CSAE Working Paper WPS/2014-01

Foreign Aid and Domestic Absorption

Jonathan R. W. Temple
University of Bristol

Nicolas Van de Sijpe
University of Oxford

9th January 2014

ABSTRACT

We introduce a new 'supply-push' instrument for foreign aid, to be used together with an instrumental variable estimator that filters out interactive fixed effects. 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 increases consumption, and there is also some evidence that aid raises investment, but with a delayed effect.

We are grateful to Chris Adam, Paddy Carter, Gabriela Cohen-Freue, Markus Eberhardt, Gerdie Everaert, Benedikt Goderis, Carlos Lamarche, Anita Ratcliffe, Francis Teal and Frank Windmeijer for various contributions and advice. The first author is grateful for financial support from the British Academy.

UNIVERSITY OF OXFORD

Centre for the Study of African Economies
Department of Economics · University of Oxford · Manor Road Building · Oxford OX1 3UQ
T: +44 (0)1865 271084 · F: +44 (0)1865 281447 · E: csae.enquiries@economics.ox.ac.uk · W: www.csae.ox.ac.uk

1 Introduction

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

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

Second, we shift the focus to how foreign aid is absorbed by the domestic economy. In particular, we are interested in the causal effects of aid on consumption, investment and net imports, in the short run and the long run. For example, net imports must respond if the aid is to be absorbed, rather than spurring capital flight or building up foreign exchange reserves. If aid improves the investment climate, we would expect to see an increase in investment relative to GDP, so that aid is at least partly absorbed by higher investment. We will argue that the effects of aid on macroeconomic ratios are inherently easier to study than the effects on steady-state GDP levels or growth rates.

As is well known, the basic 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. 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 at some initial date t_0, multiplied by the current donor budget at date t. The basic idea is that, when there are many aid recipients, donor budgets vary largely independently of conditions in any given recipient. We make this idea more precise below, and show that it can be used to isolate exogenous variation in aid under a relatively plausible exclusion restriction.

Further intuition can be gained from the case of two donors – say, Britain and France. When Britain's total aid budget increases relative to that of France, former British colonies are likely to see an increase in aid received, relative to former French colonies. It is this form of variation that our synthetic measure is intended to capture, isolating it from variation driven by the particular circumstances of individual aid recipients. In other words, we look for exogenous variation driven by changes to total donor budgets. If the evolution of total budgets is not influenced by the domestic circumstances of specific recipient countries, we have a candidate instrument. We call this a supply-push instrument, since it is closely related to the work of Card (2001) on the labour market effects of immigration.

An important objection to the supply-push instrument is that, in practice, donor budgets are influenced by economic forces that have independent effects on aid recipients. For example, the state of world economic conditions is likely to affect donor generosity, and also the economic outcomes of poor countries. Drawing on recent work in the panel time series literature, these forces can often be seen in terms of interactive fixed effects: latent common factors with loadings that differ across countries. These represent a substantial generalization of conventional country and time fixed effects. We filter out these interactive fixed effects using an instrumental-variable version of a common correlated effects (CCE) estimator. This class of estimators was introduced by Pesaran (2006) and extended to instrumental variables by Harding and Lamarche (2011). Once the factors have been filtered out, an argument that our instrument could be (statistically) endogenous is harder to construct.

Using this approach, we seek to recover the causal relationships between aid and various macroeconomic ratios. Aid, as a capital transfer, is not part of measured GDP. In principle, it could be offset by a corresponding capital outflow. It could also be used to accumulate foreign exchange reserves. Alternatively, aid may lead to changes in at least two of the components of GDP: household consumption, government consumption, gross investment, exports and imports. We study the effects of aid on the ratios of these components to GDP, in the short run and long run. We cannot reject the hypothesis that aid leads to a one-for-one increase in net imports. Absorption occurs mainly through an increase in imports rather than a decline in exports, and hence we do not find evidence for Dutch Disease effects. There is clear evidence that aid raises total consumption. The effect on investment is less clear, but in dynamic models we find that aid may raise investment, with a delayed effect.

The rest of the paper has the following structure. In section 2, we describe the basic framework and the approach to estimation and how it relates to the previous literature. Section 3 describes the data. In section 4, we present the main results, while section 5 considers Dutch Disease effects. Section 6 presents a number of robustness checks, before section 7 concludes.

2 Basic framework

In this section we set out the main ideas of the paper. We are interested in whether aid flows are absorbed and, if so, whether the absorption is reflected primarily in higher household consumption, gross investment, or government consumption. We start from the basic GDP identity, in standard notation:

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

where Y is GDP, C is household consumption, I is gross investment (private and public), G is government consumption, X is exports and M is imports. To see how aid could be absorbed, we take the ratios of these components to GDP to be the outcomes of interest, and sometimes aggregate the last two to give net imports as a share of GDP, $(M - X)/Y$.[1]

From a national accounts perspective, foreign aid is a capital transfer which does not contribute directly to GDP, but in principle allows greater expenditure relative to domestic production of goods and services. For the aid to be absorbed in this way, at least one of C, I or G must increase, along with their total. A corollary is that net imports, $(M - X)$, must increase. In other words, the non-aid current account balance must deteriorate. There is nothing problematic about this; it is what must happen if aid is to permit greater expenditure relative to GDP.[2] If aid is entirely devoted to greater domestic expenditure, net imports will

[1]Since a linear combination of these dependent variables equals unity by construction, the model for one of the dependent variables will be statistically redundant when the covariates are the same across the regressions. But for ease of interpretation, we report results for each of the dependent variables. In the terminology of Aiyar et al. (2006), the quantity $M - X$ is the non-aid current account deficit.

[2]For background on aid absorption see Aiyar et al. (2006), Aiyar and Ruthbah (2008) and

rise one-for-one with aid. The increase may be smaller in certain circumstances. Some aid may be devoted to technical assistance which primarily funds consultants from the donor country, and hence will not be available to be absorbed by the recipient. Alternatively, aid may be used to build foreign exchange reserves, or exit the country as capital flight.

As well as studying whether aid is fully absorbed, we study its effects on the ratios of consumption and investment to GDP. Simple predictions can be obtained from growth models, notably the one-sector Ramsey model without government. If aid takes the form of grants made direct to households, as in Obstfeld (1999), then a permanent increase in aid raises the investment ratio in the short run, but not in the long run. Aid promotes faster convergence to the steady-state, but the long-run level of GDP is invariant to aid. Along the balanced growth path, all aid is consumed. From a national accounts perspective, consumption is higher while investment and GDP are unchanged, and the increase in steady-state consumption is permitted by imports of the final good. This implies that when the ratio of aid to GDP increases, the long-run C/Y and $(M-X)/Y$ ratios increase by the same absolute number of percentage points, leaving the other ratios unchanged.

