



This is a repository copy of *Foreign Aid and Domestic Absorption*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/120412/>

Version: Accepted Version

Article:

Temple, J. and Van de Sijpe, N.M. orcid.org/0000-0003-0442-7020 (2017) Foreign Aid and Domestic Absorption. *Journal of International Economics*. ISSN 0022-1996

<https://doi.org/10.1016/j.jinteco.2017.07.010>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Foreign Aid and Domestic Absorption

Jonathan 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

30th May 2017

ABSTRACT

This paper introduces 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.

JEL Classifications: F35

Keywords: Foreign aid, absorption, Dutch Disease, common correlated effects

For various contributions and advice that have improved the paper, we are grateful to the editor and anonymous referees, Chris Adam, Erwin Bulte, Paddy Carter, Gabriela Cohen Freue, Sarah Dykstra, 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, Leicester, 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, the Northeast Universities Development Consortium, and the Royal Economic Society. The first author is also grateful to the British Academy for financial support from a Mid-Career Fellowship, award MD120067.

1 Introduction

Although a large literature studies the effects of aid using cross-country regressions, its problems are well-known. Researchers 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 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 estimator that filters out unobserved common factors, even when their effects differ across countries. In principle, this combination of instrument and estimator will identify the causal effect of aid under more general conditions than existing approaches. It 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. It might be offset by a corresponding capital outflow, or used to accumulate foreign exchange reserves. Some of the aid flows recorded by donors will not correspond to international transfers: for example, some forms of donor-sponsored technical assistance will not have a direct effect on the recipient’s domestic expenditure. 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. 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, and this has been a central concern of the empirical aid literature since its inception. We will argue that the effects of aid on investment, and other macroeconomic ratios, are easier to study than the effects on growth. Further, in our empirical work, these effects will be identified from relatively persistent changes in aid receipts, rather than transitory and endogenous changes.

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 .

As an example, consider what happens if the British aid budget increases relative to the French aid budget. 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. Our instrument uses changes in total donor budgets, weighted by the initial shares of recipients in those budgets, to isolate exogenous changes in aid receipts that are not driven by the conditions of individual aid recipients. We call this a supply-push instrument; it is related to the work of Bartik (1991) on regional economics and Card (2001) on the labor market effects of immigration. As in the immigration setting, the origins and destinations of flows of aid are large in number, and this makes it unlikely that the instrument — as a weighted average of many donor budgets — will be correlated with recipient-specific conditions. We investigate this further below.

A possible objection is that donor budgets may be influenced by forces common to many recipients. For example, world economic conditions are likely to affect donor generosity, and also the outcomes of poor countries. Drawing on recent work in the panel time series literature, global forces can be seen as unobserved common factors with loadings that differ across countries. We filter these out 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). Our paper is the first to apply this approach to the study of foreign aid. The combination of the new instrument and estimator should mean that we identify causal effects of aid under more general conditions than the existing literature.

We find that aid is at least partially absorbed, reflected in net imports. 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 symptoms of Dutch Disease. The find-

ings 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 by increases in household consumption, but those estimates are less precise unless we exclude outliers. The evidence that aid promotes investment is weaker. In some models, aid has a delayed effect on investment, but these results are sensitive to the estimation method and the exclusion of outliers.

The next section will sketch possible relationships between aid and macroeconomic ratios. Section 3 explains the approach to estimation and its relation to the literature. Section 4 describes the data. In section 5, we analyze whether and how aid is absorbed, and the possibility of Dutch Disease. Section 6 presents robustness checks, before section 7 concludes. An appendix describes the CCE IV estimator.

2 Aid and macroeconomic ratios

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. Alternatively, aid might be used to accumulate foreign reserves, or lead to a capital outflow. Some aid may be spent on consultants who work exclusively in the donor country, with no direct effect on the aid recipient's domestic expenditure. It is therefore interesting to ask whether aid is absorbed. Domestic absorption is typically defined as the sum of household consumption, gross investment, and government consumption. We are interested in (1) whether aid is reflected in higher domestic expenditure on final goods and services, and (2) which expenditure components are most affected. This helps to clarify what is at stake in the paper. We show that aid is generally absorbed — it increases expenditure relative to output — but also find that consumption responds more strongly than investment. We do not uncover any symptoms of Dutch Disease. These results do not establish whether aid is 'effective', a hard task for a single paper, but they do contribute new evidence to the relevant debates.

We first consider what it means for aid to be fully absorbed. Take Y as GDP, equal to the sum of household consumption C , gross investment I , and government consumption G , minus net imports $M - X$. For 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 GDP 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 permits greater domestic expenditure relative to domestic production.¹ In the short run, if aid is

¹For more on absorption see Adam (2013), Aiyar and Ruthbah (2008), Berg et al. (2010), Hansen

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.

To study absorption, we take the ratios $C/Y, I/Y, G/Y$ and $(M - X)/Y$ as our dependent variables. The way aid is absorbed might differ between the short and the long run. 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 temporarily postponed.² More generally, the relationships between macroeconomic ratios and aid could be complicated over longer time horizons. If aid is spent in ways that improve the investment climate, the long-run effect of aid on investment could be much larger than the short-run effect. Or consider what happens when donor funds are spent on consultants working in the donor country: short-run absorption will be zero, but technical advice may later be reflected in economic policies and hence in macroeconomic ratios. Our empirical analysis will distinguish between short-run and long-run effects, by estimating dynamic models, sometimes with a role for lagged aid.

The models we estimate can be related to macroeconomic theories of the aggregate effects of aid. In the one-sector Ramsey model, a permanent increase in aid raises the investment ratio in the short run, but not 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 (Obstfeld, 1999). 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. When the ratio of aid to GDP increases permanently, the long-run C/Y and $(M - X)/Y$ ratios increase to the same extent, leaving the other ratios unchanged. In a two-sector model of a small open economy, with traded and non-traded goods, the effects are more complicated. Aid may increase or decrease the long-run capital stock and gross investment, depending on whether traded production is relatively capital-intensive. Section B.1 of the online appendix discusses this in more detail.

Models with balanced growth paths typically imply that the long-run consumption and investment ratios are stable functions of structural parameters. In a model with shocks, there would be a common stochastic trend in consumption, investment and output, while the long-run ratios would be mean stationary. The long-run ratio of consumption to output would be linear in the ratio of aid to GDP, where the intercept

and Headey (2010) and Hussain et al. (2009).

²Berg et al. (2010) and Hussain et al. (2009) analyze these decisions in detail, emphasizing that absorption 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. Rodrik (2006) discusses the welfare costs associated with holding foreign reserves.

depends on structural parameters and can be treated as a fixed effect. By working in terms of ratios to GDP, we stay close to these predictions and avoid the non-stationarity that would arise with alternative explanatory variables, such as aid per capita.

One drawback is that scaling our variables by GDP risks inducing correlations between variables that were originally unrelated. We examine this possibility in detail, including the pattern of results across the different ratios. We consistently find effects of aid on some ratios and not others, where the pattern conforms with the predictions of theoretical models, and where the effect sizes have plausible magnitudes. To generate this pattern, a story based on spurious correlations might have to be somewhat contrived. Other reasons to be wary of that explanation include additional results from first-differenced models, and models with the dependent variable in logarithms and which also include the logarithm of GDP as an explanatory variable.

There are some advantages to studying absorption rather than growth. First, the relationships between macroeconomic ratios and aid intensity are more likely to be linear, for the reasons just discussed. Second, if aid improves the conditions for domestic investment, this effect might be relatively easy to detect. If we see consumption and investment as jump variables, they can respond quickly to changes in aid. In contrast, GDP is a function of state variables such as the capital stock: the relevant effects of aid could take time to emerge, and empirical researchers have to contend with the high degree of persistence of GDP. With these points in mind, it should be easier to establish reliable findings for absorption than for growth.

Our approach allows us to make progress on some fronts, but not others. Although our instrument has some major advantages, it would be difficult for us to adapt the current approach to allow for parameter heterogeneity or non-linearities. And given the limitations of the available data, we cannot disaggregate aid while retaining a sample large enough for the methods that we adopt. The estimates we obtain might be best interpreted as the average effects of typical or business-as-usual aid. The paper is therefore complementary to previous work, given its narrower focus and the trade-offs that inevitably arise in addressing some econometric issues and not others.³

3 Methods

As we noted earlier, when studying the effects of aid, a central problem is that aid is not randomly assigned across countries. Even in a model that controls for country and time fixed effects, it is likely that aid flows and outcome variables are jointly influenced

³For discussion of econometric issues, see Roodman (2007a, 2007b) and Temple (2010); on diminishing returns to aid, see Carter et al. (2015).

by time-varying variables that are not readily measured. To address this, we adopt a supply-push instrument which is a weighted average of donor budgets, where the 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 fixed shares of common aggregates, such as donor budgets. 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 $A_{it}^S/Y_{it} \equiv (\sum_{d=1}^{N_D} a_{i0}^d D_{dt})/Y_{it}$, where a_{i0}^d is the share of donor d 's total aid disbursements that recipient i receives, over an initial period that is excluded from estimation, and D_{dt} is the total aid disbursement made by donor d in period t .

One advantage of the new instrument is especially worth discussion. It is likely that much empirical work on aid conflates the effects of permanent and temporary variation, just as early work on consumption conflated the effects of permanent and transitory income (Carter, 2015). From a policy perspective, a researcher might be more interested in determining the effects of a permanent change in aid. One solution is to use an instrument that is correlated with the permanent component and not with the transitory component. Since our instrument is a weighted average of donor budgets, and individual donor budgets are persistent, it should come closer than some precursors to identifying the effects of permanent changes in aid.

In using this instrument, we are assuming that the total aid budgets of most donors are not greatly influenced by the time-varying conditions of many individual aid recipients. As background, aid flows are increasingly fragmented. The number of significant donors has increased, and most donors provide aid to a large number of countries. This is documented in Annen and Moers (2017), Djankov et al. (2009), Easterly (2007) and Knack and Rahman (2007), among others. Even as early as the 1970s, the US accounted for less than a quarter of total aid flows. According to Annen and Moers, the average bilateral donor provided aid to about 20 recipients in 1960, rising to 87 recipients by 2011.

For identification, we rely on within-donor and within-recipient fragmentation. We want to avoid correlations between the instrument and time-varying conditions in a recipient country. If aid is fragmented, endogeneity will arise only when the total aid budgets of multiple donors respond simultaneously to a recipient's conditions, and to a large degree. This does not seem especially plausible. Total aid budgets are likely to emerge from a medium-run political process, and the hypothesized strong responses

would imply greater volatility in total aid budgets than we see in the data. If aid is not fragmented, the most serious problem for the instrument would arise when a donor gives to a small number of recipients, and those recipients receive most of their aid from that donor. There are few such cases in the data. The observed fragmentation suggests that the instrument will be uncorrelated with time-varying recipient conditions in most cases.⁴

To support this claim, we first look at within-donor fragmentation. We construct Herfindahl-Hirschman (HH) indices for the 29 donors that contribute to the supply-push instrument for at least one aid recipient in our main sample; low values correspond to aid that is distributed across many recipients, or high fragmentation.⁵ Figure 1 shows box-plots for the distribution of the HH indices across donors, for each three-year period in our sample. If we take the box-plot for 1971-73, the upper border of the box indicates that 75% of the donors had HH indices below about 0.30, indicating high fragmentation even in the first period.

Reading across the figure, the long-run trend is towards greater fragmentation (lower median HH indices). The dots indicate isolated exceptions, which relate to two minor donors and a small number of recipients. Even the exceptions are unlikely to be a problem, because these recipients will typically receive aid from multiple sources. Since the instrument is a weighted average of donor budgets, this should weaken the correlation between the instrument and recipient conditions even when a subset of the donors are highly specialized. But to check this, in the later analysis, we will consider versions of the instrument which exclude specialized donors. The results continue to indicate strong effects of aid on total consumption and net imports.


The high degree of within-donor fragmentation supports our approach. It suggests that the total aid budgets of these donors are unlikely to be strongly driven by the domestic conditions of individual recipients. We next examine within-recipient fragmentation. As indicated previously, this will promote identification when the value of the supply-push instrument for a given recipient draws on aid from multiple donors. This is because, even if some donors to a given recipient have total aid budgets which are correlated with that country's domestic conditions, summing over a larger set of donors will tend to weaken the correlation between those conditions and the instrument.

Note that, for identification, it will be fragmentation *within the instrument* at the recipient level that matters, rather than within a recipient's aid. We construct HH indices

⁴Section B.3 of the online appendix discusses the relationship between fragmentation and identification in more detail.

⁵The HH indices are based on the shares of each recipient in that donor's total aid disbursement, ignoring a very small fraction of negative aid entries arising from the repayment of loans. The shares are squared and summed over recipients to give an index between 0 and 1, where larger numbers indicate low fragmentation.

