, and there are only limited further labour supply and output responses when the spending is actually implemented in subsequent years. To investigate this possibility further, in Table 8 I document the effects of changes in government spending on the major expenditure components of GDP. I do this by re-estimating Equation (1), but replacing the dependent variable in turn with changes in private consumption and investment, government consumption and investment, and net exports. As before, I use the four alternative measures of disbursements on World Bank loans as instruments for total government spending.¹¹

Not surprisingly, I find that government consumption expenditures and government investment increase when total government spending increases – this is true almost as a matter of arithmetic. Of more interest are the responses of private consumption and private investment. These are much less precisely estimated than the effects on output, and so it is difficult to draw very firm conclusions. However, it is noteworthy that in three out of four specifications, the response of private consumption to changes in government spending is negative, while the response of private investment is positive in three of four cases. This pattern of consumption and investment responses to changes in public spending looks much more like the patterns the theory would predict for an *unanticipated* change in spending rather than an anticipated change. And this in turn casts some doubt on the importance of anticipation effects in accounting for the small estimated multipliers in my benchmark specifications.¹²

¹¹ Data on private and government consumption, total investment, and net exports are taken from the national accounts as reported in the World Development Indicators. I use data on total public investment from the IMF's WEO database to separate total investment into public and private investment. Data on the expenditure components of GDP and/or public investment are missing for a handful of observations in my sample. However, the fit of the first-stage regressions of changes in government spending on changes in disbursements on World Bank loans are nearly the same as those reported in Table 3 for the full sample, and so are not repeated for this subsample.

¹² A possible explanation for this pattern can be found in the work of Leeper, Walker and Yang (2010), who use a calibrated model of the US economy to investigate the short and long-run effects of productive government

Another way to assess the importance of anticipation effects is to examine directly the correlation between World Bank project *approvals* and changes in output. If anticipation effects are important, one would expect that consumption and output are more likely to react at the time that the project is approved, rather than when the spending is actually implemented. The difficulty in doing this is that, as discussed earlier, project approvals are potentially endogenous responses to contemporaneous macroeconomic conditions. In order to identify the impact effect of project approvals on output, it is necessary to somehow distinguish between projects that are approved for cyclical reasons and those that are not.

I do this by consulting the documentation for individual projects to assess whether or not their stated objectives suggest that the projects were a response to a contemporaneous macroeconomic shock. Doing so exhaustively for all projects in my sample is difficult given the large number of projects involved (in the 29 countries I consider, there are 1516 projects approved since 1985). Instead, I selected a random sample of 10 projects in each of 10 countries, and consulted the statement of project objectives for each of these 100 projects.¹³ Based on these, I identified 10 projects where the statement of project objectives clearly suggested cyclical rationale for the approval of the project. Not surprisingly, these 10 projects approved for countercyclical reasons also disburse much faster than typical World Bank projects: several disburse in just one or two years, and the longest takes only four years after project approval to disburse fully. Based on this regularity, I isolate all those projects that required more than 4 years to fully disburse, and assume that these are unlikely to have been approved for cyclical reasons. In my full sample of 1516 projects, 79 percent of projects accounting for 67 percent of total approved value fall in this category of slow-disbursing projects that plausibly are not cyclically-motivated. Moreover, this is a conservative way of identifying such non-cyclical projects, because it is

investments on output when there is time-to-build in public capital. In their model, public investment spending plans are announced in advance and are implemented over time, but public capital only becomes productive when the project is fully completed. In this environment, the initial neoclassical labour supply and investment responses to the announcement of the spending plan are muted because private agents would prefer to postpone investments and labour until the future when public capital becomes productive and raises the marginal products of private capital and labour. On the other hand, when the spending is complete, there are positive responses of labour supply, investment, and output due to the complementary effects of public capital.