Given the difficulties of using cross-country data to study aid effectiveness, there are at least three good reasons to start with the effects of aid on macroeconomic ratios, rather than growth. First, their relationships with aid are more likely to be linear, as in the case of the Ramsey model just discussed. In contrast, the study of growth has to contend with the possibility of diminishing returns to aid. Second, aid is sometimes justified precisely as improving the conditions for domestic investment. If we see investment as a jump variable, it can respond relatively quickly to changes in aid. This fast response may be easier to detect in the data than a longer-lived effect on GDP. Third, growth researchers have to contend with the possibility of slow convergence, and hence a high degree of persistence of GDP. In contrast, the macroeconomic ratios we study are likely to be less persistent.

With these arguments in mind, we study the effects of aid (as a share of GDP) on macroeconomic ratios. The endogeneity problem, or the non-random allocation of aid, is central. Even in a model that controls for country and period fixed effects, it is likely that aid flows and outcome variables are jointly influenced by one or more variables that are not readily measured, and hence will be omitted

Hansen and Headey (2010). The results in Table 2 of Werker et al. (2009) can also be interpreted as findings on absorption.

from the regression. A conventional instrumental variable approach can address this, but only if the instrument is uncorrelated with the error term. We are interested in whether we can achieve identification even when the error term may include a number of latent common factors, with factor loadings that vary across countries. This is a natural structure in the current context, where aid flows and macroeconomic outcomes are likely to be jointly influenced by hard-to-measure variables such as world economic conditions.

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

This form of presentation makes clear the need to allow for latent factors with heterogeneous effects. Imagine the process generating the macroeconomic outcome of interest M_{it} is given by:

$$M_{it} = \eta + \beta\left(D_{it}/Y_{it}\right) + \phi_i f_t + \varepsilon_{it} \tag{1}$$

where f_t is a vector of unobserved common factors (including, say, world economic conditions) and ϕ_i is a set of factor loadings which may vary across countries. This interactive fixed effects specification nests both conventional fixed effects (where one common factor is time-invariant) and conventional period effects (where loadings on one time-varying factor are the same across countries) as special cases, but is more general than either.

We proceed on the assumption that we do not have observable proxies for the various possible components of the interactive fixed effects. For a conventional fixed-effects IV estimator to achieve identification, we would then need the fixed-share instrument D^S_{it}/Y_{it} to be uncorrelated with the interactive fixed effects that appear in the data generating process, and hence in the error term. This could easily be questioned. For example, it is plausible that donor budgets will be correlated with world economic conditions which also influence macroeconomic ratios, in which case our supply-push instrument would be correlated with the error term.

This suggests the need to go beyond conventional IV estimation. We emphasize results which filter out the interactive fixed effects using the approach of Pesaran (2006). That paper introduced common correlated effect (CCE) estimators for panel data, which proxy for the combined effect of common factors using cross-section averages of the dependent variable and all the explanatory variables, all with country-specific coefficients. This idea has been extended to the case of instrumental variables by Harding and Lamarche (2011), yielding a CCE IV estimator. We report the results from several methods, but give most emphasis to the CCE IV estimator, as the one most likely to yield consistent estimates under a broad range of scenarios.

The CCE class of estimators has been widely adopted in recent empirical work. The approach can accommodate various forms of cross-section dependence, and has been shown to perform relatively well even in small samples.[3] A remaining limitation of standard CCE approaches is that factor loadings which are correlated with the regressors can lead to inconsistent estimates. This problem does not arise when a suitable instrument is available, as Harding and Lamarche (2011) discuss.[4] They also present simulation evidence in which the CCE IV estimator performs well even when the factor loadings are correlated with the regressors.

It is important to note that the interactive fixed effects structure is a major generalization of conventional fixed effects and the usual approaches to cross-section dependence in the cross-country literature. We could have proxied for common factors by using time dummies, and used time dummies interacted with observed, country-specific variables to proxy for some common factors with differential effects.[5] Such approaches, although more general than a single set of time dummies, heavily restrict the structure of the unobserved, heterogeneous factor loadings. Hence, they are less general than the approach we adopt.

One possible concern raised by our approach is that our 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. At first glance, this might

[3]Relevant papers include Chudik, Pesaran and Tosetti (2011), Kapetanios, Pesaran and Yamagata (2011), and Pesaran and Tosetti (2011).

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

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

suggest that the instrument would be eliminated from the first stage of 2SLS, when filtering out interactive fixed effects. It is easy to show, however, that with one endogenous variable and one instrument, the instrument is only eliminated from the first stage when there is either a single donor or the initial shares of aid recipients in donor budgets are the same across donors.[6] 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.

Our supply-push instrument was first used by Van de Sijpe (2010) to study the effect of aid on governance, but without filtering out interactive fixed effects. Related synthetic measures of aid, but based on average shares in donor budgets rather than initial shares, have been used in Hodler and Raschky (2011) and Nunn and Qian (forthcoming). These average shares in donor budgets may be affected by developments within a recipient economy over the sample period, which weakens the case for exogeneity. To see this, 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. 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 $(1/T) \cdot (a_{i1} + ... + a_{iT}) \cdot D_t / Y_{it}$, 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.[7]

So far, we have said nothing about dynamic aspects of the specification. Pesaran (2006, p. 975) notes that common feature dynamics across the units (here, countries) are captured through the serial correlation structure of the common factors. But a remaining concern with our initial-share instrument is that circumstances specific to individual aid recipients, such as their domestic political developments, may be serially correlated. For each country, the initial share in a donor's budget may then be correlated with shocks in some of the subsequent periods, which undermines the exogeneity of the instrument. This po-

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

[7]In addition, at least some of the aid shares in other periods, a_{is} for $s \neq t$, are likely to be a function of the transient shock at date t, and this could also lead to a failure of exogeneity. Nunn and Qian experiment with calculating shares over the recent past, which avoids the problem if the shocks are not serially correlated.

tential limitation of the supply-push approach is acknowledged by Card (2001, footnote 23). It is likely to be an especial concern for the earlier time periods of the panel, and when there are relatively few time periods overall. Later in the paper, we investigate the problem by dropping some of the early time periods from the estimated model. This means that the initial share is measured some years before the first time period used for estimation. When we do this, we find no warning signs that our main results are substantially affected by underlying serial correlation in country-specific circumstances.