Figure 1: Within-donor Herfindahl-Hirschman indices: distribution across donors




Note: These are box-plots that summarize the distribution of HH indices across donors, in each sample period, for donors that contribute to the instrument values of at least one aid recipient in our sample in that period. The lower and upper borders of the box correspond to the 25th and 75th percentiles respectively, and the line inside the box shows the median. Dots highlight atypical observations more than 1.5 times the interquartile range from either the 25th or the 75th percentile. The whiskers indicate the range for the observations excluding these dots.

for each recipient and year, based on the contribution of each donor to that recipient's supply-push instrument in that year. Figure 2 shows box-plots for the distributions of these within-instrument HH indices across recipients, drawing on the 1099 recipient-period observations in our main sample. The lines in the middle of the boxes show the degree of fragmentation within the instrument for the median recipient: for many recipients, the instrument is relatively fragmented, supporting identification. There are some exceptions where recipients have HH indices close to one, and hence where the instrument places a high weight on a single donor for that particular recipient. These are isolated cases which are unlikely to dominate the variation, and should not threaten identification given the within-donor fragmentation documented above.⁶ In summary, the analysis suggests that a supply-push approach can be applied to the study of foreign aid.

⁶Note that, for identification, dependence on a single donor is not a problem in itself. Concerns arise when that donor gives to only a small number of recipients, since then it becomes more likely that the donor's total budget is endogenous to conditions in those recipients in a way that will not be weakened by summing over donors.

Figure 2: Within-instrument Herfindahl-Hirschman indices: distribution across recipients



Note: The box-plots summarize the distribution of HH indices across recipients, based on within-recipient instrument fragmentation in each period. For interpretation, see the main text and the notes to Figure 1.

We now turn to a different objection, which is that total aid budgets might be influenced by disaster and emergency relief. But even broadly defined, humanitarian assistance accounts for a small share of global aid flows: for 1995-2013, Qian (2015) finds that it ranged between 5% and 9% of official development aid. Some humanitarian assistance is long-term rather than emergency-related, and several of the major recipients are not in our data set.⁷ Finally, since emergency relief may be funded by reallocations within existing budgets, and disasters do not appear to have major effects on aid receipts (Qian, 2015), we do not see strong grounds to reject the supply-push approach on this basis.

The paper's contribution goes beyond the supply-push instrument for aid. This is the first paper to estimate the causal effect of aid while allowing for latent common factors with heterogeneous effects, such as global economic conditions. Imagine the

⁷For example, Afghanistan, Haiti, Iraq, Somalia, and the West Bank and Gaza Strip. Major recipients of humanitarian assistance in recent years are listed in Development Initiatives (2014); on their definition, humanitarian aid accounted for roughly 10% of aid from OECD DAC donors in 2004-2013.

process generating the outcome of interest Q_{it}/Y_{it} is given by:

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

$$\varepsilon_{it} = \lambda_i' F_t + u_{it} \quad (2)$$

where F_t is a vector of unobserved common factors (including, say, world economic conditions) and λ_i is a vector 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 time effects (where loadings on one time-varying factor are the same across countries) as special cases. The generality of this structure has made it a focus of recent econometric research, and applications have spanned a range of fields, as we note below.

Importantly, allowing for heterogeneous effects of common factors will help to ensure that our supply-push instrument is exogenous. In contrast, conventional panel estimators with time fixed effects assume that common factors, such as global shocks, have exactly the same effects on all the countries in the sample.⁸ If the data generating process is more complicated, identification could fail, because the instrument might be correlated with the effects of the common factors. The combination of a supply-push instrument and a flexible approach to common factors is new to this paper, and should achieve identification under a wider range of circumstances than previous work.

We now explore this in more detail, and 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 $\lambda_i' F_t$, or with the country-specific shock u_{it} . For a 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} , at all dates. This could be questioned. For example, donor budgets may be correlated with world economic conditions which also influence macroeconomic ratios in individual aid recipients. In that case a supply-push instrument could be correlated with $\lambda_i' F_t$ even when there is no correlation with u_{it} .

To address this, we filter out common factors using the approach of Pesaran (2006). His paper introduced common correlated effects (CCE) estimators for panel data. This class of estimators proxies for the combined effects of common factors using linear combinations of the cross-section means of the observable variables. To allow the effects of the factors to differ across countries, the combinations are estimated from

⁸When this assumption seems too restrictive, researchers sometimes interact time dummies with a small set of country-specific observables; Breinlich et al. (2014) call this latter approach *proportional time effects*. But this is also restrictive, because it treats the unobserved factor loadings as simple in structure and known to the researcher.

the data and vary across countries. This is achieved by augmenting the regression with cross-section means of the dependent variable and the explanatory variables, all with country-specific coefficients. Under a condition on the number of linearly independent common factors, their combined effect will be captured by the cross-section means. The method thereby addresses an important class of omitted variables in a way that is straightforward to implement. It can accommodate various forms of cross-section dependence, and simulations suggest that it can perform well even in small samples and when the factors are non-stationary.⁹

The CCE method has been extended to the case of instrumental variables by Harding and Lamarche (2011), yielding a CCE IV estimator that we describe in Appendix A. This estimator is again easy to implement. The first and second stages of 2SLS are augmented with cross-section means of the observable variables, including the instrument, with country-specific coefficients. This approach has costs and benefits. It builds in robustness, helping to ensure that the instrument is exogenous. However, filtering out the common factors is parameter-intensive and could ask a lot of the data. The key point here is that, although the benefits of CCE IV may be offset by higher standard errors, our estimates are often precise enough to be informative.¹⁰ Although we report the results from several methods, we give most emphasis to CCE IV, as the estimator most likely to yield consistent estimates of the parameters of interest.

We now discuss how the paper relates to previous work. The instrument we adopt was first used by Van de Sijpe (2010) to study aid and governance, but without allowing for a multifactor error structure.¹¹ Work using other supply-push instruments includes Nunn and Qian (2014) and Werker et al. (2009). The instrument in the latter paper interacts the world oil price with a dummy for Muslim countries, since aid to Muslim aid recipients may be sensitive to the oil price. Werker et al. use this to study the effects of aid on a range of outcomes, including macroeconomic ratios. Their findings tally closely with ours. They find a significant effect of aid on consumption, where the IV estimate is much larger; no evidence that aid leads to higher government consumption; some evidence that aid promotes gross investment, but this is not robust; no evidence

⁹For references to the literature, see chapter 29 of Pesaran (2015). Other textbook presentations include Hsiao (2014) and Söderbom et al. (2015), while previous applications include Baltagi and Li (2014), Bond et al. (2010), Eberhardt et al. (2013), Holly et al. (2010) and Imbs et al. (2011).

¹⁰To get a sense of whether the degrees of freedom are sufficient for over-fitting to be avoided, the balance of parameters and observations in our models is comparable to a hypothetical fixed effects, static panel data model with $N = 350$ and $T = 3$. See also Section B.5 of the online appendix, and the simulation evidence in Harding and Lamarche (2011), Table 1.

¹¹A related synthetic measure of aid, based on average shares in donor budgets rather than initial shares, was used in Hodler and Raschky (2014). Average shares may be affected by developments within recipients, which weakens the case for exogeneity, as we discuss in Section B.2 of the online appendix.

that aid affects exports; and strong evidence that aid leads to higher imports. Although the two papers differ in significant respects, the findings on absorption are remarkably similar.

Our instrument can be used to study broadly-defined aid, whereas natural experiments tend to be informative about 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 Nunn and Qian (2014) is focused on US food aid, using lagged US wheat production as the supply-side push variable. Another advantage of our approach is that we can allow for common factors. Although some papers introduce observable proxies for $\lambda'_i F_t$, this is necessarily restrictive. In contrast, the CCE approach does not require either the common factors or their heterogeneous loadings to be observable.

Our approach is related to other work on aid using instrumental variables, including Galiani et al. (2017), Jarotschkin and Kraay (2016) and Tavares (2003). The latter paper used the distance between recipients and donors, and whether they share a common border, language or religion, to instrument for aid. In our study, the initial shares in donor budgets can proxy for many possible connections between donors and recipients while remaining agnostic about their sources. Put differently, we infer connections from the data, rather than restricting them to take specific forms.

In summary, the combination of a supply-push instrument and the CCE IV estimator is new to this paper. The approach has several benefits. First, by using the total aid budgets of many donors to construct the instrument, we are exploiting the fragmentation of global aid flows to lessen the risk that the instrument is correlated with the conditions of individual aid recipients. This will be the case even if those conditions are driven by country-level trends, for example. Second, we use the CCE IV approach to address a remaining concern, that total donor budgets could be influenced by common factors, such as world economic conditions, that are correlated with conditions in aid recipients. The combination of instrument and estimator allows us to go beyond the natural experiments studied in the literature to date, and study the effects of broadly-defined aid from a wide range of donors.

4 Data

Our models will be estimated using three-year averages over 1971-2012.¹² 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 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, and our final sample comprises 88 aid recipients.

Our synthetic measure for aid is constructed 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).¹³ 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 will be measured relative to GDP. The data on GDP and its components are taken from online World Bank data using `wbopendata` (Azevedo, 2011).

The dependent variables considered will include household consumption, government consumption, gross capital formation, imports and exports, again relative to GDP.¹⁴ Net imports are defined as imports minus exports. In the recipient-year data, before

¹²The 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, given the need to estimate models with a more complicated dynamic structure.

¹³For a small percentage of observations, the numerator or denominator 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. This is likely to reflect repayments of the principal of loans, which are treated as negative flows in the net ODA data. 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.

¹⁴Since a linear combination of the ratios equals unity by construction, the model for one of the dependent variables will be statistically redundant when the covariates are the same across regressions, but the covariates differ in the case of CCE estimation.

collapsing to three-year averages, observations for these variables are turned to missing whenever at least one of the other components of the GDP identity is missing. This keeps the sample consistent across the different outcomes we consider. In our final data set, the expenditure components sum to total GDP, or very close to GDP, for each country-period observation.¹⁵ We exclude countries with populations fewer than 500,000 people in the first period of the sample. These steps leave us with a panel with 1099 observations, based on 88 countries.¹⁶

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 models that include a lagged dependent variable.¹⁷ In each case, the coefficient estimates indicate the effect of the aid-to-GDP ratio on the dependent variable, also measured as a ratio. For example, in the case of net imports, a point estimate of one implies that the ratio of net imports to GDP increases one-for-one with the ratio of aid to GDP, corresponding to full absorption.

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 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 estimator in Harding and Lamarche (2011) is not known. The bootstrapped standard errors are noticeably larger than the conventional standard errors in some cases, but our main findings obtain under either approach to inference.¹⁸

¹⁵Only 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.

¹⁶These are Algeria, Angola, Argentina, Benin, Bolivia, Botswana, Brazil, Burkina Faso, Burundi, Cambodia, Cameroon, Central African Rep., Chad, Chile, Colombia, Costa Rica, Cote d'Ivoire, Cuba, Cyprus, Dem. Rep. of the Congo, Dominican Rep., Ecuador, Egypt, El Salvador, Ethiopia, Fiji, Gabon, Ghana, Guatemala, Guinea, Guinea-Bissau, Guyana, Honduras, Hong Kong, India, Indonesia, Iran, Israel, Jamaica, Jordan, Kenya, Kuwait, Lebanon, Lesotho, Libya, Madagascar, Malawi, Malaysia, Mali, Mauritania, Mauritius, Mexico, Morocco, Mozambique, Nepal, Nicaragua, Niger, Nigeria, Oman, Pakistan, Panama, Papua New Guinea, Paraguay, Peru, Philippines, Rep. of the Congo, Rwanda, Saudi Arabia, Senegal, Sierra Leone, Singapore, South Korea, Sri Lanka, Sudan, Syria, Tanzania, Thailand, Togo, Trinidad and Tobago, Tunisia, Turkey, Uganda, Uruguay, Venezuela, Vietnam, Yemen, Zambia, and Zimbabwe.

¹⁷The sample for each of the eight regressions consists of the observations included in the CCE IV estimation of the dynamic model. In our main estimates, we have 13 or 14 time series observations for many countries, with a maximum of 14.

¹⁸In the case of CCE IV estimates, we also report bootstrapped, bias-corrected 90% confidence intervals for the long-run effect, based on the BC_a method (see Davison and Hinkley, 1997, pp.

