

Foreign Aid and Domestic Absorption

Jonathan R. W. Temple

Department of Economics, University of Bristol

jon.temple@bristol.ac.uk

Nicolas Van de Sijpe

Department of Economics, University of Sheffield

n.vandesijpe@sheffield.ac.uk

6th October 2014

ABSTRACT

We introduce a new ‘supply-push’ instrument for foreign aid, to be used together with an instrumental variable estimator that filters out unobserved common factors. We use this instrument to study the effects of aid on macroeconomic ratios, and especially the ratios of consumption, investment, imports and exports to GDP. We cannot reject the hypothesis that aid is fully absorbed rather than used to build foreign reserves or exiting as capital flight, nor do we find evidence of Dutch Disease effects. Aid leads to higher consumption, while the evidence that it promotes investment is less robust.

This is a revised version of Temple and Van de Sijpe (2014), using newer data. For various contributions and advice, we are grateful to Chris Adam, Erwin Bulte, Paddy Carter, Gabriela Cohen-Freue, Markus Eberhardt, Gerdie Everaert, Benedikt Goderis, Doug Gollin, Arne Risa Hole, Carlos Lamarche, Oliver Morrissey, Anita Ratcliffe, Abhijeet Singh, Francis Teal, Frank Windmeijer, Adrian Wood, seminar participants at the universities of Bristol, Heriot-Watt, Oxford, Royal Holloway and Sheffield, and participants at the 2014 conferences of the Centre for the Study of African Economies, the Econometric Society European Meeting, the International Association for Applied Econometrics, and the Royal Economic Society. The first author is grateful to the British Academy for financial support from a Mid-Career Fellowship.

1 Introduction

A large empirical literature studies the effects of aid on growth using cross-country data. It is fair to say that even its strongest adherents recognize the difficulties of interpreting the results. Researchers studying the effect of aid on growth must contend with the endogeneity of aid, the high persistence of output, the uncertain determinants of growth rates, nonlinear effects of aid, biases from measurement error, and the likelihood of substantial cross-country heterogeneity in the effects of aid. Moreover, since aid is given in many different forms and with a variety of motives, these regressions invite concerns that are not purely statistical. For its detractors, this literature uses unreliable data to arrive at fragile answers to the wrong question.

These criticisms may seem decisive, but some important questions are hard to answer without cross-country data. In this paper, we seek to advance the literature in two ways. First, we introduce a new ‘supply-push’ instrument for aid, to be used together with an instrumental variables estimator that filters out unobserved common factors, even when their effects differ across countries. In principle, this combination of instrument and estimator could be applied to a wide range of aid-related questions in future research.

Second, we shift the focus to whether and how foreign aid is absorbed by the domestic economy. Aid, as a capital transfer, is not part of measured GDP. The aid could be absorbed, by allowing increased domestic expenditure, but this is not the only possibility. The capital transfer may be offset by a corresponding capital outflow, or used to accumulate foreign exchange reserves. Some of the aid flows recorded by donors may not correspond to genuine transfers, since they may have been spent on technical assistance provided by foreign consultants, without ever reaching the recipient. In all these cases, aid is not absorbed by the domestic economy. For absorption to take place, domestic expenditure must increase relative to domestic production, implying an increase in net imports. Hence, we begin by examining the causal effect of aid on net imports.

We are also interested in how absorption takes place. Absorption requires an increase in at least one of the components of domestic final expenditure: household consumption, government consumption, and gross investment. We study the effects of aid on the ratios of these components to GDP, in the short run and the long run. This should help us to understand the potential effects of

aid. For example, if aid improves the investment climate, we would expect to see an increase in investment relative to GDP. We will argue that the effects of aid on macroeconomic ratios are inherently easier to study than the effects on steady-state GDP levels or growth rates.

The well-known identification problem in the cross-country literature is that aid is not randomly assigned. To address this problem, we introduce a supply-push instrument. It is based on the idea that the exposure of recipients to changes in donor budgets varies across recipients. Consider two aid recipients, A and B, and a single donor. Country A accounts for a larger share of aid from the donor, and this greater exposure persists over time. In that case, when the donor's budget increases for some exogenous reason, the movement in aid is larger for country A than for country B, driven solely by the changing supply of aid. This suggests the following instrument: we can construct a synthetic measure of aid at each date t , based on each country's share of aid in a donor budget at some initial date t_0 , multiplied by the current donor budget at date t .

To give a specific example, consider the two donors Britain and France. When Britain's total aid budget increases relative to that of France, former British colonies are likely to see an increase in aid received, relative to former French colonies. More generally, there will often be long-term connections between particular donors and recipients, so that recipients are more exposed to variation in some donor budgets than others. It is this form of variation that our synthetic measure will capture, isolating it from variation driven by the particular circumstances of individual aid recipients. In other words, we look for exogenous changes in aid receipts that are driven by changes to total donor budgets. We call this a supply-push instrument, since it is closely related to the work of Card (2001) on the labor market effects of immigration. As in the immigration setting, the destinations of flows from most 'sending' countries are likely to be sufficiently diverse that the instrument will not be correlated with time-varying recipient-specific circumstances.

A remaining objection to the supply-push instrument is that, in practice, donor budgets will be influenced by forces that are common to many recipients. For example, the state of world economic conditions is likely to affect donor generosity, and also the economic outcomes of poor countries. This could also be true of other global events or trends, ranging from climate conditions to political developments, such as the democratizations of the 1990s. Drawing on

recent work in the panel time series literature, these forces can be seen as unobserved common factors with loadings that differ across countries. We filter out the common factors using an instrumental-variable version of a common correlated effects (CCE) estimator. This class of estimators was introduced by Pesaran (2006) and extended to instrumental variables by Harding and Lamarche (2011). Once the factors have been filtered out, an argument that our instrument could be (statistically) endogenous is harder to construct.

The combination of the instrument and CCE estimators proves informative in our application. We find that aid is at least partially absorbed, reflected in significant increases in net imports. In fact, we cannot reject the hypothesis that aid leads to a one-for-one increase in net imports, corresponding to full absorption. This occurs mainly through an increase in imports rather than a decline in exports, and hence we do not find any symptoms of Dutch Disease. These findings hold across a range of estimators and robustness checks. There is similarly robust evidence that aid leads to increases in total consumption. This appears to be driven primarily by increases in household consumption, but those estimates are less precise unless we exclude outliers. The evidence that aid promotes investment is weaker. Aid sometimes appears to have a delayed effect on investment, but these results are not robust to alternative estimation methods or the exclusion of outliers. Hence the balance of evidence is that aid is absorbed at least partially, and this is achieved by increases in consumption and imports relative to output.

The rest of the paper has the following structure. In section 2, we sketch various possible relationships between aid and macroeconomic ratios. Section 3 explains the approach to estimation and how it relates to the previous literature. Section 4 describes the data. In section 5, we first analyze whether aid is absorbed, and whether there are Dutch Disease effects, before studying how aid is absorbed. Section 6 presents a number of robustness checks, before section 7 concludes.

2 Aid and macroeconomic ratios

In this section we set out the main ideas of the paper. From a national accounts perspective, foreign aid is a capital transfer which does not contribute directly to GDP, but in principle allows an increase in domestic expenditure on final

goods and services, relative to domestic production. As we noted above, this is not the only possibility, since aid may alternatively be used to accumulate foreign reserves, or lead to a capital outflow. Some aid may be devoted to forms of technical assistance which fund consultants from the donor country, with no direct effect on the aid recipient's domestic expenditure.

Domestic absorption is typically defined as the sum of household consumption, gross investment, and government consumption. We are interested in whether aid is absorbed, or put differently, whether aid is reflected in higher domestic expenditure on final good and services. We are also interested in the effect of aid on the separate components relative to GDP, in the short run and the long run. Investigating these links should contribute to a better understanding of the effects of aid.

We start with the question of what it means for aid to be fully absorbed. It helps to note the basic GDP identity:

$$Y \equiv C + I + G + X - M$$

where Y is GDP, C is household consumption, I is gross investment (private and public), G is government consumption, X is exports and M is imports. At the risk of stating the obvious, the identity should not be interpreted as a theory of how GDP is determined. It says only that the distinct uses of output (C, I, G, X) must sum to the total amount of output available: GDP plus imports, or $Y + M$. Our paper does not examine how output or productivity are determined, but how aid influences the different types of expenditure.

For the aid to be absorbed, at least one of C, I or G must increase, along with their total. If they increase relative to GDP, the identity implies that the ratio of net imports to GDP, $(M - X)/Y$, must also increase. There is nothing problematic about this; it is what must happen if aid is to permit greater domestic expenditure relative to the domestic production of goods and services.¹ In the short run, if aid received by a country is entirely devoted to higher domestic expenditure on final goods and services, net imports will rise one-for-one with aid. If the response of net imports is smaller than this, aid absorption is only partial.

¹In the terminology of Aiyar et al. (2006), the quantity $M - X$ is the non-aid current account deficit. For more background on aid absorption see Aiyar et al. (2006), Aiyar and Ruthbah (2008), Berg et al. (2010), Hansen and Headey (2010) and Hussain et al. (2009).

To study absorption, we take macroeconomic ratios, such as C/Y , I/Y , G/Y and $(M - X)/Y$, as our dependent variables.² The way aid is absorbed might differ between the short and the long run. For instance, in the short run, aid might be used to build foreign exchange reserves which are used to finance higher expenditure only later, so that full absorption is delayed.³ More generally, the relationships between aid and macroeconomic ratios could be complicated over longer time horizons. As the aid is spent, this will have indirect effects on the evolution of expenditure components relative to output. For example, if aid is spent in ways that greatly improve the investment climate, the long-run effect of aid on the investment rate could be much larger than the short-run effect. Or consider what happens if donor funds are spent entirely on foreign consultants: short-run absorption will be zero, but technical advice may be associated with subsequent changes in economic policies and hence in macroeconomic ratios. Our empirical analysis will take these possibilities into account by distinguishing between short-run and long-run effects, where the latter are estimated using dynamic models.

Some macroeconomic models make sharp predictions, notably the one-sector Ramsey model. If aid takes the form of grants made direct to households, as in Obstfeld (1999), then a permanent increase in aid raises the investment ratio in the short run, but not in the long run. Aid promotes faster convergence to the steady-state, but the long-run levels of the capital stock and GDP are invariant to aid. Along the balanced growth path, all aid is consumed. From a national accounts perspective, consumption is higher while investment and GDP are unchanged, and the increase in steady-state consumption is permitted by imports of the final good. This implies that when the ratio of aid to GDP increases permanently, the long-run C/Y and $(M - X)/Y$ ratios increase by the same absolute number of percentage points, leaving the other ratios unchanged.

The effects are more complicated in a two-sector model, such as a dynamic version of the dependent economy model, with traded and non-traded goods. In standard versions of that model, the relative price of non-traded goods is in-

²Since a linear combination of these dependent variables equals unity by construction, the model for one of the dependent variables will be statistically redundant when the covariates are the same across the regressions. But for ease of interpretation, we report results for each of the dependent variables.

³Berg et al. (2010) and Hussain et al. (2009) analyze these decisions in detail, emphasizing that absorption outcomes will typically be influenced by the actions of both the fiscal authority and the central bank, with scope for these to pull in different directions.

variant to aid in the long run.⁴ The effects on gross investment and exports are more complicated. Brock and Turnovsky (1994) and Brock (1996) showed that a small open economy's long-run adjustment to aid will depend on the relative capital intensities of the traded and non-traded sectors. If aid increases demand for traded and non-traded goods, domestic output of the latter must expand; this attracts capital and labour into the non-traded sector. With a long-run relative price that is invariant to the transfer, restoring equilibrium will require the steady-state capital stock (and hence gross investment) to be either higher or lower, depending on whether non-traded or traded production is most capital intensive. This suggests that aid could be associated with higher or lower gross investment in the long run.