Although our work is related to that using average-share instruments, our use of CCE estimation should achieve identification under more general conditions than previous work. Some influential recent studies based on cross-country data, such as Clemens et al. (2012) and Rajan and Subramanian (2008), do not address cross-sectional dependence in any detail. Nunn and Qian (forthcoming) and Werker et al. (2009) use instruments that are similar in spirit to ours, and carry out robustness tests which correspond to searching for observable proxies for the elements of $\phi_i f_t$ in equation (1) that seem especially likely to threaten identification.[8] This increases the credibility of identification in these papers, but also highlights the potential benefit of the more general approach we adopt in this paper. A further advantage of our approach is that we can study the effects of total development aid, whereas studies based on natural experiments are currently limited to specific forms of aid, rather than aid as a whole. For example, the Werker et al. findings are most informative about the effects of unconditional grants from Gulf oil exporters to Muslim recipients, while the Nunn and Qian findings relate to US food aid.

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

[8] Werker et al. construct an instrument for aid by interacting the world oil price with a dummy for Muslim countries. Nunn and Qian instrument wheat aid received from the US with an interaction of one-period-lagged US wheat production and the fraction of years in the sample in which a country has received food aid from the US.

3 Data

Our estimation sample consists of three-year averaged data for the period 1975-2010.[9] For the reasons set out in the previous section, the initial shares of recipients in donor budgets on which the synthetic instrument is based are calculated before the start of the sample used in estimation, namely over the period 1960-1974.

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

Our synthetic measure for aid is constructed in the following way from the DAC's recipient-donor-year data. For each donor, we calculate the average of the annual shares of a given recipient country in a donor's aid for the years 1960-1974 (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).[10] We then sum these numbers across donors to get $D_{it}^{S} = \sum_{d=1}^{N_D} a_{i0}^{d} D_{dt}$. 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 1960-1974 average shares. It is this time-varying, synthetic measure of aid that we use to instrument for aid in panel data regressions. Both the endogenous aid variable and the instrument in the regressions below are expressed as a percentage of GDP. Our GDP data, and the other macroeconomic variables used in the regressions, are extracted from online World Bank data files using `wbopendata` for Stata (Azevedo, 2011).

The dependent variables considered below include household consumption, government consumption, gross capital formation, imports and exports, all ex-

[9] The 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.

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

pressed as a percentage of GDP. Net imports are defined as imports minus exports. In the recipient-year data, before collapsing to three-year averages, observations for these variables are turned to missing whenever at least one of the other variables of the GDP identity is missing. This keeps the sample consistent across the different dependent variables we consider below. We exclude countries with small populations (fewer than 500,000 people in the first period of the sample, 1975-77). In our final data set, the available macroeconomic aggregates sum to total GDP, or very close to GDP, for each observation.[11]

4 Results

For each dependent variable, we report eight regressions. For reference purposes, we report FE and pooled CCE results that do not instrument for aid. We report estimates for static models, and dynamic models that include a lagged dependent variable.[12] Whenever the estimated model is dynamic, we report the long-run effect of aid, with a standard error approximated by the delta method. We note that CCE-type estimators are consistent in dynamic panel data models only under more restrictive assumptions than in the static case (Chudik and Pesaran, 2013, Everaert and De Groote, forthcoming) and so we consider both types of model. The standard errors that we report are heteroskedasticity-robust and clustered by country, and we make a small-sample adjustment to take into account the large number of estimated parameters. In experiments, we compared the adjusted standard errors for balanced panels for the CCE IV estimator to those obtained from a non-parametric bootstrap, given that the asymptotic distribution of pooled CCE-type estimators depends on nuisance parameters (Pesaran, 2006) and the asymptotic variance of the CCE IV estimator introduced in Harding and Lamarche (2011) has not yet been studied. The adjusted standard errors are typically close to bootstrapped versions so, for simplicity, we report the former throughout the paper.

Before we turn to the results, one point of interest is that the effects of instrumenting for aid will differ across dependent variables.[13] The effect of in-

[11]Three country-year observations show discrepancies larger than 1% of GDP. Dropping these observations makes little difference to our main results.

[12]The sample for each of these 8 regressions consists of the observations that are included in the CCE IV estimation of the dynamic model. Since CCE IV estimation of the dynamic model requires estimating 5 country-specific parameters (1 country fixed effect and 4 factor loadings for the cross-section means) only countries with at least 6 observations are retained in the sample. In our main sets of estimates, we have 11 or 12 time series observations for many countries, with a maximum of 12.

[13]In models with the same covariates, the effects of aid on our dependent variables must sum

strumenting will depend on the extent to which aid allocation is non-random. Macroeconomic ratios are likely to differ in their sensitivity to particular types of aid-relevant shocks, and so the effect of instrumenting is likely to differ across the different outcome variables that we consider.

We first study the effect of aid on household consumption (Table 1). Some of the estimated effects are imprecise but, comparing the upper row of estimates with the lower row, the difference made by instrumental variables can be seen clearly. Compared to the FE and CCE estimates, the point estimates from IV estimators suggest much larger effects of aid on household consumption. We should avoid over-interpreting this, because the differences are not statistically significant. But it is easy to see how the pattern could arise. In terms of the within variation, aid may respond positively to circumstances which lower consumption, such as adverse shocks. By instrumenting aid we alleviate this source of bias, and find much larger effects of aid on consumption. We find much less evidence that aid influences government consumption, as Table 2 shows. This pattern is similar to that found by Werker et al. (2009, Table 2).

Aid has often been characterized as primarily government-to-government transfers. Our finding that a substantial fraction of aid is reflected in higher household consumption, but not in higher government consumption, may therefore appear surprising. One mechanism could be lower taxes: standard reasoning suggests that increased aid to governments will not only be used to increase government purchases, but also to reduce taxes (Kimbrough, 1986). It is also possible that recipient governments use aid to finance transfers for political ends. This assumption has been common in the literature, as in Boone (1996), Adam and O'Connell (1999) and Hodler and Raschky (2011), among others. Finally, some aid is given in ways which bypass domestic governments, such as off-budget aid projects, or support for NGOs.[14] In all of these cases, household consumption is where the effect of aid is most likely to be manifested in the national accounts.

Table 3 shows results for total consumption, defined as the sum of household and government consumption. While household and government consumption are distinct concepts, 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. Examining the effect on total consumption is also a natural counterpart for our later results on investment. When we study the effects of

to zero, by construction. Hence, not all the estimates can move in the same direction when an instrument is used.

[14]Van de Sijpe (2013) discusses off-budget aid in more detail.

aid on total consumption, the estimates are generally similar to those we find for household consumption, but more precise. Again, the use of an instrument increases the estimated effect of aid.

For gross capital formation, the results in Table 4 suggest the opposite pattern. Now the estimated effects weaken once we instrument for aid. In the lower set of estimates in this table, it is clear that we cannot reject the hypothesis that increases in the ratio of aid to GDP have no effect on the ratio of investment to GDP. Again, there is a contrast between the uninstrumented and instrumented estimates, visible by comparing the upper and lower panels. A potential explanation is that unobserved changes in the climate for investment may influence both aid and investment simultaneously. Once we eliminate this source of spurious correlation, the evidence that aid influences investment is weaker. This is not necessarily a surprise: as discussed earlier, in the one-sector Ramsey model, aid has no effect on the steady-state ratio of investment to GDP.

