HOW LARGE IS THE GOVERNMENT SPENDING MULTIPLIER? EVIDENCE FROM WORLD BANK LENDING

Aart Kraay
The World Bank

First Draft: June 2010
This Draft: September 2011

Abstract: This paper proposes a novel approach to empirically identifying government spending multipliers that relies on two features unique to many low-income countries: (1) borrowing from the World Bank finances a substantial fraction of government spending, and (2) spending on World Bank-financed projects is typically spread out over several years following the original approval of the project. The first fact means that fluctuations in spending on World Bank-financed projects are a significant source of fluctuations in overall government 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 so are unlikely to be correlated with shocks to output in the current year. I use World Bank project-level disbursement data to isolate the component of World Bank-financed government spending in a given year that is associated with past project approval decisions. I then use this as an instrument for total government spending to estimate multipliers in a sample of 29 primarily low-income countries where variation in government spending from this source is large relative to the size of the economy. The resulting spending multipliers are small, and reasonably precisely estimated to be in the vicinity of 0.5.

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 Ilzetzki, Justin Lin, Claudio Raddatz, Martin Ravallion, Luis Serven, Susan Yang, and four anonymous referees for helpful comments. 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 government spending multipliers requires a strategy to isolate changes in government spending that plausibly are uncorrelated with contemporaneous economic shocks. In this paper, I propose a novel method of identifying such fluctuations in government spending that relies on two features unique to many low-income countries: (1) borrowing from the World Bank finances a substantial fraction of government spending, and (2) spending on World Bank-financed projects is typically spread out over several years following the original approval of the project. The first fact means that fluctuations in spending on World Bank-financed projects are a significant source of fluctuations in overall government 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 so are unlikely to be correlated with shocks to output in the current year. I use World Bank project-level disbursement data to isolate the component of World Bank-financed government spending in a given year that is associated with past project approval decisions. I then use this as an instrument for total government spending, in order to estimate spending multipliers in a sample of 29 primarily low-income countries where variation in government spending from this source is large relative to the size of the economy.

The recent global 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 consists of vector auto-regression (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 for analyzing 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, Ilzetzki and Végh (2008) and Ilzetzki, Mendoza and Végh (2010) are able to assemble quarterly data in order to analyze the cyclical effects of fiscal policy in these

The second identification strategy consists of isolating a subcomponent of spending or taxes that is likely to be uncorrelated with contemporaneous economic shocks. For example, many authors have followed Barro (1981) in arguing that fluctuations in military spending in the United States during major wars can be thought of in this way, because these conflicts occurred outside the United States (so that there was no direct effect of the conflict on the US economy), and their timing was determined by geopolitical factors unrelated to US macroeconomic fluctuations.[2] The third strategy consists of finding some external source of variation in government spending that is unlikely to be correlated with contemporaneous macroeconomic events. This approach has most often been employed in the context of spending by sub-national governments, often using US state and local government spending.[3]

---

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.

[2] Ramey and Shapiro (1998), Hall (2009), Fisher and Peters (2010), Ramey (2011b), and Barro and Redlick (2011) all follow variants on this basic approach. Nakamura and Steinsson (2011) also focus on military spending in the United States, but exploit cross-state variation in the intensity of defense spending. This permits a weaker identifying assumption that military spending buildups are unrelated to *differences* in macroeconomic conditions across US states. A common drawback of these military spending-based studies is that it is difficult to control for the macroeconomic effects of other key features of wartime economies, such as price controls or mandatory military service. 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 United States, and use this to distinguish between those changes that were made 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.

[3] 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. On the other hand, Chodorow-Reich, Feiveson, Liscow, and Woolston (2011) study the effects of federal Medicaid transfers to US states during the 2009 fiscal stimulus in the United States and find significant positive effects on state-level employment. 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 government spending and find output multipliers ranging from 0.9 to 1.7 depending on the type of spending, although no appreciable impact on employment. Clemens and Miran (2010) exploit variation across US states in the stringency of balanced-budget rules to isolate changes in state government spending that plausibly are unrelated to state-level economic conditions, and find a spending multiplier equal to 1.7. Shoag (2010) uses idiosyncratic variation in returns in US state government-run pension funds to isolate "windfalls" that strongly predict state government spending, and finds a spending multiplier of 2.1 . Serrato and Wingender (2011) observe that significant portions of federal spending are allocated to localities based on population estimates, and then exploit sharp changes in population estimates due to methodological changes in census years to isolate exogenous variations in federal spending at local levels. They also find large estimated multipliers in the vicinity of 1.9. Acconcia, Corsetti, and Simonelli (2011) exploit interruptions in municipal public works projects on evidence of mafia involvements as a source of variation in government spending across Italian regions, and find estimated spending multipliers ranging from 1.4 to 2.

The identification strategy in this paper is a hybrid of the second and third approaches. As in the second approach, I begin by identifying a subcomponent of government spending whose fluctuations are plausibly uncorrelated with contemporaneous macroeconomic events. 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. The basic identifying assumption is that these project approval decisions made in previous years do not react to shocks to growth in the current year, and so current-year disbursements attributable to project approval decisions made in previous years are unlikely to be correlated with current-year shocks to growth. As in the third approach, I use fluctuations in this component of World Bank-financed spending as an instrument for changes in total government spending in order to estimate overall government spending multipliers.

I apply this approach in a sample of 29 aid-dependent developing countries in which World Bank lending is an important source of financing of government expenditures, and for which well-identified estimates of spending multipliers have until now not been available. My baseline estimate of the one-year government spending multiplier is small, at 0.48, with some variation around this value depending on the specification and methodology used. This is well below existing estimates of the multiplier for the United States.[4] Moreover, the spending multiplier is reasonably precisely estimated, with standard errors that are comparable to those in leading papers using US data.

There are numerous possible objections to this basic identification strategy, two of which deserve immediate mention. First, although past project approval decisions are made before current macroeconomic shocks are known, they may nevertheless be correlated with current shocks if these shocks are persistent over time, or are 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. Second, even if past project approval decisions are uncorrelated with contemporaneous shocks, the timing of subsequent disbursements on those projects may nevertheless be contemporaneously correlated with macroeconomic shocks. I address this concern by relying throughout the paper on a synthetic measure of project-level disbursements based on typical disbursement rates for projects in the same economic sector and geographical region, rather than using actual disbursements. By construction, these artificial disbursement rates will not reflect country-specific macroeconomic shocks.

[4] Ramey (2011a) surveys this literature and summarizes a consensus view that spending multipliers for the United States are between 0.8 and 1.5.

Several limitations of the evidence in this paper are also worth acknowledging at the outset. First, this paper shares with much of the literature the difficulty that econometrically-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 burden of the eventual taxes that finance the spending. For this reason, the term "multiplier" as I use it in this paper is perhaps best understood simply as short-hand for the empirical correlation between a plausibly exogenous component of changes in government spending and changes in output. As such, the estimates of the multiplier in this paper may not be a good guide to the effects of particular types of spending increases in particular situations, even within my sample of countries. Second, out of necessity the empirical results in this paper are based on data for a particular set of poor and/or small economies in which World Bank lending is an important source of financing for government expenditure. The short-run effects of government spending on output in other countries outside this particular sample may very well be different.

Third, I emphasize that the empirical work here is designed only to assess the short-run output effects of changes in total government spending, using changes in disbursements on World Bank projects as an instrument. This question is related to -- but distinct from -- that of the long-run growth impacts of foreign assistance more generally, which has been debated extensively in the vast empirical aid-growth literature. This paper differs from the aid-growth literature in two main respects. First, the objective of this paper is to estimate the short-run stimulative effects of overall government spending, using annual fluctuations in total government spending and per capita GDP. In contrast, the aid-growth literature has primarily focused on estimating the medium- to long-run growth effects of aid alone. Second, the identification strategy in this paper differs importantly from that of the aid-growth literature. As discussed above, in this context the challenge of identification consists of finding a source of short-run within-country variation in total government spending that plausibly is uncorrelated with contemporaneous macroeconomic shocks at an annual frequency. The solution consists of isolating the component of spending that is associated with World Bank projects approved in previous years. In contrast, the identification challenge in the aid-growth literature consists of finding a source of medium- to long-run variation in aid that is likely to be uncorrelated with medium- to long-run growth, usually across countries. The solution in this literature has typically consisted of proposing a range of

instruments of varying degrees of plausibility, none of which have been based on disbursement lags associated with aid-financed projects, as is done in this paper.[5]

Section 2 of the paper presents the empirical framework I use to estimate the government spending multiplier. Section 3 describes the World Bank project-level disbursement data I use to construct the instrument for total government spending. Section 4 contains the 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 in my sample, and Section 6 investigates whether the spending multiplier varies across different types of government spending. Section 7 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:

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

Here, y_t and g_t denote GDP and total government spending, both measured in constant local currency units; 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. All data are measured at annual frequency out of necessity, given the unavailability of quarterly data in the sample of primarily low-income countries used in this paper. 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 unit of government spending does not lead to reductions in any of the other expenditure components of GDP, the multiplier would be one. As noted in the introduction, β is difficult to interpret as a deep structural parameter. Rather, it should simply be

[5] Clemens, Radelet, Bhavnani and Gozzi (2011) document how the most common identifying assumption in the empirical aid-growth literature is that country size is exogenous to growth, and discuss the weaknesses of other external instruments used in this literature. As an alternative, they investigate the relationship between lagged aid and contemporaneous growth, using panels of 4-year averages. Roodman (2008) extensively documents the shortcomings of internal instruments in the smaller literature that uses dynamic panel techniques to estimate aid-growth regressions, most commonly using panels of 4- or 5-year averages.

thought of as a reduced-form empirical summary of the contemporaneous relationship between annual fluctuations in government spending and output.

Although later I will be combining information from multiple countries, for notational convenience I suppress country subscripts. In all the empirical work that follows, I will also restrict β to be the same across countries, and I will include a full set of country and year dummies. Let Δx_t denote the deviation of $\frac{x_t - x_{t-1}}{y_{t-1}}$ from its country- and year-average for any variable x, so that the version of Equation (1) that I estimate can be re-written more compactly as:

$$\Delta y_t = \beta \Delta g_t + \varepsilon_t \tag{2}$$

The standard difficulty in identifying β 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 $E[\Delta g_t \varepsilon_t] < 0$. In this case, 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, a different concern might be that government spending is procyclical, i.e. $E[\Delta g_t \varepsilon_t] > 0$. In this case, OLS estimates of the multiplier would be biased upwards.[6] Yet another possibility is that, in aid-dependent countries such as those studied here, any procyclical tendencies in domestically-financed government spending are offset by countercyclical tendencies in aid-financed government spending, so that total spending could be either procyclical or countercyclical.

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

[6] See Ilzetzki and Végh (2008) Fatás and Mihov (2002) for evidence on the cyclicality of government spending in developing countries.

implement that they cannot react to economic activity within a quarter.[7] The military spending approach restricts attention to US defense spending, and argues that fluctuations in this component of spending are uncorrelated with the error term in Equation (2). As long as changes in military and non-military spending are also uncorrelated, a consistent estimate of the multiplier can be obtained from an OLS regression of changes in output on changes in military spending. Finally, the instrumental variables approach to identification involves finding some external source of variation in government spending that arguably is uncorrelated with the error term in Equation (2).

My approach to identification is a hybrid of these second and third approaches. As in the military spending literature, I begin by focusing on one particular component of government spending: the portion that is associated with disbursements on World Bank-financed projects. As noted in the introduction, the World Bank finances a significant share of total government spending in my sample of aid-dependent low-income countries. Crucially for my purposes, individual World Bank projects typically take many years to implement. As a result, most of World Bank-financed spending in a typical country-year is associated with project approval decisions made in many previous years, and before contemporaneous macroeconomic shocks are known. I exploit these disbursement lags, together with data on the time profiles of disbursements over the life of individual World Bank projects, to construct a measure of World Bank-financed spending in a given year that is associated with projects approved in previous years. Since my primary interest is in estimating an overall government spending multiplier, and not simply an aid-financed spending multiplier, I use changes in this measure of spending associated with previous-approved projects as an instrument for changes in total government spending when estimating Equation (2).

To better understand this identification strategy, 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 government 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 piece of infrastructure, or to fund teacher training, or to support a particular health intervention, or a myriad of other development-oriented government activities 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.

[7] VAR-based identification also relies on the further strong assumptions that policymakers do not react quickly to other information not included in the VAR, and that the cyclically-adjusted spending and taxes used in the VAR are fully purged of all automatic responses of fiscal policy to the business cycle.