It might be asked why we focus on macroeconomic ratios, rather than more conventional outcomes like growth. Given the difficulties of using cross-country data to study aid, there are at least three good reasons to start with the effects on macroeconomic ratios.⁵ First, their relationships with aid are more likely to be linear, as in the case of the Ramsey model just discussed. In contrast, the study of growth has to contend with the possibility of diminishing returns to aid. Second, aid is sometimes justified precisely as improving the conditions for domestic investment, an effect that might be relatively visible in the data. If we see consumption and investment as jump variables, they can respond quickly to changes in aid. These short-run responses of macroeconomic ratios to aid should be easier to detect than longer-term development benefits. Third, growth researchers have to contend with the possibility of slow convergence, and hence a high degree of persistence of GDP, perhaps to the extent that GDP contains a unit root. In contrast, the macroeconomic ratios we study are more likely to be stationary, and should be less persistent than GDP. Hence, it may be easier to establish reliable findings in our setting than in the case of aid and growth.

⁴In models with two sectors and two factors, with sectoral factor mobility and international capital mobility, the long-run relative price depends solely on supply conditions; see, for example, Obstfeld and Rogoff (1996, section 4.2).

⁵Some of the relevant econometric issues are discussed in Temple (2010).

3 Methods

With these arguments in mind, we study the effects of aid on macroeconomic ratios. The endogeneity problem arising from the non-random allocation of aid is a central issue. Even in a model that controls for country and time fixed effects, it is likely that aid flows and outcome variables are jointly influenced by one or more variables that are not readily measured, and hence will be omitted from the regression. A conventional instrumental variable approach can address this, but only if the instrument is uncorrelated with the error term. We are interested in achieving identification even when the error term may include a number of latent common factors, with factor loadings that vary across countries. This is a natural structure in the current context, where aid flows and macroeconomic outcomes are likely to be jointly influenced by hard-to-measure variables such as world economic conditions and other global events or trends.

Our chosen instrument has a supply-push form. We instrument aid using a weighted average of donor budgets, where the sets of weights are fixed over time but vary across aid recipients. To make this more precise, we are interested in the case where a country-specific time-varying variable A_{it}/Y_{it} (aid received by country i at time t divided by GDP) is instrumented by a synthetic predictor based on a fixed share of a common aggregate. In the case of aid with one donor, for instance, we have $(a_{i0}D_t)/Y_{it}$, where D_t is the donor budget and a_{i0} is the share of recipient i in that donor's aid budget at time zero. In the case of two donors, we have $(a_{i0}^1D_{1t} + a_{i0}^2D_{2t})/Y_{it}$, and so on. In the general case of N_D donors, the synthetic aid measure is therefore $A_{it}^S/Y_{it} \equiv \left(\sum_{d=1}^{N_D} a_{i0}^d D_{dt}\right)/Y_{it}$, where a_{i0}^d is the share of donor d 's total aid disbursements that recipient i receives, calculated over an initial period that is excluded from estimation, and D_{dt} is the total aid disbursement made by donor d in period t .

This form of presentation hints at the possible importance of latent factors with heterogeneous effects. Imagine the process generating the macroeconomic outcome of interest Q_{it}/Y_{it} is given by:

$$Q_{it}/Y_{it} = \beta (A_{it}/Y_{it}) + \varepsilon_{it} \quad (1)$$

$$\varepsilon_{it} = \phi_i f_t + u_{it} \quad (2)$$

where f_t is a vector of unobserved common factors (including, say, world eco-

economic conditions) and ϕ_i is a set of factor loadings which may vary across countries. This multifactor error structure nests both conventional fixed effects (where one common factor is time-invariant) and conventional period effects (where loadings on one time-varying factor are the same across countries) as special cases. The greater generality of this structure has made it a focus of recent research in econometric theory, and applications have begun to emerge in a variety of fields, as we note below.

We will assume that we do not have observable proxies for the common factors or their loadings. This means there are two possible sources of endogeneity: aid might be correlated with the effects of the omitted factors, the $\phi_i f_t$, or with the country-specific shock u_{it} . For a conventional fixed-effects IV estimator to be consistent, we would need our supply-push instrument A_{it}^S/Y_{it} to be uncorrelated with both, and hence with ε_{it} . This could easily be questioned. For example, it is plausible that donor budgets will be correlated with world economic conditions and trends which also influence macroeconomic ratios. In that case a supply-push instrument will be correlated with $\phi_i f_t$ even if there is no correlation with u_{it} .

This suggests the need to go beyond conventional IV estimation. We will present results which filter out the common factors using the approach of Pesaran (2006). His paper introduced common correlated effect (CCE) estimators for panel data. This class of estimators proxies for the combined effects of common factors using weighted averages of the cross-section averages of the observable variables, where the weights are estimated from the data and vary across countries. This is done by augmenting the regression with cross-section averages of the dependent variable and of the explanatory variables, all with country-specific coefficients. The CCE approach has been extended to the case of instrumental variables by Harding and Lamarche (2011), yielding a CCE IV estimator. We report the results from several methods, but give most emphasis to the CCE IV estimator, as the one most likely to yield consistent estimates. Recent applications of CCE estimators include Baltagi and Li (2014), Bond et al. (2010), Eberhardt et al. (2013), Holly et al. (2010) and Imbs et al. (2011).

The CCE approach can accommodate various forms of cross-section dependence, and has been shown to perform relatively well even in small samples.⁶

⁶Relevant papers include Chudik et al. (2011), Kapetanios et al. (2011), and Pesaran and Tosetti (2011).

A remaining limitation of standard CCE approaches is that factor loadings which are correlated with the regressors can lead to inconsistent estimates. This problem does not arise when a suitable instrument is available: see Harding and Lamarche (2011).⁷ They also present simulation evidence in which the CCE IV estimator performs well even when the factor loadings are correlated with the regressors. If we denote our instrument at time s by Z_{is} , our maintained assumption is that $E(Z_{is}u_{it}|\phi, f) = 0$, but Z_{is} may be correlated with ϕ or with f .

It is useful to contrast the multifactor error structure with the assumptions of conventional fixed effects estimators. If the common factors had homogeneous effects, we could have proxied for them using time fixed effects, the standard approach. Sometimes, common factors with heterogeneous effects can be proxied by interacting time fixed effects with observed, country-specific variables.⁸ Although that approach is more general than time fixed effects, it heavily restricts the structure of the unobserved, heterogeneous factor loadings. Hence, it is less general than the approach we adopt.

The CCE IV estimator of Harding and Lamarche (2011) can be implemented using 2SLS. The difference from conventional 2SLS is that cross-section averages of the observable variables are included in the first and second stage, with country-specific coefficients. One possible concern with the supply-push instrument is that it has a factor structure: it is a weighted average of donor budgets, with sets of weights (initial budget shares) that vary across aid recipients. This raises a concern: perhaps the instrument would be eliminated from the first stage of 2SLS estimation when filtering out the common factors. It is easy to show, however, that with one endogenous variable and one instrument, the instrument is only eliminated from the first stage in two unlikely cases: when there is either a single donor or the initial shares of aid recipients in donor budgets are the same across donors.⁹ Since in practice there are multiple donors and budget shares differ across donors, in principle our instrument will retain explanatory power in the first stage, even conditional on the inclusion of cross-section means

⁷For a theoretical analysis of the pooled CCE estimator when factor loadings are correlated with the regressors, see Westerlund and Urbain (2013).

⁸This latter approach, which Breinlich et al. (2014) call *proportional time effects*, has often been used for sub-national data, but less often in cross-country research.

⁹Under the stated conditions, the instrument would be perfectly collinear within each country with the cross-section mean of the instrument, and then identification fails.

with country-specific coefficients. Moreover, this is a testable assumption.

So far, we have said nothing about dynamic aspects of the specification. Pesaran (2006, p. 975) notes that common feature dynamics across the units (here, countries) are captured through the serial correlation structure of the common factors. But a remaining concern with our initial-share instrument is that circumstances specific to individual aid recipients, such as their domestic political developments, may be serially correlated. For each country, the initial share in a donor's budget may then be correlated with shocks in some of the subsequent periods, which undermines the exogeneity of the instrument. This potential limitation of the supply-push approach is acknowledged by Card (2001, footnote 23). It is likely to be an especial concern for the earlier time periods of the panel, and when there are relatively few time periods overall. Later in the paper, we investigate the problem by dropping some of the early time periods from the estimated model. This means that the initial share is measured some years before the first time period used for estimation. When we do this, we find no warning signs that our main results are substantially affected by underlying serial correlation in country-specific circumstances.

We now discuss the relation of our paper to previous work. Among recent cross-country studies, Clemens et al. (2012) and Rajan and Subramanian (2008) are especially well known. Both studies investigate the effect of aid on growth, but Clemens et al. (2012) argue that the literature has failed to identify credible instruments. Since aid is unlikely to be randomly assigned, either across countries or over time, the lack of an instrument has posed a major problem. The solution that we adopt, the supply-push instrument, was first used by Van de Sijpe (2010) to study the effect of aid on governance, but without allowing for a multifactor error structure. A related synthetic measure of aid, but based on average shares in donor budgets rather than initial shares, was used in Hodler and Raschky (2014). Average shares in donor budgets may be affected by developments within a recipient economy over the sample period, which weakens the case for exogeneity. We sketch this argument in appendix A.

Perhaps the closest precursors to our paper are Nunn and Qian (2014) and, especially, Werker et al. (2009). The latter paper uses a simpler form of supply-push instrument, based on interacting the world oil price with a dummy for Muslim countries. The argument is that aid flows from oil-rich countries (in particular, from Gulf states) to Muslim aid recipients will be positively correl-

ated with the world oil price. Werker et al. study the effects of aid on a range of outcomes, and the findings in their Table 2 relate to how aid is absorbed. Those findings tally closely with ours. As in our paper, they find a significant effect of aid on household consumption, where the estimated effect is much larger in IV estimates. There is no evidence that aid leads to higher government consumption; some evidence that aid promotes gross investment, but this is not robust; no evidence that aid affects exports; but strong evidence that aid leads to higher imports. Hence, although the two papers differ in significant respects, the findings on how aid is absorbed are remarkably similar.

The papers differ in the choice of instrument, and the approach to estimation. One benefit of our instrument is that we can study the effects of broadly-defined aid, whereas studies based on natural experiments are informative about relatively narrow categories of aid. For example, the Werker et al. findings are most informative about the effects of unconditional grants from Gulf oil exporters to Muslim aid recipients, while the Nunn and Qian approach is currently restricted to US food aid.¹⁰ As for the estimation method, the CCE estimators that we adopt should be consistent under more general conditions than the standard approaches in the literature. Although some papers include robustness tests which correspond to searching for observable proxies for $\phi_i f_i$, the CCE approach is more general, because it does not require the common factors or their heterogeneous loadings to be observable.

Our approach is also related to earlier work on aid using instrumental variables, especially Tavares (2003). He used the geographic distance between recipients and donors, and whether or not they share a common border, language or religion, to instrument for aid, since bilateral connections between donors and recipients are likely to influence aid flows. In our approach, the initial shares in donor budgets can be interpreted as proxying for connections between donors and recipients, while remaining more agnostic than Tavares about the potential sources of these connections. Put differently, we infer connections from the aid data itself, rather than relying on sets of connections that are already known to the econometrician.

¹⁰Nunn and Qian (2014) study the effect of US food aid on conflict in recipient countries. They instrument wheat aid received from the US with an interaction of one-period-lagged US wheat production and the fraction of years in the sample in which a country has received food aid from the US.

4 Data

Our models will be estimated using three-year averaged data for the period 1971-2012.¹¹ To construct the synthetic aid measure that we use as an instrument, we need the initial shares of aid recipients in donor budgets. These initial shares will be based on the period 1960-1970.