Finally, we study the effect of aid on net imports. Recall that net imports must increase if aid is to be absorbed. If the net import share rises one-for-one with the aid share, this should assuage concerns that aid is diverted abroad (capital flight) or primarily used to accumulate foreign exchange reserves. The relevant results are shown in Table 5. 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 very close to unity, both in static models and in the long run derived from dynamic models. Again, the contrast with the upper row of estimates is instructive. If we estimate a static model that does not instrument for aid, the evidence that aid is fully absorbed is much weaker.

We have not yet discussed the strength of our instrument. The tables report the first stage F-statistic as a heuristic guide.[15] This approach has been widely used, but it should be noted that the conventional Stock and Yogo (2005) benchmarks for first-stage F-statistics do not apply directly to panel data models. This is partly because the benchmarks assume i.i.d. errors. By considering a model for a single time series, Bun and de Haan (2010) show that the standard benchmarks for the first-stage F-statistic do not apply when the errors are serially correlated. In their Monte Carlo simulations, a robust F-statistic tends to underestimate instrument strength. It is not clear whether this is general or would extend to panels, and hence our use of F-statistics is best seen as heuristic; this

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

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

5 Dutch Disease?

We can also use this approach to investigate whether Dutch Disease effects are at work. To see this, note that aid absorption may be reflected in a one-for-one increase in imports, or a combination of higher imports and lower exports. The former implies limited scope for Dutch Disease, because the extra expenditure made possible by aid has been spent entirely on externally-produced goods and services.[16] But if the increase in imports is less than one-for-one, expenditure on domestically-produced goods will have increased, which will tend to imply an increase in the relative price of non-traded goods. This relative price change will tend to attract labour into the sector, raising labour costs in the traded sector: this shifts the supply curve upwards in the traded sector, leading to a decline in exports. We do not investigate directly whether or not such a change in relative prices is harmful. Instead, we examine whether increased domestic expenditure is met mainly through imports.

Tables 6 and 7 show the effects of the aid share on import and export shares respectively. The results for the import share are similar to those for the net import share. A strong positive effect of aid is found in the CCE IV regressions in particular. In contrast, we do not find a clear-cut effect of aid on the export share. Some of the point estimates are negative, but we cannot reject the hypothesis that aid has no causal effect on the export share. This does not rule out Dutch Disease – for that, we would need to find zeroes estimated with greater precision – but neither do we find any persuasive evidence that would be consistent with Dutch Disease effects.

6 Robustness

In this section, we consider various robustness tests. Most of these involve reducing the sample in some way, and hence often lead to less precise estimates. Nevertheless, our main conclusions turn out to be reasonably robust, even when

[16]If some imports are of capital goods, one effect of aid will be to augment productive capacity, but the general equilibrium effect will depend on which sectors are investing.

we make adjustments to the instrument that weaken its explanatory power in the first stage.

In our main results, we followed much of the cross-country literature and excluded countries with small populations. If, instead, we include these countries, we generally find that the instrument becomes weaker in the first stage of 2SLS.[17] These results are shown in row 2 of Table 8; row 1 repeats the main results from Tables 1-7 for ease of comparison. A potential explanation is that, for countries which account for small and possibly volatile shares of donor budgets, the share of a budget at an initial date may be relatively uninformative about that country's long-term exposure to changes in that budget. In that case, we would expect our supply-push instrument to have more explanatory power for aid for larger countries than for small ones. Qualitatively, the second stage results are similar: we continue to find large effects of aid on total consumption, household consumption (though less clearly than before), net imports, and imports, and no clear effects on the other variables.

We next examine the implications of transitions from colonial rule to independence. In some cases, the DAC dataset includes reports of aid flows before a recipient country became independent. In our case, this implies that, for some countries, we have constructed an instrument based on initial shares in donor budgets in the period 1960-74 even though the country only became independent later. To the extent that recorded aid flows before independence are incomplete or measured less accurately, this may affect our results. Hence, as an alternative, we calculate the aid variable (and the initial shares in donor budgets needed to construct an instrument) only after discarding all aid data in the years before a recipient's independence.[18] The results, shown in row 3 of Table 8, are similar to those found before. In the static model, the CCE IV estimate of the effect of aid on imports is slightly smaller than before, but there is no evidence that aid undermines exports. In the dynamic model, the long-run effect of aid on imports is still large, significantly different from zero, and not significantly different from one.

One objection to much cross-country research on aid is that it does not consider delayed effects of aid in sufficient detail (this point is especially associated

[17]The countries that we return to the sample are Antigua and Barbuda, the Bahamas, Bahrain, Barbados, Belize, Bhutan, Cape Verde, Comoros, Dominica, Equatorial Guinea, Grenada, Kiribati, Macao, Malta, St. Kitts and Nevis, St. Lucia, St. Vincent and the Grenadines, Suriname, Tonga, and Vanuatu.

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

with Clemens et al., 2012). Our next set of robustness tests addresses this point in various ways. Row 4 in Table 8 replaces current aid by its one period lag, instrumented by the first lag of the synthetic instrument. Row 5 includes both current and one period lagged aid, instrumented with current and one period lagged values of the instrument.[19] Row 6 replaces aid and the instrument by the unweighted averages of their current and one period lagged values.[20] Row 7 measures the dependent variable in the final year of each three-year period instead of taking the average over all three years. Overall, we find some evidence that aid has a delayed effect on investment, while the effects on consumption are sometimes more muted. The estimated effects of aid on investment are generally larger than before, and are significant in the specifications with only lagged aid or with the final year values of the dependent variable. Aid's effects on exports, imports and net imports are similar to those reported previously.

We now turn to various potential criticisms of our instrument. In terms of the robustness of our main results, there are two especially important concerns. One is that serial correlation in country-specific circumstances might undermine exogeneity. The second is that the strength of the instrument could decline over time. Our IV strategy relies on the idea that shares in donor budgets in 1960-74 are informative about exposure to later changes in total donor budgets. If, for example, strategic or economic connections between countries change over time, then the instrument is likely to have less explanatory power for aid in later periods of the sample. We can address both of these concerns with a single robustness test. Rows 8 to 10 in Table 8 repeat the main analysis but exclude, respectively, the first, the first two, and the first three periods from the panel data sample. As additional checks, row 11 uses an instrument based on initial shares calculated over the period 1960-71 and an estimation sample that starts with the period 1972-74, while rows 12 to 14 exclude the first, the first two and the first three periods from this sample.