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. These projects typically are designed to be carried out over several years: in the sample of primarily low-income countries I consider, the median length of a World Bank project from approval to completion is six 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.

My core identifying assumption is that, while project approval decisions in a given year may be correlated with current and past macroeconomic shocks, they are not correlated with future macroeconomic shocks that are unknown at the time of project approval. If, in addition, the subsequent disbursements on the project follow the schedule initially planned at the time of project approval, then fluctuations in aggregate disbursements (across all active projects) on projects approved in previous years will be uncorrelated with contemporaneous macroeconomic shocks.

Importantly for my purposes, individual disbursements do not always unfold as originally planned in the initial project proposal. Disbursements over the life of individual projects are made primarily at the discretion of World Bank staff responsible for the project, and reflect the pace of actual project implementation. Departures from initially-planned disbursements can reflect a wide range of factors, including unforeseen technical problems in the implementation of the project, procurement delays, failure by the recipient country to meet conditions specified in the project agreement, unexpected delays to investigate possible financial irregularities in the project, and many other considerations.[8] The particular concern relevant for my identification strategy is the possibility that departures from initially-planned disbursements may be correlated with contemporaneous macroeconomic shocks. For example, disbursements on projects approved in previous years might be 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 shocks, such as

[8] In extraordinary cases of non-repayment of World Bank loans, the World Bank will suspend disbursements on all active projects, as well as approvals of new projects, in the defaulting country. While a large literature has studied the allocation of aid projects across countries, relatively little attention has been paid to the determinants of the rate of disbursement on approved aid projects. An exception is Kilby (2010) who empirically documents a correlation between disbursement rates and measures of US political influence over the recipient country.

the onset of civil conflict, could make project implementation more difficult, and thus lead to a reduction in disbursements relative to original plans at the same time that output declines.

To address this concern, I replace actual disbursements on World Bank-financed projects with a synthetic measure of predicted disbursements based on typical disbursement profiles observed on all projects in the same economic sector and geographical region. I apply these typical disbursement profiles to the total size of the project to obtain a series of predicted disbursements for each year of the life of each project. Finally, I aggregate these across projects, again excluding disbursements on projects approved in the same year, to arrive at a measure of total predicted disbursements. By construction, the only country-specific information reflected in this measure consists of project approval decisions made in previous years. Conditional on my identifying assumption that project approval decisions do not anticipate future macroeconomic shocks, fluctuations in this predicted disbursement measure are uncorrelated with the error term in Equation (2).

My primary interest is in estimating an overall government spending multiplier, and not simply a World Bank-financed spending multiplier. In principle, the two could be very different, depending on the nature of the response of the rest of government spending to World Bank disbursements. An extreme example might be that World Bank disbursements fully crowd out other types of government spending, or other aid flows, so that total government spending is unchanged. Conversely, to the extent that World Bank-financed projects require co-financing from other donors, or from the recipient government itself, an additional dollar of World Bank disbursements might lead to a greater increase in total government spending. I therefore use changes in predicted disbursements as an instrument for changes in total government spending, implying the following first-stage regression:

$$\Delta g_t = \gamma \Delta d_t + u_t \tag{3}$$

where Δd_t denotes the change in predicted disbursements. In aid-dependent countries in which World Bank-financed spending is an important component of total government spending, and given the long disbursement lags on many World Bank projects, Δd_t will be an important source of variation in Δg_t. This will result in a strong estimated first-stage relationship between World Bank disbursements and total government spending. Moreover, since changes in predicted disbursements only reflect changes in project approval decisions from previous years, my identifying assumption that past project approval

decisions do not anticipate future shocks to growth implies that Δd_t will be uncorrelated with the error term in Equation (2).[9] I can therefore use changes in predicted disbursements as an instrument for changes in total government spending to generate two-stage least squares (2SLS) estimates of the multiplier based on Equation (2).

3. Data

In order to implement the empirical strategy described in the previous section, 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).[10] For each project, I sum the quarterly disbursements within each calendar year to arrive at annual disbursement flows for each project, in order to match the annual frequency of macroeconomic data available for the countries in my sample. A detailed description of the project-level disbursement data, as well as the other data used in the paper, is provided in the Data Appendix.

The first step in constructing my instrument is to discard disbursements in a given country-year that are associated with projects approved in the same year, as these may be reactions to contemporaneous macroeconomic shocks. In order for this not to eliminate all of the variation in disbursements, it is crucial that there are substantial lags between project approval and eventual full disbursement for most projects. I document that this is indeed the case in Figure I, which reports the average disbursement profile, i.e. fraction of total spending that is disbursed in year t of the project (t=0,...,10 with t=0 indicating the year in which the project was approved), pooling all 7,443 World Bank projects for which I have disbursement data. For the average World Bank project, just over 12 percent

[9] It is tempting to interpret the slope coefficient in the first-stage regression as a measure of the extent to which World Bank-financed aid crowds in, or alternatively supplants, other forms of government spending. However, the interpretation of the slope coefficient is complicated by the fact that the instrument, Δd_t, measures *predicted* rather than *actual* disbursements. A more natural way to investigate any crowding in/out of government spending due to aid would be to regress changes in government spending on changes in actual aid, and using the predicted disbursements measure as an instrument for actual aid disbursements. While this issue is interesting and highly relevant to aid policy discussions, it is tangential to the emphasis in this paper on estimating government spending multipliers.

[10] The dataset in principle covers all 10,631 projects approved between 1948 and early 2010. 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 7,443 projects that were active since 1985, and for which electronic data on disbursement flows are available.

of the total value of the project 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 of this ratio across country-years is 99 percent. This means that, for a typical country-year observation in my sample, 99 percent of total disbursements are associated with project approval decisions made in previous years.[11]

The second step in the construction of my instrument addresses concerns about the potential endogeneity of actual disbursements in a given year on projects approved in previous years. I do this by assigning a typical disbursement profile to each project, based on a simple average of the actual disbursement profiles across all projects in the same economic sector and geographical region.[12] In particular, this involves constructing average disbursement profiles like the one shown in Figure I, but separately for each region/sector bin. I apply these region/sector-specific average disbursement profiles to the total value of each individual project to arrive at a time series of annual predicted disbursements for each project. Finally, I aggregate these predicted project-level disbursements to the country-year level, excluding projects approved in the same year, to obtain my measure of aggregate predicted disbursements that I use as an instrument.[13] By construction, the only country-specific information included in fluctuations in this instrument are fluctuations in past project approval decisions, which I have assumed to be uncorrelated with contemporaneous shocks.

Figure II illustrates these two steps in the construction of my instrument, for one country in my sample, Zambia. The overall height of the bars shows annual actual total disbursements on World Bank projects in Zambia. World Bank disbursements are large as a fraction of GDP in Zambia: total

[11] The difference between the previous figure of 88 percent and this figure of 99 percent is that for the former, the unit of observation is the project, while for the latter, the unit of observation is the country-year. Since country-year observations vary in the number and size of active projects, these two figures need not be the same.

[12] The region-sector averages are based on 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 on the 10-sector classification of projects shown in Appendix Table A1.

[13] Rather than rely on these predicted disbursement profiles, another approach would be to use information on the planned disbursement profiles at the time of project approval. There are however two considerations which make this infeasible. First, initially-planned disbursement profiles are electronically available only for a subset of projects in the second half of the time period covered in the sample. Given the large number of projects involved, manually retrieving these initially-planned disbursement profiles is infeasible. Second, planned future disbursement profiles are frequently revised over the life of the project, calling into question the reliability of the initial disbursement profiles set at the time of project approval.

disbursements average 3.4 percent of GDP over the sample period 1985-2009, and in several years exceed five percent of GDP. The dark-shaded lower portion of each bar isolates disbursements on projects approved in previous years, but not the current year, while the remainder of the bar shows disbursements on projects approved in the same year. As discussed above, given the length of a typical World bank project, in most years the bulk of World Bank disbursements are associated with projects approved in previous years. However, in a few years such as 1991 and 1999, there are also cases of large disbursements on projects approved in the same year. Finally, the solid line graphs my measure of predicted disbursements on previously-approved projects. As described above, this series is constructed by applying region/sector average disbursement profiles to the actual amount of each project, and then aggregating across projects to the country-year level, excluding projects approved in the same year.

Two episodes in the data for Zambia in Figure II are useful as illustrations of my identification strategy. First, in 1991 actual disbursements are very large at 6.3 percent of GDP, and nearly all of this reflects disbursements on projects approved in the same year. Much of this was due to one large (for Zambia) project approved in 1991, worth $210 million, of which approximately three-quarters was disbursed in the same year. This project was an "Economic Recovery Credit" and was intended to "support economic reforms aimed at macroeconomic stabilization". Both the approval of this project, and the decision to disburse a large fraction of it in the approval year, clearly are countercyclical responses to contemporaneous macroeconomic events. For this reason, when constructing the instrument in a given year, it is important to exclude disbursements on all projects approved in the same year. Failure to do so would, in this case, result in downward-biased 2SLS estimates of the government spending multiplier, since the instrument would be negatively correlated with macroeconomic shocks. At the same time, however, this strategy of eliminating approval-year disbursements on all projects is conservative because it also excludes approval-year disbursements on many projects that were very likely undertaken for non-cyclical reasons.

A second interesting episode is the late 1980s, when disbursements fall from a high of nearly 7 percent of GDP in 1986 to nearly zero in 1988-1990. The very low disbursements during the late 1980s occurred during a time of significant political unrest in Zambia, including a coup attempt against independence-era President Kaunda, and culminating in a transition to multi-party democracy in 1991. In 1987 the World Bank suspended most activities in Zambia, including new project approvals and disbursements on previously-approved projects, due to non-repayment of debt service due. World Bank disbursements did not resume until 1991. This example illustrates how both approvals, as well as

disbursements in a given year on projects approved in previous years, may decline in response to an adverse contemporaneous macroeconomic shock. For this reason, I rely on typical predicted disbursement profiles, rather than actual realized disbursement profiles, in constructing my instrument. Not surprisingly, predicted aggregate disbursements fall much less than actual disbursements during this period since they do not reflect the Zambia-specific suspension of disbursements. It is also important to notice that, in this case, the gap between actual and predicted disbursements is *procyclical*, with actual disbursements declining more than predicted disbursements during a negative economic shock. In this case, failure to correct for this potentially-cyclical component of disbursements on previously-approved projects would result in upward-biased 2SLS estimates of the government spending multiplier, since my instrument would be positively correlated with macroeconomic shocks.

I limit my sample of countries to those in which disbursements on World Bank projects are large relative to borrowing-country GDP and government spending, so that fluctuations in World Bank disbursements contribute significantly to fluctuations in total government spending. In particular, I focus on a set of aid-dependent countries where (a) annual disbursements on World Bank projects are available for at least 20 of the 25 years between 1985 and 2009, and (b) annual disbursements on World Bank projects as a share of GDP, averaged over the sample period 1985-2009, exceed one percent. This results in a set of 42 primarily low-income countries where World Bank lending has been an important source of financing for government spending over the past 25 years.

My country sample is further limited by the availability of data on total government spending required to construct Δg_t. My primary source 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 13 countries for which there are remaining gaps in the government spending data or national accounts data. This results in a final sample of 29 countries reported in Table I. 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 poor, and are located in Sub-Saharan Africa. For these 29 countries, World Bank-financed 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 as a share of GDP averages 2.2 percent across countries, and ranges from 1 percent in Morocco to 4.3 percent in Burundi.

The share of World Bank disbursements in total government spending ranges from a low of 2.8 percent in Morocco to a high of 18.6 percent in Uganda, and averages 9.4 percent.

Table II documents summary statistics for the key variables of interest: changes in output, Δy_t, changes in total government spending, Δg_t, and changes in predicted disbursements on previously-approved World Bank projects, Δd_t, pooling all country-years in my sample, and taking deviations from country and year averages. For comparison purposes, I also report changes in actual disbursements. Fluctuations in actual World Bank disbursements are quite substantial, with a standard deviation of 1.4 percent of GDP. Unsurprisingly, fluctuations in predicted disbursements on previously-approved projects are smaller, with a standard deviation of 0.5 percent of GDP, but are non-trivial. By way of comparison, the standard deviation of GDP growth rates is 3.5 percent in this sample, and the standard deviation of changes in total government spending is 2.7 percent.