The long-run responses of macroeconomic ratios to aid may differ from those in the short run. The inclusion of a lagged dependent variable is one way to capture this. Whenever the estimated model is dynamic, the long-run effect of aid is estimated using the ratio of the short-run effect to one minus the coefficient of the lagged dependent variable, with a standard error approximated by the delta method. Since the long-run effect is a ratio, the estimates are likely to be less precise than the estimates of short-run effects.¹⁹ Dynamic models also have the drawback that CCE-type estimators will be consistent under more restrictive assumptions than in the static case.²⁰

With this in mind, our discussion will give more emphasis to static models; the use of three-year averages implies that, even in these models, absorption can extend over several years. Before we turn to the results, note that the effects of instrumenting for aid and allowing for common factors are likely to vary across the dependent variables. Macroeconomic ratios are likely to differ in their sensitivity to particular aid-relevant shocks, and in their sensitivity to common factors.

We first study the effects of aid on trade-related variables, starting with net imports. Recall that net imports must increase if aid is 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), used to accumulate foreign exchange reserves, or spent on forms of technical assistance that do not have a direct effect on expenditure beyond the donor country. The 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 different from unity, both in static models and in the long run derived from dynamic models. Note the contrast with the upper row of estimates: in the absence of an instrument, the evidence that aid is fully absorbed is weaker.

We can also investigate whether there are symptoms of aid-driven Dutch Disease. An increase in domestic expenditure will often fall partly on non-traded goods, increasing both their relative price and the costs facing the export sector.²¹ Tables 2 and 3 show the effects of aid intensity on import and export shares respectively. The results for the import share are similar to those for net imports, with the exception of the static model estimated by FE IV. A strong positive effect of aid is found in the two CCE IV

203-211) and using 1999 replications.

¹⁹As is well known, the estimation of dynamic panel data models with fixed effects is subject to the Nickell (1981) bias. That bias is typically less serious when the data span a long time, and also for a long-run effect, because the biases in the individual parameter estimates can offset each other.

²⁰See Chudik and Pesaran (2015) and Everaert and De Groot (2016). The development of estimators for this case is an active area of research.

²¹In 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, with imports partly financed by the transfer. See Cerra et al. (2009, p. 149).

Table 1: Aid and net imports

	(1) FE	(2) FE	(3) CCE	(4) 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
	(5) FE IV	(6) FE IV	(7) CCE IV	(8) 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			[0.73,1.86]	[0.97,2.80]
First stage F-statistic	19.10	19.27	18.14	13.45
Underidentification	0.001	0.001	0.009	0.014

Note: Dependent variable is net imports. All variables measured 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 averages, 1971-2012 ($N = 88$, $NT = 1099$). IV regressions carried out using `xtivreg2` for Stata (Schaffer, 2010). FE and FE IV regressions allow for country and time fixed effects, coefficients not reported. Country fixed effects and country-specific coefficients on cross-section means in CCE and CCE IV regressions not reported. Heteroskedasticity-robust standard errors, clustered by country, in brackets. Standard errors (SE) for long-run effects based on the delta method. *, **, and *** denote significance at 10, 5 and 1% respectively. BC_a shows a bias-corrected-and-accelerated 90% confidence interval from a non-parametric block bootstrap. Underidentification shows the p-value of the Kleibergen and Paap (2006) LM test for underidentification.

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 – that finding would require zeroes estimated with greater precision – but nor is there robust evidence that aid adversely affects exports.²²

²²This is consistent with Jarotschkin and Kraay (2016), who find little evidence that aid leads to real exchange rate appreciations. Previous evidence on aid-induced Dutch Disease is mixed. Rajan and Subramanian (2011) use variation across sectors and countries, and find that aid lowers the relative growth of tradable sectors. Nsor-Ambala (2015) finds less evidence of this. The review by Adam (2013) discusses the evidence in more detail.

Table 2: Aid and imports

	(1) FE	(2) FE	(3) CCE	(4) 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
	(5) FE IV	(6) FE IV	(7) CCE IV	(8) 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			[0.33,1.12]	[0.54,2.39]
First stage F-statistic	19.10	21.00	29.27	19.52
Underidentification	0.001	0.001	0.019	0.006

Note: Dependent variable is imports. $N = 88, NT = 1099$. For other notes, see Table 1.

Next, we study how aid is absorbed. Note that, since aid leads to an increase in net imports, the GDP identity implies that the sum of household consumption, government consumption and total investment must have also increased. The question is whether we can reliably identify the components of GDP which respond most strongly to aid. We first look at 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, there are sectors such as education 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 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. Donors may respond to country-specific adverse shocks by allocating countries more aid, and this form of endogeneity will weaken the correlation between aid and total consumption. By instrumenting aid we alleviate this source of bias, and find larger effects of aid.

Table 3: Aid and exports

	(1) FE	(2) FE	(3) CCE	(4) 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
	(5) FE IV	(6) FE IV	(7) CCE IV	(8) 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			[-0.37,0.58]	[-0.31,1.21]
First stage F-statistic	19.10	18.21	30.69	30.97
Underidentification	0.001	0.001	0.018	0.012

Note: Dependent variable is exports. $N = 88$, $NT = 1099$. For other notes, see Table 1.

For the effects of aid on household consumption, the estimates 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 a different 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 be surprising. One mechanism could be lower taxes: increased aid to governments may not be used to increase government purchases, but to reduce taxes (Kimbrough, 1986). Alternatively, recipient governments may use aid to finance transfers for political ends, as in Adam and O'Connell (1999), Boone (1996) and Hodler and Raschky (2014), among others. Finally, some aid is given in ways which bypass domestic governments, such as off-budget aid projects and support for NGOs.²³ In these cases, household consumption is where the effect of aid is most likely to be manifested in the national accounts.

Finally, we look at the effect of aid on the investment rate, in Table 7. When

²³Van de Sijpe (2013) discusses off-budget aid in more detail.

Table 4: Aid and total consumption

	(1) FE	(2) FE	(3) CCE	(4) 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
	(5) FE IV	(6) FE IV	(7) CCE IV	(8) 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			[0.30,1.19]	[0.24,1.61]
First stage F-statistic	19.10	18.10	19.84	12.12
Underidentification	0.001	0.001	0.010	0.007

Note: Dependent variable is total consumption, the sum of household and government consumption. $N = 88$, $NT = 1099$. For other notes, see Table 1.

aid is instrumented, the estimated long-run effect is significant at the 5% (FE IV) or 10% (CCE IV) level. But to anticipate our later discussion, the investment results are less robust than the consumption results. Overall, our findings are in line with Boone (1996), who finds that aid translates mainly into consumption rather than investment and growth. Clemens et al. (2012) present stronger evidence that aid has modest effects on investment and growth, but note that their results are 'not incompatible' with Boone's suggestion that aid is often consumed rather than invested. The 2SLS estimates of Werker et al. (2009, Table 2) suggest that aid has stronger effects on consumption than investment. They find no evidence that aid affects growth when using four-year averages, which tallies with the lack of robustness of an investment effect in this paper.

Considering the magnitudes of the effects, the specification allows these to be interpreted easily. In Table 1, for example, a one percentage point increase in aid's share of GDP leads to a 0.772 percentage point increase in net imports as a share of GDP (in FE IV estimates) or a 1.085 increase (in CCE IV estimates). Although these two estimates are individually statistically significant, the difference between them may not be significant. To investigate this, we use a bootstrap-based test described in Section B.4 of the online appendix. For the results in Tables 1-7, we typically cannot reject the

Table 5: Aid and household consumption

	(1) FE	(2) FE	(3) CCE	(4) 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
	(5) FE IV	(6) FE IV	(7) CCE IV	(8) 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			[0.20,1.48]	[0.10,2.01]
First stage F-statistic	19.10	17.86	18.18	11.84
Underidentification	0.001	0.002	0.006	0.006

Note: Dependent variable is household consumption. $N = 88$, $NT = 1099$. For other notes, see Table 1.

null that the CCE IV point estimate is not significantly different from the FE IV estimate, given their high standard errors. In choosing between them, there is a trade-off between robustness and efficiency. The CCE IV estimator should be consistent under more general conditions than FE IV, and hence more robust, but will usually lead to higher standard errors. Readers may legitimately differ in which set of estimates they prefer.

We have not yet discussed the strength of our instrument. The tables report the first-stage F-statistic as a guide, indicating the significance of the single excluded instrument.²⁴ 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. In keeping with other papers, our application of the F-statistic in panel 2SLS is best seen as heuristic. In most cases, and especially in the static CCE IV regressions, the first-stage robust F-statistic is reasonably high, and the Kleibergen-Paap LM test always rejects the null of under-identification at the 5% level.

²⁴Note that the first-stage F-statistic differs across tables for the CCE IV regressions. This is because the CCE IV estimator includes the cross-sectional mean of the dependent variable with country-specific coefficients in both the first and second stage, so the first-stage models differ across tables in this case.

Table 6: Aid and government consumption

	(1) FE	(2) FE	(3) CCE	(4) 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
	(5) FE IV	(6) FE IV	(7) CCE IV	(8) 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			[-0.49,0.15]	[-0.49,0.22]
First stage F-statistic	19.10	22.01	14.99	13.89
Underidentification	0.001	0.001	0.005	0.004

Note: Dependent variable is government consumption. $N = 88, NT = 1099$. For other notes, see Table 1.

6 Robustness

We now consider several alternative models and estimators. These will tend to increase robustness – in the sense of reducing likely biases – 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. Rows 2-9 then correspond to the robustness tests listed in the notes to the table, and that we discuss below. In each row, we report the estimated effects of aid in static and dynamic models, and the first-stage F statistics.

In our main results, we followed much of the literature and excluded countries with small populations. If we include these countries, the instrument becomes weaker in the first stage of 2SLS: see row 2 of the table.²⁵ 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 uninformative about that recipient's long-term exposure to changes in that budget. Hence, we would expect our supply-push instrument to have

²⁵The 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.

Table 7: Aid and gross capital formation

	(1) FE	(2) FE	(3) CCE	(4) 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
	(5) FE IV	(6) FE IV	(7) CCE IV	(8) 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			[0.02,0.94]	[0.28,1.91]
First stage F-statistic	19.10	18.70	14.67	11.21
Underidentification	0.001	0.001	0.009	0.011

Note: Dependent variable is gross capital formation. $N = 88$, $NT = 1099$. For other notes, see Table 1.

less explanatory power for aid to smaller countries. But despite the weakening of the instrument, the results are qualitatively unchanged.

This is also true when we exclude aid observations for colonies prior to independence, as discussed in Section B.5 of the online appendix. Also in that appendix, we examine whether the estimated effects of aid are driven by distinct subgroups of countries, namely economies where natural resources play a major role, and countries with unusually strong or weak institutions. Estimates based on subsamples continue to suggest that aid increases consumption and net imports, while the evidence that aid affects investment varies across samples. As before, aid appears to increase imports, while there is no evidence of a negative effect on exports.

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	0.303 (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
4	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
5	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
6	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
7	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
8	Static	1.533** (0.717)	-0.223 (0.230)	1.389** (0.548)	0.201 (0.403)	-0.542 (0.691)	0.571 (0.484)	1.835*** (0.547)
	F	9.141	9.387	10.43	9.050	10.79	8.912	10.08
	Dynamic	1.566* (0.911)	-0.349 (0.288)	1.408** (0.673)	0.745 (0.479)	-0.0609 (1.069)	1.307* (0.792)	2.485*** (0.896)
	F	8.690	8.690	9.017	7.855	8.758	8.156	7.961
9	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 averages (1971-2012) using an instrument based on initial shares in donor budgets over 1960-70, unless noted below. Heteroskedasticity-robust standard errors, clustered by country, in brackets. Standard errors (SE) for long-run effects based on 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 ($N = 88$, $NT = 1099$).

Row 2 includes small countries ($N = 103$, $NT = 1248$).

Row 3 replaces aid and the instrument by their first lag; sample starts with 1974-76 ($N = 88$, $NT = 1039$).

Row 4 replaces aid and the instrument by averages of their current and one-period-lagged values; sample starts with 1974-76 ($N = 88$, $NT = 1032$).

Row 5 excludes the first period (1971-73) from estimation ($N = 88$, $NT = 1032$).

Row 6 uses an instrument based on initial shares over 1960-73; sample starts with 1974-76 ($N = 91$, $NT = 1061$).

Row 7 uses an instrument based on initial shares over 1960-73; sample starts with 1977-79 ($N = 91$, $NT = 993$).

Row 8 uses only donors whose Herfindahl-Hirschman index never exceeds 0.25 in any sample period ($N = 88$, $NT = 1099$).

Row 9 excludes outliers ($N = 81$, $NT = 1004$).

Responses to aid may take time to emerge, as Clemens et al. (2012) argue. In Table 8, we estimate models with lagged aid rather than current aid (row 3) and using six-year averages of aid and the instrument (row 4).²⁶ These point to stronger effects of aid on investment, but we interpret this cautiously. As we discuss later, our findings on investment are more sensitive to modelling choices than our findings on consumption and net imports. Section B.5 of the online appendix reports on some additional variations on these robustness tests.