¹³ For completed projects I relied on the Implementation Completion Report (ICR), while for currently ongoing projects I relied on either Project Approval Documents (PAD) or Project Information Documents (PID). These documents are publicly available for individual projects at <http://www.data.worldbank.org>. A summary of the statements of project objectives and my classification based on them for all 100 projects is available upon request.

likely that at least some of the projects that eventually disburse in less than four years were also not approved for cyclical reasons.¹⁴

I next measure total disbursements over the life of slow-disbursing projects, aggregate this across all projects for each country-year, and express this total as a fraction of GDP in the year that the project was approved. I then include this as a control variable to the benchmark specifications, and assume that such approvals on slow-disbursing projects can be thought of as plausibly exogenous to contemporaneous macroeconomic shocks. The results can be found in Table 9. The coefficient on the announcement of slow-disbursing projects is very small, ranging from 0.04 to 0.05, and is very precisely estimated with standard errors around 0.08. As a result, I cannot reject the null hypothesis that there are no output responses to announcements of World Bank projects that are approved for non-cyclical reasons. The estimated spending multipliers are also quite similar to those in the benchmark specification, and moreover remain strongly identified. Taken together, both the results on the responses of consumption and investment, and the estimates of output responses to project approvals, support the idea that anticipation effects are at most small and likely cannot account for the low estimates of the government spending multiplier in my benchmark specifications.¹⁵

5.5 Is the Multiplier for World Bank-Financed Spending Different?

Thus far, my empirical evidence has focused on the size of the overall government spending multiplier, using changes in disbursements on World Bank loans as an instrument for changes in total government spending. A key underlying assumption has been that the multiplier is the same for all types of public spending. Yet one might reasonably wonder whether such an assumption is valid. For example, it could be the case that World Bank-financed spending has a larger stimulative effect on output than other forms of government spending, perhaps because it is less prone to mismanagement or corruption than other forms of spending. And so the estimates of the overall government spending multiplier I have reported thus far may not be a good guide to the likely stimulative effects of World Bank-financed projects, which could be larger.

¹⁴ In my random subsample of 100 projects, 8 of those that I classify as non-cyclical nevertheless disburse in four years or less.

¹⁵ Of course, another interpretation is that even project approvals themselves are anticipated, and the output responses to the projects occurred even prior to approval. Absent proxies for the expectation of project approvals, I cannot convincingly address this concern. Rather my argument is simply that project approvals are less likely to be anticipated than the eventual disbursements on the projects, since the latter can be predicted with considerable accuracy from project approvals alone.

To empirically investigate this possibility, I consider this simple generalization of Equation (2):

$$(5) \quad \Delta y_t = \beta \Delta g_t + \theta \Delta g_{1t} + \varepsilon_t$$

As before, β is the overall government spending multiplier, and now in addition θ is the differential impact of World Bank-financed spending on output, i.e. since total government spending includes World Bank and non-World Bank financed spending, the impact of the former on output is $\beta + \theta$. The difficulty in estimating Equation (5) is that while changes in the World Bank-financed component of spending, Δg_{1t} , are plausibly exogenous to contemporaneous macroeconomic shocks, changes in total public spending, Δg_t , are not. Moreover, since I now want to estimate the differential effects of World Bank-financed spending on output, I can no longer use it as an instrument for total spending. And finally, since we have already seen that Δg_t and Δg_{1t} are correlated, this endogeneity problem will bias OLS estimates of both β and θ .

While I cannot solve this identification problem, it is straightforward to explore the sensitivity of estimates of the differential impact of World Bank-financed spending, θ , to alternative prior assumptions about the size of the overall spending multiplier, β . In particular define $\Delta \tilde{y}_t = \Delta y_t - \beta \Delta g_t$, for a given fixed prior value for β , so that we can re-write Equation (5) as

$$(6) \quad \Delta \tilde{y}_t = \theta \Delta g_{1t} + \varepsilon_t$$

The probability limit of the OLS estimator of θ based on Equation (6) is:

$$(7) \quad \hat{\theta}(\beta) = \frac{COV(\Delta \tilde{y}_t, \Delta g_{1t})}{V(\Delta g_{1t})} = \frac{COV(\Delta y_t, \Delta g_{1t})}{V(\Delta g_{1t})} - \beta \frac{COV(\Delta g_t, \Delta g_{1t})}{V(\Delta g_{1t})}$$