Our aid variable is taken from Table 2a of the standard OECD Development Assistance Committee (DAC) data tables. We follow Arndt, Jones and Tarp (2010) in our treatment of some missing values: they argue that some apparently missing values in fact correspond to zeroes. In each year, we turn missing recipient-donor-year aid to zero for combinations of recipients that receive aid from at least one donor in that year and donors that disburse aid to at least one recipient in that year. Aid in recipient-year format is found by keeping the entries that list ‘All donors, total’ as a donor. Our focus is on net aid disbursements.

Our synthetic measure for aid is constructed in the following way from the DAC’s recipient-donor-year data. For each donor, we calculate the average of the annual shares of a given recipient country in a donor’s aid for the years 1960-1970 (this yields a_{i0}^d), and multiply this by the donor’s current budget (D_{dt} , the sum of the donor’s aid disbursements over all recipient countries in period t).¹² We then sum these numbers across donors to get $A_{it}^S = \sum_{d=1}^{N_D} a_{i0}^d D_{dt}$. For each recipient country, this yields the aid that the recipient would have received at each date, had its shares in the various donor budgets remained constant, and hence equal to the 1960-1970 average shares. It is this time-varying, synthetic measure of aid that we use to instrument for aid in panel data regressions. Both the endogenous aid variable and the instrument in the regressions will be measured relative to GDP. Our GDP data, and the other macroeconomic variables used in the regressions, are extracted from online World Bank data files using

¹¹The cross-country literature often uses four-year or five-year averages, but those choices would leave us with a relatively short time dimension, given that the CCE estimators require country-specific coefficients for each cross-section mean. Moving in the opposite direction, to annual data, would also have disadvantages: we would need to estimate models with a more complicated dynamic structure, sitting uneasily with the use of both our instrument and CCE estimators.

¹²For a small percentage of observations the numerator in these annual shares (aid received by country i from donor d in each year) or the denominator (total aid disbursed by donor d in each year) are negative. Hence, before we calculate the annual shares, negative values for the numerator are changed to zero, and the denominator is recalculated by summing the non-negative numerators over all recipients.

the `wbopendata` software for Stata (Azevedo, 2011).

The dependent variables considered will include household consumption, government consumption, gross capital formation, imports and exports, again relative to GDP. Net imports are defined as imports minus exports. In the recipient-year data, before collapsing to three-year averages, observations for these variables are turned to missing whenever at least one of the other variables of the GDP identity is missing. This keeps the sample consistent across the different dependent variables we consider below. We exclude countries with small populations (fewer than 500,000 people in the first period of the sample). In our final data set, the available expenditure components sum to total GDP, or very close to GDP, for each country-period observation.¹³

A final data issue is that, in a small number of cases, the distinctions between recipients and donors are blurred. Countries such as Cyprus, Israel, Saudi Arabia, South Korea, Thailand and Turkey have taken both roles at one time or another. Since these are relatively small donors in global terms, we do not investigate this in detail, but a later robustness check will restrict attention to the ten largest donors.

5 Results

For each dependent variable, we report eight regressions. For reference purposes, we report FE and pooled CCE results that do not instrument for aid. We report estimates for static models, and dynamic models that include a lagged dependent variable.¹⁴ The standard errors that we report are heteroskedasticity-robust and clustered by country, and we make a small-sample adjustment to take into account the large number of estimated parameters. In experiments, we compared adjusted standard errors to those obtained from a non-parametric block bootstrap, given that the asymptotic distribution of pooled CCE-type estimators is non-standard (Pesaran, 2006) and the asymptotic variance of the CCE IV es-

¹³Only one country-period observation, for Mali in 2004-6, shows a discrepancy larger than 1% of GDP. Dropping this observation makes little difference to our results.

¹⁴The sample for each of the eight regressions consists of the observations that are included in the CCE IV estimation of the dynamic model. Since CCE IV estimation of the dynamic model requires estimating five country-specific parameters (one country fixed effect and four factor loadings for the cross-section means) only countries with at least six observations are retained in the sample. In our main sets of estimates, we have 13 or 14 time series observations for many countries, with a maximum of 14.

estimator introduced in Harding and Lamarche (2011) has not yet been studied. The bootstrapped standard errors are noticeably larger than the conventional standard errors in some cases, and smaller in others. Hence, in the case of CCE IV estimates, we also report 90% bootstrapped, bias-corrected confidence intervals based on the BC_a method (see, for example, Davison and Hinkley, 1997, pp. 203-211). Our main findings obtain under either approach to inference.

We noted earlier that the short-run response of macroeconomic ratios to aid could be modest, but the long-run effects larger. The inclusion of a lagged dependent variable is one way to capture this, with the usual partial adjustment interpretation. Whenever the estimated model is dynamic, we report the long-run effect of aid on the ratio of interest, with a standard error approximated by the delta method; but since long-run effects correspond to ratios of parameters, their estimates are likely to be imprecise. We note that CCE-type estimators are consistent in dynamic panel data models under more restrictive assumptions than in the static case (Chudik and Pesaran, 2013, and Everaert and De Groote, forthcoming) and we give more emphasis to the results from static models.¹⁵ The use of three-year averages implies that, even in these models, we allow the absorption process to extend over several years, as seems likely.

Before we turn to the results, note that the effects of instrumenting for aid are likely to vary across the dependent variables. The effect of instrumenting will depend on the extent to which aid allocation is non-random. Macroeconomic ratios are likely to differ in their sensitivity to particular types of aid-relevant shocks, and hence there is scope for the effects of instrumenting to change with the dependent variable.

We first study the effects of aid on trade-related variables, starting with net imports. Recall that net imports must increase if aid is to be absorbed. If the net import share rises one-for-one with the aid share, this should assuage concerns that aid is diverted abroad (capital flight) or primarily used to accumulate foreign exchange reserves. The relevant results are shown in Table 1. In our IV estimates, we cannot reject the hypothesis that aid is fully absorbed domestically. The coefficient on aid is large, significantly different from zero, and not significantly

¹⁵Hayakawa et al. (2014) introduce an estimator for dynamic panel data models with a multifactor error structure, but it requires the regressors other than the lagged dependent variable to be strictly exogenous. Moon and Weidner (2013) develop an estimator for dynamic panels that can accommodate interactive fixed effects and endogenous regressors, but it treats the number of factors as known, and would be more complicated to implement than the approach we adopt.

Table 1: Aid and net imports

	(1)	(2)	(3)	(4)
	FE	FE	CCE	CCE
Aid	0.545*** (0.113)	0.315*** (0.0745)	0.628*** (0.123)	0.479*** (0.117)
Lagged dep. variable		0.575*** (0.0399)		0.440*** (0.0842)
Long-run effect aid		0.741***		0.856***
Long-run effect SE		0.167		0.200
Countries	88	88	88	88
Observations	1099	1099	1099	1099
	(1)	(2)	(3)	(4)
	FE IV	FE IV	CCE IV	CCE IV
Aid	0.772*** (0.217)	0.555*** (0.169)	1.085*** (0.285)	0.991** (0.390)
Lagged dep. variable		0.548*** (0.0516)		0.358*** (0.102)
Long-run effect aid		1.227***		1.543***
Long-run effect SE		0.289		0.513
BC_a			[0.73,1.83]	[0.99, 2.92]
First stage F-statistic	19.10	19.27	18.14	13.45
Underidentification	0.001	0.001	0.009	0.014
Countries	88	88	88	88
Observations	1099	1099	1099	1099

Note: Dependent variable is net imports. All variables expressed relative to GDP. Fixed effects (FE), fixed effects IV (FE IV), common correlated effects (CCE) and common correlated effects IV (CCE IV) results, three-year averaged data, 1971-2012. IV regressions carried out using xtivreg2 for Stata (Schaffer, 2010). FE and FE IV regressions include time dummies, coefficients not reported. Country-specific coefficients on cross-section means in CCE and CCE IV regressions not reported. Heteroskedasticity-robust standard errors, clustered by country and corrected for degrees of freedom, in brackets. Standard errors (SE) for the long-run effects obtained via the delta method. *, **, and *** denote significance at 10, 5 and 1%, respectively. BC_a shows a 90% bias-corrected-and-accelerated confidence interval obtained from a non-parametric block bootstrap. Underidentification shows the p-value of the Kleibergen and Paap (2006) LM test for underidentification.

different from unity, both in static models and in the long run derived from dynamic models. The contrast with the upper row of estimates is instructive: in the absence of an instrument, the evidence that aid is fully absorbed is much weaker.

We can also investigate whether there are symptoms of aid-driven Dutch Disease. The underlying idea is that an increase in domestic expenditure will often fall partly on non-traded goods. At least in the short run, this will increase their price relative to traded goods, given that the prices of the latter are determined by world markets and hence exogenous. As the relative price of non-traded goods increases, the domestic traded sector contracts in relative terms. The effect arises because aid augments income without directly augmenting productive capacity. We do not investigate directly whether or not these changes are harmful for long-run growth and development, which remains a hard question to resolve. Instead, we examine whether aid leads to higher imports and/or lower exports, both measured relative to GDP. If we find only a modest effect on imports, or a large negative effect on exports, these would be symptoms of a possible Dutch Disease.

Whether the Dutch Disease can arise partly depends on how the aid is spent. In principle, aid could lead to a one-for-one increase in imports. With no increase in domestic expenditure on non-traded goods, there is no scope for Dutch Disease.¹⁶ Alternatively, if aid is partly spent on domestic non-traded goods, their relative price will increase, at least in the short run. The price change will tend to attract labour and other mobile factors into the non-traded sector, raising costs in the traded sector: this shifts the supply curve inwards in the traded sector, leading to a decline in exports. Hence, by studying the effects of aid on imports and exports, we can get a sense of whether Dutch Disease effects may be at work.¹⁷

Tables 2 and 3 show the effects of the aid share on import and export shares respectively. The results for the import share are similar to those for the net import share, with the exception of the static model estimated by FE IV. A strong

¹⁶We should qualify this slightly: if some imports are of capital goods, one effect of aid will be to augment productive capacity, but the general equilibrium effect will depend on which sectors see increased investment.

¹⁷In models in which the capital stock is endogenous, a pure transfer may have no long-run relative price effects, but exports will be lower in equilibrium, because imports are financed partly by the transfer. See Cerra et al. (2009, p. 149).

Table 2: Aid and imports

	(1)	(2)	(3)	(4)
	FE	FE	CCE	CCE
Aid	0.491*** (0.109)	0.343*** (0.0692)	0.609*** (0.100)	0.407*** (0.0879)
Lagged dep. variable		0.682*** (0.0481)		0.607*** (0.0642)
Long-run effect aid		1.080***		1.037***
Long-run effect SE		0.232		0.243
Countries	88	88	88	88
Observations	1099	1099	1099	1099
	(1)	(2)	(3)	(4)
	FE IV	FE IV	CCE IV	CCE IV
Aid	0.123 (0.336)	0.451*** (0.145)	0.622*** (0.210)	0.655*** (0.187)
Lagged dep. variable		0.677*** (0.0491)		0.413*** (0.0877)
Long-run effect aid		1.395***		1.115***
Long-run effect SE		0.441		0.385
BC_a			[0.33, 1.18]	[0.52, 2.44]
First stage F-statistic	19.10	21.00	29.27	19.52
Underidentification	0.001	0.001	0.019	0.006
Countries	88	88	88	88
Observations	1099	1099	1099	1099

Note: See note Table 1. Dependent variable is imports.

positive effect of aid is found in the two CCE IV regressions in particular. In contrast, we do not find a clear-cut effect of aid on the export share. In the FE IV estimates of a static model, aid has a negative effect on the export share which is significant at the 10% level, but this finding is not robust to alternative models and estimators. In the dynamic FE IV estimates, and the two sets of CCE IV estimates, we cannot reject the hypothesis that aid has no effect on the export share. This does not rule out Dutch Disease – for that, we would need to find zeroes estimated with greater precision – but neither do we find robust evidence pointing to Dutch Disease effects.