We now turn to potential criticisms of our instrument. Serial correlation in country-specific conditions might undermine exogeneity (Card, 2001). Another concern is that the relevance 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 strategic or economic connections between countries evolve, the instrument may have less explanatory power for aid in later periods. We address these concerns as follows. Row 5 in Table 8 repeats the main analysis but excludes the first period. Row 6 uses an instrument based on initial shares calculated over the period 1960-73 and an estimation sample that starts with 1974-76. Row 7 excludes the first period from this sample. Section B.5 of the online appendix reports some further variations, including ones with the initial shares calculated over 1960-76.

We might expect instrument strength to weaken when early time periods are excluded, and this is what we find. The estimated second-stage coefficients are fairly stable, however, when dropping early time periods. In rows 5-7, we continue to find that aid is absorbed via higher consumption and higher imports, without much effect on exports. In experiments that drop up to three periods, the point estimates are quite stable, but the results for total consumption become a little less precise once we drop three periods; see Section B.5 of the online appendix. We also note that, if serially-correlated shocks were a major problem in our static models, we would have expected a greater contrast between static and dynamic models in Tables 1-7.

It could be argued that our instrument makes use of too many donors. By considering all DAC donors, we have included some whose budgets could be dominated by a few recipients, which risks endogeneity. To investigate this, row 8 in Table 8 shows results using an instrument based only on less specialized donors, whose HH index never exceeds 0.25.²⁷ The instrument remains informative and the results are in line with our main findings, but yield larger point estimates for the effects of aid on consumption and on

²⁶In rows 3 and 4, 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.

²⁷This leaves eighteen donors. In descending order by their share of world aid, averaged over 1960-2012, these are the USA, Japan, France, Germany, IDA, EU institutions, UK, Netherlands, Canada, Italy, WFP, Sweden, UNDP, Norway, Denmark, UNICEF, Switzerland, and UNTA.

net imports. We cannot reject full absorption achieved through higher consumption.²⁸

Next, we consider whether donors have been sorting themselves across recipient countries in a way that could undermine identification. We examine this using an approach developed by Greenstone et al. (2015). We implement two versions of their test and, in both cases, find little evidence that sorting has taken place. But there are some grounds for caution over the applicability of these tests in our setting, and Section B.3 of the online appendix examines these issues in greater detail.

We also investigate the possibility of outliers. Given that we use 2SLS, outliers could arise in 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 variable estimator of Cohen Freue et al. (2013), after partialling out fixed effects and cross-section means. As they discuss, robust parameter estimates can then be used to identify multivariate outliers. Across our dependent variables, seven countries regularly give rise to one or more outlying country-period observations: Burundi, the Central African Republic, Chad, the Democratic Republic of Congo, Jordan, Madagascar, and Mauritania. The results when we exclude them are shown in Row 9 of Table 8. In line with our main findings, the instrument retains explanatory power, and we cannot reject strong effects of aid on net imports, household consumption and total consumption. The effects on net imports are sufficiently strong that we cannot reject the null hypothesis of full absorption. In the dynamic model, the increase in net imports appears to arise from higher imports rather than lower exports. The estimated effect of aid on investment is too imprecise to draw conclusions, and this remains the case in models (not reported) that allow for delayed effects.

We now consider an alternative estimation method. Thus far, we have emphasized estimators that use a within transformation. For the static models, another way to eliminate country-specific effects is to first difference the model.²⁹ A comparison of the two approaches should be informative about the validity of our assumptions, and help to address 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 the models we estimate: in that case, the instrument would be correlated with the error term. But, given the persistence of GDP, first differencing would weaken the correlation between the instrument and the error term in differences. Hence, we investigate the

²⁸With a lower cutoff, 0.15 or 0.20, we get similar results to the 0.25 cutoff used in the table. With a higher cutoff of 0.30 we get similar results again, while a cutoff of 0.35 leads to long-run effects on consumption and net imports closer to our main findings.

²⁹Under our maintained assumptions, an estimator based on first differences should have the same probability limit as a within groups estimator, but is likely to be more efficient if the error term is highly persistent; see Wooldridge (2010, pp. 321-326). The result is likely to apply even for CCE estimators.

results from this approach, and compare them to our earlier findings.

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

$$\Delta(Q_{it}/Y_{it}) = \beta\Delta(A_{it}/Y_{it}) + \lambda'_i\Delta F_t + \Delta\varepsilon_{it} \quad (3)$$

and if we define a new set of factors $F_t^2 \equiv \Delta F_t$, 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 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. These estimates suggest that absorption may be less than complete but, 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 show the effects 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 measured relative to GDP. CCE IV estimation on three-year averages (1974-2012) using an instrument based on initial shares in donor budgets over 1960-70. Heteroskedasticity-robust standard errors, clustered by country, in brackets. *, **, and *** denote significance at 10, 5 and 1%, respectively. F shows the first-stage F-statistic.

Row 1 is a first-differences version of row 1 in Table 8 ($N = 88$, $NT = 1032$).

Row 2 removes outliers and is a first-differences version of row 9 in Table 8 ($N = 81$, $NT = 943$).

The first-differenced results suggest that our earlier findings are not spurious. Nevertheless, a sceptical observer might still be worried about a 'false positive', because the variables in our regressions take the form of ratios to nominal GDP. In principle, this could lead to an observed relationship between two ratios even when their numerators

are unrelated.³⁰ One reason to doubt this interpretation is the pattern of the results. We consistently find that the ratios of consumption and net imports to GDP are increasing in aid intensity, in line with theoretical predictions, while the effects of aid on the other ratios are weaker, again in line with theory. A convincing story based on ‘false positives’ would have to explain why spurious correlations emerge for some ratios and not others, and why the magnitudes of the estimated effects remain plausible. But to investigate this further, Section B.5 of the online appendix reports on some additional analysis, including models with the dependent variable in logarithms and the logarithm of nominal GDP included as an explanatory variable. The results on absorption and its channels are in line with our main findings.

In summary, we have carried out a range of robustness tests. These include demanding versions of CCE and CCE IV estimators with many additional parameters, and go beyond those typically used in the literature. It is not surprising that the effects sometimes become imprecise, but the point estimates are quite stable, and rarely change sign. Across a range of models and approaches, we continue to find that aid is absorbed in line with theoretical predictions, and with effect sizes that have plausible magnitudes. More precise estimates are likely to require longer spans of data or the use of additional instruments, such as those of Galiani et al. (2017) or Jarotschkin and Kraay (2016). In the meantime, we note that our findings are consistent with Werker et al. (2009), while adopting a new approach to identification and filtering out common factors. Moreover, in contrast to previous studies based on distinct natural experiments, the results we present relate to a broad concept of aid, for the full set of DAC donors.

7 Conclusions

Using cross-country data to study aid 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 instrument that can be used to identify persistent changes in 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 is an advance on much of the existing literature. Second, we use the instrument to investigate the effects of aid on domestic absorption, consumption and investment, and whether aid is associated with Dutch Disease.

The evidence suggests that aid is absorbed at least partially, and 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, and we find no evidence that aid lowers

³⁰See Kronmal (1993). In the aid literature, the problem was noted by Arndt et al. (2010).

exports through Dutch Disease effects. We also investigate the relationship between aid and the separate components of domestic expenditure. Our estimates suggest that aid is absorbed primarily by increased consumption rather than investment. Although we do sometimes find a significant effect of aid on investment, it seems less robust than the effect on consumption.

Overall, 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. The instrument is strong enough to generate some informative findings, and as future years of data become available, the prospects for robust and precise estimates should improve further. The instrument could have many possible applications in the future study of aid.

Appendix A The CCE IV estimator

Here, we outline how the Common Correlated Effects (CCE) approach of Pesaran (2006) can be applied in a 2SLS setting, based on Harding and Lamarche (2011). Using their notation, we consider a model for $i = 1, \dots, N$ and $t = 1, \dots, T$:

$$\begin{aligned} y_{it} &= \beta x_{it} + \varepsilon_{it} \\ \varepsilon_{it} &= \lambda_i' F_t + u_{it} \\ x_{it} &= \pi w_{it} + \Lambda_i' F_t + v_{it} \end{aligned}$$

The second equation represents the multifactor error structure, where λ_i is an $r \times 1$ vector of factor loadings, and F_t are the r common factors, where λ_i and F_t are unobserved. We assume a single endogenous variable x_{it} and a single instrument w_{it} which satisfies the usual requirements for identification. For simplicity, we do not cover deterministic time components (such as time trends) or factor loadings that are correlated with the regressors; for a more general treatment, see Harding and Lamarche (2011).

We gather y and x in the vector $z_{it} = (y_{it}, x_{it})'$. The above expressions can be combined to obtain:

$$z_{it} = C_1 w_{it} + C_{3i}' F_t + \xi_{it}$$

where $\xi_{it} = (\beta v_{it} + u_{it}, v_{it})'$ and $C_1 = (\pi\beta, \pi)'$, $C_{3i} = ((\Lambda_i\beta + \lambda_i)', \Lambda_i)'$, where our notation again follows Harding and Lamarche (2011).

The next step is to take cross-section means, to obtain:

$$\bar{z}_t = C_1 \bar{w}_t + \bar{C}_3' F_t + \bar{\xi}_t$$

where bars over variables indicate cross-section means. Pre-multiplying by \bar{C}_3 and rearranging for F_t , we have:

$$F_t = (\bar{C}_3 \bar{C}_3')^{-1} \bar{C}_3 (\bar{z}_t - C_1 \bar{w}_t - \bar{\xi}_t)$$

We first consider the case of OLS estimation of the equation for y_{it} when x_{it} is strictly exogenous and there is no instrument w_{it} . If we substitute the above expression for F_t into the model $y_{it} = \beta x_{it} + \lambda_i' F_t + u_{it}$, this suggests that we can proxy for the unobserved F_t by including cross-section means of y_{it} and x_{it} . These cross-section means will require country-specific coefficients, to allow the factor loadings λ_i to differ across countries. This approach is due to Pesaran (2006), and the estimator we have just described corresponds to his pooled CCE (CCEP) estimator. He showed that, under regularity conditions, $\bar{\xi}_t \rightarrow 0$ and $\bar{C}_3 \rightarrow C_3$ (constant) as the cross-section dimension tends to infinity. The CCEP estimator of β is then asymptotically (large N) unbiased and normally distributed as (N, T) jointly tend to infinity.

One limitation is that, for the estimator of β to be asymptotically unbiased, the number r of unobserved factors F_t can be no greater than the number of observable variables — typically three or four in the estimates of this paper. Intuitively, this is because the unobserved factors F_t are only eliminated completely if they lie in the space spanned by the cross-section means; see Sarafidis and Wansbeek (2012, p. 496). Stating this more formally, the properties of CCEP summarized above rely on a rank condition on \bar{C}_3 . But the CCEP estimator may improve on more conventional estimators, and lead to bias reduction, even when the rank condition is not met; see Pesaran (2006) for discussion and simulations.

Harding and Lamarche (2011) extend the CCE approach to 2SLS panel estimation, when x_{it} is endogenous and a suitable instrument w_{it} is available. Again using the above expression for F_t , the 2SLS analogue to the CCEP estimator introduces cross-section means into the first and second stage of 2SLS, with country-specific coefficients in each case. In simulations, Harding and Lamarche show that this estimator outperforms conventional estimators in cases with unobserved factors and endogeneity (Design III in their Table 1). It is easy to implement: we augment an otherwise conventional FE IV model with the cross-section means of the observable variables (y, x, w) in the first and second stage, interacted with country dummies in each case.

As noted in the main text, the interactive fixed effects specification nests conventional fixed effects (corresponding to a constant factor $F_t = 1$) and time dummies (corresponding to $\lambda_i = \lambda$ for one of the factors F_t).³¹ As Harding and Lamarche (2014,

³¹A technical aside is that, when augmenting models with cross-section means in a CCE approach, time dummies should be excluded; see Sarafidis and Wansbeek (2012, p. 502).

fn. 3) note, a CCE approach can be combined with fixed effects — in our case, country dummies — which in our notation corresponds to introducing a factor $F_t = 1$. In simulations for their quantile regression setting, an estimator which uses conventional fixed effects as well as cross-section means has the best small-sample performance. In this paper, we also include fixed effects, and use CCE IV to denote this modified version of the Harding and Lamarche (2011) estimator. But our main findings are not sensitive to this choice, as discussed in Section B.5 of the online appendix.