The first term is simply the slope of the reduced-form regression of changes in output on changes in World Bank-financed spending, reported in the top panel of Table 3. The second term is the overall spending multiplier, β , multiplied by slope of the first-stage regression of changes in total government spending on changes in World Bank-financed spending, reported in the second panel of Table 3.

Revisiting the results in Table 3 with the help of Equation (7), it is immediately apparent that the data do not provide much evidence in support of the hypothesis that World Bank-financed spending itself has a very different stimulative impact on output in the short run than overall government

spending. Consider the results in the first column, corresponding to the total disbursements measure. The reduced-form slope from a regression of output changes on changes in World Bank-finance spending delivers a slope coefficient of 0.14, while the corresponding first-stage slope of 0.31. Inserting these into Equation (7), the estimated differential impact of World Bank spending on output is $0.14 - \beta 0.31$. Even if the overall government spending multiplier were $\beta = 0$, the multiplier for World Bank-financed spending would only be 0.14. For larger assumed values of the overall spending multiplier, the differential impact of World Bank-financed spending on output would be even smaller, and even negative. Based on this evidence, it seems difficult to argue that my estimates of the overall government spending multiplier are too low because they conflate differential short-run effects of World Bank-financed and non-World Bank-financed spending on output.

6. Conclusions

In this paper, I have proposed a novel way to identify fluctuations in public spending that are likely to be uncorrelated with contemporaneous macroeconomic shocks. My identification strategy is based on two key features of many low-income countries: (1) borrowing from the World Bank is an important source of financing for public spending, and (2) projects financed by the World Bank typically take several years to implement following the initial approval of the project. While project approval decisions are potentially endogenous to contemporaneous macroeconomic shocks, I have argued that various measures of disbursements on projects approved in previous years are unlikely to be correlated with shocks to output in the current year. Under this assumption, they can be used as an instrument for changes in total public spending in order to estimate spending multipliers, in a set of 29 low-income countries where systematic evidence on the cyclical effects of government spending does not exist.

The multipliers I have estimated are for the most part very small, reasonably precisely estimated, and rarely significantly different from zero. These findings survive a range of robustness checks designed to address concerns about the noisiness of the macroeconomic data for these countries, as well as concerns about potential violations of the exclusion restriction. In fact, the most puzzling feature of my basic results is that the estimated spending multipliers are so small. I investigate a range of potential explanations for this, including the role of anticipation effects, the composition of World Bank-financed spending and concessionality of World Bank loans, and characteristics of recipient countries. While some of the empirical results around these hypotheses are suggestive, none provide a definitive explanation for the small size of the estimated multipliers. This indicates a need for more

work to understand better why government spending seems to have so little short-run impact on output in these countries.

The basic estimates of the government spending multiplier presented here suggest a limited role for countercyclical fiscal expansions during economic downturns as a way of stimulating aggregate economic activity, in the sample of 29 low-income countries that I consider. This is of course not to say that there is no role for public spending in response to adverse macroeconomic shocks. In many of these countries there is a strong rationale -- and considerable scope -- for expanding social safety net programs to aid the most vulnerable during economic downturns. However the rationale for such programs is better understood as one of providing social protection, rather than as one of providing short-term macroeconomic stimulus. And of course, as noted in the introduction, the small impact effects of government spending on output that I estimate are potentially consistent with very different longer-run effects of public spending in general, and World Bank-financed spending in particular, on growth.