Inevitably, since these robustness checks involve a significant loss of observa-

[19] This places heavy demands on the data for a number of reasons. CCE IV estimation of the dynamic model now requires estimating 7 country-specific parameters. This also implies that only countries with at least 8 observations can be included in estimation, reducing the sample by more than 100 observations. Moreover, high collinearity between aid and its lag, and between the instrument and its lag, weakens the instruments, further contributing to large second-stage standard errors. Reassuringly for the validity of our instrument, in the first stage only the instrument in period t is significant for aid in period t, and only the instrument in period $t-1$ matters for aid in $t-1$.

[20] In all three cases, the first period of the sample is dropped so that the sample starts in the period 1978-80. This is done to avoid any overlap between the period over which the initial shares in donor budgets, on which the instrument is based, are calculated (1960-74) and the periods over which the aid variables on the right hand side of the structural equation are measured.

tions, we would expect the standard errors to be larger, at least on average, and the first-stage F statistic to gradually weaken as more periods are excluded, and this is what we find. The estimated second-stage coefficients are fairly stable, however, when dropping early time periods. As before, the results suggest that aid is primarily consumed rather than invested. We continue to find that aid is absorbed domestically, and that this happens via an increase in imports rather than a reduction in exports.

Another possible criticism is that, by considering all donors, we have included some donors whose budgets are dominated by a small number of recipient countries. In that case, the domestic circumstances of these aid recipients may influence the evolution of the donor's aid budget over time, which risks endogeneity. This is likely to be an especial concern for smaller donors, who may concentrate most of their aid in a few recipient countries. To investigate this issue, row 15 in Table 8 shows results using an instrument based on the top ten donors. These are defined as those with the largest average annual share in global aid over the period 1960-2010.[21] As expected, this instrument is slightly weaker than the one based on all donors, and the second-stage coefficients are estimated less precisely. Nevertheless, we continue to find large effects of aid on consumption, imports and net imports.

Our next robustness tests address various related concerns. The first is that some of the identifying within variation could arise from large-scale events, in particular output collapses associated with economic crisis or civil war. These events could generate large swings in the ratio of aid to GDP and the macroeconomic ratios, and form a significant component of the within variation. To the extent that this happens, the effects of aid might be identified mainly from the variation generated by 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', this is an important concern.

To address this, we investigate what happens when we gradually eliminate the countries with the most substantial declines in real GDP. For each 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 1 shows the evolution of the estimated effect of aid as we progressively drop the countries with the largest percentage declines in real GDP, while Figure 2 does

[21]These donors are the United States, Japan, France, Germany, the International Development Association (IDA), the EU institutions, the United Kingdom, the United Arab Emirates, the Netherlands, and Canada. Over the period 1960-2010, these ten donors accounted for more than three-quarters of total aid, on average.

the same for the long-run effect of aid estimated from a model that includes a lagged dependent variable. If the impact of aid is different in times of crisis, we would expect the magnitudes of the estimated coefficients to move substantially. We find little evidence of this. In the static model, the effect of aid on net imports decreases somewhat when 12 or more countries are omitted, but even then the estimated effect remains large, and the confidence interval continues to include unity and exclude zero. In the dynamic model, the effect of aid on net imports remains close to unity throughout.

A somewhat related set of concerns is that movements in GDP may influence our results. In particular, our dependent and independent variables both contain nominal GDP in the denominator. Our maintained assumption is that we can estimate structural relationships between macroeconomic ratios and aid intensity, of the type suggested by the neoclassical growth model, but an observer sceptical about such relationships could argue that there is a risk of a 'false positive'. Movements in GDP could lead to spurious positive correlations between the dependent variables and aid intensity, so that the models could appear to have explanatory power even in the absence of a structural relationship. Moreover, where the data generating process involves a negative relationship, perhaps between exports and aid, this may be hidden by the role of GDP in the denominator of the ratios. Discussions of potentially spurious relationships between ratios date back to Pearson (1897) and Yule (1910), while Kronmal (1993) is a more recent treatment. In the aid literature, the problem is raised by Arndt, Jones and Tarp (2010) in particular.

One way to investigate whether we have found 'false positives' is to add the reciprocal of the denominator as an explanatory variable (Kronmal, 1993). Before we undertake this, note that an estimated relationship between a macroeconomic ratio and 1/GDP lacks a clear structural interpretation. First, nominal GDP is trended; at least in the long run, this cannot be true of the bounded ratios that form our dependent variables.[22] A second argument against including 1/GDP is that GDP will be determined jointly with the ratios we analyze, and hence 1/GDP is an endogenous variable. These considerations make the results hard to interpret, and we therefore add 1/GDP to the models only as an experiment, to gauge whether there is any risk of a false positive. Considering non-CCE results first, the inclusion of 1/GDP makes little difference. The estimated effects of aid are similar, instrument strength is barely affected, and 1/GDP is only

[22]The 'Great Ratios' are typically modelled as mean stationary processes, as in King et al. (1991).

significant in the second stage in four out of 28 FE and FE IV regressions. The latter is mainly the case in the model for imports, where 1/GDP is positive and significant in three out of four regressions. But overall, these results raise few concerns about false positives. The case of CCE estimators is more complicated, and we consider it in the appendix.

A related argument is that, in principle, at least some of the macroeconomic ratios could be functions of the level of development. In that case, the error term in our models would contain a nonlinear function of real GDP per capita, and this omitted component could be correlated with the instrument, which has nominal GDP in its denominator. To the extent that this effect drives our results, we would expect it to have been exposed by several of the above robustness tests, in particular when we exclude countries that experienced large declines in real GDP, and when we control for the reciprocal of nominal GDP. An alternative approach would be to control directly for some function of real GDP per capita. This is not especially attractive, since real GDP per capita is likely to be itself endogenous and/or I(1), and its use to explain a bounded ratio would need attention to the functional form. Instead, we have proceeded under the maintained assumption that our instrument is uncorrelated with omitted determinants of the macroeconomic ratios, once we have filtered out interactive fixed effects.

7 Conclusions

As many authors have emphasized, using cross-country data to study aid effectiveness is fraught with difficulties, and yet some of the relevant research questions are hard to answer any other way. This paper has aimed to make progress on two fronts. First, we have analyzed the effects of aid on various macroeconomic ratios. This is informative about the extent of domestic absorption, the relationship between aid and the investment ratio, and the extent to which the data are consistent with Dutch Disease concerns. Second, and more important for future research, we have introduced a new 'supply push' instrument for aid, and shown how to combine it with a panel time series estimator that filters out interactive fixed effects. Given this combination, we can make a stronger claim to have recovered causal effects of aid than the majority of the existing literature.