We now study the relationships between aid and other macroeconomic ratios, starting with the effect of aid on total consumption. We define this as the sum of household and government consumption ($C + G$). While household and government consumption are distinct concepts, there are sectors such as educa-

Table 3: Aid and exports

	(1)	(2)	(3)	(4)
	FE	FE	CCE	CCE
Aid	-0.0541 (0.106)	0.0852* (0.0465)	0.0396 (0.0946)	0.0416 (0.0805)
Lagged dep. variable		0.760*** (0.0409)		0.520*** (0.0537)
Long-run effect aid		0.356*		0.0867
Long-run effect SE		0.207		0.169
Countries	88	88	88	88
Observations	1099	1099	1099	1099
	(1)	(2)	(3)	(4)
	FE IV	FE IV	CCE IV	CCE IV
Aid	-0.649* (0.357)	0.0176 (0.130)	-0.0464 (0.267)	0.196 (0.216)
Lagged dep. variable		0.757*** (0.0418)		0.433*** (0.0837)
Long-run effect aid		0.0727		0.346
Long-run effect SE		0.538		0.402
BC_a			[-0.40,0.53]	[-0.37,1.16]
First stage F-statistic	19.10	18.21	30.69	30.97
Underidentification	0.001	0.001	0.018	0.012
Countries	88	88	88	88
Observations	1099	1099	1099	1099

Note: See note Table 1. Dependent variable is exports.

tion and health where the distinction is somewhat artificial for welfare purposes, given a mix of public and private provision. The results are shown in Table 4 and suggest that aid has a large positive effect on total consumption. The difference made by instrumental variables can be seen clearly, by comparing the upper row of estimates with the lower row. Compared to the FE and CCE estimates, the point estimates from IV estimators suggest much larger effects of aid on total consumption. We should avoid over-interpreting this, because the differences are not statistically significant. Nevertheless, it is easy to see how this pattern could arise in the within variation. Aid may respond positively to country-specific circumstances which lower total consumption, such as adverse shocks. By instrumenting aid we alleviate this source of bias, and find much larger effects of aid on total consumption.

When we study the effects of aid on household consumption, the estimates

Table 4: Aid and total consumption

	(1)	(2)	(3)	(4)
	FE	FE	CCE	CCE
Aid	0.391*** (0.131)	0.137* (0.0764)	0.489*** (0.130)	0.283*** (0.102)
Lagged dep. variable		0.624*** (0.0341)		0.516*** (0.0517)
Long-run effect aid		0.364*		0.585***
Long-run effect SE		0.185		0.193
Countries	88	88	88	88
Observations	1099	1099	1099	1099
	(1)	(2)	(3)	(4)
	FE IV	FE IV	CCE IV	CCE IV
Aid	0.695*** (0.245)	0.272** (0.133)	0.662** (0.284)	0.429* (0.248)
Lagged dep. variable		0.606*** (0.0393)		0.447*** (0.0607)
Long-run effect aid		0.690**		0.775*
Long-run effect SE		0.300		0.436
BC_a			[0.29,1.22]	[0.26,1.62]
First stage F-statistic	19.10	18.10	19.84	12.12
Underidentification	0.001	0.001	0.010	0.007
Countries	88	88	88	88
Observations	1099	1099	1099	1099

Note: See note Table 1. Dependent variable is total consumption, the sum of household and government consumption.

are generally similar to those we find for total consumption, but less precise. These results are shown in Table 5. The use of an instrument again increases the estimated effect of aid. We find much less evidence that aid influences government consumption, as Table 6 shows. These results are similar to those found by Werker et al. (2009, Table 2) using an alternative instrument.

Aid has often been characterized as primarily government-to-government transfers. Our finding that a substantial fraction of aid is reflected in higher household consumption, but not in higher government consumption, may therefore appear surprising. One mechanism could be lower taxes: standard reasoning suggests that increased aid to governments will not only be used to increase government purchases, but also to reduce taxes (Kimbrough, 1986). It is also possible that recipient governments use aid to finance transfers for political ends. This assumption has been common in the literature, as in Adam and O'Connell

Table 5: Aid and household consumption

	(1)	(2)	(3)	(4)
	FE	FE	CCE	CCE
Aid	0.337** (0.142)	0.125* (0.0694)	0.381* (0.211)	0.226* (0.123)
Lagged dep. variable		0.664*** (0.0282)		0.535*** (0.0477)
Long-run effect aid		0.373*		0.486*
Long-run effect SE		0.197		0.264
Countries	88	88	88	88
Observations	1099	1099	1099	1099
	(1)	(2)	(3)	(4)
	FE IV	FE IV	CCE IV	CCE IV
Aid	0.844** (0.370)	0.309* (0.173)	0.708 (0.435)	0.462 (0.358)
Lagged dep. variable		0.645*** (0.0330)		0.460*** (0.0615)
Long-run effect aid		0.869*		0.856
Long-run effect SE		0.450		0.655
BC_a			[0.17,1.45]	[0.12,2.02]
First stage F-statistic	19.10	17.86	18.18	11.84
Underidentification	0.001	0.002	0.006	0.006
Countries	88	88	88	88
Observations	1099	1099	1099	1099

Note: See note Table 1. Dependent variable is household consumption.

(1999), Boone (1996) and Hodler and Raschky (2014), among others. Finally, a significant share of aid is given in ways which bypass domestic governments, such as off-budget aid projects, or support for NGOs.¹⁸ In all of these cases, household consumption is where the effect of aid is most likely to be manifested in the national accounts.

Finally, we turn to the relationship between aid and gross investment as a share of GDP. These results are shown in Table 7. The long-run effect of aid on the investment rate is noticeably larger than the short-run effect and, when the instrument is used, the estimated long-run effect is significant at the 5% (FE IV) or 10% (CCE IV) level. To anticipate some of our later discussion, however, the investment results will turn out to be less robust than the consumption results.¹⁹

¹⁸Van de Sijpe (2013) discusses off-budget aid in more detail.

¹⁹This instability also obtained in the earlier working paper (Temple and Van de Sijpe, 2014). There, we found some evidence of lagged effects of aid, but in our baseline results, we could

Table 6: Aid and government consumption

	(1)	(2)	(3)	(4)
	FE	FE	CCE	CCE
Aid	0.0541 (0.0884)	-0.00931 (0.0406)	0.105 (0.0686)	0.0836 (0.0695)
Lagged dep. variable		0.716*** (0.0325)		0.426*** (0.0744)
Long-run effect aid		-0.0328		0.146
Long-run effect SE		0.144		0.120
Countries	88	88	88	88
Observations	1099	1099	1099	1099
	(1)	(2)	(3)	(4)
	FE IV	FE IV	CCE IV	CCE IV
Aid	-0.149 (0.209)	-0.0538 (0.0703)	-0.0583 (0.150)	-0.0339 (0.102)
Lagged dep. variable		0.721*** (0.0328)		0.400*** (0.0843)
Long-run effect aid		-0.193		-0.0566
Long-run effect SE		0.260		0.172
BC_a			[-0.46,0.15]	[-0.56,0.19]
First stage F-statistic	19.10	22.01	14.99	13.89
Underidentification	0.001	0.001	0.005	0.004
Countries	88	88	88	88
Observations	1099	1099	1099	1099

Note: See note Table 1. Dependent variable is government consumption.

We have not yet discussed the strength of our instrument. The tables report the first stage F-statistic as a heuristic guide.²⁰ This approach has been widely used, but the conventional Stock and Yogo (2005) benchmarks for first-stage F-statistics do not apply directly to panel data models. This is partly because the benchmarks assume i.i.d. errors. By considering a model for a single time series, Bun and de Haan (2010) show that the standard benchmarks for the first-stage F-statistic do not apply when the errors are serially correlated. In their Monte Carlo simulations, a robust F-statistic tends to underestimate instrument strength. It is not clear whether this is general or would extend to panels, and

not reject the hypothesis that aid had no effect on investment. Those results were based on a smaller sample, 992 observations compared to 1099 observations here.

²⁰While the sample and first stage model are the same across the different tables reported below, the first stage estimates for CCE IV are not. The CCE IV estimator includes the cross-sectional mean of the dependent variable with country-specific coefficients in both the first and second stage, and hence the first-stage F-statistic varies across tables for the CCE IV regressions.

Table 7: Aid and gross capital formation

	(1)	(2)	(3)	(4)
	FE	FE	CCE	CCE
Aid	0.154*** (0.0580)	0.158*** (0.0369)	0.123 (0.0749)	0.172*** (0.0599)
Lagged dep. variable		0.579*** (0.0459)		0.462*** (0.0584)
Long-run effect aid		0.375***		0.320**
Long-run effect SE		0.110		0.125
Countries	88	88	88	88
Observations	1099	1099	1099	1099
	(1)	(2)	(3)	(4)
	FE IV	FE IV	CCE IV	CCE IV
Aid	0.0763 (0.194)	0.251** (0.122)	0.279 (0.252)	0.442* (0.237)
Lagged dep. variable		0.579*** (0.0474)		0.368*** (0.0795)
Long-run effect aid		0.596**		0.700*
Long-run effect SE		0.288		0.394
BC_a			[-0.02,0.88]	[0.29, 1.74]
First stage F-statistic	19.10	18.70	14.67	11.21
Underidentification	0.001	0.001	0.009	0.011
Countries	88	88	88	88
Observations	1099	1099	1099	1099

Note: See note Table 1. Dependent variable is gross capital formation.

hence our use of F-statistics is best seen as heuristic; this is also true of the Kleibergen and Paap (2006) LM test for underidentification. In most cases below, and especially in the static CCE IV regressions, the first-stage robust F-statistic is reasonably high, and we do not have grounds for concern about the strength of the instrument.

6 Robustness

In this section, we consider several alternative models and estimators. Most of these can be interpreted as measures that will increase robustness – in the sense of making biases less likely – at the expense of reduced efficiency. Our main conclusions continue to find support, even when we make adjustments to the instrument that weaken its explanatory power in the first stage. The estimates

are summarized in Table 8, where row 1 shows the main results from Tables 1-7 for ease of comparison.

In our main results, we followed much of the cross-country literature and excluded countries with small populations. If, instead, we include these countries, we generally find that the instrument becomes weaker in the first stage of 2SLS: see row 2 of the table.²¹ A potential explanation is that, for aid recipients which account for small and volatile shares of donor budgets, the share of a budget at an initial date may be relatively uninformative about that recipient's long-term exposure to changes in that budget. Hence, we would expect our supply-push instrument to have more explanatory power for aid for larger countries than for small ones. But despite the weakening of the instrument, the second stage results are qualitatively similar to those we presented earlier.

We next examine the implications of transitions from colonial rule to independence. In some cases, the DAC dataset includes reports of aid flows before an aid recipient became independent. This implies that, for some countries, we have constructed an instrument based on initial shares in donor budgets in the period 1960-70 even though the country only became independent later. To the extent that recorded aid flows before independence are incomplete or measured less accurately, this may affect our results. Hence, as an alternative, we calculate the aid variable (and the initial shares in donor budgets needed to construct an instrument) only after discarding aid data in the years before a recipient's independence.²² The results, shown in row 3 of Table 8, are again similar to those found before.

²¹The countries that we return to the sample are the Bahamas, Bahrain, Barbados, Belize, Bhutan, Comoros, Djibouti, Gambia, Kiribati, Macao, Malta, Suriname, Swaziland, Tonga, and Vanuatu.

²²The year of independence is taken to be the first year that a country is listed in the Polity IV dataset (Marshall, Gurr and Jagers, 2013). For countries not included in Polity IV, we use the CIA World Factbook (<https://www.cia.gov/library/publications/the-world-factbook/>).