Finally, it is worth reiterating that the work here shares an important weakness with much of the broader empirical literature on estimating government spending multipliers: the difficulty in assigning structural interpretations to reduced-form empirical estimates of multipliers. As emphasized by Leeper (2010), government spending multipliers based on calibrated theoretical models summarize a complex array of factors, including the type of spending involved, the time path of spending, and the nature of the taxes that ultimately will be used to finance the spending increases. As a result it is difficult and even misleading to talk about “the” spending multiplier when its magnitude depends on so many factors, many of which are difficult to control for empirically. Rather, it is better to view empirical measures of multipliers such as those developed in this paper as contributing to a body of stylized facts on the partial correlation between government spending changes and output changes, that can be used to discriminate between alternative theoretical models of the short-run impacts of fiscal policy on output.

References

- Barro, Robert (1981). "Output Effects of Government Purchases". *Journal of Political Economy*. 89:1086-1121.
- Barro, Robert and Charles Redlick (2010). "Macroeconomic Effects from Government Purchases and Taxes. Manuscript, Harvard University.
- Belsely, D.A., E. Kuh, and R.E. Welsch (1980). *Regression Diagnostics: Identifying Influential Data and Sources of Collinearity*. New York, Wiley.
- Blanchard, Olivier and Roberto Perotti (2002). "An Empirical Characterization of the Dynamic Effects of Changes in Government Spending". *Quarterly Journal of Economics*. 117(4): 1329-1368.
- Cohen, Lauren, Joshua Covall, and Christopher Malloy (2010). "Do Powerful Politicians Cause Corporate Downsizing?". NBER Working Paper No. 15839.
- Fatás, Antonio and Ilian Mihov (2003). "The Case for Restricting Fiscal Policy Discretion". *Quarterly Journal of Economics*. 118(4):1419-1448.
- Fisher, Jonas and Ryan Peters (2010). "Using Stock Returns to Identify Government Spending Shocks". *The Economic Journal*. 120:414-436.
- Fishback, Price and Valentina Kachanovskaya (2010). "In Search of the Multiplier for Federal Spending on the States During the New Deal". Manuscript, University of Arizona.
- Hall, Robert (2009). "By How Much Does GDP Rise if the Government Buys More Output?". *Brookings Papers on Economic Activity*.
- Ilzetzki, Ethan and Carlos Végh (2008). "Procyclical Fiscal Policy in Developing Countries: Truth or Fiction?". Manuscript, University of Maryland.
- Ilzetzki, Ethan, Enrique Mendoza and Carlos Végh (2010). "How Big (Small?) Are Fiscal Multipliers?". Manuscript, University of Maryland and LSE.
- Mertens, Karel and Morten Ravn (2010). "Measuring the Impact of Fiscal Policy in the Face of Anticipation: A Structural VAR Approach". *The Economic Journal*. 120(544):393-413.
- Leeper, Eric, Todd Walker and Susan Yang (2010). "Government Investment and Fiscal Stimulus in the Short and Long Runs". Manuscript, Congressional Budget Office and Indiana University.
- Leeper, Eric (2010). "Monetary Science, Fiscal Alchemy". Paper prepared for Jackson Hole Conference.
- Ramey, Valerie and Matthew Shapiro (1998). "Costly Capital Reallocation and the Effects of Government Spending". *Carnegie-Rochester Conference Series on Public Policy*. 48(1): 145-194.

Ramey, Valerie (2009). "Identifying Government Spending Shocks: It's All in the Timing". NBER Working Paper No. 15464.

Romer, David and Christina Romer (2010). "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks". *American Economic Review*. 100:763-801.

Staiger, Douglas and James Stock (1997). "Instrumental Variables Regression with Weak Instruments". *Econometrica*. 65(3):557-586.

Wooldridge, Jeffrey (2002). *Econometric Analysis of Cross Section and Panel Data*. Cambridge: MIT Press.