In terms of the substantive findings, we derive a number of results. We find some evidence that aid increases household consumption, obtaining larger point estimates when we instrument for aid, consistent with the idea that aid flows respond positively to adverse shocks that would otherwise lower household consumption. In the case of investment, there is a positive association between aid and the investment ratio which disappears once we instrument for aid. This is consistent with the possibility that aid flows respond positively to improvements in the investment climate. If we use an instrument to alleviate this source of bias, we find that the association between investment and contemporaneous aid is weak, but there is some evidence that aid has delayed benefits for investment.

We also study absorption by using movements in import and export shares. Our results suggest that aid increases net imports one-for-one, assuaging concerns that aid exits as capital flight or is used primarily to accumulate foreign reserves. The domestic absorption of aid comes about mainly via higher imports. Once we instrument for aid, we find no evidence that aid lowers exports, and hence no grounds on which to infer that Dutch Disease effects play an important role. Most of these results demonstrate the importance of instrumenting for aid in cross-country empirical research. Overall, our work suggests that the supply-push instrument has the potential to be informative, and could have many possible applications in the study of aid.

8 Appendix

In the main text, we discussed the potential concern that our partial correlations are driven by the presence of GDP in the denominator of the macroeconomic ratios and aid intensity. We found little evidence for this in the non-CCE case. In this appendix, we discuss the issue for the case of the CCE estimators. As well as adding 1/GDP to our model, we must add its cross-section mean with country-specific coefficients. The comparison with our main results is then complicated by a loss of degrees of freedom, and a reduction in sample size.[23] Identification is weaker in the first stages of the 2SLS estimates, with lower F-statistics. In the second stages, the effects of aid intensity are generally insignificant. This is also true, however, of the new 1/GDP variable, which is insignificant in all 14 models estimated by CCE IV and has a negative sign in half the cases.

The evidence for, or against, false positives in CCE estimates is therefore rather inconclusive. Given the number of parameters now estimated with CCE IV — 501 in the case of the dynamic models, from 944 observations — it is not

[23] Only countries with at least seven observations are included in the estimation, which means we lose eight countries, or 48 observations, as well as having more country-specific parameters to estimate.

surprising that the estimates are imprecise. There are some good theoretical and empirical reasons to exclude 1/GDP. It is not significant in the second stages of 2SLS estimates, and cannot readily be given a structural interpretation for the reasons explained in the main text. In the particular case of the CCE estimators, since 1/GDP is a function of a variable that is likely to be I(1) or I(2), introducing its cross-section mean with country-specific coefficients may be an exercise of doubtful validity. Nevertheless, for readers who are sceptical about the existence of structural relationships between macroeconomic outcomes and aid intensity, these results are likely to be a concern. Since it is hard to separate the effect of the problem from the effect of a smaller sample size, a satisfactory resolution may require longer spans of data, as in the wider cross-country literature on aid.

References

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Hodler, Roland and Raschky, Paul A. (2011). Foreign aid and enlightened leaders. Mimeo, January 11.

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

Kapetanios, George, Pesaran, M. Hashem and Yamagata, Takashi (2011). Panels with non-stationary multifactor error structures. *Journal of Econometrics*, 160(2), 326-348.

Kimbrough, Kent P. (1986). Foreign aid and optimal fiscal policy. *Canadian Journal of Economics*, 19(1), 35-61.

King, Robert G., Plosser, Charles I., Stock, James H. and Watson, Mark W. (1991). Stochastic trends and economic fluctuations. *American Economic Review*, 81(4), 819-40.

Kleibergen, Frank and Paap, Richard (2006). Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics*, 133(1), 97-126.

Kronmal, Richard A. (1993). Spurious correlation and the fallacy of the ratio standard revisited. *Journal of the Royal Statistical Society: Series A*, 156(3), 379-392

Marshall, Monty G., Gurr, Ted Robert and Jaggers, Keith (2013). Polity IV project: dataset users' manual. Center for Systemic Peace, mimeo, April 21.

Nunn, Nathan and Qian, Nancy (forthcoming). U.S. food aid and civil conflict. *American Economic Review*.

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

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

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

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

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

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

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

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

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

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

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

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

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

Table 1: Aid and household consumption

	(1) FE	(2) FE	(3) CCE	(4) CCE
Aid	0.233	0.140**	0.260	0.168
	(0.148)	(0.0667)	(0.220)	(0.112)
Lagged dep. variable		0.662***		0.564***
		(0.0285)		(0.0719)
Long-run effect aid		0.413**		0.386
Long-run effect SE		0.195		0.287
Countries	91	91	91	91
Observations	992	992	992	992
	(5) FE IV	(6) FE IV	(7) CCE IV	(8) CCE IV
Aid	0.578*	0.283*	0.722*	0.677
	(0.321)	(0.152)	(0.432)	(0.464)
Lagged dep. variable		0.654***		0.461***
		(0.0296)		(0.0855)
Long-run effect aid		0.820*		1.256
Long-run effect SE		0.427		0.946
First stage F-statistic	20.31	20.18	13.28	8.646
Underidentification	0.000	0.000	0.003	0.007
Countries	91	91	91	91
Observations	992	992	992	992

Note: Dependent variable is household consumption. All variables expressed as a % of GDP. Fixed effects (FE), fixed effects IV (FE IV), common correlated effects (CCE) and common correlated effects IV (CCE IV) results, three-year averaged data, 1975-2010. IV regressions carried out using `xtivreg2` for Stata (Schaffer, 2010). FE and FE IV regressions include time dummies, coefficients not reported. Country-specific coefficients on cross-section means in CCE and CCE IV regressions not reported. Heteroskedasticity-robust standard errors, clustered by country and corrected for degrees of freedom, in brackets. Standard errors (SE) for the long-run effects obtained via the delta method. $*$, $**$, and $***$ denote significance at 10, 5 and 1%, respectively. Underidentification shows the p-value of the Kleibergen and Paap (2006) LM test for underidentification.

Table 2: Aid and government consumption

	(1) FE	(2) FE	(3) CCE	(4) CCE
Aid	0.109	0.0174	0.0576	0.0571
	(0.0882)	(0.0372)	(0.0495)	(0.0518)
Lagged dep. variable		0.696***		0.275**
		(0.0416)		(0.115)
Long-run effect aid		0.0573		0.0788
Long-run effect SE		0.121		0.0767
Countries	91	91	91	91
Observations	992	992	992	992

	(5) FE IV	(6) FE IV	(7) CCE IV	(8) CCE IV
Aid	0.123	0.0488	-0.0154	-0.00676
	(0.173)	(0.0784)	(0.151)	(0.127)
Lagged dep. variable		0.691***		0.239**
		(0.0400)		(0.120)
Long-run effect aid		0.158		-0.00888
Long-run effect SE		0.252		0.166
First stage F-statistic	20.31	20.34	12.53	10.46
Underidentification	0.000	0.000	0.001	0.001
Countries	91	91	91	91
Observations	992	992	992	992