References

- Adam, C. S. (2013). Dutch Disease and foreign aid. In *The New Palgrave Dictionary of Economics*, Online Edition. Palgrave Macmillan.
- Adam, C. S. and O'Connell, S. A. (1999). Aid, taxation and development in Sub-Saharan Africa. *Economics and Politics*, 11(3), 225-253.
- Aiyar, S. and Ruthbah, U. H. (2008). Where did all the aid go? An empirical analysis of absorption and spending. IMF Working Paper no. 08/34. Washington DC: IMF.
- Annen, K. and Moers, L. (2017). Donor competition for aid impact, and aid fragmentation. *World Bank Economic Review*, forthcoming.
- Arndt, C., Jones, S. and Tarp, F. (2010). Aid, growth and development: have we come full circle? *Journal of Globalization and Development*, 1(2), Article 5.
- Azevedo, J. P. (2011). wbopendata: Stata module to access World Bank databases. <http://ideas.repec.org/c/boc/bocode/s457234.html>
- Baltagi, B. H. and Li, J. (2014). Further evidence on the spatio-temporal model of house prices in the United States. *Journal of Applied Econometrics*, 29(3), 515-522.
- Bartik, T. (1991). *Who benefits from state and local economic development policies?* W.E. Upjohn Institute for Employment Research: Kalamazoo, MI.
- Berg, A., Mirzoev, T., Portillo, R. and Zanna, L.-F. (2010). The short-run macroeconomics of aid inflows: understanding the interaction of fiscal and reserve policy. IMF Working Paper no. 10/65. Washington DC: IMF.
- Bond, S., Leblebicioglu, A. and Schiantarelli, F. (2010). Capital accumulation and growth: a new look at the empirical evidence. *Journal of Applied Econometrics*, 25(7), 1073-1099.

- Boone, P. (1996). Politics and the effectiveness of foreign aid. *European Economic Review*, 40(2), 289-329.
- Breinlich, H., Ottaviano, G. I. P. and Temple, J. R. W. (2014). Regional growth and regional decline. In P. Aghion and S. Durlauf (eds.), *Handbook of Economic Growth, Volume 2*. The Netherlands: North-Holland.
- Card, D. (2001). Immigrant inflows, native outflows, and the local market impacts of higher immigration. *Journal of Labor Economics*, 19(1), 22-64.
- Carter, P. (2015). Aid econometrics: lessons from a stochastic growth model. University of Bristol Discussion Paper no. 15/659, May.
- Carter, P., Postel-Vinay, F. and Temple, J. R. W. (2015). Dynamic aid allocation. *Journal of International Economics*, 95(2), 291-304.
- 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, A. and Pesaran, M. H. (2015). Common correlated effects estimation of heterogeneous dynamic panel data models with weakly exogenous regressors. *Journal of Econometrics*, 188, 393-420.
- Clemens, M. A., Radelet, S., Bhavnani, R. R. and Bazzi, S. (2012). Counting chickens when they hatch: timing and the effects of aid on growth. *Economic Journal*, 122(561), 590-617.
- Cohen Freue, G. V., Ortiz-Molina, H. and Zamar, R. 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.
- Development Initiatives (2014). *The Global Humanitarian Assistance Report 2014*.
- Djankov, S., Montalvo, J. G. and Reynal-Querol, M. (2009). Aid with multiple personalities. *Journal of Comparative Economics*, 37, 217-229.
- Easterly, W. (2007). Are aid agencies improving? *Economic Policy*, October, 633-678.
- Eberhardt, M., Helmers, C. and Strauss, H. (2013). Do spillovers matter when estimating private returns to R&D? *Review of Economics and Statistics*, 95(2), 436-448.

- Everaert, G. and De Groote, T. (2016). Common correlated effects estimation of dynamic panels with cross-sectional dependence. *Econometric Reviews*, 35(3), 428-463.
- Galiani, S., Knack, S., Xu, L. C., and Zou, B. (2017). The effect of aid on growth: evidence from a quasi-experiment. *Journal of Economic Growth*, 22(1), 1-33.
- Greenstone, M., Mas, A., and Nguyen, H.-L. (2015). Do credit market shocks affect the real economy? Quasi-experimental evidence from the Great Recession and 'Normal' economic times. Working paper, University of Chicago, September.
- Hansen, H. and Headey, D. (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, M. and Lamarche, C. (2011). Least squares estimation of a panel data model with multifactor error structure and endogenous covariates. *Economics Letters*, 111(3), 197-199.
- Harding, M. and Lamarche, C. (2014). Estimating and testing a quantile regression model with interactive effects. *Journal of Econometrics*, 178, 101-113.
- Hodler, R. and Raschky, P. A. (2014). Regional favoritism. *Quarterly Journal of Economics*, 129(2), 995-1033.
- Holly, S., Pesaran, M. H. and Yamagata, T. (2010). A spatio-temporal model of house prices in the USA. *Journal of Econometrics*, 158(1), 160-173.
- Hsiao, C. (2014). *Analysis of Panel Data* (third edition). Cambridge University Press, Cambridge.
- Hussain, M., Berg, A. and Aiyar, S. (2009). The macroeconomic management of increased aid: policy lessons from recent experience. *Review of Development Economics*, 13(S1), 491-509.
- Imbs, J., Jondeau, E. and Pelgrin, F. (2011). Sectoral Phillips curves and the aggregate Phillips curve. *Journal of Monetary Economics*, 58(4), 328-344.
- Jarotschkin, A. and Kraay, A. (2016). Aid, disbursement delays, and the real exchange rate. *IMF Economic Review*, 64(2), 217-238.
- Kimbrough, K. P. (1986). Foreign aid and optimal fiscal policy. *Canadian Journal of Economics*, 19(1), 35-61.

- Kleibergen, F. and Paap, R. (2006). Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics*, 133(1), 97-126.
- Knack, S. and Rahman, A. (2007). Donor fragmentation and bureaucratic quality in aid recipients. *Journal of Development Economics*, 83, 176-197.
- Kronmal, R. A. (1993). Spurious correlation and the fallacy of the ratio standard revisited. *Journal of the Royal Statistical Society: Series A*, 156(3), 379-392.
- Nickell, S. J. (1981). Biases in dynamic models with fixed effects. *Econometrica*, 49(6), 1417-1426.
- Nsor-Ambala, A. A. (2015). Foreign transfers, manufacturing growth and the Dutch Disease revisited. University of Bristol, Discussion Paper no. 15/663.
- Nunn, N. and Qian, N. (2014). U.S. food aid and civil conflict. *American Economic Review*, 104(6), 1630-1666.
- Obstfeld, M. (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.
- Pesaran, M. H. (2006). Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica*, 74(4), 967-1012.
- Pesaran, M. H. (2015). *Time Series and Panel Data Econometrics*. Oxford University Press, Oxford.
- Qian, N. (2015). Making progress on foreign aid. *Annual Review of Economics*, 7, 277-308.
- Rajan, R. G. and Subramanian, A. (2011). Aid, Dutch Disease, and manufacturing growth. *Journal of Development Economics*, 94, 106-118.
- Rodrik, D. (2006). The social cost of foreign exchange reserves. *International Economic Journal*, 20(3), 253-266.
- Roodman, D. (2007a). The anarchy of numbers: aid, development, and cross-country empirics. *World Bank Economic Review*, 21(2), 255-277.
- Roodman, D. (2007b). Macro aid effectiveness research: a guide for the perplexed. CGD Working Paper 135. Washington DC: Center For Global Development.

- Sarafidis, V. and Wansbeek, T. (2012). Cross-sectional dependence in panel data analysis. *Econometric Reviews*, 31(5), 483-531.
- Schaffer, M. E. (2010). xtvreg2: 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>.
- Söderbom, M., Teal, F., Eberhardt, M., Quinn, S. and Zeitlin, A. (2015). *Empirical Development Economics*. Routledge, London.
- Stock, J. H. and Yogo, M. (2005). Testing for weak instruments in linear IV regression. In D. W. K. Andrews and J. H. Stock (eds.), *Identification and inference for econometric models: essays in honor of Thomas Rothenberg*. Cambridge, UK: Cambridge University Press.
- Tavares, J. (2003). Does foreign aid corrupt? *Economics Letters*, 79(1), 99-106.
- Temple, J. R. W. (2010). Aid and conditionality. In Dani Rodrik and Mark Rosenzweig (eds.), *Handbook of Development Economics, Volume 5*, 4415-4523.
- Van de Sijpe, N. (2010). Foreign aid and government behaviour. DPhil Thesis, Department of Economics, University of Oxford.
- Van de Sijpe, N. (2013). Is foreign aid fungible? Evidence from the education and health sectors. *World Bank Economic Review*, 27(2), 320-356.
- Werker, E., Ahmed, F. Z. and Cohen, C. (2009). How is foreign aid spent? Evidence from a natural experiment. *American Economic Journal: Macroeconomics*, 1(2), 225-244.
- Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data* (second edition). MIT Press, Cambridge, MA.

Appendix B Online appendix

B.1 Aid and absorption in theory

We briefly consider the aggregate effects of aid in two-sector models. When aid increases demand for traded and non-traded goods, it might be anticipated that the price of non-traded goods would increase. But in standard versions of the dependent economy model with traded and non-traded goods, the relative price of non-traded goods is invariant to aid in the long run. In 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. For more on dynamic dependent economy models, see chapter 4 of Turnovsky (1997).

Brock and Turnovsky (1994) and Brock (1996) showed that long-run adjustment depends on the relative capital intensities of the traded and non-traded sectors. In the long-run equilibrium, imports of traded goods can be financed partly by aid as well as exports, and hence aid is associated with a smaller traded sector in the long run. Since the long-run relative price is invariant to the transfer, aid can increase or decrease the steady-state capital stock, depending on the relative capital intensities of non-traded and traded production. This implies that aid could be associated with higher or lower gross investment in the long run.

B.2 Further notes on the instrument

In this appendix, we first sketch 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 are likely to be a function of the transient errors at earlier dates, and this could reinforce the probable failure of exogeneity when using average shares.

The use of budget shares, whether initial or average, might raise another concern. The supply-push instrument itself has a factor structure: it is a weighted average of donor

budgets, with sets of weights (initial budget shares) that vary across aid recipients. It might be thought that this will lead the instrument to 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: either when there is a single donor, or when the initial shares of aid recipients in donor budgets are the same across donors.³² 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 with country-specific coefficients. Moreover, this is a testable assumption.

But a remaining concern with our initial-share instrument is that conditions 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. We investigate the problem by dropping some of the early time periods from the estimated models. 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 conditions. The results are reported in Table 8 in the main text and, especially, in Table B.1 below.

B.3 Sorting and identification

This subsection explains the assumptions we use for identification in more detail, and summarizes the results of a test of sorting, based on the Greenstone et al. (2015) working paper. The test makes assumptions that may not be reasonable in our setting but, when implemented, does not provide any evidence against our approach to identification.

We begin by recalling that our instrument for recipient country i at time t can be written as:

$$\sum_d a_{i0}^d (D_{dt}/Y_{it})$$

where d indexes donors, a_{i0}^d is the initial share of recipient i in donor d 's budget, D_{dt} is the donor's nominal budget and Y_{it} is the nominal GDP of the recipient. From the construction of the instrument, it is clear that the instrument will not be influenced by

³²Under the stated conditions, the instrument would be perfectly collinear within each country with the cross-section mean of the instrument, and then identification fails.

donors reallocating aid within their total budget in response to time-varying recipient conditions. Indeed, this is one way to understand what the instrument achieves.

But a remaining threat to identification is that total donor budgets may be functions of the time-varying conditions of individual aid recipients. Informally, this may not be a persuasive account of how total donor aid budgets are determined. In most cases, they seem likely to evolve through a longer-term political process, such as the UK's commitment to give 0.7% of GDP. Moreover, if total donor budgets respond strongly to the conditions of multiple recipients, this would tend to suggest that total aid budgets will be highly volatile — and it is not clear this is true in the data. Nevertheless, we can study when this source of endogeneity might be a major problem for our instrument.

Our analysis runs parallel with the study of bank lending by Greenstone et al. (2015). One difference is that our results use multiple periods of a panel, whereas Greenstone et al. work with first-differenced cross-section regressions, which are then estimated separately for different time periods. Nevertheless, we can draw on their work to clarify the assumptions we rely on in our own setting.

We denote time-varying recipient-specific conditions by c_{it} . A natural example in our case would be the political conditions within aid recipient i at time t . We can think of the recipient-specific conditions as the counterpart to the county-level demand shock in Greenstone et al. In the following discussion, we assume that c_{it} cannot be observed by the econometrician but influences aid flows and also the dependent variable. It will then be a component of the transitory error term in the models that we estimate, and the source of the overall endogeneity problem we are seeking to address.

We now write the total donor budget for donor d as a function of recipient conditions, and a supply-side variable S_{dt} which captures exogenous changes in the size of a donor's total aid budget. Since aid budgets typically increase exponentially over time, whether in nominal or real terms, we use the following formulation:

$$D_{dt} = \exp \left(\pi_{id} c_{it} + \sum_{k \neq i} \pi_{kd} c_{kt} + S_{dt} \right)$$

where the new π_{id} parameters capture the responsiveness of the total aid budget for donor d to country i conditions. Note that the π_{id} will not be the same as the overall responsiveness of the donor's aid to recipient i , because thus far, we are modelling only the response of the total donor budget, and not reallocations within the budget.