Figure 1: Fluctuations in Disbursements on World Bank Loans: Example of Zambia




Figure 2: Disbursement Rates on World Bank Projects
(Fraction of original approval disbursed per year)




Figure 3: Core Results






Table 1: Sample of Countries

Code	Country	Number of Observations	Average Disbursements on World Bank Loans (Percent of GDP)	Total Government Spending (Percent of GDP)	World Bank Disbursements as Share of Total Government Spending (Percent)
BDI	Burundi	24	4.3%	29.0%	15.0%
BEN	Benin	24	1.5%	18.9%	8.1%
BFA	Burkina Faso	24	2.0%	19.4%	10.4%
BOL	Bolivia	24	1.1%	28.2%	3.7%
CAF	Central African Republic	24	1.6%	17.9%	8.9%
CIV	Cote d'Ivoire	24	1.3%	25.2%	5.3%
COM	Comoros	24	2.0%	27.7%	7.2%
CPV	Cape Verde	23	1.7%	37.4%	4.5%
ETH	Ethiopia	24	1.9%	21.2%	8.9%
GHA	Ghana	24	3.0%	26.5%	11.3%
GIN	Guinea	24	1.7%	18.7%	9.1%
GMB	Gambia	24	3.1%	27.0%	11.3%
JOR	Jordan	24	1.2%	36.7%	3.3%
KEN	Kenya	24	1.3%	22.9%	5.8%
LSO	Lesotho	24	2.0%	56.3%	3.6%
MAR	Morocco	24	1.0%	24.8%	4.0%
MDG	Madagascar	24	2.9%	18.4%	15.8%
MLI	Mali	24	2.5%	24.4%	10.1%
MWI	Malawi	24	4.6%	31.0%	14.7%
NER	Niger	24	2.2%	19.1%	11.6%
RWA	Rwanda	24	2.6%	21.6%	11.8%
SEN	Senegal	24	1.7%	31.0%	5.3%
SLE	Sierra Leone	24	3.0%	23.8%	12.7%
TCD	Chad	24	2.0%	18.7%	10.9%
TGO	Togo	24	1.9%	22.0%	8.4%
TUN	Tunisia	24	1.1%	29.2%	3.8%
TZA	Tanzania	21	2.8%	18.7%	15.1%
UGA	Uganda	24	3.1%	17.0%	18.5%
ZMB	Zambia	24	3.4%	29.9%	11.3%
	Average	24	2.2%	25.6%	9.3%

Table 2: Summary Statistics

	<u>Changes in Total</u>	<u>Changes in Disbursements on</u>				
	<u>Disbursements</u>	<u>Previously-Approved Projects</u>				
		<u>Excluding</u>	<u>Excluding</u>	<u>Excluding</u>	<u>Change in Total</u>	<u>Change in Real</u>
		<u>Current Year</u>	<u>Current and</u>	<u>Current Year and</u>	<u>Government</u>	<u>GDP</u>
			<u>Previous Year</u>	<u>Using Predicted</u>	<u>Expenditure</u>	
				<u>Disbursements</u>		
<i>Percent of Lagged GDP</i>						
mean	0.06	0.03	0.02	0.04	1.01	3.67
sd	1.43	1.10	0.89	0.55	2.93	3.93
p10	-1.38	-1.13	-0.94	-0.53	-2.59	-0.83
p25	-0.61	-0.45	-0.34	-0.24	-0.70	1.60
p50	-0.04	0.00	-0.02	-0.02	1.03	4.08
p75	0.57	0.45	0.36	0.29	2.82	6.04
p90	1.55	1.22	0.96	0.65	4.57	7.93
N	610	610	610	610	610	610