4. Estimates of the Government Spending Multiplier

4.1 Basic Results

My basic estimates of the government spending multiplier are reported in Table III. I first report the OLS estimate of the multiplier, in order to provide a benchmark for comparison with the 2SLS results that follow. The OLS estimator in the first column delivers a slope of 0.23 that is very strongly significantly different from zero. The 2SLS estimator in the second column, using changes in predicted disbursements as an instrument, delivers an estimated multiplier of 0.48 with a standard error that implies a 95 percent confidence interval ranging from -0.09 to 1.04. While the estimated 2SLS standard error of 0.29 is inevitably larger than the OLS standard error, it is not unreasonably large.[14] By way of comparison, the standard error associated with this estimate of the multiplier is not much larger than those reported in Barro and Redlick (2011), 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 II, 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.

[14] The predicted disbursements measure is a generated instrument (consisting of actual project approvals multiplied by estimated average disbursement rates). However, 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.

The respectable precision of this first basic estimate of the multiplier in part reflects the fact that changes in the predicted disbursement instrument are a good predictor of changes in total government spending. This can be seen in the corresponding first-stage regression reported in the bottom panel of Table III. The first-stage regression delivers a slope of 1.03 and a standard error of 0.27, indicating a first-stage relationship that is highly significant. The first-stage F-statistic is nearly 15 and comfortably exceeds the Staiger and Stock (1997) rule of thumb of 10, indicating that weak instrument pathologies are unlikely to be a concern in this specification. Nevertheless, I also report a weak instrument-consistent confidence interval for the 2SLS estimate of the multiplier based on the Moreira (2003) conditional likelihood ratio statistic. This confidence interval is very similar to the one based on the usual asymptotic normal approximation, extending from -0.06 to 1.12.

Although the difference between the two is not statistically significant, qualitatively it is interesting to note that the 2SLS estimate of the multiplier is larger than the OLS estimate. This could reflect a combination of (a) attenuation bias in the OLS estimates due to classical measurement error in the data on government spending changes, and (b) a countercyclical response of total government spending to contemporaneous output shocks. As long as my instrument is also uncorrelated with measurement error in government spending, the 2SLS estimator will correct both sources of downward bias in the OLS estimates.[15]

The last two columns of Table III report 2SLS estimates of the multiplier, using two alternative measures of World Bank disbursements as instruments. These alternative specifications are useful for illustrating the rationale behind my predicted disbursements instrument built up from project-level disbursement flows. In the third column, I use changes in actual total disbursements on World Bank-financed projects, rather than predicted disbursements excluding those associated with projects approved in the same year, as the instrument. In this case, the 2SLS estimate of the multiplier of 0.67 is somewhat larger than in the case where predicted disbursements are used as the instrument, although the difference between the two is small relative to estimated standard errors. To understand this

[15] Unfortunately it is not possible to definitively isolate the relative importance of these two sources of downward bias in the OLS estimates. However, a suggestive exercise is to ask how much classical measurement error alone would be required to rationalize the observed difference between the OLS and 2SLS estimates. Since the 2SLS estimate is roughly double the OLS estimate, standard textbook calculations imply that the signal-to-noise ratio in government spending would have to be approximately one in order to account for this difference, i.e. the variance of measurement error would have to be equal to the variance of actual government spending changes. If measurement error were less (more) severe than this benchmark, this would imply that in addition changes in government spending would have to be counter- (pro-) cyclical in order to account for the observed gap between the OLS and 2SLS results.

difference, recall that actual disbursements in a given country-year include two potentially-endogenous responses to current shocks: (a) disbursements on projects approved in the current year, as well as (b) deviations from typical disbursement rates for projects approved in previous years. Both of these have been eliminated from the predicted disbursement measure, which excludes disbursements on projects approved in the same year, and relies on typical rather than actual disbursement rates. The fact that the 2SLS estimates of the multiplier are larger when using actual disbursements is consistent with the idea that these two potentially-endogenous components of disbursements are procyclical. This might be the case if there are difficulties in project implementation when countries experience negative growth shocks, leading to a decline in disbursements. An extreme example might be a country that falls into civil conflict, triggering both lower growth and a suspension of World Bank activity, including both new project approvals as well as disbursements on existing projects.

The last column sheds some light on which of these two potentially-endogenous components of World Bank disbursements is more important. In this column, I again use actual disbursements as the instrument, but now excluding disbursements that are associated with projects approved in the same year. This eliminates the first of the two potentially endogenous components of World Bank disbursements, i.e. the portion associated with new project approvals in the same year. In this case, the 2SLS estimate of 0.63 is only slightly smaller than in the previous column, where total disbursements were used as an instrument. This evidence, while admittedly indirect, suggests that cyclically-motivated World Bank project approvals are the less important of the two potentially endogenous components of total disbursements. Instead, deviations from predicted disbursements are more likely to be an important source of correlation between actual disbursements and contemporaneous macroeconomic shocks. This is why I eliminate them from my instrument based on predicted disbursements.

The picture that emerges from these first basic results in Table III is clear: the government spending multiplier is reasonably precisely estimated to be small, at 0.48, with a 95 percent confidence interval extending from just below zero to just above one. This suggests a rather modest 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 multiplier is so small in the sample of primarily low-income countries I consider.

4.2 Sensitivity of Results to Influential Observations

Figure III provides a visual summary of the relationships in the data underlying the benchmark results in Table III. The three panels of Figure III report the reduced-form relationship between changes in output and changes in predicted disbursements on World Bank projects; the first-stage relationship between changes in total government spending and changes in predicted disbursements; and the structural relationship between changes in output and changes in total spending, with the OLS and 2SLS regression lines superimposed. A striking feature of these graphs is that there are very large fluctuations in all three variables in this sample of primarily low-income countries. In light of this, a natural concern is that the results in Table III could be sensitive to a small number of influential observations.

I explore this possibility in Table IV, by means of four robustness checks. First, I re-estimate the OLS and 2SLS regressions in the first and second columns of Table III 29 times, dropping one country at a time from the sample. In the first panel of Table IV, I report the minimum and maximum across these 29 samples of the OLS and 2SLS slope coefficients and t-statistics, as well as the minimum and maximum first-stage F-statistics. Dropping individual countries has relatively little impact on my results. The OLS estimates are very stable across specifications, varying between 0.20 and 0.25. The 2SLS estimates of the multiplier vary a bit more, between 0.34 and 0.59. However, the basic conclusion that the estimated multipliers are insignificantly different from zero is preserved across nearly all specifications: the maximum t-statistic for the null hypothesis that the multiplier is zero is 2.09, and it is below two in all other samples. Finally, it is worth noting that my first-stage regressions are reasonably strong across all subsamples, with first-stage F-statistics ranging from 10.8 to 19.0.

In the remaining three panels of Table IV, I consider the robustness of my results to dropping potentially influential individual data points rather than entire countries. I do this by using three standard rules to identify potentially influential observations in the first-stage and reduced-form OLS regressions underlying the 2SLS estimates (the slope in the latter being the ratio of the two former slopes). In the second panel, I use the covariance ratio statistic, which measures changes in the precision of OLS estimates as individual observations are dropped from the sample, while in the third panel, I use the DFITS measure which captures changes in OLS slopes as individual observations are dropped from the sample (see Belsely, Kuh and Welsch (1980) for details). Finally, in the fourth panel, I use the Hadi (1992) procedure for identifying potentially influential observations. I then re-estimate the OLS, 2SLS, and first-stage regressions eliminating the corresponding set of possibly-influential data points.

Removing influential observations in this way has minimal impact on the OLS estimates of the multiplier, which range from 0.13 to 0.22. The 2SLS estimates are somewhat more unstable: using the covariance ratio statistic to eliminate influential observations reduces the estimated multiplier sharply to 0.09, while using the DFITS statistic reduces it to 0.18. On the other hand, using the Hadi (1992) procedure increases the estimated multiplier to 0.82, but at the same time substantially increases the standard error as well. In all three cases, the basic conclusion that the estimated multiplier is insignificantly different from zero is preserved. For all three methods, the first-stage F-statistics remain above the Staiger and Stock (1997) rule of 10, and are much larger using the DFITS and covariance ratio statistics. While there is some variability in the 2SLS estimates, overall, this first set of robustness checks based on eliminating potentially influential observations is broadly consistent with the benchmark conclusion that the estimated multiplier is small and reasonably precisely estimated.

4.3 Robustness to Concerns about the Identifying Assumption

I next consider several further robustness checks designed to address potential concerns with my basic identifying assumption that project approvals in previous years are uncorrelated with shocks to growth in the current year. A first potential objection is that, while project approval decisions in previous years were made prior to the realization of current macroeconomic shocks, these shocks may have been anticipated 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 projects are approved in response to shocks that are persistent over time, then subsequent disbursements on those projects will be correlated with subsequent shocks, in violation of my exclusion restriction.

A natural way to address this first concern is impose a longer lag between project approvals and disbursements. In my benchmark measure of predicted disbursements, I do this by eliminating all disbursements on projects approved in the same year. In the first panel of Table V, I construct an alternative measure of predicted disbursements that discards disbursements on projects approved in the current *and* previous year, and use it as an instrument for changes in government spending. Unfortunately, this substantially weakens the first-stage relationship between changes in predicted disbursements and changes in total government spending, with a first-stage F-statistic of only 5.52. The

point estimate of the multiplier falls to 0.13, and a weak instrument-consistent 95% confidence interval is very wide, ranging from -0.69 to 0.86 and including the benchmark estimate of 0.48.[16]

Another way of addressing the same concern is simply to control for lagged growth. I do this in the second panel of Table V. In this sample of countries, annual growth rates exhibit little persistence, as lagged growth is not significantly correlated with contemporaneous growth (conditional on the included country and year fixed effects). In light of this, controlling for lagged growth has virtually no impact on the estimated multiplier, which changes marginally, from 0.48 in Table III to 0.44 in Table V. Controlling for lagged growth also does not appreciably weaken identification: the first stage F-statistic declines only slightly from 15.0 to 13.9 when I control for lagged growth. These findings suggest that persistence in growth rates, combined with a possible correlation between project approvals and shocks to growth, is not seriously undermining my identification strategy.

Another possible concern with the identifying assumption is that lagged changes in government spending themselves might matter for contemporaneous growth. This might be the case if the stimulative effects of government spending take more than a year to fully materialize. If, in addition, changes in disbursements on World Bank-financed projects are persistent over time, then my instrument may be correlated with the omitted variable of lagged spending changes, in violation of my exclusion restriction. In the third panel of Table V, I add the lagged change in government spending to address this possibility, and use the lagged change in predicted disbursements as an additional instrument. I do find that the coefficient on lagged government spending is positive, but it is small at 0.18, and it is not significantly different from zero. The coefficient on contemporaneous changes in government spending also increases a little relative to the benchmark specification, from 0.48 to 0.58, and now is significantly different from zero at the 10 percent level. In this specification, the sum of the estimated coefficients on current and lagged changes in government spending captures the cumulative

[16] Yet another strategy for dealing with this concern would be to eliminate all disbursements on projects that were approved for cyclical reasons in response to current shocks. As discussed in more detail in Section 5, it is difficult to come up with a fully-convincing rule for distinguishing between cyclically-motivated and non-cyclically-motivated projects. One crude possibility is simply to eliminate projects corresponding World Bank adjustment lending operations (known as Structural Adjustment Loans or Development Policy Loans). These typically involve budget support to countries experiencing adverse economic shocks, and so are more likely to be cyclically motivated than investment lending operations which generally finance long-term development projects. Eliminating all such projects from the construction of my predicted disbursements instrument and re-estimating the core specification in Table III delivers a slightly smaller estimate of the multiplier of 0.38 with a standard error of 0.33, which is reasonably strongly-identified with a first-stage F-statistic of 12.8.

effect of spending on output after two years, and is equal to 0.76 with a standard error of 0.41, implying a 95% confidence interval extending from -0.04 to 1.56.

A third potential concern is that World Bank project approvals might be correlated with improvements in policy performance which themselves have direct effects on growth, rather than the government spending itself. This could be because disbursements on World Bank projects trigger reforms through some form of policy conditionality.[17] Or, World Bank project approvals themselves could be triggered by policy improvements that are persistent over time. In either case, this would undermine my basic identifying assumption by creating a correlation between changes in predicted disbursements and changes in policy performance that are not included in my core specification. Moreover, to the extent that World Bank projects are associated with growth-enhancing policy changes, this would lead to an upward bias in my estimated multipliers.