We can now substitute this expression into that for the instrument for country i at time t :

$$\sum_d a_{i0}^d \exp(\pi_{id} c_{it}) \exp \left(\sum_{k \neq i} \pi_{kd} c_{kt} \right) \exp(S_{dt}) (1/Y_{it})$$

This expression highlights three forces that influence the extent of a correlation between the instrument for a given recipient, and that recipient's conditions. These are the three terms in exponentials, written separately for clarity.

The first term is the direct effect of recipient conditions, working through the total budget — it is this direct effect which risks an endogeneity problem. The second term reflects the dependence of the total budget on the conditions of other countries, which will typically weaken the association between the instrument for this recipient and the recipient's own conditions. This is especially likely when donors spread their aid relatively evenly over many recipients. The last exponential term, in S_{dt} , reflects the supply-side measure of the donor's generosity, determined independently of recipient conditions.

Note also that these exponentials are summed across donors. Hence, if the π_{id} parameters vary across donors, the summation across donors should also weaken the association between recipient conditions and the final instrument. Problems could arise if the π_{id} parameters are simultaneously high for several large donors (those with large values of S_{dt}). Large donors matter more, because they receive more weight in the sum over donors, through the S_{dt} term.

But even then, note the role of the second exponential. The sensitivity of donor budgets to the conditions of *other* recipients, in the context of many recipients, should ensure that the final instrument for aid to country i does not vary strongly with the conditions of country i in most cases. There may be a few exceptions — for example, when donors give a large share of their budget to a recipient, and that recipient receives most of their aid from that donor — but in general it seems likely that the instrument will be uncorrelated with recipient conditions for most recipients.

More formally, our approach assumes that donors spread their aid relatively evenly over many recipient countries, which is likely to weaken the association between donor budgets and the conditions of a given recipient. In terms of our expression above, if aid is spread across many recipients relatively evenly, this will reduce the contribution of the first exponential relative to the second. This indicates why it is relevant to consider within-donor fragmentation, using Herfindahl-Hirschman indices. If donors give to multiple countries, in a way that does not concentrate their aid on a small number of countries, it becomes less plausible to imagine that total aid budgets are responding strongly to the conditions of a given recipient.

There could still be a concern that the π_{id} parameters, capturing the responsiveness of total aid budgets, could be simultaneously high for larger donors, risking a correlation between the instrument and recipient conditions for at least a subset of recipients. But, as noted before, strong budgetary responses of that type, across many recipients, would be likely to generate greater overall volatility of total aid budgets than is observed

empirically, given that donor budgets are often fairly stable as shares of donor GDP.

Our supply-push instrument has parallels in the analysis of bank lending by Greenstone et al. (2015). In developing their own supply-push instrument, they allow lending outcomes at the bank-county level to depend on a county effect (demand) and a bank effect (supply), assuming that demand shocks do not have bank-specific effects; this is equation (5) in their paper. They use this partly to study whether individual banks may be sorting into activity at particular locations, which could threaten identification.

We implement this test below, but note that an assumption used in Greenstone et al. (2015) is unlikely to be reasonable in our setting. To examine this further, we describe a model for aid flows at the recipient-donor-time (i, d, t) level. A natural model in the aid case, with links to the aid allocation literature, is:

$$\begin{aligned} aid_{idt} &= \left[\frac{\exp(\theta_{id}c_{it})}{\sum_k \exp(\theta_{kd}c_{kt})} \right] D_{dt} \exp(u_{idt}) \\ &= \left[\frac{\exp(\theta_{id}c_{it})}{\sum_k \exp(\theta_{kd}c_{kt})} \right] \exp\left(\pi_{id}c_{it} + \sum_{k \neq i} \pi_{kd}c_{kt} + S_{dt}\right) \exp(u_{idt}) \end{aligned}$$

where the new θ_{id} parameters reflect aid allocation, and index the responsiveness of a donor to recipient conditions through reallocations within the donor's overall budget. For simplicity, we interpret u_{idt} as measurement error.

As this expression makes clear, without further assumptions, there is no clean way in which to use donor-time fixed effects to extract estimates of the supply-push parameters S_{dt} needed for the Greenstone et al. (2015) test. This is because several of the other terms in this expression will also contribute to variation in the dimensions (d, t) and so the various effects are likely to be confounded in a regression making use of donor-time fixed effects. Unless we observe the θ_{id} and π_{id} parameters directly, or assume they are the same across donors to eliminate their dependence on the index d , we cannot separate the supply-push shocks S_{dt} cleanly from the responses of donors to recipient conditions.

This point runs parallel to Greenstone et al. (2015), who explicitly assume that different banks respond in exactly the same way to a given demand shock, as they note on page 12 of the September 2015 version. The restriction could be plausible in their case, but the corresponding assumption seems implausible in our setting. This is partly because of the role of connections between specific donor-recipient pairs, and partly because the large empirical literature on aid allocation has documented that donors differ in their responses to observable variables.

To make the issues clearer, drop the t subscript from the previous equation, and rewrite it in log-differences to get:

$$\Delta \log aid_{id} = (\theta_{id} + \pi_{id})\Delta c_i + \sum_{k \neq i} \pi_{kd}\Delta c_k - \Delta \log \left(\sum_k \exp[\theta_{kd}c_k] \right) + \Delta S_d + \Delta u_{id}$$

If we assume $\theta_{id} = \theta_i$ and $\pi_{id} = \pi_i$, then the above can be written as the sum of a recipient effect and a donor effect:

$$\Delta \log aid_{id} = c'_i + S'_d + \Delta u_{id}$$

which is a decomposition parallel to Greenstone et al. (2015). We now describe what happens when we implement this decomposition, estimate the donor fixed effects S_d , and run sorting tests similar to those of Greenstone et al. (2015). In brief, we find no evidence of sorting.

In more detail, to run the sorting test, we estimate an equation for the first-difference of log aid at the recipient-donor level for each time period, with recipient and donor fixed effects. We extract the estimated donor effects. Across recipients, we then run a regression of the fixed effect of a recipient's most important donor on the fixed effect of the recipient's second most important donor. An alternative version of the test, also taken from Greenstone et al. (2015), uses a share-weighted average of the fixed effects of donors other than the largest donor.

If there is sorting of donors across recipients, these regressions should reveal a significant positive correlation for at least some time periods. In practice, for the fourteen time periods we study, the correlation is positive and significant at the 10% level in only one case out of fourteen, a result that could easily arise by chance. We obtain a similar result using the alternative version of the test, based on the share-weighted average: again, there is one correlation which is positive and (close to) significant at the 10% level, but the rest are not. In summary, we find little evidence of sorting that could threaten identification.

Overall, we consider that the effect of recipient conditions on total donor budgets will not be a major threat to identification given: (1) the relatively plausible scenario that total aid budgets are determined by a medium-run political process which does not respond strongly to the conditions of many individual recipients (see, for example, the UK's formal commitment to give 0.7% of GDP); and (2) we also present evidence, based on Herfindahl-Hirschman indices, that aid budgets are dispersed widely (within-donor fragmentation). Moreover, many recipients receive aid from many sources (within-recipient fragmentation). This suggests that the instrument and recipient conditions will not be strongly correlated for most aid recipients, once the various summations

have been worked through in the construction of the instrument for each recipient.

Finally, we note that the arguments above have made no particular assumptions about the process driving recipient conditions. The aid literature has sometimes used instruments based on interacting a single driving variable with a country characteristic, an approach that might be sensitive to assumptions about recipient conditions: for example, those conditions might be driven by country-level trends that are correlated with the driving variable. Two aspects of our approach help us to avoid that problem. First, since we are averaging across the total budgets of many donors, we use the fragmentation of global aid flows to lessen the risk that the instrument is correlated with the conditions of individual aid recipients. Second, we use the CCE approach to address a remaining concern, that total donor budgets could be influenced by common factors (such as world economic conditions) that are correlated with conditions in aid recipients. Our combination of a supply-push instrument and CCE IV estimation is new, and should ensure identification in a wider range of circumstances than has been possible to date.

B.4 Testing equality of estimates

One question of interest is whether the findings from CCE IV differ from more conventional approaches, such as FE IV. The discussion in the main text emphasizes the economic magnitude of the differences in point estimates. In this section, we describe a test of whether the differences in parameter estimates across estimators are *statistically* significant, using the distribution of their difference. Our model for the macroeconomic outcome of interest Q_{it}/Y_{it} is given by:

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

$$\varepsilon_{it} = \lambda_i' F_t + u_{it} \quad (\text{B.2})$$

We want to test whether allowing for common factors has a material effect on the parameter estimate β , the effect of aid. This is a non-standard testing problem, but we can test whether β changes to a significant extent using our bootstrap estimates. We call one estimate $\hat{\beta}_1$ (say, FE IV) and the other $\hat{\beta}_2$ (say, CCE IV). Since our bootstrap yields a simulated joint distribution for $\hat{\beta}_1$ and $\hat{\beta}_2$, we can construct a test statistic:

$$\frac{(\hat{\beta}_1 - \hat{\beta}_2)^2}{\text{Var}(\hat{\beta}_1) + \text{Var}(\hat{\beta}_2) - 2\text{Cov}(\hat{\beta}_1, \hat{\beta}_2)}$$

where the variances and the covariance are taken from our bootstrapped estimates. For example, to compare FE IV with CCE IV, we estimate the coefficient under FE IV and

the coefficient under CCE IV for each of the bootstrap samples. We can then calculate the test statistic above. If the parameter estimates are asymptotically normal under the null, the test statistic will have a χ^2 distribution with one degree of freedom.

Compared to a standard Hausman test, this approach does not require one estimator to be asymptotically efficient. This is because we have simulated the joint distribution of the parameter estimates rather than (as in the Hausman test) using a theoretical result about their covariance. The formal interpretation of the test is heuristic, in that we proceed as if both $\hat{\beta}_1$ and $\hat{\beta}_2$ correspond to well-defined parameters in a population, which is open to question in this setting. Given this qualification, the results should be seen as only indicative. Since the testing problem is non-standard, there is no better approach readily available.

B.5 Additional robustness checks

In this section we discuss a wide range of additional robustness checks. These are reported in Table B.1 and, similarly to Table 8, the first row presents our main results for ease of comparison. In row 2, we report results which omit conventional fixed effects. In this case, time-invariant country effects will be captured by a linear combination of the cross-section means, given that they have country-specific coefficients. This approach, since it reduces the number of parameters, provides a useful check that our findings are not driven by over-fitting. As can be seen by comparing rows 1 and 2, our findings are not sensitive to this change. This exercise is informative partly because it helps to confirm that our main results do not somehow reflect an extreme over-fitting problem that has not been communicated by the standard errors. If that was the case, the results should have changed markedly in moving from row 1 to row 2. In fact, the two sets of results are very close to each other.

The other robustness tests we present generally move in the other direction — to greater robustness, at the expense of reduced efficiency. We first 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) after discarding aid data in the years before a recipient's independence.³³

³³The 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

The results, shown in row 3 of Table B.1, are again similar to those found before.

Table B.1: Additional 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	0.808* (0.473)	-0.0552 (0.202)	0.718** (0.298)	0.313 (0.246)	-0.299 (0.339)	0.608** (0.300)	1.016*** (0.247)
	F	19.80	18.36	19.32	17.45	21.86	20.18	16.15
	Dyn.	0.800 (0.679)	-0.0216 (0.247)	0.724* (0.425)	0.678* (0.355)	0.494 (0.509)	1.396*** (0.503)	1.582*** (0.566)
	F	16.25	18.35	15.06	13.87	18.21	18.09	12.70
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	0.719 (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
5	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
6	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

Continued on next page

World Factbook (<https://www.cia.gov/library/publications/the-world-factbook/>).