Table 3: Basic Results

		<u>Changes in Total Disbursements</u>	<u>Changes in Disbursements on Previously-Approved Projects</u>		
Measure of Δg_1 :			<u>Excluding Current Year</u>	<u>Excluding Current and Previous Year</u>	<u>Excluding Current Year and Using Predicted Disbursements</u>
Panel A: Reduced-Form Regressions of Δy on Δg_1					
Slope		0.138	0.108	0.0334	-0.242
Std.Err.		(0.0978)	(0.149)	(0.182)	(0.212)
R-sq		0.003	0.001	0.000	0.001
F		1.996	0.527	0.0336	1.298
Panel B: First-Stage Regressions of Δg on Δg_1					
Slope		0.310***	0.399***	0.459***	0.791***
Std.Err.		(0.0872)	(0.123)	(0.128)	(0.218)
R-sq		0.023	0.022	0.019	0.022
F		12.67	10.48	12.80	13.11
Panel C: Regressions of Δy on Δg					
	OLS		2SLS Using Δg_1 as Instrument		
Slope	0.320***	0.445	0.270	0.0729	-0.306
Std.Err.	(0.0615)	(0.337)	(0.384)	(0.387)	(0.298)
N	610	610	610	610	610

Table 4: Robustness Checks: Dropping Influential Countries and Observations

		<u>Changes in Total</u>	<u>Changes in Disbursements on</u>		
		<u>Disbursements</u>	<u>Previously-Approved Projects</u>		
Measure of Δg_1 :			<u>Excluding</u>	<u>Excluding</u>	<u>Excluding Current</u>
			<u>Current Year</u>	<u>Current and</u>	<u>Year and Using</u>
				<u>Previous Year</u>	<u>Predicted</u>
					<u>Disbursements</u>
Panel A: Controlling for Lagged Growth					
	OLS	2SLS Using Δg_1 as instrument			
Coefficient on Δg	0.295***	0.561	0.362	0.129	-0.249
Std. Err.	(0.0569)	(0.372)	(0.436)	(0.421)	(0.333)
Coefficient on Lagged Growth	0.152**	0.130*	0.146*	0.165**	0.196**
Std. Err.	(0.0618)	(0.0673)	(0.0743)	(0.0687)	(0.0725)
First-Stage F-Statistic		14.65	10.49	11.55	13.1
N	592	592	592	592	592
Panel B: Controlling for Lagged Growth and Lagged Government Spending					
	OLS	2SLS Using Δg_1 as instrument			
Coefficient on Δg	0.298***	0.591*	0.409	0.109	-0.291
Std. Err.	(0.0534)	(0.314)	(0.321)	(0.415)	(0.413)
Coefficient on Lagged Δg	0.140***	0.212	0.140	0.0619	0.121
Std. Err.	(0.0467)	(0.302)	(0.237)	(0.317)	(0.357)
Coefficient on Lagged Growth	0.128***	0.0930*	0.119**	0.156***	0.179***
Std. Err.	(0.0365)	(0.0554)	(0.0482)	(0.0549)	(0.0649)
Cragg-Donald Statistic		5.24	6.03	3.47	4.04
<i>(Stock-Yogo critical value for 15% (25%) size distortion = 4.58 (3.63))</i>					
N		584 584	584	584	584
Panel C: Controlling for Contemporaneous Policy Changes					
	OLS	2SLS Using Δg_1 as instrument			
Coefficient on Δg	0.314***	0.285	0.128	0.0383	-0.362
Std. Err.	(0.0596)	(0.341)	(0.394)	(0.402)	(0.318)
Coefficient on Change in CPIA	0.0139**	0.0140**	0.0145**	0.0147**	0.0159**
Std. Err.	(0.00614)	(0.00637)	(0.00634)	(0.00627)	(0.00675)
First-Stage F-Statistic		10.94	10.23	12.46	12.43
N	610	610	610	610	610