Table 8: Robustness checks

Row	Model	C	G	C+G	I	X	M	M-X
1	Static	0.708 (0.435)	-0.0583 (0.150)	0.662** (0.284)	0.279 (0.252)	-0.0464 (0.267)	0.622*** (0.210)	1.085*** (0.285)
	F	18.18	14.99	19.84	14.67	30.69	29.27	18.14
	Dyn.	0.856 (0.655)	-0.0566 (0.172)	0.775* (0.436)	0.700* (0.394)	0.346 (0.402)	1.115*** (0.385)	1.543*** (0.513)
	F	11.84	13.89	12.12	11.21	30.97	19.52	13.45
2	Static	0.532 (0.326)	-0.0909 (0.152)	0.531** (0.237)	0.412* (0.232)	0.283 (0.369)	1.081*** (0.295)	1.076*** (0.250)
	F	11.27	11.85	13.50	10.72	13.14	15.65	12.58
	Dyn.	0.889 (0.606)	-0.138 (0.204)	0.925* (0.495)	0.730* (0.377)	0.611 (0.523)	1.856*** (0.575)	1.663*** (0.496)
	F	7.116	9.122	7.552	7.704	11.01	9.891	8.094
3	Static	0.795* (0.470)	-0.149 (0.191)	0.744** (0.317)	0.187 (0.267)	-0.155 (0.299)	0.554** (0.253)	1.099*** (0.291)
	F	18.00	16.03	20.02	15.49	30.89	27.57	19.34
	Dyn.	1.052 (0.690)	-0.194 (0.232)	0.885* (0.469)	0.612* (0.351)	0.161 (0.465)	1.043*** (0.387)	1.588*** (0.518)
	F	12.33	14.37	12.39	12.05	27.18	19.40	14.66
4	Static	0.303 (0.235)	-0.0756 (0.168)	0.346* (0.199)	0.561** (0.224)	-0.0650 (0.243)	0.722** (0.351)	0.988*** (0.302)
	F	26.34	22.10	38.18	28.95	51.17	76.25	24.78
	Dyn.	0.188 (0.344)	-0.0130 (0.230)	0.0931 (0.395)	0.807*** (0.302)	0.132 (0.362)	0.798 (0.496)	1.058*** (0.371)
	F	19.54	19.17	27.73	20.02	39.72	64.83	20.51
5	Static	0.719 (0.576)	-0.0654 (0.206)	0.784** (0.392)	0.558** (0.247)	-0.134 (0.432)	0.990*** (0.370)	1.334*** (0.499)
	F	7.602	5.275	6.749	5.463	6.495	5.317	5.559
	Dyn.	1.303 (0.820)	0.0131 (0.328)	1.417** (0.651)	0.845* (0.458)	0.392 (0.653)	1.566** (0.720)	1.591** (0.610)
	F	8.061	5.944	7.580	3.845	3.296	2.516	5.224

Continued on next page

Table 8 – continued from previous page

Row	Model	C	G	C+G	I	X	M	M-X
6	Static	0.541 (0.425)	-0.0326 (0.213)	0.576** (0.273)	0.561** (0.251)	0.00975 (0.329)	0.930*** (0.279)	1.205*** (0.337)
	F	19.05	17.50	23.05	18.73	51.88	50.16	19.61
	Dyn.	0.748 (0.599)	0.0249 (0.255)	0.665 (0.400)	0.837** (0.389)	0.362 (0.421)	1.087*** (0.385)	1.412*** (0.461)
	F	11.56	14.93	15.97	13.73	39.54	42.81	17.07
7	Static	0.476 (0.320)	-0.0607 (0.201)	0.557** (0.239)	0.424** (0.186)	-0.0305 (0.338)	0.828*** (0.212)	0.984*** (0.217)
	F	19.85	13.41	21.40	16.78	27.42	33.04	20.93
	Dyn.	0.617 (0.445)	-0.108 (0.233)	0.635** (0.301)	0.488** (0.228)	0.302 (0.415)	0.856*** (0.287)	1.065*** (0.245)
	F	10.21	10.87	12.58	13.40	27.31	25.62	15.26
8	Static	0.619 (0.443)	0.0197 (0.133)	0.665** (0.334)	0.208 (0.268)	0.0245 (0.259)	0.577** (0.234)	1.043*** (0.255)
	F	18.72	14.41	21.00	14.20	29.01	25.49	18.28
	Dyn.	0.926 (0.682)	0.0339 (0.150)	0.949* (0.544)	0.550 (0.382)	0.532 (0.385)	0.964** (0.395)	1.470*** (0.496)
	F	12.03	12.50	12.87	11.99	29.62	15.87	14.13
9	Static	0.613 (0.464)	0.0626 (0.111)	0.618** (0.308)	0.166 (0.270)	0.0614 (0.246)	0.621** (0.244)	1.013*** (0.233)
	F	18.28	14.34	23.80	15.08	35.05	25.79	20.80
	Dyn.	1.027 (0.718)	0.0584 (0.122)	0.949* (0.509)	0.0941 (0.391)	0.546 (0.420)	1.120*** (0.396)	1.331*** (0.461)
	F	11.23	11.86	13.24	16.15	27.40	13.28	17.81
10	Static	0.602 (0.528)	0.120 (0.0854)	0.557 (0.372)	0.217 (0.276)	0.158 (0.236)	0.625** (0.302)	1.084*** (0.323)
	F	15.34	10.90	15.81	10.81	20.15	14.09	13.07
	Dyn.	0.943 (0.846)	0.135 (0.0905)	0.765 (0.597)	0.215 (0.398)	0.656 (0.415)	1.179** (0.486)	1.323** (0.548)
	F	8.524	8.619	9.171	10.15	14.98	9.339	12.28
11	Static	0.537 (0.371)	0.0822 (0.171)	0.632** (0.301)	0.141 (0.255)	-0.0362 (0.277)	0.651*** (0.182)	0.919*** (0.228)
	F	20.19	15.82	19.94	15.06	19.02	20.65	18.05
	Dyn.	0.925 (0.643)	0.106 (0.183)	1.009* (0.513)	0.515 (0.373)	0.375 (0.455)	1.013*** (0.355)	1.426*** (0.430)
	F	14.77	14.89	13.58	13.71	22.71	17.40	13.56

Continued on next page

Table 8 – continued from previous page

Row	Model	C	G	C+G	I	X	M	M-X
12	Static	0.533 (0.380)	0.121 (0.157)	0.626** (0.294)	0.101 (0.262)	-0.0299 (0.262)	0.647*** (0.206)	0.959*** (0.210)
	F	19.29	15.10	20.59	15.25	22.03	20.93	18.45
	Dyn.	1.070 (0.665)	0.102 (0.171)	1.068** (0.496)	0.0970 (0.356)	0.406 (0.492)	1.348*** (0.333)	1.280*** (0.400)
	F	12.50	13.38	12.40	14.73	21.29	13.59	14.78
13	Static	0.534 (0.473)	0.155 (0.132)	0.594 (0.399)	0.154 (0.264)	0.110 (0.283)	0.670*** (0.250)	1.083*** (0.319)
	F	14.30	9.650	12.77	11.57	14.79	13.57	12.58
	Dyn.	1.154 (0.902)	0.103 (0.110)	0.920 (0.657)	0.257 (0.350)	0.653 (0.545)	1.454*** (0.446)	1.330** (0.540)
	F	7.046	8.070	7.177	9.970	12.49	9.778	10.64
14	Static	1.015 (0.729)	0.0775 (0.106)	0.904 (0.624)	-0.182 (0.340)	-0.0195 (0.330)	0.398 (0.370)	1.235** (0.499)
	F	9.582	6.157	7.454	5.216	9.684	6.543	7.461
	Dyn.	1.073 (0.919)	0.0957 (0.0917)	1.000 (0.835)	-0.264 (0.603)	0.237 (0.686)	0.724 (0.582)	1.335* (0.756)
	F	6.350	5.934	5.566	3.444	5.367	3.692	6.222
15	Static	1.520** (0.743)	-0.109 (0.209)	1.525** (0.620)	0.117 (0.441)	-0.584 (0.745)	0.436 (0.572)	1.884*** (0.602)
	F	8.391	9.127	9.714	7.941	9.928	7.828	9.030
	Dyn.	1.445 (0.876)	-0.235 (0.253)	1.472** (0.711)	0.652 (0.492)	-0.129 (1.137)	1.031 (0.877)	2.543** (0.971)
	F	8.274	8.398	8.622	7.070	8.076	7.083	7.539
16	Static	1.583** (0.795)	-0.0706 (0.315)	1.282** (0.626)	-0.187 (0.478)	-0.128 (0.472)	0.308 (0.390)	0.981** (0.457)
	F	9.665	12.58	9.033	8.418	11.12	13.20	8.509
	Dyn.	1.766* (1.027)	-0.193 (0.436)	1.370* (0.755)	-0.0103 (0.649)	1.117 (0.905)	1.149 (0.696)	1.347** (0.660)
	F	8.599	10.75	6.999	6.179	9.434	10.44	5.168

Note: The entries in this table show the long-run effect of aid on household consumption (C), government consumption (G), total consumption ($C + G$), gross capital formation (I), exports (X), imports (M) and net imports ($M - X$) in models with (“Dyn.”) and without (“Static”) a lagged dependent variable. All variables expressed relative to GDP. CCE IV estimation on three-year averaged data (1971-2012) using an instrument based on initial shares in donor budgets calculated over the period 1960-70, unless reported otherwise below. 1099 observations from 88 countries, unless reported otherwise below. Heteroskedasticity-robust standard errors, clustered

by country and corrected for degrees of freedom, in brackets. Standard errors (SE) for the long-run effects obtained via the delta method. *, **, and *** denote significance at 10, 5 and 1%, respectively. F shows the first-stage F-statistic.

Row 1 repeats the main results from Tables 1-7 for ease of comparison.

Row 2 includes small countries (1248 observations from 103 countries).

Row 3 constructs the endogenous aid variable in the second stage and the initial shares in donor budgets on which the instrument is based after discarding all aid data in the years before a recipient country's independence (1033 observations from 81 countries).

Row 4 replaces aid and the instrument by their first lag; sample starting with the period 1974-76 (1039 observations from 88 countries).

Row 5 includes both aid and its first lag, instrumented by the current and one period lagged values of the synthetic instrument; sample starting with the period 1974-76 (978 observations from 80 countries). F is the Kleibergen-Paap Wald rk F statistic.

Row 6 replaces aid and the instrument by the unweighted average of its current and one period lagged values; sample starting with the period 1974-76 (1032 observations from 88 countries).

Row 7 uses the final year values in each period for the dependent variables instead of the three-year averages (1065 observations from 86 countries).

Row 8 excludes the first period (1971-73) from estimation (1032 observations from 88 countries).

Row 9 excludes the first two periods (1971-73 and 1974-76) from estimation (965 observations from 88 countries).

Row 10 excludes the first three periods (1971-73, 1974-76 and 1977-79) from estimation (892 observations from 88 countries).

Row 11 uses an instrument based on initial shares calculated over the period 1960-73 and a sample that starts with the period 1974-76 (1061 observations from 91 countries).

Row 12 uses an instrument based on initial shares calculated over the period 1960-73 and a sample that starts with the period 1977-79 (993 observations from 91 countries).

Row 13 uses an instrument based on initial shares calculated over the period 1960-73 and a sample that starts with the period 1980-82 (919 observations from 91 countries).

Row 14 uses an instrument based on initial shares calculated over the period 1960-71 and a sample that starts with the period 1981-83 (846 observations from 91 countries).

Row 15 uses an instrument based on the largest ten donors only.

Row 16 removes outliers (1004 observations from 81 countries).