The natural way to address this concern is simply to control for changes in policy, which I do in the fourth panel of Table V. The specific policy 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 eligible borrowing countries. The CPIA provides a rating on a six-point scale of the quality of policies and institutions, based on a checklist of various policy areas.[18] Annual changes in the CPIA are significantly positively correlated with changes in output in both the OLS and 2SLS specifications. However, including this control variable has only small effects on the estimated multipliers. The OLS estimate is nearly identical to the benchmark specifications in Table III. The 2SLS estimate of 0.44 is only slightly smaller than in the benchmark specifications. Qualitatively, at least, this is consistent with the hypothesis of an upward bias in previous estimates due to the omission of contemporaneous policy changes . Moreover, the 2SLS estimate of the multiplier is still reasonably strongly identified, with a first-stage F-statistic of 14.4. Overall, this evidence suggests that although policy reforms as captured by

[17] As noted earlier, by relying on predicted rather than actual disbursements, I have removed the country-specific component of disbursements on existing projects which might be correlated with policy reforms: for example, a missed policy reform that might both reduce growth and delay disbursements on existing projects. However, a remaining potential source of concern is that project approvals themselves might trigger subsequent policy reforms, and for this reason controlling for policy directly is important. For a skeptical view on whether World Bank lending actually influences the policy choices of borrowers, see Easterly (2005).

[18] 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 practical advantage of the CPIA data in this context is that it has full coverage of all country-year observations used in this paper.

changes in the CPIA are in fact significantly correlated with growth, they are not sufficiently correlated with changes in disbursements on World Bank projects to undermine my identification strategy.

In the last panel of Table V, I implement all three robustness checks simultaneously, controlling for lagged growth, lagged changes in spending, and changes in policy. The results here are similar to those in the previous panels where these robustness checks were implemented separately. I find no evidence that lagged growth, or lagged spending changes, are significantly correlated with current growth, while changes in policy performance are significantly correlated with growth. Finally, the estimated multiplier increases only slightly relative to the benchmark specification in Table III, from 0.48 to 0.51, but remains insignificantly different from zero at conventional levels. Taken together, these robustness checks support the conclusion from the basic results that the spending multiplier is small, and reasonably precisely estimated to be in the vicinity of 0.5.

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

The estimated multipliers in the previous section are small, with a baseline point estimate from Table III of only 0.48. This is well below the values for the multiplier of one or more that are often assumed in policy discussions.[19] My baseline estimate is also at the low end of estimated multipliers in the empirical literature. For example, Hall (2009) summarizes the existing evidence for the United States based on VARs and military spending shocks as delivering one-year multipliers ranging from 0.5 to one, while Ramey (2011a) summarizes a plausible range of estimates from 0.8 to 1.5. My estimate of 0.48 is also considerably smaller than most of the estimates of sub-national government spending multipliers in the United States and other advanced economies -- the studies along these lines cited in the introduction of this paper deliver estimates of spending multipliers ranging from 0.9 to over 2.0.[20]

A general reason for these differences between my findings and most of the rest of the literature could simply be the very different sample of primarily low-income countries where my identification strategy based on fluctuations in disbursements on World Bank projects is feasible. In this

[19] For example, some initial estimates of the projected effects of the American Recovery and Reinvestment Act of 2009 implied a one-year impact multiplier of around 1.4. See Cogan, Cwik, Taylor, and Wieland (2010) for a discussion and critique of these projected spending impacts.

[20] One exception to this pattern of large multipliers can be found in the work of Ilzetzki, Mendoza and Vegh (2010), who find VAR-based estimates of the one-year spending multiplier in the vicinity of 0.3 in their sample of emerging markets.

section, I examine the extent to which several specific features of my data, methodology, and sample of countries could account for the low estimates of the multiplier. I first consider the consequences of the fact that my instrument is based on fluctuations in World Bank-financed spending, that, although plausibly uncorrelated with contemporaneous shocks, are also likely to be anticipated at the time of project approval. I also consider the implications for the size of the spending multiplier of the fact that much of government spending in my sample of aid-dependent countries is financed by concessional loans and outright grants. Finally, I briefly discuss the possibility that the types of nominal rigidities emphasized by New Keynesian models as contributing to larger spending multipliers might be less prevalent in my sample of countries.

5.1 Role of Anticipation Effects

One 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 such 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 (2011b). 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 (2011b) 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 (2011b) 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 anticipated increase in spending

eventually 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 happened, and now investment (consumption) are declining (increasing) back to their steady-state levels. At first glance, this provides a good candidate explanation for why the estimated multipliers are so small – it could simply be that most of the private sector response to the increase in spending occurs 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.

I empirically investigate this possibility in two ways. First, I document the estimated contemporaneous effects of changes in government spending on the major expenditure components of GDP. I do this by re-estimating Equation (2), but replacing the dependent variable in turn with changes in private consumption and investment, government consumption and investment, and net exports. Table VI reports OLS results, as well as 2SLS estimates using changes in predicted disbursements on World Bank projects as an instrument for changes in total government spending.[21] Not surprisingly, I find that government consumption and investment expenditures increase when total government spending increases – this is true almost as a matter of arithmetic. Of more interest to understanding the potential role of anticipation effects are the responses of private consumption and private investment. These are substantially smaller than the effects on government consumption and investment, and are not significantly different from zero, so that it is difficult to draw very firm conclusions. Qualitatively however, it is noteworthy than in both the OLS and 2SLS results, private consumption falls when government spending increases, while private investment increases. This pattern of consumption and investment responses to changes in government spending looks 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.[22]

[21] 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 the sample, and so the sample size is slightly smaller than in Table III.

[22] A possible explanation for this pattern of positive investment responses to actual spending changes 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 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

Another way to assess the importance of anticipation effects is to examine directly the contemporaneous correlation between World Bank project approvals and changes in output. If anticipation effects are important, one should expect that output is more likely to react at the time that projects are approved, rather than when the spending is actually implemented. The difficulty in doing this is, of course, that some project approvals are likely to be endogenous responses to contemporaneous macroeconomic conditions. As a result, I cannot simply regress changes in output on contemporaneous approvals of World Bank projects in order to capture responses to the associated anticipated future spending increases.

In order to identify the impact of project approvals on output, it is necessary to somehow distinguish between projects that were approved for cyclical reasons and those that were not. I do this by consulting the documentation for individual projects to assess whether or not their stated objectives suggest that the projects were approved in 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 1975 projects registering disbursements at some point since 1985). Instead, I selected a random sample of 10 projects in each of 10 countries in my sample, and consulted the statement of project objectives for these 100 projects.[23] In this subset of projects, I identified 10 projects in which the stated objectives clearly suggested a 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 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 sample of 29 countries, 79 percent of projects accounting for 67 percent of total project value fall in this category of slow-disbursing projects that plausibly are not cyclically-motivated. Moreover, this is a conservative way

---

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 work effort 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.

[23] For completed projects I relied on the Implementation Completion Report, while for currently ongoing projects I relied on either Project Approval Documents or Project Information Documents. These documents for individual projects are publicly available at http://www.data.worldbank.org.

of identifying such non-cyclical projects, because it is likely that at least some of the projects that eventually disburse in less than four years were not approved for cyclical reasons.[24]

I next measure total disbursements over the life of each slow-disbursing project, aggregate this across all slow-disbursing projects approved in each country-year, and express this total as a fraction of GDP in the approval year. I then include this as a control variable in the benchmark specifications, on the assumption 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 VII. The 2SLS estimate of the coefficient on the announcement of slow-disbursing projects is 0.13 and is reasonably precisely estimated, so that the estimated impact is marginally statistically significant. This suggests that there is some weak evidence of output responses to the announcement of approvals of World Bank-financed projects, consistent with the neoclassical mechanism. Interestingly, however, controlling for these anticipation effects has virtually no impact on the contemporaneous estimated spending multiplier, which is nearly identical to that in the baseline specification in Table III.

Taken together, the evidence on anticipation effects is somewhat mixed. On the one hand, the qualitative patterns of responses of private consumption and investment to fluctuations in government spending induced by changes in predicted disbursements on World Bank projects are consistent with theoretically-predicted responses to unanticipated spending changes. On the other hand, controlling directly for approvals of plausibly acyclical World Bank projects suggests a weakly significant positive output response to the project approvals themselves. However, controlling for anticipation effects does not appear to matter much for the size of the estimated multiplier once the spending actually occurs.

5.2 Role of Concessional Financing of Government Spending

A key feature of the countries in my sample is that they are heavily dependent on foreign aid, not just from the World Bank, but from other donors as well. In fact, for the median country-year observation in my sample, net official development assistance (ODA) accounts for 12 percent of GDP, and 54 percent of total government spending. This observation has implications for the relevance of the

[24] For example, in the random subsample of 100 projects, 8 of those that I classify as non-cyclical nevertheless disburse in four years or less. In light of the earlier discussion in Footnote 16, it is also interesting to note that my categorization of slow- and fast-disbursing projects corresponds very closely to the classification of projects as adjustment versus investment lending operations. Investment lending operations, which generally finance long-term development projects, account for 95 percent of projects and 87 percent of lending volumes in my sample of slow-disbursing projects. In contrast, they account for only 26 percent of projects and 11 percent of lending volumes in my sample of fast-disbursing projects.

basic neoclassical mechanism that delivers a positive government spending multiplier. According to this mechanism, when government spending increases, private wealth falls by the present value of future taxes that are required to eventually finance the spending increase. This negative wealth effect leads to a positive labour supply and output response as private agents rebuild their wealth.

To the extent that increases in government spending are financed by ODA, the associated future tax burden associated with these expenditures will naturally be much lower than if government spending were financed by borrowing on market terms. In the extreme case where ODA takes the form of outright grants to the recipient country, there is no change in the future tax burden when aid-financed spending increases, and the neoclassical mechanism would predict no output response at all (provided that this spending is unproductive). Historically, however, much of aid has been provided in the form of concessional loans which do entail future repayment obligations. A typical example would be loans provided by the International Development Association (IDA), the concessional lending facility of the World Bank. These loans are highly concessional: 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 low, and in particular 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. And if in addition private agents assign some non-zero probability to eventual future debt relief, the expected present value of future repayments could be even lower than this.

This suggests that a possible explanation for the small estimated multipliers in my sample of poor countries is that these countries are able to finance government spending cheaply through foreign aid, thereby undermining the basic neoclassical mechanism for a positive spending multiplier. One way to shed light on this hypothesis is to investigate whether spending multipliers are smaller in country-year observations in my sample where foreign aid inflows are relatively more important. I do this by constructing the ratio of net ODA to total government spending, and dividing my sample of country-year observations in half at the median value of this ratio. Since net ODA measures grants plus net disbursements on concessional loans, this is a reasonable proxy for access to low-cost financing for government spending. I then estimate spending multipliers separately in the two subsamples.

The results are reported in Table VIII. I find that both the OLS and the 2SLS estimates of the multiplier are smaller when estimated using the relatively more aid-dependent observations. For example, the 2SLS estimates are 0.66 in the low-aid sample versus 0.13 in the high-aid sample.

However, these differences are small relative to the estimated standard errors. Unsurprisingly, I find that identification is stronger in the highly aid-dependent subsample (with a first-stage F-statistic of 10.2) than in the low-aid sample (where the first-stage F-statistic is only 7.3). This reflects the fact that World Bank-financed spending is also relatively more important as a share of total government spending in this subsample. In contrast, the much larger estimate of the multiplier is far more weakly identified in the low-aid subsample. While the difference in the estimated multipliers is not strongly identified, qualitatively the pattern of results in Table VIII is consistent with the hypothesis that access to concessional finance undermines the neoclassical mechanism for a positive spending multiplier, by reducing the burden of future taxes associated with a given amount of government spending. This in turn could contribute to an explanation of why the estimated multiplier is small in my sample of aid-dependent countries.

5.3 Role of New Keynesian Rigidities

A third set of potential explanations for the small estimated multipliers has to do with the relative importance in my primarily low-income country sample of leading rigidities emphasized by New Keynesian models as key determinants of the size of the government spending multiplier. Two such rigidities are (a) the presence of "rule-of-thumb" consumers whose consumption is tied to their current income, rather than their lifetime income, and (b) sticky prices. Gali, Lopez-Salido, and Valles (2007) and Hall (2009) provide a detailed discussion of how these two rigidities, and especially their interaction, can result in larger output and consumption responses to changes in government spending. Briefly, the presence of "rule of thumb" consumers by assumption dampens the crowding out of private consumption that occurs when government spending increases in the neoclassical model, since these consumers do not reduce consumption in response to anticipated higher future taxes. The presence of sticky prices reduces the decline in real wages that occurs when labour supply increases, and so enables a more positive response of consumption, particularly for "rule-of-thumb" consumers.