Table 5: Robustness Checks, Cont'd: Possible Violations of Exclusion Restriction

		<u>Changes in Total Disbursements</u>	<u>Changes in Disbursements on Previously-Approved Projects</u>		
			<u>Excluding Current Year</u>	<u>Excluding Current and Previous Year</u>	<u>Excluding Current Year and Using Predicted Disbursements</u>
Measure of Δg_1 :					
Panel A: Controlling for Lagged Growth					
	<i>OLS</i>	<i>2SLS Using Δg_1 as instrument</i>			
Coefficient on Δg	0.295***	0.561	0.362	0.129	-0.249
Std.Err.	(0.0569)	(0.372)	(0.436)	(0.421)	(0.333)
Coefficient on Lagged Growth	0.152**	0.130*	0.146*	0.165**	0.196**
Std. Err.	(0.0618)	(0.0673)	(0.0743)	(0.0687)	(0.0725)
First-Stage F-Statistic		14.65	10.49	11.55	13.1
N	592	592	592	592	592
Panel B: Controlling for Lagged Growth and Lagged Government Spending					
	<i>OLS</i>	<i>2SLS Using Δg_1 as instrument</i>			
Coefficient on Δg	0.298***	0.591*	0.409	0.109	-0.291
Std.Err.	(0.0534)	(0.314)	(0.321)	(0.415)	(0.413)
Coefficient on Lagged Δg	0.140***	0.212	0.140	0.0619	0.121
Std. Err.	(0.0467)	(0.302)	(0.237)	(0.317)	(0.357)
Coefficient on Lagged Growth	0.128***	0.0930*	0.119**	0.156***	0.179***
Std. Err.	(0.0365)	(0.0554)	(0.0482)	(0.0549)	(0.0649)
Cragg-Donald Statistic		5.24	6.03	3.47	4.04
(Stock-Yogo critical value for 15% (25%) size distortion = 4.58 (3.63))					
N		584 584	584	584	584
Panel C: Controlling for Contemporaneous Policy Changes					
	<i>OLS</i>	<i>2SLS Using Δg_1 as instrument</i>			
Coefficient on Δg	0.314***	0.285	0.128	0.0383	-0.362
Std.Err.	(0.0596)	(0.341)	(0.394)	(0.402)	(0.318)
Coefficient on Change in CPIA	0.0139**	0.0140**	0.0145**	0.0147**	0.0159**
Std. Err.	(0.00614)	(0.00637)	(0.00634)	(0.00627)	(0.00675)
First-Stage F-Statistic		10.94	10.23	12.46	12.43
N	610	610	610	610	610

Table 6: Composition and Concessionality of World Bank-Financed Spending

		<u>Changes in Total</u>	<u>Changes in Disbursements on</u>		
		<u>Disbursements</u>	<u>Previously-Approved Projects</u>		
Measure of Δg_1 :			<u>Excluding</u> <u>Current Year</u>	<u>Excluding</u> <u>Current and</u> <u>Previous Year</u>	<u>Excluding Current</u> <u>Year and Using</u> <u>Predicted</u> <u>Disbursements</u>
Panel A: Dropping World Bank Spending on Foreign Procurement					
	OLS	2SLS Using Δg_1 as instrument			
Coefficient on Δg	0.320***	0.429	0.0165	-0.163	-0.423
Std.Err.	(0.0615)	(0.350)	(0.396)	(0.400)	(0.410)
First-Stage F-Statistic		6.705	4.958	8.101	6.880
N	610	610	610	610	610
Panel B: Post-Debt-Relief Period (1996-2008)					
	OLS	2SLS Using Δg_1 as instrument			
Coefficient on Δg	0.246***	0.468	0.308	0.519**	0.0160
Std.Err.	(0.0759)	(0.340)	(0.292)	(0.232)	(0.511)
First-Stage F-Statistic		9.956	9.698	15.09	4.566
N	368	368	368	368	368
Panel C: Pre-Debt-Relief Period (1985-1995)					
	OLS	2SLS Using Δg_1 as instrument			
Coefficient on Δg	0.336***	0.538	0.522	-1.053	0.0456
Std.Err.	(0.107)	(0.895)	(0.958)	(2.073)	(0.354)
First-Stage F-Statistic		2.139	1.638	1.260	7.104
N		242	242	242	242