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

Table 3: Aid and total consumption

	(1) FE	(2) FE	(3) CCE	(4) CCE
Aid	0.342***	0.180**	0.418***	0.242**
	(0.110)	(0.0711)	(0.138)	(0.112)
Lagged dep. variable		0.594***		0.489***
		(0.0379)		(0.0523)
Long-run effect aid		0.444***		0.473**
Long-run effect SE		0.161		0.216
Countries	91	91	91	91
Observations	992	992	992	992

	(5) FE IV	(6) FE IV	(7) CCE IV	(8) CCE IV
Aid	0.701***	0.382**	0.661**	0.548**
	(0.238)	(0.146)	(0.269)	(0.259)
Lagged dep. variable		0.572***		0.367***
		(0.0432)		(0.0769)
Long-run effect aid		0.893***		0.865**
Long-run effect SE		0.306		0.396
First stage F-statistic	20.31	19.84	21.50	12.96
Underidentification	0.000	0.000	0.001	0.002
Countries	91	91	91	91
Observations	992	992	992	992

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

Table 4: Aid and gross capital formation

	(1) FE	(2) FE	(3) CCE	(4) CCE
Aid	0.199***	0.182***	0.162**	0.172***
	(0.0543)	(0.0336)	(0.0665)	(0.0550)
Lagged dep. variable		0.535***		0.360***
		(0.0424)		(0.0731)
Long-run effect aid		0.390***		0.269***
Long-run effect SE		0.0892		0.0881
Countries	91	91	91	91
Observations	992	992	992	992

	(5) FE IV	(6) FE IV	(7) CCE IV	(8) CCE IV
Aid	-0.0241	0.157	0.0472	0.148
	(0.179)	(0.115)	(0.269)	(0.236)
Lagged dep. variable		0.535***		0.221**
		(0.0419)		(0.103)
Long-run effect aid		0.338		0.190
Long-run effect SE		0.251		0.311
First stage F-statistic	20.31	20.44	14.89	15.51
Underidentification	0.000	0.000	0.000	0.000
Countries	91	91	91	91
Observations	992	992	992	992

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

Table 5: Aid and net imports

	(1) FE	(2) FE	(3) CCE	(4) CCE
Aid	0.542***	0.392***	0.556***	0.470***
	(0.0905)	(0.0716)	(0.109)	(0.107)
Lagged dep. variable		0.482***		0.308***
		(0.0336)		(0.0856)
Long-run effect aid		0.757***		0.680***
Long-run effect SE		0.129		0.158
Countries	91	91	91	91
Observations	992	992	992	992
	(5) FE IV	(6) FE IV	(7) CCE IV	(8) CCE IV
Aid	0.690***	0.583***	0.888***	0.884***
	(0.173)	(0.158)	(0.197)	(0.271)
Lagged dep. variable		0.462***		0.239**
		(0.0416)		(0.0986)
Long-run effect aid		1.082***		1.161***
Long-run effect SE		0.244		0.317
First stage F-statistic	20.31	20.56	20.69	16.50
Underidentification	0.000	0.000	0.000	0.000
Countries	91	91	91	91
Observations	992	992	992	992

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

Table 6: Aid and imports

	(1) FE	(2) FE	(3) CCE	(4) CCE
Aid	0.600***	0.447***	0.639***	0.521***
	(0.106)	(0.0615)	(0.124)	(0.0972)
Lagged dep. variable		0.613***		0.541***
		(0.0478)		(0.0748)
Long-run effect aid		1.156***		1.134***
Long-run effect SE		0.178		0.249
Countries	91	91	91	91
Observations	992	992	992	992
	(5) FE IV	(6) FE IV	(7) CCE IV	(8) CCE IV
Aid	0.241	0.464***	0.660**	0.547**
	(0.271)	(0.144)	(0.265)	(0.270)
Lagged dep. variable		0.612***		0.410***
		(0.0489)		(0.0967)
Long-run effect aid		1.197***		0.926*
Long-run effect SE		0.352		0.490
First stage F-statistic	20.31	22.70	20.13	15.39
Underidentification	0.000	0.000	0.000	0.001
Countries	91	91	91	91
Observations	992	992	992	992

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

Table 7: Aid and exports

	(1) FE	(2) FE	(3) CCE	(4) CCE
Aid	0.0584	0.101	0.140	0.152
	(0.125)	(0.0615)	(0.111)	(0.0944)
Lagged dep. variable		0.713***		0.476***
		(0.0374)		(0.0709)
Long-run effect aid		0.352		0.291
Long-run effect SE		0.219		0.193
Countries	91	91	91	91
Observations	992	992	992	992
	(5) FE IV	(6) FE IV	(7) CCE IV	(8) CCE IV
Aid	-0.450	-0.0244	-0.0458	0.104
	(0.304)	(0.141)	(0.385)	(0.343)
Lagged dep. variable		0.711***		0.422***
		(0.0377)		(0.0870)
Long-run effect aid		-0.0842		0.180
Long-run effect SE		0.485		0.598
First stage F-statistic	20.31	20.54	17.72	15.00
Underidentification	0.000	0.000	0.000	0.000
Countries	91	91	91	91
Observations	992	992	992	992