If such rigidities are less prevalent in the primarily low-income countries that comprise my sample, in principle this could help to rationalize the low estimated multipliers that I obtain.[25] However, it seems difficult to marshal convincing empirical evidence for such an explanation. If "rule-of-thumb"

[25] Although, it is important to note that even calibrated models with these New Keynesian features do not necessarily deliver multipliers that are all that large, or even greater than one. For example, Cogan, Cwik, Taylor, and Wieland (2010) use a calibrated version of the Smets and Wouters (2007) model of the US economy, including sticky wages and rule of thumb consumers, and find multipliers associated with the American Recovery and Reinvestment Act of 2009 that are substantially smaller than one.

consumption behaviour is due to the absence of well-developed financial markets that provide households with opportunities to smooth consumption over time, if anything it would seem more plausible that "rule-of-thumb" behaviour is more --rather than less -- prevalent in low-income countries where financial market imperfections most likely are more severe. Other things equal, this should imply larger, rather than smaller, spending multipliers in my sample of primarily low-income countries. Regarding sticky prices, a growing literature has documented the frequency of disaggregated nominal price changes, primarily in industrial countries, as well as a handful of emerging markets including Israel, Poland, Brazil, and Argentina (see Alvarez, Gonzalez-Rozada, Neumeyer, and Beraja (2011) for a detailed survey of existing evidence). This literature has primarily documented a strong positive association between the frequency of price changes and average inflation rates, and has as yet not documented any systematic association between frequency of price changes and levels of development, conditional on the inflation rate. Absent such information, it is unclear whether such nominal rigidities emphasized by New Keynesian models are likely to be more or less prevalent in the set of poor countries analyzed in this paper. And for this reason it is unclear whether we should expect estimated multipliers to be systematically larger or smaller in the sample of primarily low-income countries that are considered here.

6. Is the Multiplier Different for Different Types of Spending?

Thus far, my empirical evidence has focused on the size of the overall government spending multiplier, using changes in predicted disbursements on World Bank projects 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 government spending. Yet there are various breakdowns of total spending for which one might reasonably wonder whether such an assumption is valid. A first distinction is between government purchases of goods and services and other components of spending, such as interest payments, transfers and subsidies. Most of the theoretical work on government spending multipliers focuses on the effects of government purchases alone, which may have different short-run output effects than transfers, subsidies, or interest payments. Within government purchases, a further distinction is between consumption and investment expenditures, for which theoretical models can also imply very different short-run effects. For example, Baxter and King (1993) use a calibrated neoclassical model to show how increases in government investment result in large positive output effects that increase over time as the stock of productive public capital increases. On the other hand, nonproductive

government consumption expenditures in their model imply much smaller output effects, both on impact and in the long run.

Yet another possibility highlighted by the focus of this paper is the distinction between the short-run stimulative effects of World Bank-financed and non-World Bank-financed spending. To the extent that World Bank projects focus on long-run development objectives, their short-run stimulative effects might be different from those of other types of government spending. Within World Bank-financed spending, a further distinction has to do with the extent to which it finances purchases of domestic versus foreign goods and services, as the domestic stimulative effect of the latter might be considerably less than the former. To take an extreme scenario, if a World Bank-financed project consists primarily of purchases of equipment 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.

Any such differences in multipliers across expenditure categories imply that 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 increases in any of one these particular categories of spending. To illustrate this point more specifically, consider this simple generalization of Equation (2):

$$\Delta y_t = \beta_A \Delta g_t^A + \beta_B \Delta g_t^B + \varepsilon_t \quad (4)$$

where $g_t = g_t^A + g_t^B$ is a decomposition of total government spending into two categories A and B, and the slope coefficients β_A and β_B measure the impacts of the two types of expenditures on output. Ideally, I would simply estimate these disaggregated spending multipliers directly. Unfortunately, however, two key data limitations make this direct approach infeasible. The first is simply that, as discussed in more detail below, reliable and comprehensive data on the decomposition of total government spending into the categories described above is not available for many of the low-income countries in my sample. The second is that, for most decompositions of total government spending, I cannot provide a matching decomposition of my predicted disbursements instrument. For example, while I can (imperfectly) distinguish current versus capital spending in the aggregate government spending data, I cannot do the same at the level of disbursements on individual World Bank projects on which my instrument is based. As a result, I also cannot construct predicted disbursements measures

separately for these two types of spending, which would be a natural approach to identifying the corresponding disaggregated multipliers.

Instead, my more modest objective in this section is to provide some suggestive evidence on the extent to which disaggregated spending multipliers might differ from the aggregate multipliers I have estimated, given the limited available data on the composition of spending. To do so, it is useful to first observe that the probability limit of the aggregate spending multiplier based on 2SLS estimation of Equation (2), will be a linear combination of the two disaggregated multipliers:

$$\hat{\beta} = \frac{COV(\Delta y_t, \Delta d_t)}{COV(\Delta g_t, \Delta d_t)} = \frac{\gamma_A \beta_A + \gamma_B \beta_B}{\gamma_A + \gamma_B} \tag{5}$$

where γ_A and γ_B are the slopes of the first-stage regressions of the two categories of government spending on the aggregate predicted disbursements instrument. This expression highlights that the aggregate spending multiplier will be close to the category A (category B) spending multiplier if the first-stage regression primarily reflects the relationship between the instrument and fluctuations in category A (category B) spending. Although I cannot separately identify the disaggregated spending multipliers β_A and β_B, this expression is nevertheless useful in that the weights on the two multipliers depend on observable first-stage relationships between the various categories of spending and the aggregate predicted instrument. Knowledge of these weights, together with the estimated aggregate spending multiplier, can give a sense of plausible magnitudes for the disaggregated spending multipliers.

Table IX reports results for various decompositions of total government spending. Within each panel, I report the first-stage relationship between changes in category A spending and changes in the aggregate predicted disbursements instrument, γ_A, and the ratio $\gamma_A/(\gamma_A + \gamma_B)$. I also report the corresponding estimate of the category A spending multiplier under the extreme assumption that the multiplier for category B spending is zero. From Equation (5), this is simply $\hat{\beta}_A = \hat{\beta}\frac{\gamma_A+\gamma_B}{\gamma_A}$. In panel A of Table IX, I consider the distinction between non-interest and interest expenditures, denoting the former as category A spending. As noted earlier, a basic problem is that data on the functional composition of total government spending is extremely sparse in the sample of primarily low-income countries that I consider. For example, the IMF's WEO database, which is my primary source for total government spending, reports data on total interest payments for only 25 percent of the country-year observations

in my sample of countries for which total spending is available.[26] A less comprehensive, but much more widely-available source of information on interest payments is the World Bank's Global Development Finance database, which reports data on interest payments on public- and publicly-guaranteed external debt for all of the country-years in my sample. The disadvantage of this data, however, is that it covers only interest payments to foreign, and not domestic, creditors. The first two columns of Table IX show that the first-stage slopes obtained from regressing changes in total spending, and changes in non-interest spending, on the predicted disbursement instrument deliver very similar slopes (the ratio $\gamma_A/(\gamma_A + \gamma_B)$ is equal to 0.9). Since the first-stage relationship for changes in total government spending primarily reflects fluctuations in non-interest spending, this implies that the estimated multiplier using only non-interest spending is nearly identical to the one using total spending, at 0.53 (as compared to 0.48 in the benchmark specification).

I next consider netting out transfers and subsidies from non-interest spending to arrive at government purchases. As is the case with interest expenditures, standard sources such as the IMF's Government Finance Statistics Yearbook report only very sparse data on this category of expenditures, for the sample of countries considered in this paper. Instead, I approximate government purchases by summing government consumption expenditures from the national accounts with government investment data from the IMF's World Economic Outlook database. While in principle this shoudl correspond to total government purchases, in practice the sum of national accounts data on government consumption and investment often exceeds my total government non-interest spending measure -- implying implausible negative transfers and subsidies.[27] A crude way of addressing this problem is simply to drop those observations for which implied transfers and subsidies are negative. Doing so reduces my sample from 610 to 452 country-year observations. In this much-reduced sample, the second panel of Table IX reveals that the first-stage slopes estimated using total spending and government purchases are quite similar. This suggests that, in this sample, the first-stage regression primarily reflects the relationship between predicted disbursements and government purchases.

[26] Data on interest payments can be inferred from the WEO series on the overall and primary fiscal balance. Data in the IMF's Government Finance Statistics Yearbook on the composition of spending is similarly sparse in my sample of countries.

[27] In addition to plain measurement error, one possible substantive explanation for this pattern of negative transfers is that these correspond to observations for which the total spending data is less comprehensive than the investment and consumption data. If, for example, the spending data refers to the central government while the investment data and national accounts government consumption data are based on a broader consolidated public sector concept, this could contribute to the finding that my measures of consumption and investment exceed the total spending measure.

Identification, however, is weaker in the sense that the first-stage F-statistics are only 9.8 and 9.6 using total spending and purchases, respectively. Although less precisely estimated, the multiplier of 0.49 using this imperfect measure of government purchases again turns out to be quite similar to the benchmark estimate of 0.48 from Section 3.

In the third panel of Table IX, I consider the distinction between government investment and non-investment expenditures, denoting the former as category A spending. In contrast with the previous two panels, I find that the slope of the first-stage relationship between changes in investment spending and changes in the predicted disbursements instrument is only about half as large as the first-stage relationship for total spending. This indicates that the instrument has predictive power for both investment and non-investment of spending. This in turn implies that a 2SLS regression of changes in output on changes in investment spending alone (such as that reported in the last column of Table IX) will deliver a consistent estimate of the investment multiplier only under the extreme assumption that the multiplier for non-investment spending is zero. Given this very strong identifying assumption, I obtain an estimate for the investment multiplier of 0.90 with a standard error of 0.53.[28] However, as is clear from Equation (5), assuming positive values for the non-investment spending multiplier β_B would necessarily imply a considerably smaller multiplier for investment spending, given that $\gamma_A/(\gamma_A + \gamma_B)$ is only equal to 0.54.

In the last two panels of Table IX, I focus on World-Bank-financed spending. As discussed above, simple Keynesian mechanisms for government spending multipliers operating through expansions in domestic aggregate demand depend crucially on whether spending occurs at home or abroad. To investigate whether this distinction is important empirically in my sample, ideally I would like to have information on the import content of total government spending. While such information is not available for total government spending, it is possible to roughly approximate the import content of World Bank-financed spending, and so I focus on this alone.[29] I do this by matching 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

[28] Qualitatively this larger multiplier for investment spending is consistent with the findings of Ilzetzki, Mendoza, and Vegh (2010) who document larger investment multipliers in their sample of emerging markets, using VAR-based identification techniques.

[29] While information on the cyclical variation in the import content of total government spending is not available for the countries in my sample, it is noteworthy that both the OLS and 2SLS results in Table VI indicate a sizeable decline in net exports when government spending increases. This suggests that at least some of the increase in government spending falls on imports.

the supplier to which each procurement contract was awarded, as well as 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 which is tendered to suppliers situated in the country where the project is located. I then aggregate these project-level estimates of domestically-spent annual disbursement flows up to the country-year level to arrive at an estimate of actual disbursements that are spent domestically. The domestic share of World Bank spending is substantial: pooling all country-year observations in my sample, the median share of total disbursements that is contracted domestically is 43 percent, and the 25^{th} and 75^{th} percentiles of this ratio are 31 percent and 57 percent, respectively.

In the fourth panel of Table IX , I identify the foreign component of World Bank-financed public spending as category B spending, and the balance of total spending as category A. Excluding this estimate of spending abroad has minimal impact on my estimates of the multiplier. The first-stage slopes for total spending and domestic spending are very similar. The estimated spending multiplier for domestic spending, under the polar assumption that the multiplier for World Bank-financed spending abroad is zero, is accordingly slightly larger than the multiplier for domestic spending, at 0.58.

In the final panel of Table IX, I distinguish between total World Bank-financed and non-World Bank-financed spending. As noted above, it could be the case that the short-run stimulative effects of World Bank-financed spending, which tends to focus on projects with long-run development objectives, might be different from those of other forms of government spending. To investigate this possibility, in this panel I identify non-World Bank-financed spending as category A spending. In this case, the first-stage regression of changes in non-World Bank-financed spending on changes in predicted disbursements delivers a smaller and less-precisely estimated slope than when total spending is used (0.71 versus 1.03). Under the extreme assumption that the short-run multiplier for World Bank-financed spending is zero, this implies a slightly larger multiplier for non-World Bank-financed spending of 0.69.