Table B.1 – continued from previous page

Row	Model	C	G	C+G	I	X	M	M-X
7	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
8	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
9	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
10	Static	0.330 (0.265)	0.166 (0.143)	0.552** (0.255)	0.133 (0.250)	-0.0150 (0.240)	0.690*** (0.204)	0.857*** (0.199)
	F	21.91	15.85	21.12	14.48	19.29	20.43	16.76
	Dyn.	0.784* (0.442)	0.166 (0.173)	0.935** (0.398)	0.138 (0.330)	0.351 (0.483)	1.304*** (0.316)	1.241*** (0.415)
	F	15.12	14.25	13.81	13.26	16.78	12.87	13.37
11	Static	0.291 (0.322)	0.201* (0.120)	0.482 (0.335)	0.171 (0.270)	0.125 (0.247)	0.723*** (0.233)	0.875*** (0.286)
	F	16.33	11.03	13.70	11.07	15.04	14.79	12.02
	Dyn.	0.746 (0.578)	0.154 (0.128)	0.695 (0.492)	0.318 (0.353)	0.595 (0.536)	1.420*** (0.432)	1.220** (0.558)
	F	8.965	9.430	8.251	9.307	10.98	9.911	9.787
12	Static	0.755 (0.502)	0.102 (0.0809)	0.765 (0.503)	-0.210 (0.422)	-0.0218 (0.317)	0.455 (0.378)	0.968** (0.438)
	F	10.79	6.742	8.237	5.529	9.698	7.153	7.973
	Dyn.	0.714 (0.548)	0.134 (0.110)	0.855 (0.637)	-0.119 (0.560)	0.154 (0.700)	0.717 (0.612)	1.255* (0.759)
	F	7.236	5.963	5.857	3.869	5.229	4.365	6.690

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 averages (1971-

2012) using an instrument based on initial shares in donor budgets over 1960-70, unless noted below. Heteroskedasticity-robust standard errors, clustered by country, in brackets. Standard errors (SE) for long-run effects based on 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 ($N = 88$, $NT = 1099$).

Row 2 excludes fixed effects from estimation ($N = 88$, $NT = 1099$).

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 ($N = 81$, $NT = 1033$).

Row 4 includes both aid and its first lag, instrumented by the current and one-period-lagged values of the instrument; sample starts with 1974-76 ($N = 80$, $NT = 978$). F is the Kleibergen-Paap Wald rk F statistic.

Row 5 uses the final year values in each period for the dependent variables instead of the three-year averages ($N = 86$, $NT = 1065$).

Row 6 excludes the first two periods (1971-73 and 1974-76) from estimation ($N = 88$, $NT = 965$).

Row 7 excludes the first three periods (1971-73, 1974-76 and 1977-79) from estimation ($N = 88$, $NT = 892$).

Row 8 uses an instrument based on initial shares over 1960-73; sample starts with 1980-82 ($N = 91$, $NT = 919$).

Row 9 uses an instrument based on initial shares over 1960-73; sample starts with 1983-85 ($N = 91$, $NT = 846$).

Row 10 uses an instrument based on initial shares over 1960-76; sample starts with 1977-79 ($N = 92$, $NT = 1003$).

Row 11 uses an instrument based on initial shares over 1960-76; sample starts with 1980-82 ($N = 92$, $NT = 929$).

Row 12 uses an instrument based on initial shares over 1960-76; sample starts with 1983-85 ($N = 92$, $NT = 856$).

The remaining entries in the table are variations on the treatment of delayed effects (rows 4 and 5), the exclusion of initial periods (rows 6 and 7), and the use of longer time periods over which to measure the initial budget shares (rows 8-12). The results for net imports are remarkably consistent across the various experiments, and the null hypothesis of full absorption is never rejected. The results for total consumption are a little more uneven: the point estimates are reasonably stable, but the standard errors sometimes increase, so that the results are sometimes significant at conventional levels and sometimes not. This is perhaps not surprising, given that some of these experiments are asking a lot of the data when the sample size is reduced. As before, the increase in net imports appears to come about by increased imports rather than a reduction in exports.

We now summarize some other experiments. First, we examine whether the estimated effects of aid are driven by distinct subgroups of countries, namely economies where natural resources play a major role, and countries with unusually strong or weak

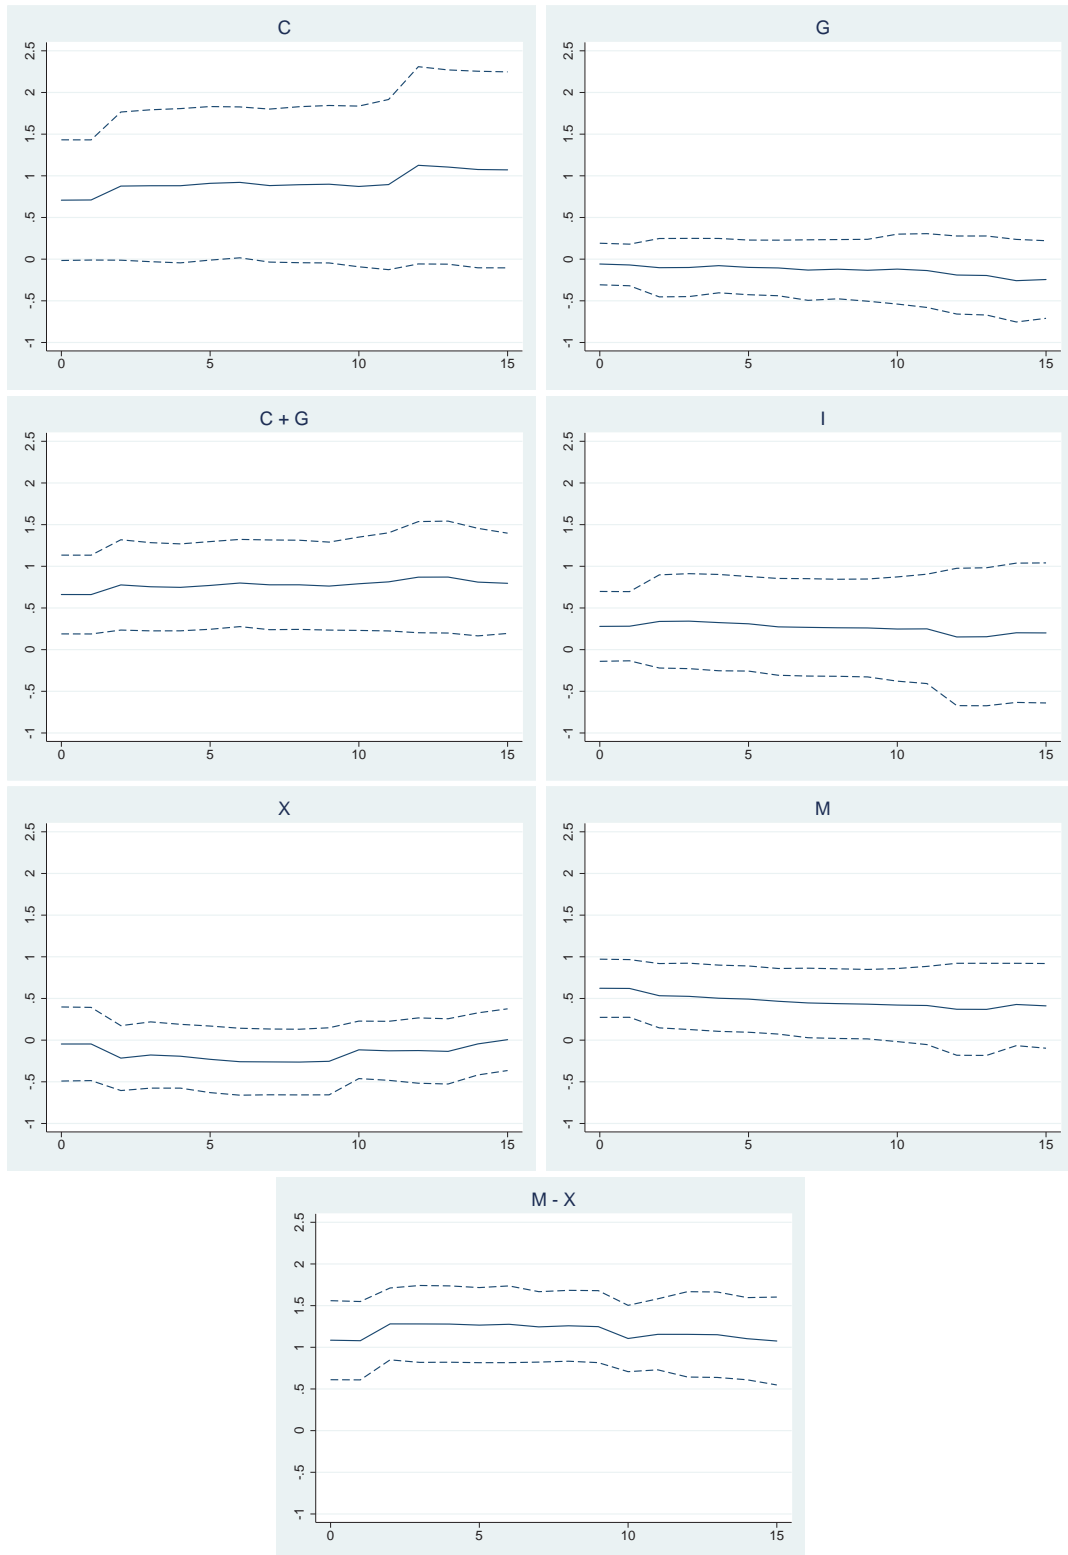
institutions. To identify the aid recipients who rely on mineral exports, we use the share in merchandise exports of fuel, ores and metals, and multiply this by the share of merchandise exports in GDP. We then exclude the eleven countries where mineral exports, averaged over the relevant time periods, exceed 30% of GDP: in descending order, these are Gabon, Oman, Kuwait, Saudi Arabia, the Republic of the Congo, Libya, Trinidad and Tobago, Nigeria, Angola, Yemen, and Zambia. In this reduced sample, we continue to find strong effects of aid on net imports and a significant effect on imports, while there is no evidence that exports are adversely affected (results not reported). As before, absorption seems to work through total consumption: the estimated effect of aid on total consumption is significant at the 5% level in the static CCE IV estimates, although only at 15% in the dynamic case. The dynamic models suggest that aid may also have a positive effect on investment, whether estimated by FE IV or CCE IV.

To analyze subgroups linked to institutional quality, we take the average of the six Worldwide Governance Indicators. These data are available bi-annually for 1996-2002, and annually thereafter. For each aid recipient, we average the governance indicator over all sample periods in which it is available and the recipient is in the sample. When we drop the ten best-governed countries, we continue to find strong effects of aid on total consumption and on net imports, while dynamic models continue to suggest that aid also promotes higher investment (results not reported). As before, there is evidence of an increase in imports, while the estimated effect on exports is insignificant. When we instead drop the ten worst-governed countries, the results are similar. We still find a strong effect on net imports, and no evidence of an adverse effect on exports. The effect on total consumption is significant at 10% in the CCE IV estimates of a static model, but only at 18% in the dynamic case. In line with the previously-acknowledged fragility of the results for investment, the effect on investment is now insignificant.

We now return to the potential for spurious correlations. One 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 influence 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').