An important objection to cross-country aid research is that it rarely considers delayed effects of aid in sufficient detail, as Clemens et al. (2012) argue. Our next set of robustness tests addresses this point in various ways. We estimate models with lagged aid rather than current aid; current and lagged aid; using six-year averages of aid on the right-hand-side; and replacing each three-year average of the dependent variable with its value at the end of each three-year period. These alternative ways of capturing delayed effects of aid tend to point

to stronger effects of aid on investment, but we interpret this cautiously: as we discuss later, our findings on investment lack robustness.

The details of these results are shown in rows 4-7 of Table 8. In row 4, we replace current aid by its one-period lag, instrumented by the first lag of the synthetic instrument. Row 5 includes both current and one-period-lagged aid, instrumented with current and one-period-lagged values of the instrument.²³ Row 6 replaces aid and the instrument by the average of their current and one period lagged values, corresponding to a six-year average.²⁴ Row 7 measures the dependent variable in the final year of each three-year period instead of taking the average over all three years. Overall, we find some evidence that aid has a delayed effect on investment, while the effects on consumption are sometimes more muted than before.

We now turn to various potential criticisms of our instrument. In terms of the robustness of our main results, there are two especially important concerns. One is that serial correlation in country-specific circumstances might undermine exogeneity, as Card (2001) notes. The second is that the strength of the instrument could decline over time. Our IV strategy relies on the idea that shares in donor budgets in 1960-70 are informative about exposure to later changes in total donor budgets. If, for example, strategic or economic connections between countries change over time, then the instrument may have less explanatory power for aid in later periods of the sample. We can address both of these concerns as follows. Rows 8 to 10 in Table 8 repeat the main analysis but exclude, respectively, the first, the first two, and the first three periods from the panel data sample. As additional checks, row 11 uses an instrument based on initial shares calculated over the period 1960-73 and an estimation sample that starts with the period 1974-76, while rows 12 to 14 exclude the first, the first two and the first three periods from this sample. All of these checks are likely to increase robustness at the expense of efficiency. In particular, we might expect

²³This places heavy demands on the data: CCE IV estimation of the dynamic model now requires estimating seven country-specific parameters for each country. This implies that only countries with at least 8 observations can be included in estimation, reducing the sample by more than 100 observations. Moreover, high correlations between aid and its lag, and between the instrument and its lag, will contribute to large standard errors. But reassuringly for the validity of our instrument, in the first stage only the instrument in period t is significant for aid in period t , and only the instrument in period $t - 1$ matters for aid in $t - 1$.

²⁴In rows 4-6, the first period of the sample is dropped and the sample starts in the period 1974-76. This avoids overlap between the period over which the initial shares in donor budgets are calculated (1960-70) and the periods over which aid receipts are measured.

instrument strength to weaken as more periods are excluded, and this is what we find. The estimated second-stage coefficients are fairly stable, however, when dropping early time periods. As before, the results suggest that aid is primarily consumed rather than invested. We continue to find that aid is absorbed domestically, and that this happens via an increase in imports rather than a reduction in exports.

Another possible criticism is that, in constructing our instrument, we have considered too many donors. By considering all DAC-affiliated donors, we have included some donors whose budgets could be dominated by a small number of recipient countries. In that case, the domestic circumstances of these aid recipients may influence the evolution of the donor's aid budget over time, which risks endogeneity. This is likely to be a particular concern for smaller donors, who may concentrate most of their aid in a few recipient countries. To investigate this issue, row 15 in Table 8 shows results using an instrument based on the top ten donors. These are defined as those with the largest average annual share in global aid over the period 1960-2012.²⁵ As expected, this instrument is slightly weaker than the one based on all donors, and the second-stage coefficients are estimated less precisely. Nevertheless, we continue to find large effects of aid on consumption, imports and net imports.

Our next robustness tests acknowledge the potential importance of outliers. Given that we emphasize 2SLS results, outliers could arise in either the first stage or the second stage. Some of our robustness checks give rise to large first-stage F statistics, which may be a warning sign of outliers. To address this, we use the robust instrumental variables estimator of Cohen Freue et al. (2013), after partialling out fixed effects and cross-section means. The estimator can be used to identify multivariate outliers using the Mahalanobis distances of individual country-period observations (on the observable variables considered jointly) from the robust IV estimates, with respect to robust covariance estimates.²⁶ Across our various dependent variables, seven countries regularly emerge as giving rise to one or more outlying country-period observations: Burundi, the Central African Republic, Chad, the Democratic Republic of

²⁵These donors are the United States, Japan, France, Germany, the International Development Association (IDA), the EU institutions, the United Kingdom, the United Arab Emirates, the Netherlands, and Canada. Over the period 1960-2012, these ten donors accounted for more than three-quarters of total aid, on average.

²⁶See Cohen Freue et al. (2013) for more details of this approach.

Congo, Jordan, Madagascar, and Mauritania. The results when we exclude these countries are shown in the last row of Table 8. The results are in line with our earlier findings: the instrument retains explanatory power in the first stage, and we cannot reject strong effects of aid on net imports, household consumption and total consumption. Nor can we reject the null hypothesis of full absorption. There is no evidence that aid raises the investment rate, and this remains the case in models (not reported) that allow for delayed effects.

We now consider a more radical departure from our earlier approach. Thus far, we have mainly emphasized estimators that use a within transformation (the CCE IV estimator can be interpreted as a generalization of fixed-effects instrumental variable approaches to panel data). For the static models, an alternative way to address country-specific effects would be to first difference the model. Under our maintained assumptions, an estimator based on first differences (FD) should have the same probability limit as a within groups (FE) estimator, but may be more efficient if the error term is highly persistent.²⁷ Perhaps more importantly, a comparison of FD and FE estimates is potentially informative about the validity of our assumptions, and helps to address some potential concerns. In particular, our dependent and independent variables both contain nominal GDP in the denominator, as does the instrument. This would be problematic if some function of nominal GDP cannot legitimately be excluded from our models for the dependent variables: in that case, the instrument would be correlated with the error term. But if this is a problem, FD estimates should look rather different to FE, since first differencing is likely to weaken the correlation between the instrument and the error term.

If we first difference the model implied by equations (1)-(2), we obtain:

$$\Delta Q_{it}/Y_{it} = \beta \Delta (A_{it}/Y_{it}) + \phi_i \Delta f_t + \Delta u_{it} \quad (3)$$

and if we define a new set of factors $f'_t \equiv \Delta f_t$, we can estimate the model in first differences using the CCE IV estimator as before. Results based on first differences are shown in rows 1 and 2 of Table 9. The first three-year period in these regressions is 1974-76, to avoid overlap with the years used to construct the instrument. In the case of row 1, the effect of aid on total consumption is

²⁷See, for example, Wooldridge (2010, pp. 321-326). The result is likely to apply even for CCE estimators.

still present, but the effect on net imports is less precisely estimated than before, and the effect on imports is smaller. In row 2, we have dropped the same seven outliers identified previously. This greatly strengthens the results: now aid has significant effects on total consumption, household consumption, imports and net imports. The effect of aid on net imports is somewhat lower in magnitude than in our baseline results. Hence, these estimates suggest that absorption may be less than complete, although given the high standard error, it remains true that we cannot reject the null of full absorption.

Table 9: First-differenced models

Row	Sample	C	G	C+G	I	X	M	M-X
1	Full	0.618 (0.418)	0.0597 (0.107)	0.653* (0.347)	-0.217 (0.272)	-0.0685 (0.329)	0.385* (0.195)	0.689 (0.431)
	F	21.12	22.01	22.83	17.31	18.45	16.44	23.95
2	No outliers	1.273** (0.498)	0.157 (0.164)	1.308*** (0.460)	-0.856* (0.497)	0.250 (0.407)	0.750* (0.407)	0.607** (0.252)
	F	9.049	10.42	9.760	8.054	9.381	7.888	9.080

Note: The entries in this table show the effect of aid on household consumption (C), government consumption (G), total consumption ($C + G$), gross capital formation (I), exports (X), imports (M) and net imports ($M - X$) in static first-differenced models. All variables expressed relative to GDP. CCE IV estimation on three-year averaged data (1974-2012) using an instrument based on initial shares in donor budgets calculated over the period 1960-70. Heteroskedasticity-robust standard errors, clustered by country and corrected for degrees of freedom, in brackets. *, **, and *** denote significance at 10, 5 and 1%, respectively. F shows the first-stage F-statistic. Row 1 is the first-differenced equivalent of row 1 in Table 8 (1032 observations from 88 countries).

Row 2 drops outliers. It is the first-differenced equivalent of row 16 in Table 8 (943 observations from 81 countries).

There is a further reason that the first-differenced results are reassuring. One maintained assumption is that we can estimate structural relationships between macroeconomic ratios and aid intensity, such as those suggested by the neoclassical growth model. An observer sceptical about such relationships might be worried about a ‘false positive’, because all the variables in our regressions take the form of ratios to nominal GDP. In principle, movements in nominal GDP could lead to spurious correlations between the dependent variables and aid intensity: the models might appear to have explanatory power even in the absence

of a structural relationship.²⁸ This argument is less applicable once the variables have been first-differenced, and hence the similarity between our baseline results and row 2 of Table 9 is reassuring.

There are other ways to respond to the potential concern about spurious correlations, such as including the reciprocal of nominal GDP as an additional explanatory variable. The inclusion of the denominator of ratios is the approach suggested by Kronmal (1993) in a cross-section setting. In our case, this will tend to increase the standard errors, particularly in CCE estimates given the additional cross-section mean and associated country-specific coefficients, and the inclusion of the reciprocal of nominal GDP in the first stage of 2SLS. But the results, shown in Table B.1 in appendix B, are generally quite similar to our baseline estimates. We continue to find significant effects on net imports and consumption, and little evidence for effects on government consumption, investment and exports. In some cases, the point estimates for the consumption effects become larger, and the associated confidence intervals are wider than before; in a few cases, the instrument noticeably weakens. But the reciprocal of nominal GDP is often insignificant, especially in CCE estimates. These results suggest that our main findings are unlikely to be driven by spurious correlations.

A final way to see if output plays a confounding role is to examine large changes in *real* GDP, perhaps arising from economic crisis or civil war. If such events generate large swings in aid and expenditure components relative to GDP, they could make a major contribution to the within variation. In that case, the effects of aid might be identified mainly from extreme events, but responses to aid may also be different at those times. Since our interest is primarily in the effects of aid in ‘normal times’, we investigate what happens when we gradually eliminate countries which sometimes exhibit rapid declines in real GDP (‘output collapses’). For each country-period observation, we calculate the percentage change in real GDP from the previous to the current three-year period. For a static model estimated by CCE IV, Figure B.1 in appendix B shows the evolution of the estimated effect of aid as we progressively drop the countries with the largest output collapses. Figure B.2 does the same for the long-run effect of aid estimated from a model with a lagged dependent variable. If the impact of aid is

²⁸Discussions of potentially spurious relationships between ratios with common denominators date back to Pearson (1897) and Yule (1910), while Kronmal (1993) is a systematic treatment. In the aid literature, the problem is noted by Arndt et al. (2010) in particular.

different in times of crisis, we would expect the estimated coefficients to move substantially. We find little evidence of this. The effects on net imports are fairly stable in the static and dynamic models, and the confidence intervals always exclude zero and include unity. Arguably, compared to our earlier findings, the main differences are the wider confidence intervals for the effects of aid on household consumption, total consumption and imports. The point estimates are reasonably stable, however.

7 Conclusions

Using cross-country data to study aid effectiveness is fraught with difficulties, and yet some research questions are hard to answer any other way. This paper has aimed to make progress on two fronts. First, we have introduced a new ‘supply push’ instrument for aid, and combined it with a panel time series estimator that filters out unobserved common factors even when their effects differ across countries. This approach represents an advance on much of the existing applied literature. Second, we use the instrument to investigate the effects of aid on various macroeconomic ratios. This is informative about the extent of domestic absorption, the effects of aid on consumption and investment, and whether aid might be associated with Dutch Disease.