Summing up, it is theoretically plausible that different types of government spending will have different short-run effects on output. Here I have considered the possibility that multipliers might be different for different subcomponents of total government spending, including government purchases, government investment, and domestic and total World Bank-financed spending. My ability to identify disaggregated multipliers is limited by imperfect data on the composition of spending for the primarily low-income countries that are the focus of this paper, as well as the fact that it is not possible to decompose my predicted disbursements instrument in the same way as I can (imperfectly) decompose

total spending. There is, however, suggestive evidence that multipliers for non-interest spending, government purchases, and government spending excluding World Bank-financed spending abroad are probably not very different from the aggregate spending multipliers reported in the rest of the paper.

6. Conclusions

In this paper, I have proposed a novel way to identify fluctuations in government 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 government 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 responses to contemporaneous macroeconomic shocks, they are unlikely to be correlated with future macroeconomic shocks. Under this assumption, changes in disbursements in a given year that are associated with project approvals in previous years can be used as an instrument for changes in total government spending in order to estimate spending multipliers. I implement this approach in a set of 29 primarily low-income countries for which systematic evidence on the cyclical effects of government spending has until now not been available. The resulting 2SLS estimates of the government spending multiplier are for the most part small, and reasonably precisely estimated to be in the vicinity of 0.5. 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.

A puzzling feature of my basic results is that the estimated spending multipliers are so small. My baseline estimates in Table III suggest an impact multiplier of 0.48, which is at the low end of existing estimates of the multiplier. Most of these estimates in the existing literature are based on data from advanced economies and rely on very different identification strategies. It is therefore possible that the small estimates of the multiplier that I find are due to unique characteristics of the set of countries I work with in this paper, as well as the specifics of my identification strategy.

I investigate a range of potential explanations along these lines. A key feature of the spending shocks I rely on for identification is that they are anticipated in advance. However, while there is some modest evidence that output reacts to the announcement of World Bank projects, this does not undermine my conclusion that the contemporaneous spending multiplier, once the spending occurs, is small. A somewhat more promising explanation is based on the fact that the countries in my sample on

average finance a significant portion of government spending with foreign aid. To the extent that this reduces the burden of future taxation for a given amount of spending, this weakens the standard neoclassical mechanism operating through wealth effects. Consistent with this, I do find even smaller estimates of the multiplier in the most aid-dependent half of my sample. Another possibility is that my results are due to a lower prevalence of key rigidities contributing to larger multipliers in New Keynesian models. Absent direct evidence of a lower prevalence of rule-of-thumb consumption behaviour and/or a lesser degree of price stickiness in my primarily low-income country sample, this explanation for small estimated multipliers is necessarily speculative. Overall, the absence of a definitive explanation for the small size of the estimated multipliers 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.

Although the basic estimates of the government spending multiplier presented here suggest a rather modest role for countercyclical fiscal expansions during economic downturns as a way of stimulating aggregate economic activity, this does not mean that there is no role for government spending in response to adverse macroeconomic shocks. For example, in many of the countries in my sample, 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 government spending in general, and World Bank-financed spending in particular, on growth. However, my empirical work is focused on estimating the short-run relationship between annual fluctuations in government spending and output within countries, and does not address these longer-term growth effects of aid and/or government spending.

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 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 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 plausibly exogenous government spending changes and output changes, that can be used to discriminate between alternative theoretical models of the short-run impacts of fiscal policy.

Data Appendix

World Bank Project-Level Disbursement Data

I retrieve project-level disbursement data from the World Bank's internal Business Warehouse Database. This source provides electronic records on project-level disbursements beginning in 1985. Disbursements for earlier years are available only as paper records. My sample covers 7443 projects that were active at some point between 1985 and 2009. For each project I have the full time series of quarterly disbursements until project closing (for closed projects) or until the fourth quarter of 2009 (for active projects as of end-2009). I first aggregate the quarterly flows to calendar year frequency, resulting in a series of annual disbursement flows for all projects, beginning in 1985.[30] Next, for each country-year observation, I construct total disbursements by aggregating the project-level disbursements on active projects in that country approved in the current year and the previous 10 years. By construction, this excludes a small fraction of total disbursements that are associated with projects with disbursement profiles longer than 10 years. I do this for consistency with the measure of predicted disbursements, described next. Nevertheless, this measure of disbursements aggregated up from individual projects coincides almost perfectly with the sum of total disbursements on World Bank loans and IDA grants as reported in the World Bank's Global Development Finance at the country-year level.

To construct predicted disbursements, I first calculate actual annual disbursement rates for each project, as the fraction of the total amount disbursed in each year of the life of the project. I measure total disbursements in two different ways. For projects approved and completed between 1985 and 2009, I have the full series of annual disbursements, and I simply sum these to obtain the total amount disbursed on the project. Note that these total realized disbursements for the project need not coincide with the original amount initially committed to the project, which can be larger or smaller than eventual total disbursements. For projects that are still active as of 2009, the total amount that will eventually be disbursed is not yet known, and so instead I use the total initial commitment to the project as the denominator when calculating disbursement rates. Similarly, for projects approved prior to 1985, I do not have the full annual disbursement flows required to calculate total realized disbursements for the project. For these projects I also use the total initial commitment as the denominator when calculating disbursement rates.

With these annual project-level disbursement rates in hand, I next construct a 10-year typical disbursement profile for all projects in the same economic sector and geographical region. The

[30] In a handful of projects, disbursements are negative in some years. This can occur if, for example, a procurement contract issued by the recipient government and financed by World Bank funds was subsequently found not to adhere to relevant Bank policies, and the funds are then repaid to the World Bank. I drop these cases of negative disbursements.

distribution of projects across sectors in my sample is indicated in Appendix Table A1.[31] I calculate the typical disbursement profile as the simple unweighted average disbursement rate for all projects in the same sector/region bin. For example, I take all Transport sector projects approved in Africa between 1985 and 2009, and then compute the average disbursement rate across all such projects, for the year of project approval as well as the 10 subsequent years.[32] I then calculate a time series of predicted total disbursements for each individual project within each sector/region bin by applying the estimated average disbursement profile to the total amount disbursed on the project.

Data on the distinction between domestic and foreign World Bank spending is based on contract-level data, also from the Business Warehouse database. Electronic records on individual World Bank procurement contracts contain information on the fraction of the contract value that is sourced by the country of origin of the supplier. I first match each individual procurement contract to its corresponding project using World Bank project identification codes. I then retrieve the portion of each contract's value that is awarded to suppliers in the country where the project takes place, and express it as a fraction of total procurement associated with the project. This is only an imperfect indicator of where disbursements on World Bank projects 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.

Other Data

Total general government spending is obtained from the IMF's World Economic Outlook database (series code GGX). For countries in Africa, I augment this with data from the World Bank's African Development Indicators, which reports data on total government expenditure, extending back to the 1980s for many countries. Data on public investment are taken from the internal unpublished version of the World Economic Outlook database, which contains unofficial estimates of public investment gathered by IMF country economists from a variety of sources. In most cases, this data corresponds closely to similar data on public investment flows reported in the African Development Indicators. However, I rely on the WEO database because of its greater country-year coverage.

[31] Many World Bank projects span multiple sectors. Data is available on the percent of project value assigned to up to five different sectors. I assign each project to the largest single sector in which the project is active.

[32] An alternative would be to omit the project in question from each sector-region average. However, with 7443 projects spread over 60 sector-region bins, the number of projects in each bin is sufficiently large that including or excluding the project in question has negligible effects on the results.

National accounts data on GDP and its expenditure components are taken from the World Bank's World Development Indicators database, available at www.data.worldbank.org.

Real Changes in Variables

Data on total government spending, public investment, and the expenditure components of GDP are all retrieved in current local currency units. For each such variable X_t, I construct its constant price change as: $\left(\frac{X_t}{P_t} - \frac{X_{t-1}}{P_{t-1}}\right) / \left(\frac{Y_{t-1}}{P_{t-1}}\right)$, where Y_t is GDP in current local currency units, and P_t is the overall GDP deflator. World Bank disbursements are in current US dollars. I construct real changes in these variables analogously as $\left(\frac{X_t e_t}{P_t} - \frac{X_{t-1} e_{t-1}}{P_{t-1}}\right) / \left(\frac{Y_{t-1}}{P_{t-1}}\right)$, where e_t is the annual average exchange rate expressed in local currency units per US dollar.

Sample of Countries

The sample of countries used in the paper is based on data availability and the importance of World Bank lending. I begin by identifying 42 countries for which (a) at least 20 annual observations on World Bank disbursements are available, and (b) the average over time of World Bank disbursements as a share of GDP is greater than one percent. Of these, I discard 13 countries for which fewer than 20 years of government spending data and World Bank disbursements are available. This results in the sample of 29 countries reported in Table I. In the regression sample, I further eliminate a small number of observations corresponding to implausibly large fluctuations in GDP or total government spending.

Data Availability

The Business Warehouse database from which the project-level data for this paper is obtained is not publicly available. However, World Bank project-level commitments are publicly available through the OECD Creditor Reporting System database since 1995, as are project-level disbursement flows since 2002 (see http://stats.oecd.org/Index.aspx?DatasetCode=CRSNEW). Basic data on all World Bank projects, including commitment amounts and selected project documentation, is available through www.data.worldbank.org. Finally, data on individual procurement contracts since 2000 are publicly available at www.worldbank.org/projects. Country-level aggregate actual and predicted disbursement series, but not the underlying project-level data, are available in the replication dataset for the paper.

In Panels D and E of Table V, I use data from the World Bank's Country Policy and Institutional Assessments (CPIA). CPIA ratings are publicly available, but only since 2005, and only for countries eligible to borrow from the International Development Association (IDA), the concessional lending window of the World Bank. CPIA data for these countries can be found at www.worldbank.org/ida. However, the full CPIA series used to generate these results is confidential.

In columns 2 and 4 of Table VII, and panels B and C of Table IX, I use data on total public investment taken from the internal World Economic Outlook database of the International Monetary Fund. As noted above, this data is not reported publicly in the World Economic Outlook database.

DEVELOPMENT RESEARCH GROUP

THE WORLD BANK

Appendix Table A1: Distribution of World Bank Projects by Sector

	All Countries		29 Country Sample	
	Number of Projects	*Percent of Total*	*Number of Projects*	*Percent of Total*
Agriculture	1,313	17.64	336	17.01
Public Administration	1,279	17.18	414	20.96
Transport	862	11.58	207	10.48
Energy	772	10.37	166	8.41
Eduation	696	9.35	189	9.57
Health	690	9.27	198	10.03
Water	587	7.89	133	6.73
Industry	494	6.64	117	5.92
Finance	420	5.64	104	5.27
Other	330	4.43	111	5.62
Total	**7,443**	**100**	**1,975**	**100**

This table reports the sectoral distribution of all World Bank projects active between 1985 and 2009 (first panel), and in my 29 country sample over the same period (second panel). Projects are assigned to sectors using the "major sector" data available for individual projects in the World Bank's project portal. Each project has a fraction assigned to up to five sectors, and I identify the sector of the project as the largest single major sector share.

References

Acconcia, Antonio, Giancarlo Corsetti, and Saverio Simonelli, "Mafia and Public Spending: Evidence on the Fiscal Multiplier from a Quasi Experiment," CEPR Discussion Paper No. 8305, 2011.

Alvarez, Fernando, Martin Gonzalez-Rozada, Andrew Neumeyer, and Martin Beraja, "From Hyperinflation to Stable Prices: Argentina's Evidence on Menu Cost Models," unpublished, University of Chicago, 2011.

Barro, Robert, "Output Effects of Government Purchases," *Journal of Political Economy*, 89 (1981), 1086-1121.

Barro, Robert, and Charles Redlick, "Macroeconomic Effects from Government Purchases and Taxes," *Quarterly Journal of Economics,* 126 (2011), 51-102.

Baxter, Marianne, and Robert King, "Fiscal Policy in General Equilibrium," *American Economic Review*, 83 (1993), 315-334.

Belsely, D.A., E. Kuh, and R.E. Welsch, *Regression Diagnostics: Identifying Influential Data and Sources of Collinearity* (New York: Wiley, 1980).