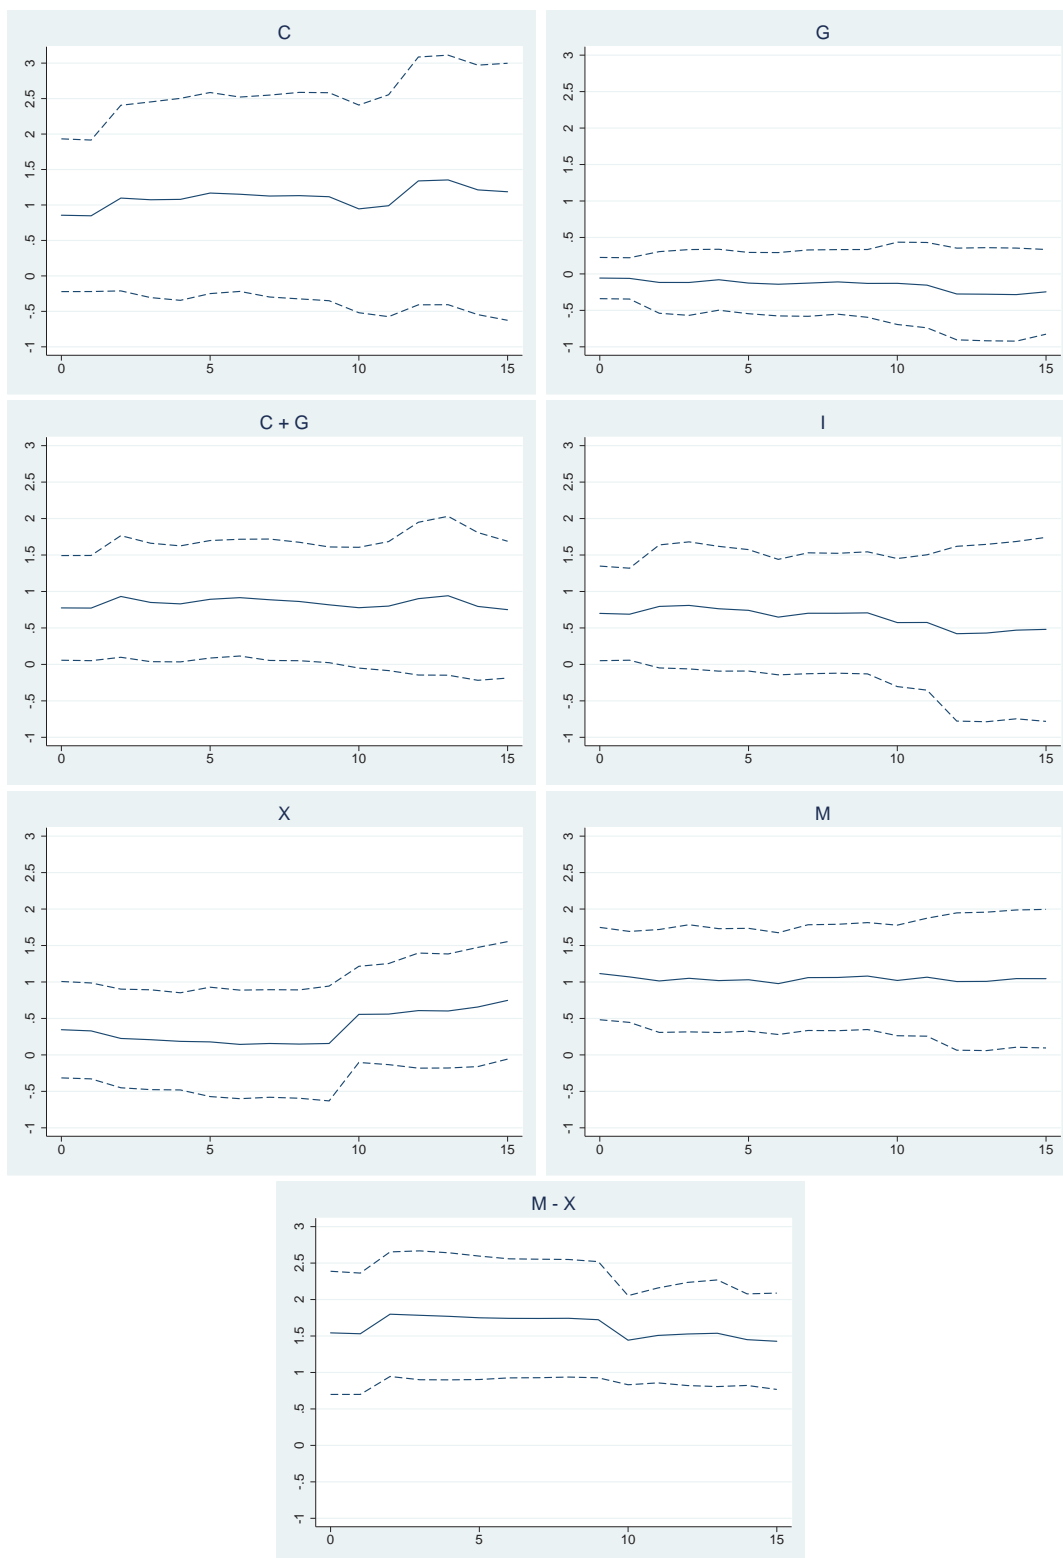
For each country-period observation, we calculate the percentage change in real GDP from the previous to the current three-year period. For a static model estimated by CCE IV, Figure B.1 shows the evolution of the estimated effect of aid as we progressively drop the countries with the largest output collapses. Figure B.2 does the same for the

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.

long-run effect of aid estimated from a model with a lagged dependent variable. If the impact of aid differed in times of crisis, or earlier results were driven by large movements in the denominator of the ratios, 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.

The spurious correlation argument is less applicable once the variables have been first-differenced. But we now investigate another approach, which is to work with an approximation to (1) and (2). This will allow us to separate the effects of GDP from those of aid, by expressing the dependent variable in logarithms. Consider a version with fixed effects rather than common factors and a multiplicative error term:

$$Q_{it}/Y_{it} = [\eta_i + \beta (A_{it}/Y_{it})] \exp(\varepsilon_{it}) \quad (\text{B.3})$$

We can assume both sides of this equation are strictly positive, the empirically relevant case in what follows. If we take logarithms, we can write:

$$\log(Q_{it}/Y_{it}) = \log[1 + \beta (A_{it}/Y_{it}) + (\eta_i - 1)] + \varepsilon_{it} \quad (\text{B.4})$$

Consider cases where $\beta (A_{it}/Y_{it}) + (\eta_i - 1) \approx 0$. Then we can use $\log(1 + x) \approx x$ for small x to arrive at an approximation:

$$\log(Q_{it}/Y_{it}) \approx \beta (A_{it}/Y_{it}) + (\eta_i - 1) + \varepsilon_{it} \quad (\text{B.5})$$

This model can be estimated as before, now with redefined fixed effects $\eta'_i = \eta_i - 1$. Given that the approximation is likely to work best if η_i is close to unity, we focus on two dependent variables in particular: the ratio of total consumption to GDP and the ratio of domestic absorption ($C + I + G$) to GDP. The latter can be seen as the mirror image of the earlier results for net imports.

Table B.2: Log dependent variables

Row	Model	Without cross-sectional mean of $\log Y$		With cross-sectional mean of $\log Y$	
		$\log(C + G)$	$\log(C + I + G)$	$\log(C + G)$	$\log(C + I + G)$
1	Static	0.00693** (0.00324)	0.00923*** (0.00257)	0.00243 (0.00311)	0.00516 (0.00429)
	F	19.21	17.71	15.74	15.08
	Dyn.	0.00794* (0.00471)	0.0125*** (0.00445)	0.00446 (0.00418)	0.00534 (0.00540)
	F	11.95	13.04	9.350	9.262
2	Static	0.0136* (0.00746)	0.00800* (0.00475)	0.0169 (0.0110)	0.00321 (0.00926)
	F	10.30	9.561	3.937	3.171
	Dyn.	0.0141 (0.00871)	0.00977* (0.00570)	0.0183 (0.0126)	0.00325 (0.0106)
	F	7.981	6.045	3.730	1.592

Note: The entries in this table show the long-run effect of aid on the log of total consumption ($\log(C + G)$) and the log of domestic absorption ($\log(C + I + G)$) in models with (“Dyn.”) and without (“Static”) a lagged dependent variable. All variables expressed relative to GDP. CCE IV estimation on three-year averages (1971-2012) using an instrument based on initial shares in donor budgets over 1960-70. The first two columns control for $\log(Y)$. The final two columns in addition control for the cross-sectional mean of $\log(Y)$, with country-specific coefficients, and only include countries with at least 7 time series observations. Heteroskedasticity-robust standard errors, clustered by country, in brackets. Standard errors (SE) for long-run effects based on the delta method. *, **, and *** denote significance at 10, 5 and 1%, respectively. F shows the first-stage F-statistic.

Row 1 uses the full sample ($N = 88, NT = 1099$; when controlling for the cross-sectional mean of $\log(Y)$: $N = 86, NT = 1087$).

Row 2 removes outliers ($N = 81, NT = 1004$; when controlling for the cross-sectional mean of $\log(Y)$: $N = 79, NT = 992$).

This approach has a significant advantage. As a test of whether the correlations are spurious, we can introduce a role for the logarithm of Y_{it} as an explanatory variable. This can be interpreted as a regression of, say, the logarithm of total consumption on the logarithm of GDP, and the aid/GDP ratio in levels, in which case the concern about spurious correlations with aid does not apply.³⁴ In CCE and CCE IV estimates, we could also include the cross-section mean of the logarithm of Y_{it} with country-specific coefficients. For a number of reasons, we should expect this to work less well than our main approach. It is based on an approximation, which in itself will weaken the estimated models under our maintained assumptions, while the CCE case implies a large increase

³⁴It might be asked why we do not work with a model which includes aid in logarithms as well. That log-linear model would imply a multiplicative relationship in levels, which is hard to justify in terms of economic theory. It would also have unconvincing implications: for example, if aid is close to zero, then consumption will be predicted to be close to zero.

in the number of parameters to estimate.

If our earlier results arose from spurious correlations, we should find that expressing the dependent variable in logarithms, and controlling for the logarithm of GDP on the right-hand-side, leads to different findings. But it is clear from the first two columns of estimates in Table B.2 that we continue to find effects of aid on total consumption and total absorption. The magnitudes of the coefficients cannot be compared directly to our earlier models, now that the dependent variable is in logarithms, but a simple calculation indicates that the average marginal effects remain similar to those previously reported. The main qualification to the findings is that, if we also add the cross-section mean of the logarithm of nominal GDP, with country-specific coefficients, the instrument weakens, and the estimated effects become imprecise; this is natural given the large number of parameters, but the signs of the effects do not change. This is a far more stringent test than is usual in the cross-country literature. To investigate it further might require a switch to annual data, and a model with richer dynamics, but we leave this for future work.

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, although it is tied to a particular data generating process. In our case, including the reciprocal of nominal GDP will again 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 are generally quite similar to our baseline estimates; see Table B.3. 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.

Table B.3: Controlling for $1/Y$

Row	Model	C	G	C+G	I	X	M	M-X
1	Static	0.708	-0.0583	0.662**	0.279	-0.0464	0.622***	1.085***
		(0.435)	(0.150)	(0.284)	(0.252)	(0.267)	(0.210)	(0.285)
	F	18.18	14.99	19.84	14.67	30.69	29.27	18.14
	Dyn.	0.856	-0.0566	0.775*	0.700*	0.346	1.115***	1.543***
(0.655)		(0.172)	(0.436)	(0.394)	(0.402)	(0.385)	(0.513)	
	F	11.84	13.89	12.12	11.21	30.97	19.52	13.45

Continued on next page

Table B.3 – continued from previous page

Row	Model	C	G	C+G	I	X	M	M-X
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
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 averages (1971-2012) using an instrument based on initial shares in donor budgets over 1960-70. Heteroskedasticity-robust standard errors, clustered by country, 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 ($N = 88, NT = 1099$).

Row 2 removes outliers, repeating row 9 of Table 8 ($N = 81, NT = 1004$).

Row 3 controls for $1/Y$ in the full sample ($N = 88, NT = 1099$).

Row 4 controls for $1/Y$ in the sample without outliers ($N = 81, NT = 1004$).

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 ($N = 86, NT = 1087$).

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 ($N = 79, NT = 992$).

In summary, although we have implemented a battery of tests that are far more stringent than is usual in the literature, our overall findings seem to hold up well. The result that aid increases net imports around one-for-one emerges repeatedly, so that the null of full absorption is not rejected. In terms of channels, there is some variation in the results, which is not surprising when the number of parameters is increased or the sample size reduced. Nevertheless, most of the estimated models indicate that absorption is taking place primarily through increases in total consumption. Also in line with our main findings, the evidence tends to suggest that the increase in net imports is achieved by increased imports rather than a reduction in exports, so we do not uncover any evidence of Dutch Disease.

B.6 List of donors

Table B.4 lists the 29 donors that contribute to the instrument values of recipient countries in the sample, alongside their average annual share in world aid over the period 1960-2012 and their median Herfindahl-Hirschman (HH) index over the sample period (1971-2012). Donors are listed in descending order of their share in world aid.

References

Brock, Philip L. (1996). International transfers, the relative price of nontraded 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.

Greenstone, M., Mas, A., and Nguyen, H.-L. (2015). Do credit market shocks affect the real economy? Quasi-experimental evidence from the Great Recession and 'Normal' economic times. Working paper, University of Chicago, September.

Table B.4: List of donors that contribute to the instrument

Donor	Share in world aid	Median HH index
United States	25.78	0.010
Japan	9.00	0.066
France	7.99	0.042
Germany	7.46	0.038
IDA	7.43	0.071
EU Institutions	6.52	0.031
United Kingdom	4.72	0.052
Netherlands	2.58	0.042
Canada	2.29	0.045
Australia	2.02	0.165
Italy	1.72	0.112
Other Multilaterals	1.63	0.525
Kuwait	1.60	0.132
WFP	1.59	0.044
Sweden	1.58	0.058
UNDP	1.52	0.017
Other donor countries	1.24	0.344
Belgium	1.22	0.139
AsDB Special Funds	1.19	0.176
IDB Sp.Fund	1.14	0.156
Norway	1.03	0.057
Denmark	0.94	0.061
UNICEF	0.82	0.037
UNHCR	0.79	0.080
Switzerland	0.65	0.034
Austria	0.53	0.204
UNTA	0.40	0.013
Portugal	0.34	0.284
UNRWA	0.30	0.369

Note: Share in world aid is the average annual share of the donor in total world aid over the period 1960-2012. Median HH index is the median Herfindahl-Hirschman index of the donor over the sample period (1971-2012). For details on the construction of the HH index, see Figure 1 and the surrounding text in the paper.

Jann, Ben (2013). coefplot: Stata module to plot regression coefficients and other results. <http://ideas.repec.org/c/boc/bocode/s457686.html>.

Turnovsky, Stephen J. (1997). *International Macroeconomic Dynamics*. MIT Press, Cambridge, MA.