The balance of evidence can be summarized as follows. For the trade variables, we find robust effects of aid on the ratio of net imports to GDP, across a wide range of estimators and robustness tests. This suggests that aid is absorbed at least partially, assuaging concerns that most aid exits as capital flight or is used primarily to accumulate foreign reserves. In fact, we cannot reject the hypothesis of full absorption, in which aid increases net imports one-for-one. Absorption seems to arise mainly via increased imports. We find little support for the idea that aid lowers exports through Dutch Disease effects, although the relevant estimates are imprecise.

We also investigate the relationship between aid and the separate components of domestic expenditure. Our estimates indicate that aid increases total consumption, across a wide range of estimators and robustness tests. The effect of aid on household consumption is less precisely estimated, unless we exclude multivariate outliers identified by a robust estimator. The evidence that aid promotes investment is noticeably weaker. Although we sometimes find significant

effects on investment, especially in dynamic models, they are generally not robust to changes in specification or the exclusion of outliers.

In terms of implications for future research, the supply-push instrument appears to work well. For many of the dependent variables studied, instrumenting for aid has a substantial effect on the results, confirming that aid should not be treated as exogenous. The instrument is strong enough to generate some informative findings, and as future years of data become available, the prospects for robust, precise estimates will improve still further. With all of this in mind, the instrument could have many possible applications in the future study of aid.

References

Adam, Christopher S. and O'Connell, Stephen A. (1999). Aid, taxation and development in Sub-Saharan Africa. *Economics and Politics*, 11(3), 225-253.

Aiyar, Shekhar, Berg, Andrew, Hussain, Mumtaz, Mahone, Amber and Roache, Shaun (2006). High aid inflows: the case of Ghana. In Peter Isard, Leslie Lipschitz, Alexandros Mourmouras, and Boriana Yontcheva (eds.), *The macroeconomic management of foreign aid*. Washington DC: International Monetary Fund.

Aiyar, Shekhar and Ruthbah, Ummul Hasanath (2008). Where did all the aid go? An empirical analysis of absorption and spending. IMF Working Paper no. 08/34. Washington DC: International Monetary Fund.

Arndt, Channing, Jones, Sam and Tarp, Finn (2010). Aid, growth and development: have we come full circle? *Journal of Globalization and Development*, 1(2), Article 5.

Azevedo, João Pedro (2011). wbopendata: Stata module to access World Bank databases. <http://ideas.repec.org/c/boc/bocode/s457234.html>

Baltagi, Badi H. and Li, Jing (2014). Further evidence on the spatio-temporal model of house prices in the United States. *Journal of Applied Econometrics*, 29(3), 515-522.

Berg, Andrew, Mirzoev, Tokhir, Portillo, Rafael and Zanna, Luis-Felipe (2010). The Short-Run Macroeconomics of Aid Inflows: Understanding the Interaction of Fiscal and Reserve Policy. IMF working paper no. 10/65.

Bond, Steve, Leblebicioglu, Asli and Schiantarelli, Fabio (2010). Capital accumulation and growth: a new look at the empirical evidence. *Journal of Applied Econometrics*, 25(7), 1073-1099.

Boone, Peter (1996). Politics and the effectiveness of foreign aid. *European Economic Review*, 40(2), 289-329.

Breinlich, Holger, Ottaviano, Gianmarco I. P. and Temple, Jonathan R. W. (2014). Regional growth and regional decline. In Philippe Aghion and Steven Durlauf (eds.), *Handbook of Economic Growth, Volume 2*. The Netherlands: North-Holland.

Brock, Philip L. (1996). International Transfers, the Relative Price of Non-traded Goods, and the Current Account. *Canadian Journal of Economics*, 29, 163-80.

Brock, Philip L. and Turnovsky, Stephen J. (1994). The Dependent-Economy Model with Both Traded and Nontraded Capital Goods. *Review of International Economics*, 2(3), 306-25.

Bun, Maurice J. G. and de Haan, Monique (2010). Weak instruments and the first stage F-statistic in IV models with a nonscalar error covariance structure. University of Amsterdam, School of Economics, Department of Quantitative Economics, Discussion Paper 2010/02.

Card, David (2001). Immigrant inflows, native outflows, and the local market impacts of higher immigration. *Journal of Labor Economics*, 19(1), 22-64.

Cerra, V., Tekin, S. and Turnovsky, S. J. (2009). Foreign transfers and real exchange rate adjustments in a financially constrained dependent economy. *Open Economies Review*, 20, 147-181.

Chudik, Alexander and Pesaran, M. Hashem (2013). Common correlated effects estimation of heterogeneous dynamic panel data models with weakly exogenous regressors. CESifo Working Paper Series no. 4232. Munich: CESifo.

Chudik, Alexander, Pesaran, M. Hashem and Tosetti, Elisa (2011). Weak and strong cross-section dependence and estimation of large panels. *Econometrics Journal*, 14(1), C45-C90.

Clemens, Michael A., Radelet, Steven, Bhavnani, Rikhil R. and Bazzi, Samuel (2012). Counting chickens when they hatch: timing and the effects of aid on growth. *Economic Journal*, 122(561), 590-617.

Cohen Freue, Gabriela V., Ortiz-Molina, Hernan and Zamar, Ruben H. (2013). A Natural Robustification of the Ordinary Instrumental Variables Estimator. *Biometrics*, 69, 641-650.

Davison, A. C. and Hinkley, D. V. (1997). *Bootstrap Methods and their Application*. Cambridge University Press, Cambridge, UK.

Eberhart, Markus, Helmers, Christian and Strauss, Hubert (2013). Do Spillovers Matter When Estimating Private Returns to R&D? *Review of Economics and Statistics*, 95(2), 436-448.

Everaert, Gerdie and De Groote, Tom (forthcoming). Common correlated effects estimation of dynamic panels with cross-sectional dependence. *Econometric Reviews*.

Hansen, Henrik and Headey, Derek (2010). The short-run macroeconomic impact of foreign aid to small states: an agnostic time series analysis. *Journal of Development Studies*, 46(5), 877-896.

Harding, Matthew and Lamarche, Carlos (2011). Least squares estimation of a panel data model with multifactor error structure and endogenous covariates. *Economics Letters*, 111(3), 197-199.

Hayakawa, Kazuhiko, Pesaran, M. Hashem and Smith, L. Vanessa (2014). Transformed maximum likelihood estimation of short dynamic panel data models with interactive effects. Cambridge Working Papers in Economics, CWPE 1412.

Hodler, Roland and Raschky, Paul A. (2014). Regional favoritism. *Quarterly Journal of Economics*, 129(2), 995-1033.

- Holly, Sean, Pesaran, M. Hashem and Yamagata, Takashi (2010). A spatio-temporal model of house prices in the USA. *Journal of Econometrics*, 158(1), 160-173.
- Hussain, Mumtaz, Berg, Andrew Berg and Aiyar, Shekhar (2009). The Macroeconomic Management of Increased Aid: Policy Lessons from Recent Experience. *Review of Development Economics*, 13(S1), 491-509.
- Imbs, Jean, Jondeau, Eric and Pelgrin, Florian (2011). Sectoral Phillips curves and the aggregate Phillips curve. *Journal of Monetary Economics*, 58(4), 328-344.
- Jann, Ben (2013). coefplot: Stata module to plot regression coefficients and other results. <http://ideas.repec.org/c/boc/bocode/s457686.html>.
- Kapetanios, George, Pesaran, M. Hashem and Yamagata, Takashi (2011). Panels with non-stationary multifactor error structures. *Journal of Econometrics*, 160(2), 326-348.
- Kimbrough, Kent P. (1986). Foreign aid and optimal fiscal policy. *Canadian Journal of Economics*, 19(1), 35-61.
- Kleibergen, Frank and Paap, Richard (2006). Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics*, 133(1), 97-126.
- Kronmal, Richard A. (1993). Spurious correlation and the fallacy of the ratio standard revisited. *Journal of the Royal Statistical Society: Series A*, 156(3), 379-392
- Marshall, Monty G., Gurr, Ted Robert and Jagers, Keith (2013). Polity IV project: dataset users' manual. Center for Systemic Peace, mimeo, April 21.
- Moon, Hyungsik Roger and Weidner, Martin (2013). Dynamic linear panel regression models with interactive fixed effects. cemmap working paper no. CWP63/13.
- Nunn, Nathan and Qian, Nancy (2014). U. S. food aid and civil conflict. *American Economic Review*, 104(6), 1630-1666.

Obstfeld, Maurice (1999). Foreign resource inflows, saving, and growth. In Klaus Schmidt-Hebbel and Luis Servén (eds.), *The economics of saving and growth*. Cambridge, UK: Cambridge University Press.

Obstfeld, Maurice and Rogoff, Kenneth (1996). *Foundations of International Macroeconomics*. MIT Press, Cambridge, MA.

Pearson, Karl (1897). Mathematical contributions to the theory of evolution – on a form of spurious correlation which may arise when indices are used in the measurement of organs. *Proceedings of the Royal Society of London*, 60(359-367), 489-498.

Pesaran, M. Hashem (2006). Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica*, 74(4), 967-1012.

Pesaran, M. Hashem and Tosetti, Elisa (2011). Large panels with common factors and spatial correlation. *Journal of Econometrics*, 161(2), 182-202.

Rajan, Raghuram G. and Subramanian, Arvind (2008). Aid and growth: what does the cross-country evidence really show? *Review of Economics and Statistics*, 90(4), 643-665.

Schaffer, Mark E. (2010). xtivreg2: Stata module to perform extended IV/2SLS, GMM and AC/HAC, LIML and k-class regression for panel data models. <http://ideas.repec.org/c/boc/bocode/s456501.html>.

Stock, James H. and Yogo, Motohiro (2005). Testing for weak instruments in linear IV regression. In Donald W. K. Andrews and James H. Stock (eds.), *Identification and inference for econometric models: essays in honor of Thomas Rothenberg*. Cambridge, UK: Cambridge University Press.

Tavares, José (2003). Does foreign aid corrupt? *Economics Letters*, 79(1), 99-106.

Temple, Jonathan R. W. (2010). Aid and conditionality. In Dani Rodrik and Mark Rosenzweig (eds.), *Handbook of Development Economics, Volume 5*, 4415-4523.

Temple, Jonathan R. W. and Van de Sijpe, Nicolas (2014). Foreign aid and domestic absorption. CSAE Working Paper Series 2014-01, Centre for the Study of African Economies, University of Oxford.

Van de Sijpe, Nicolas (2010). Foreign aid and government behaviour. DPhil Thesis, Department of Economics, University of Oxford.

Van de Sijpe, Nicolas (2013). Is foreign aid fungible? Evidence from the education and health sectors. *World Bank Economic Review*, 27(2), 320-356.

Werker, Eric, Ahmed, Faisal Z. and Cohen, Charles (2009). How is foreign aid spent? Evidence from a natural experiment. *American Economic Journal: Macroeconomics*, 1(2), 225-244.

Westerlund, Joakim and Urbain, Jean-Pierre (2013). On the estimation and inference in factor-augmented panel regressions with correlated loadings. *Economics Letters*, 119(3), 247-250.

Wooldridge, Jeffrey M. (2010). *Econometric analysis of cross section and panel data* (second edition). MIT Press, Cambridge, MA.

Yule, G. U. (1910). On the interpretation of correlations between indices or ratios. *Journal of the Royal Statistical Society*, 67(6/7), 644-647.