Table 7: Distinguishing Between Strong and Weak Policy Environments

		<u>Changes in Total Disbursements</u>	<u>Changes in Disbursements on Previously-Approved Projects</u>		
Measure of $\Delta g1$:			<u>Excluding Current Year</u>	<u>Excluding Current and Previous Year</u>	<u>Excluding Current Year and Using Predicted Disbursements</u>
Panel A: Strong Policy Observations (CPIA score of 3.5 or better)					
	OLS	2SLS Using $\Delta g1$ as instrument			
Coefficient on Δg	0.297***	0.0445	-0.623	-0.397	-1.142
Std.Err.	(0.0623)	(0.501)	(0.422)	(0.653)	(0.785)
First-Stage F-Statistic		10.65	5.808	6.349	3.182
N	313	313	313	313	313
Panel B: Weak Policy Observations (CPIA score below 3.5)					
	OLS	2SLS Using $\Delta g1$ as instrument			
Coefficient on Δg	0.282***	0.655	0.882	0.316	-0.0874
Std.Err.	(0.0796)	(0.519)	(0.811)	(0.386)	(0.524)
First-Stage F-Statistic		5.590	5.206	8.256	12.66
N	297	297	297	297	297

Table 8: Effects on Expenditure Components of GDP

Measure of Δg_1 :		<u>Changes in Total</u>	<u>Changes in Disbursements on</u>		
		<u>Disbursements</u>	<u>Previously-Approved Projects</u>		
			<u>Excluding</u>	<u>Excluding</u>	<u>Excluding Current</u>
			<u>Current Year</u>	<u>Current and</u>	<u>Year and Using</u>
				<u>Previous Year</u>	<u>Predicted</u>
					<u>Disbursements</u>
	<i>OLS</i>	<i>2SLS Using Δg_1 as instrument</i>			
Private Consumption					
Coefficient on Δg	-0.0683	-0.403	0.259	-0.773	-0.999*
Std.Err.	(0.114)	(0.514)	(0.758)	(0.703)	(0.571)
Private Investment					
Coefficient on Δg	0.146*	0.160	0.151	-0.0211	0.246
Std.Err.	(0.0730)	(0.313)	(0.370)	(0.389)	(0.334)
Government Consumption					
Coefficient on Δg	0.288***	0.627*	0.662**	0.611**	0.414**
Std.Err.	(0.0327)	(0.342)	(0.243)	(0.286)	(0.151)
Government Investment					
Coefficient on Δg	0.276***	0.579**	0.436**	0.604***	0.266
Std.Err.	(0.0372)	(0.235)	(0.210)	(0.165)	(0.201)
Net Exports					
Coefficient on Δg	-0.329***	-0.549	-1.340	-0.384	-0.327
Std.Err.	(0.0646)	(0.408)	(0.911)	(0.418)	(0.433)
First-Stage F-Statistic		11.65	5.649	12.12	12.53
N	590	590	590	590	590

Table 9: Anticipation Effects

		<u>Changes in Total</u>	<u>Changes in Disbursements on</u>		
		<u>Disbursements</u>	<u>Previously-Approved Projects</u>		
Measure of Δg_1 :			<u>Excluding</u> <u>Current Year</u>	<u>Excluding</u> <u>Current and</u> <u>Previous Year</u>	<u>Excluding Current</u> <u>Year and Using</u> <u>Predicted</u> <u>Disbursements</u>
	<i>OLS</i>	<i>2SLS Using Δg_1 as instrument</i>			
<i>Coefficient on Δg</i>	0.319***	0.435	0.271	0.0545	-0.332
<i>Std.Err.</i>	(0.0613)	(0.343)	(0.383)	(0.380)	(0.295)
<i>Coefficient on Approvals of</i> <i>Slow-Disbursing Projects</i>	0.0402	0.0375	0.0412	0.0461	0.0549
<i>Std.Err.</i>	(0.0669)	(0.0696)	(0.0687)	(0.0685)	(0.0815)
<i>First-Stage F-Statistic</i>		12.11	10.43	13.21	14.66
<i>N</i>	610	610	610	610	610