Blanchard, Olivier, and Roberto Perotti, "An Empirical Characterization of the Dynamic Effects of Changes in Government Spending," *Quarterly Journal of Economics,* 117 (2002), 1329-1368.

Chodorow-Reich, Gabriel, Laura Feiveson, Zachary Liscow, and William Gui Woolston, "Does State Fiscal Relief During Recessions Raise Employment? Evidence from the American Recovery and Reinvestment Act," unpublished, Stanford University, 2011.

Clemens, Michael, Steven Radelet, Rikhil Bhavnani, and Samuel Gozzi, "Counting Chickens When They Hatch: Timing and the Effects of Aid on Growth," *Economic Journal*, forthcoming, 2010.

Clemens, Jeffrey, and Stephen Miran, "The Effects of State Budget Cuts on Employment and Income," unpublished, Harvard University, 2010.

Cogan, John, Tobias Cwik, John Taylor, and Volker Wieland, "New Keynesian Versus Old Keynesian Government Spending Multipliers," Journal *of Economic Dynamics and Control,* 34 (2010), 281-295.

Cohen, Lauren, Joshua Covall, and Christopher Malloy, "Do Powerful Politicians Cause Corporate Downsizing?" NBER Working Paper No. w15839, 2010.

Easterly, William, "What Did Structural Adjustment Adjust? The Association of Policies and Growth with Repeated IMF and World Bank Adjustment Loans," *Journal of Development Economics*, 76 (2005), 1-22.

Fatás, Antonio, and Ilian Mihov, “The Case for Restricting Fiscal Policy Discretion,” *Quarterly Journal of Economics*, 118 (2003), 1419-1448.

Fisher, Jonas, and Ryan Peters, “Using Stock Returns to Identify Government Spending Shocks,” *The Economic Journal*, 120 (2010), 414-436.

Fishback, Price, and Valentina Kachanovskaya, “In Search of the Multiplier for Federal Spending on the States During the New Deal,” NBER Working Paper No. w16561, 2010.

Hadi, A. S., "Identifying Multiple Outliers in Multivariate Data," *Journal of the Royal Statistical Society,* 54 (1992), 761–771.

Hall, Robert, “By How Much Does GDP Rise if the Government Buys More Output?,” Brookings Papers on Economic Activity, 2009.

Ilzetzki, Ethan, and Carlos Végh, “Procyclical Fiscal Policy in Developing Countries: Truth or Fiction?,” unpublished, University of Maryland and LSE, 2008.

Ilzetzki, Ethan, Enrique Mendoza, and Carlos Végh, “How Big (Small?) Are Fiscal Multipliers?,” unpublished, University of Maryland and LSE, 2010.

Kilby, Christopher, "An Empirical Assessment of Informal Influence in the World Bank," unpublished, Villanova University, 2010.

Leeper, Eric, Todd Walker, and Susan Yang, “Government Investment and Fiscal Stimulus,” *Journal of Monetary Economics*, 57(2010), 1000-1012.

Leeper, Eric, “Monetary Science, Fiscal Alchemy,” prepared for Jackson Hole Conference, 2010.

Nakamura, Emi, and Jon Steinsson, "Fiscal Stimulus in a Monetary Union," unpublished, Columbia University, 2011.

Ramey, Valerie, and Matthew Shapiro, “Costly Capital Reallocation and the Effects of Government Spending,” Carnegie-Rochester Conference Series on Public Policy, 48 (1998), 145-194.

Ramey, Valerie, "Can Government Purchases Stimulate the Economy," *Journal of Economic Literature*, forthcoming, 2011a.

Ramey, Valerie, “Identifying Government Spending Shocks: It’s All in the Timing,” *Quarterly Journal of Economics*, 126 (2011b), 1-50.

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

Roodman, David, “Through the Looking-Glass, and What OLS Found There: On Growth, Foreign Aid, and Reverse Causality,” Center for Global Development Working Paper No. w137, 2008.

Serrato, Juan Carlos Suarez, and Philippe Wingender, “Estimating Local Fiscal Multipliers," unpublished, University of California Berkeley, 2011.

Shoag, Daniel, “The Impact of Government Spending Shocks: Evidence on the Multiplier from State Pension Plans," unpublished, Harvard University, 2010.

Smets, Frank, and Rafael Wouters, "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97 (2007), 586-606.

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

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

Table I: Sample of Countries

Code	Country	Number of Observations	Average Disbursements on World Bank Projects (Share of GDP)	Total Government Spending (Share of GDP)	World Bank Disbursements (Share of Total Government Spending)
BDI	Burundi	24	4.3%	29.3%	14.7%
BEN	Benin	24	1.5%	18.9%	7.4%
BFA	Burkina Faso	24	2.0%	19.8%	10.2%
BOL	Bolivia	24	1.1%	28.5%	4.0%
CAF	Central African Republic	24	1.6%	17.2%	7.3%
CIV	Cote d'Ivoire	24	1.3%	24.9%	5.2%
COM	Comoros	24	2.0%	27.1%	7.1%
CPV	Cape Verde	23	1.7%	38.0%	5.1%
ETH	Ethiopia	24	1.9%	21.5%	9.4%
GHA	Ghana	24	3.0%	27.1%	12.5%
GIN	Guinea	24	1.7%	18.6%	9.4%
GMB	Gambia	24	3.1%	26.6%	11.0%
JOR	Jordan	24	1.2%	36.4%	2.9%
KEN	Kenya	24	1.3%	22.9%	5.7%
LSO	Lesotho	24	2.0%	54.9%	3.7%
MAR	Morocco	24	1.0%	24.8%	2.8%
MDG	Madagascar	24	2.9%	18.5%	16.1%
MLI	Mali	24	2.5%	24.5%	10.1%
MWI	Malawi	24	4.6%	29.5%	15.8%
NER	Niger	24	2.2%	19.1%	11.8%
RWA	Rwanda	24	2.6%	22.0%	13.0%
SEN	Senegal	24	1.7%	31.2%	5.7%
SLE	Sierra Leone	24	3.0%	22.1%	12.7%
TCD	Chad	24	2.0%	19.5%	11.4%
TGO	Togo	24	1.9%	21.4%	8.0%
TUN	Tunisia	24	1.1%	29.0%	3.7%
TZA	Tanzania	21	2.8%	18.9%	15.4%
UGA	Uganda	24	3.1%	17.1%	18.6%
ZMB	Zambia	24	3.4%	29.1%	10.7%
	Average	23.9	2.2%	25.5%	9.4%

This table reports the sample of countries included in the empirical analysis. The sample consists of 29 countries for which (a) World Bank disbursements average more than one percent of GDP over the period 1985-2009, and (b) at least 20 annual observations on GDP, total government spending, and World Bank disbursements are available. Details on the selection of countries is available in the Data Appendix.

Table II: Summary Statistics

	Standard Deviation	*Correlation With Change In:*			
		GDP	Government Spending	Actual Disbursements	Predicted Disbursements
Change in:					
GDP	0.035	1			
Government Spending	0.027	0.182	1		
Actual Disbursements	0.014	0.076	0.143	1	
Predicted Disbursements	0.005	0.071	0.187	0.113	1

This table reports summary statistics on the indicated variables, for the 610 observation regression sample covering 29 countries over 1985-2009. All variables are scaled by lagged GDP and are expressed as deviations from country and year averages.

Table III: Basic Results

Panel A: Estimates of Spending Multiplier				
(Dependent variable is Change in Real GDP)				
	OLS	**2SLS**	**2SLS**	**2SLS**
Change in Total Government Spending	0.228***	0.475	0.665	0.627
	(0.0547)	(0.287)	(0.434)	(0.450)
Weak Instrument Consistent 95% CI		[-0.064 , 1.121]	[-0.037 , 1.757]	[-0.054 , 1.630]
Panel B: First-Stage Regressions				
(Dependent variable is Change in Total Government Spending)				
		OLS	**OLS**	**OLS**
Change in Predicted Disbursements		1.029***		
		(0.266)		
Change in Actual Disbursements			0.281***	
			(0.0936)	
Change in Actual Disbursements				0.380***
Excluding Projects Approved in Same Year				(0.133)
F-statistic on Excluded Instrument		14.97	9.03	8.15
Number of Observations	610	610	610	610

*** (**) (*) denotes significance at the 1 (5) (10) percent level. Heteroskedasticity-consistent standard errors are clustered at the country level. All regressions are estimated using pooled country-year data and include a full set of country and year fixed effects. Changes in GDP, government spending, and disbursements are all scaled by lagged GDP. Weak instrument consistent confidence intervals computed using the Moreira (2003) conditional likelihood ratio statistic. Panel A reports OLS and 2SLS estimates of Equation (2). Panel B reports OLS estimates of first-stage regression in Equation (3).

Table IV: Robustness Checks: Dropping Influential Countries and Observations

		Panel A: Dropping Individual Countries		*Panel B: Using Covariance Ratio Statistic to Drop Influential Observations*		*Panel C: Using DFITS Statistic to Drop Influential Observations*		*Panel D: Using Hadi Procedure to Drop Influential Observations*	
(Dependent variable is Change in Real GDP)									
Estimation Method:		**OLS**	**2SLS**	**OLS**	**2SLS**	**OLS**	**2SLS**	**OLS**	**2SLS**
Change in Total Government Spending	Min	0.202	0.336	0.134**	0.0851	0.172***	0.184	0.217***	0.815
	Max	0.249	0.591	(0.0550)	(0.180)	(0.0497)	(0.181)	(0.0575)	(0.487)
T-statistic	Min	3.823	1.228						
	Max	4.575	2.088						
F-statistic on Excluded Instrument	Min		10.75		23.82		34.71		10.37
	Max		18.98						
Number of Observations				569	569	527	527	605	605

*** (**) (*) denotes significance at the 1 (5) (10) percent level. Heteroskedasticity-consistent standard errors are clustered at the country level. All regressions are estimated using pooled country-year data and include a full set of country and year fixed effects. Changes in GDP, government spending, and predicted disbursements are all scaled by lagged GDP. All four panels report OLS and 2SLS estimates of Equation (2), using changes in predicted disbursements as the instrument for changes in government spending in the 2SLS results. Panel A reports minimum and maximum estimated slope coefficients, t-statistics, and first-stage F-statistics across a set of 29 samples dropping one country at a time. Panel B reports OLS and 2SLS results for a subsample of observations excluding influential observations in the corresponding first-stage, and reduced form regressions, using the covariance ratio statistic to identify influential observations. Panel C does the same as Panel B but using the DFITS statistic to identify influential observations. Panel D does the same as Panel B but using the Hadi procedure to identify influential observations.

Table V: Robustness Checks, Cont'd: Possible Violations of Exclusion Restrictions

	Panel A: Excluding Current and Previous Year Approvals		*Panel B: Controlling for Lagged Growth*		*Panel C: Controlling for Lagged Spending*		*Panel D: Controlling for Policy*		*Panel E: Combining Panels B-D*	
(Dependent Variable is Change in Real GDP)										
Estimation Method:	**OLS**	**2SLS**	**OLS**	**2SLS**	**OLS**	**2SLS**	**OLS**	**2SLS**	**OLS**	**2SLS**
Change in Total Government Spending	0.228***	0.126	0.221***	0.437	0.232***	0.581*	0.225***	0.441	0.229***	0.511
	(0.0547)	(0.347)	(0.0517)	(0.296)	(0.0519)	(0.304)	(0.0537)	(0.284)	(0.0485)	(0.307)
Lagged Change in Real GDP			0.0304	0.0205					0.00434	-0.0495
			(0.0739)	(0.0682)					(0.0836)	(0.0886)
Lagged Change in Total Government Spending					0.0955	0.180			0.0844	0.336
					(0.0618)	(0.283)			(0.0707)	(0.266)
Change in CPIA							0.0121**	0.0116**	0.0123**	0.0131*
							(0.00517)	(0.00536)	(0.00591)	(0.00655)
F-statistic on Excluded Instrument		5.52		13.91				14.37		
Stock-Yogo Weak Identification Test (critical value for 10% (15%) maximal size distortion is 7.03 (4.58))						8.93				5.05
Number of Observations	610	610	584	584	576	576	610	610	568	568

*** (**) (*) denotes significance at the 1 (5) (10) percent level. Heteroskedasticity-consistent standard errors are clustered at the country level. All regressions are estimated using pooled country-year data and include a full set of country and year fixed effects. Changes in GDP, government spending, and disbursements are all scaled by lagged GDP. All five panels report OLS and 2SLS estimates of Equation (2), using changes in predicted disbursements as the instrument for changes in government spending in the 2SLS results. Panel A estimates Equation (2) using an alternative measure of changes in predicted disbursements as the instrument, which excludes disbursements on projects approved in the same year and the previous year. Panel B extends Equation (2) by adding lagged GDP growth and treating it as exogenous. Panel C extends Equation (2) by adding the lagged change in government spending, and uses current and lagged changes in predicted disbursements as instruments. Panel D extends Equation (2) by adding the contemporaneous change in the CPIA measure of policy and treating it as exogenous. Panel E implements the three extensions of Equation (2) in Panels B-D simultaneously.