For online publication

Appendix A The average shares instrument

This appendix sketches an argument that a supply-push instrument should be based on initial shares rather than average shares. First note that, if the initial share were instead the current share, the synthetic instrument would be equal to the variable it is instrumenting (current aid receipts) and hence endogenous. But the current share is one component of the average share: in the simple case of one donor with a total budget D_t at date t , taking the average of the shares over time means that the instrument for recipient i at date t is $(1/T) \cdot (a_{i1} + \dots + a_{iT}) \cdot D_t / Y_{it}$, in which one component of the sum is therefore $a_{it} D_t / Y_{it}$, or current aid at date t . Hence, using average shares implies the value of the instrument at each date is a function of the endogenous variable at that date: this will typically imply some degree of endogeneity, although it may achieve bias reduction. In addition, at least some of the aid shares in other periods, a_{is} for $s \neq t$, are likely to be a function of the transient error at date t , and this could reinforce the likely failure of exogeneity when using average shares.

Appendix B Additional robustness checks

Table B.1: Controlling for $1/Y$

Row	Model	C	G	C+G	I	X	M	M-X
1	Static	0.708 (0.435)	-0.0583 (0.150)	0.662** (0.284)	0.279 (0.252)	-0.0464 (0.267)	0.622*** (0.210)	1.085*** (0.285)
	F	18.18	14.99	19.84	14.67	30.69	29.27	18.14
	Dyn.	0.856 (0.655)	-0.0566 (0.172)	0.775* (0.436)	0.700* (0.394)	0.346 (0.402)	1.115*** (0.385)	1.543*** (0.513)
	F	11.84	13.89	12.12	11.21	30.97	19.52	13.45
2	Static	1.583** (0.795)	-0.0706 (0.315)	1.282** (0.626)	-0.187 (0.478)	-0.128 (0.472)	0.308 (0.390)	0.981** (0.457)
	F	9.665	12.58	9.033	8.418	11.12	13.20	8.509
	Dyn.	1.766* (1.027)	-0.193 (0.436)	1.370* (0.755)	-0.0103 (0.649)	1.117 (0.905)	1.149 (0.696)	1.347** (0.660)
	F	8.599	10.75	6.999	6.179	9.434	10.44	5.168
3	Static	0.702 (0.431)	-0.0485 (0.147)	0.660** (0.281)	0.279 (0.246)	0.00702 (0.277)	0.663*** (0.216)	1.085*** (0.286)
	F	18.54	14.79	21.04	15.19	30.01	29.17	18.91
	Dyn.	0.851 (0.639)	-0.0481 (0.168)	0.773* (0.425)	0.691* (0.377)	0.387 (0.420)	1.124*** (0.394)	1.546*** (0.512)
	F	12.23	13.68	12.92	12.65	31.42	19.58	14.33
4	Static	1.570* (0.813)	-0.0485 (0.289)	1.273* (0.653)	-0.188 (0.482)	0.0363 (0.493)	0.370 (0.424)	0.963* (0.493)
	F	12.61	13.23	11.51	10.85	10.37	13.71	10.14
	Dyn.	1.730 (1.045)	-0.150 (0.401)	1.351* (0.785)	-0.0224 (0.683)	1.336 (0.947)	1.126 (0.718)	1.306* (0.704)
	F	11.08	10.88	8.895	9.467	9.438	12.07	7.209
5	Static	0.692 (0.560)	0.0233 (0.113)	0.643 (0.465)	-0.0509 (0.392)	0.0958 (0.300)	0.499 (0.504)	1.007** (0.451)
	F	12.65	9.983	15.12	7.590	16.25	6.817	13.06
	Dyn.	0.838 (0.654)	0.00645 (0.126)	0.934 (0.642)	-0.117 (0.515)	0.409 (0.356)	0.829 (0.580)	1.318* (0.711)
	F	7.699	8.707	8.491	8.250	13.37	5.812	9.979

Continued on next page

Table B.1 – continued from previous page

Row	Model	C	G	C+G	I	X	M	M-X
6	Static	2.190** (0.915)	0.0620 (0.262)	2.078** (0.889)	-1.021 (0.698)	-0.194 (0.530)	0.904 (0.801)	0.933* (0.493)
	F	6.029	5.572	6.246	5.355	6.559	5.854	6.239
	Dyn.	2.477* (1.332)	-0.0153 (0.290)	2.313** (1.120)	-1.108 (1.027)	0.795 (0.782)	1.377 (1.011)	1.169 (0.724)
	F	3.757	7.970	3.802	3.210	6.181	5.508	3.878

Note: The entries in this table show the long-run effect of aid on household consumption (C), government consumption (G), total consumption ($C + G$), gross capital formation (I), exports (X), imports (M) and net imports ($M - X$) in models with (“Dyn.”) and without (“Static”) a lagged dependent variable. All variables expressed relative to GDP. CCE IV estimation on three-year averaged data (1971-2012) using an instrument based on initial shares in donor budgets calculated over the period 1960-70. 1099 observations from 88 countries, unless reported otherwise below. Heteroskedasticity-robust standard errors, clustered by country and corrected for degrees of freedom, in brackets. Standard errors (SE) for the long-run effects obtained via the delta method. *, **, and *** denote significance at 10, 5 and 1%, respectively. F shows the first-stage F-statistic.

Row 1 repeats the main results from Tables 1-7 for ease of comparison.

Row 2 removes outliers, repeating row 16 of Table 8 (1004 observations from 81 countries).


Row 3 controls for $1/Y$ in the full sample.

Row 4 controls for $1/Y$ in the sample without outliers (1004 observations from 81 countries).

Row 5 controls for $1/Y$ and its cross-sectional mean, allowing the latter to enter with country-specific coefficients, in the full sample. Only countries with at least 7 time series observations are included (1087 observations from 86 countries).

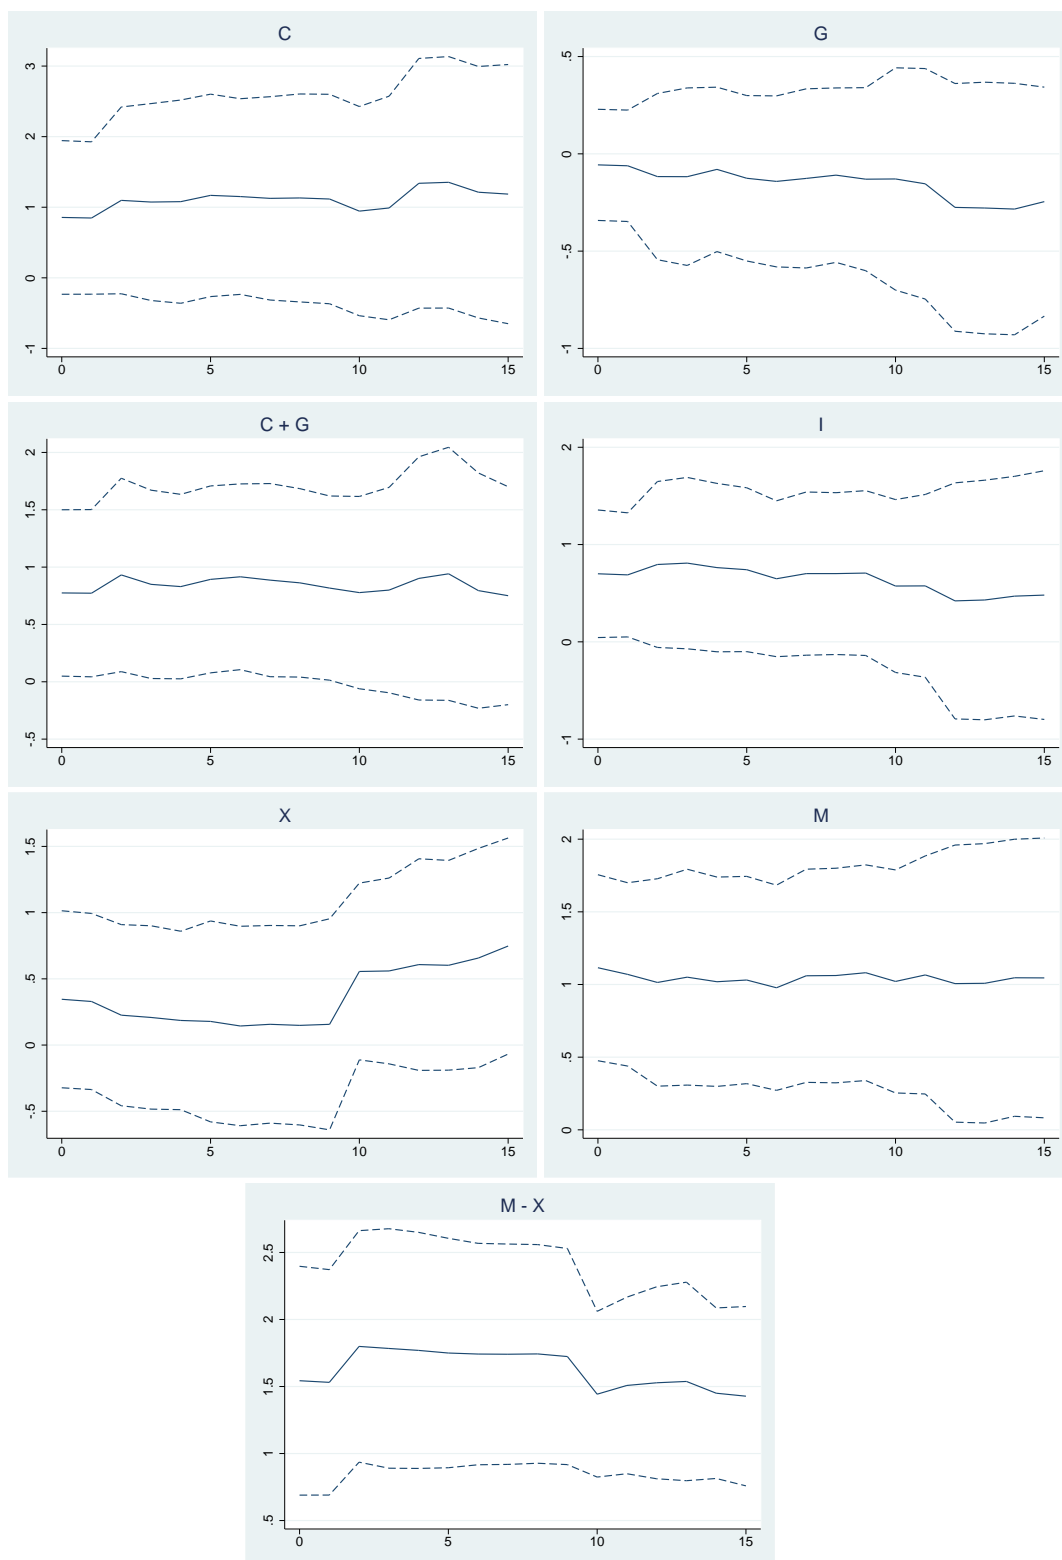
Row 6 controls for $1/Y$ and its cross-sectional mean, allowing the latter to enter with country-specific coefficients, in the sample without outliers. Only countries with at least 7 time series observations are included (992 observations from 79 countries).

Figure B.1: Dropping countries with the largest relative fall in real GDP (static model)



Note: graphs show how the estimated effects of aid (solid line) on household consumption (C), government consumption (G), total consumption ($C + G$), gross capital formation (I), exports (X), imports (M) and net imports ($M - X$) change when progressively dropping the countries with the largest percentage declines in real GDP, based on CCE IV estimation of a static model. Dashed lines indicate the 90% confidence interval. Horizontal axis shows the number of countries dropped. Graphs constructed with coefplot for Stata (Jann, 2013).

Figure B.2: Dropping countries with the largest relative fall in real GDP (dynamic model)



Note: see note Figure B.1. These graphs show the long-run effects of aid (solid line) with 90% confidence intervals (dashed lines) when progressively dropping the countries with the largest percentage declines in real GDP, based on CCE IV estimation of a model that includes a lagged dependent variable.