Table VI: Effects on Expenditure Components of GDP

Dependent Variable is Change in:	**Private Consumption**	**Private Investment**	**Government Consumption**	**Government Investment**	**Net Exports**
Panel A: OLS Estimates					
Change in Total Government Spending	-0.0462	0.0811*	0.236***	0.277***	-0.317***
	(0.0870)	(0.0439)	(0.0365)	(0.0412)	(0.0620)
Panel B: 2SLS Estimates					
Change in Total Government Spending	-0.195	0.160	0.402**	0.498***	-0.514
	(0.213)	(0.182)	(0.160)	(0.169)	(0.311)
F-statistic on excluded instrument	33.8	33.8	33.8	33.8	33.8
Number of observations	590	590	590	590	590

*** (**) (*) denotes significance at the 1 (5) (10) percent level. Heteroskedasticity-consistent standard errors are clustered at the country level. All regressions are estimated using pooled country-year data and include a full set of country and year fixed effects. Changes in all variables are scaled by lagged GDP. Table reports OLS and 2SLS estimates of Equation (2), but replacing the dependent variable with changes in the expenditure components of GDP as indicated in each column, and using changes in predicted disbursements as the instrument for changes in government spending.

Table VII: Anticipation Effects

(Dependent variable is Change in Real GDP)		
Estimation Method:	**OLS**	**2SLS**
Change in Total Government Spending	0.223***	0.438
	(0.0531)	(0.285)
Approvals on Slow-Disbursing Projects	0.140*	0.125*
	(0.0760)	(0.0725)
F-statistic on Excluded Instrument		15.06
Number of Observations	610	610

*** (**) (*) denotes significance at the 1 (5) (10) percent level. Heteroskedasticity-consistent standard errors are clustered at the country level. All regressions are estimated using pooled country-year data and include a full set of country and year fixed effects. Table reports OLS and 2SLS estimates of Equation (2), extended to include approvals on slow-disbursing projects as a fraction of GDP, as described in the main text. Changes in predicted disbursements are used as the instrument for changes in government spending in the 2SLS results, and approvals on slow-disbursing projects are treated as exogenous.

Table VIII: Effects of Access to Concessional Financing of Government Spending

	Panel A: High Aid Dependence (ODA/Gov>0.54)			***Panel B: Low Aid Dependence (ODA/Gov<0.54)***	
(Dependent Variable is Change in Real GDP)					
Estimation Method:	**OLS**	**2SLS**		**OLS**	**2SLS**
Change in Total Government Spending	0.196**	0.130		0.299**	0.656
	(0.0876)	(0.269)		(0.109)	(0.498)
F-statistic on excluded instrument		10.19			7.31
Number of observations	305	305		305	305

*** (**) (*) denotes significance at the 1 (5) (10) percent level. Heteroskedasticity-consistent standard errors are clustered at the country level. All regressions are estimated using pooled country-year data and include a full set of country and year fixed effects. Changes in GDP, government spending, and disbursements are all scaled by lagged GDP. Both panels report OLS and 2SLS estimates of generalizations of Equation (2), using changes in predicted disbursements as the instrument for changes in government spending in the 2SLS results. Panel A reports results for a subset of highly aid-dependent observations with ODA greater than 54 percent of government spending. Panel B reports results for the complementary subset of less aid-dependent observations.

Table IX: Differential Effects of Different Types of Spending

	First-Stage Slopes:				
Category A Spending Defined As:	**Total Spending ($\gamma=\gamma A+\gamma B$)**	**Category A Spending (γA)**	**Ratio ($\gamma A/\gamma$)**		**Multiplier for Category A Spending Assuming Zero Multiplier for Category B**
Panel A: Non-Interest Spending					
Coef	1.029***	0.930***	0.90		0.525
S.E.	(0.266)	(0.246)			(0.322)
First-stage F-Statistic on Excluded Instrument	14.97	14.29			
Panel B: Government Purchases (Positive Implied Transfers Sample Only)					
Coef	0.996***	1.115***	1.12		0.489
S.E.	(0.319)	(0.360)			(0.311)
First-stage F-Statistic on Excluded Instrument	9.75	9.60			
Panel C: Investment Spending					
Coef	1.029***	0.542***	0.53		0.902*
S.E.	(0.266)	(0.177)			(0.526)
First-stage F-Statistic on Excluded Instrument	14.97	9.34			
Panel D: Excluding World Bank Spending Abroad					
Coef	1.029***	0.849***	0.83		0.575
S.E.	(0.266)	(0.222)			(0.369)
First-stage F-Statistic on Excluded Instrument	14.97	14.65			
Panel E: Excluding All World Bank Spending					
Coef	1.029***	0.713***	0.69		0.686
S.E.	(0.266)	(0.229)			(0.428)
First-stage F-Statistic on Excluded Instrument	14.97	9.72			

*** (**) (*) denotes significance at the 1 (5) (10) percent level. Heteroskedasticity-consistent standard errors are clustered at the country level. All regressions are estimated using pooled country-year data and include a full set of country and year fixed effects. Each panel of the Table corresponds to a different decomposition of total government spending into category A and category B spending (with category A spending as identified and category B spending as the remainder of total government spending). The first two columns report first-stage regressions of changes in total spending and changes in category A spending on the same changes in the predicted disbursements instrument. The final column reports the results of a 2SLS regression of changes in output on changes in category A spending, using changes in predicted disbursements as the instrument.

Figure I: Disbursement Rates on World Bank Projects

(Fraction of total project spending disbursed per year of life of project, average across all projects)

This graph reports the average fraction of spending on individual World Bank projects that is disbursed in year t of the project, with t=0 corresponding to the year of project approval. Averages are taken across all 7,443 World Bank projects active at some point between 1985 and 2009. A detailed description of the project-level disbursement data can be found in the Data Appendix.

Figure II: Fluctuations in World Bank Disbursements: Example of Zambia

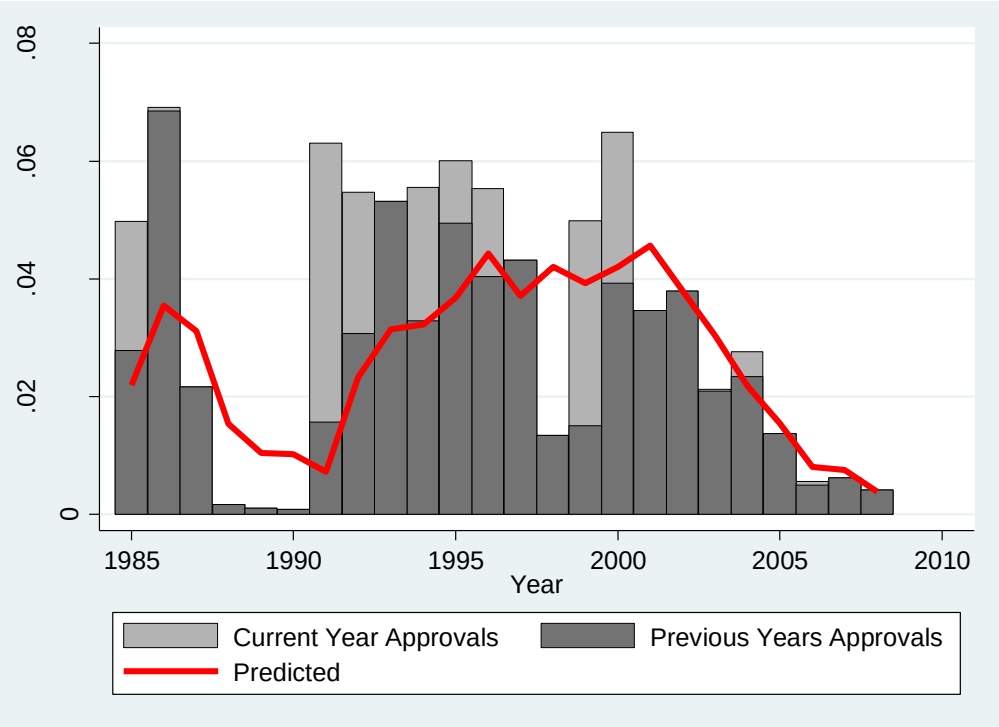

This figure plots three measures of World Bank disbursements to Zambia. The overall height of the bar corresponds to total actual disbursements on all IBRD and IDA projects. This consists of "Current Year Approvals", i.e. disbursements on projects approved in the same year, plus "Previous Years Approvals" i.e. disbursements on projects approved in previous years. "Predicted" refers to disbursements constructed as the product of total project size times typical disbursement rates for projects in the same region and sector, and then aggregated to the country-year level but excluding current-year approvals, as described in the main text and in the Data Appendix.

Figure III: Benchmark Results

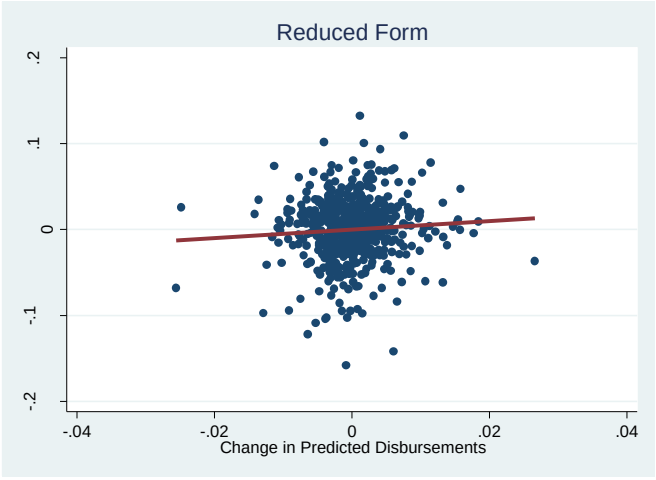

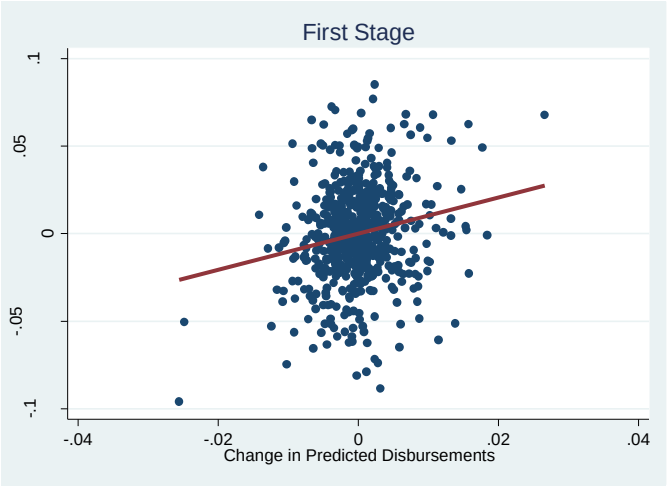

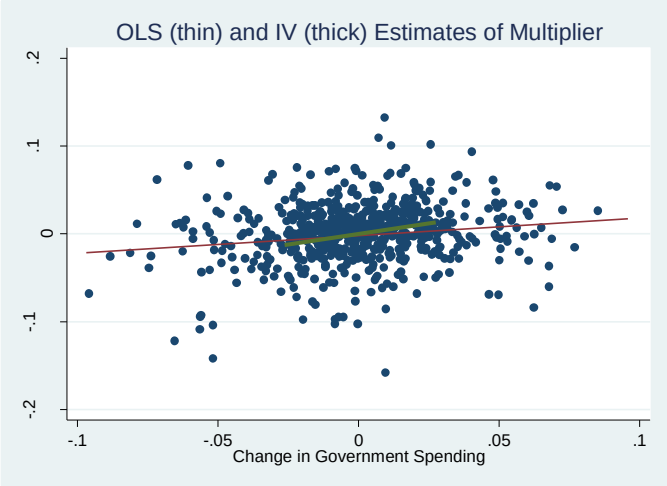

Note: All variables are scaled by lagged GDP and are expressed as deviations from country- and year-averages.