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

Table 8: Robustness checks

Row	Model	C	G	C+G	I	X	M	M-X
1	Static	0.722*	-0.0154	0.661**	0.0472	-0.0458	0.660**	0.888***
		(0.432)	(0.151)	(0.269)	(0.269)	(0.385)	(0.265)	(0.197)
	F	13.28	12.53	21.50	14.89	17.72	20.13	20.69
	Dynamic	1.256	-0.00888	0.865**	0.190	0.180	0.926*	1.161***
		(0.946)	(0.166)	(0.396)	(0.311)	(0.598)	(0.490)	(0.317)
	F	8.646	10.46	12.96	15.51	15.00	15.39	16.50
2	Static	0.506	0.283	0.853***	0.0225	-0.303	0.806***	1.008***
		(0.316)	(0.287)	(0.300)	(0.253)	(0.421)	(0.270)	(0.253)
	F	9.800	8.484	11.85	15.57	8.955	12.10	12.68
	Dynamic	0.804	0.455	1.016**	0.376	-0.104	1.684**	1.512***
		(0.740)	(0.451)	(0.405)	(0.299)	(0.601)	(0.678)	(0.533)
	F	4.970	5.218	7.211	12.84	7.403	9.757	9.094
3	Static	0.677	-0.0417	0.665**	0.0537	-0.0690	0.492*	0.918***
		(0.419)	(0.119)	(0.260)	(0.275)	(0.336)	(0.252)	(0.202)
	F	14.85	13.06	21.00	15.47	13.21	14.14	15.06
	Dynamic	1.326	-0.00744	0.893**	0.217	0.0701	0.771*	1.204***
		(0.887)	(0.102)	(0.364)	(0.319)	(0.519)	(0.430)	(0.309)
	F	10.80	10.88	16.88	17.17	9.472	18.15	18.48
4	Static	0.368	-0.261	0.480*	0.278*	-0.104	0.734**	0.951*
		(0.301)	(0.168)	(0.251)	(0.158)	(0.334)	(0.348)	(0.480)
	F	22.10	16.45	26.90	17.80	21.22	22.88	20.48
	Dynamic	0.470	-0.366	0.259	0.291*	-0.0576	0.946	0.917*
		(0.589)	(0.278)	(0.293)	(0.166)	(0.410)	(0.765)	(0.491)
	F	14.37	14.06	25.72	15.56	15.67	17.38	24.31
5	Static	1.058	-0.313	0.869*	0.520	-0.191	1.023**	1.479**
		(0.839)	(0.245)	(0.451)	(0.578)	(0.808)	(0.407)	(0.662)
	F	3.339	3.058	3.253	3.412	2.317	2.506	2.866
	Dynamic	2.478	-0.214	1.205*	0.345	-0.0963	1.804	1.369*
		(2.008)	(0.239)	(0.647)	(0.767)	(1.369)	(1.116)	(0.718)
	F	3.053	1.919	2.496	2.128	1.837	2.169	1.695
6	Static	0.710	-0.150	0.656**	0.300	-0.302	0.846***	1.194***
		(0.529)	(0.175)	(0.287)	(0.289)	(0.626)	(0.270)	(0.435)
	F	16.27	16.19	24.72	13.56	16.40	21.85	14.20
	Dynamic	1.297	-0.178	0.662**	0.301	-0.363	1.223**	1.080**
		(1.202)	(0.215)	(0.328)	(0.346)	(0.970)	(0.577)	(0.443)
	F	9.159	15.45	19.00	17.66	14.21	19.14	16.22

Continued on next page

Table 8 – continued from previous page

Row	Model	C	G	C+G	I	X	M	M-X
7	Static	0.561	-0.000403	0.493*	0.338**	0.0556	0.725***	0.892***
		(0.417)	(0.152)	(0.275)	(0.135)	(0.380)	(0.266)	(0.267)
	F	11.60	17.00	20.14	17.97	24.45	23.34	22.08
	Dynamic	0.721	0.0266	0.497	0.374**	0.284	0.877**	0.921***
		(0.581)	(0.189)	(0.333)	(0.153)	(0.596)	(0.340)	(0.307)
	F	8.722	14.11	13.63	18.62	17.46	18.13	19.54
8	Static	0.701	0.0336	0.613**	-0.00339	-0.0674	0.697**	0.896***
		(0.435)	(0.169)	(0.249)	(0.279)	(0.358)	(0.270)	(0.226)
	F	12.80	11.47	20.10	13.01	18.10	18.02	18.04
	Dynamic	1.324	0.00231	0.710**	-0.157	0.125	1.157**	0.852***
		(1.085)	(0.208)	(0.340)	(0.305)	(0.604)	(0.517)	(0.287)
	F	8.021	9.348	12.40	12.79	13.93	13.64	15.97
9	Static	0.935	0.0111	0.620*	0.0314	0.232	0.977**	1.220***
		(0.609)	(0.0831)	(0.324)	(0.294)	(0.499)	(0.443)	(0.444)
	F	11.04	7.304	10.87	7.557	11.31	8.470	9.578
	Dynamic	1.619	-0.0259	0.434	-0.00542	0.418	1.576**	1.223*
		(1.452)	(0.0955)	(0.469)	(0.311)	(0.888)	(0.759)	(0.642)
	F	5.833	5.895	5.942	7.245	7.128	5.906	8.179
10	Static	1.128*	0.0233	0.657*	-0.147	0.132	0.850	1.187**
		(0.652)	(0.0930)	(0.364)	(0.286)	(0.562)	(0.553)	(0.546)
	F	8.430	5.959	7.629	6.024	8.340	7.191	7.065
	Dynamic	1.739	-0.00268	0.532	-0.201	0.115	0.764	0.991
		(1.379)	(0.109)	(0.513)	(0.372)	(0.882)	(0.828)	(0.682)
	F	5.836	4.421	4.854	5.061	4.969	4.557	4.872
11	Static	0.773*	-0.125	0.686**	0.0699	-0.0748	0.670**	0.902***
		(0.454)	(0.164)	(0.281)	(0.280)	(0.345)	(0.261)	(0.234)
	F	12.81	12.10	22.04	14.42	19.59	22.02	19.56
	Dynamic	1.146	-0.161	0.849**	0.289	0.128	0.832*	1.230***
		(0.847)	(0.195)	(0.419)	(0.381)	(0.467)	(0.484)	(0.380)
	F	8.967	10.98	13.44	13.23	21.88	15.97	14.81
12	Static	0.754	-0.0310	0.666**	-0.00820	-0.0300	0.658**	0.887***
		(0.478)	(0.133)	(0.279)	(0.284)	(0.316)	(0.266)	(0.213)
	F	11.96	11.62	22.40	14.53	21.51	22.07	21.23
	Dynamic	1.301	-0.0518	0.865**	0.113	0.232	0.824	1.132***
		(1.024)	(0.150)	(0.420)	(0.329)	(0.507)	(0.516)	(0.343)
	F	7.770	9.986	12.89	15.24	19.80	15.39	16.94

Continued on next page

Table 8 – continued from previous page

Row	Model	C	G	C+G	I	X	M	M-X
13	Static	0.737	0.000995	0.583**	-0.0317	-0.0541	0.695**	0.894***
		(0.488)	(0.143)	(0.245)	(0.293)	(0.306)	(0.275)	(0.237)
	F	11.53	10.53	22.28	12.79	22.16	19.61	19.01
	Dynamic	1.382	-0.0439	0.679*	-0.191	0.167	0.970*	0.817***
		(1.202)	(0.180)	(0.356)	(0.301)	(0.521)	(0.584)	(0.277)
	F	7.205	8.794	12.63	14.51	18.63	14.04	17.66
14	Static	0.999	0.0231	0.606*	0.0126	0.272	0.995**	1.251***
		(0.673)	(0.0910)	(0.335)	(0.325)	(0.494)	(0.484)	(0.468)
	F	9.738	6.563	10.76	6.773	12.93	7.901	8.738
	Dynamic	1.726	-0.0141	0.390	-0.0151	0.459	1.500*	1.201*
		(1.619)	(0.100)	(0.509)	(0.318)	(0.881)	(0.795)	(0.654)
	F	5.292	5.418	6.057	7.407	8.470	5.304	8.169
15	Static	1.168*	-0.0164	1.010**	-0.0733	0.120	0.759	1.306***
		(0.647)	(0.151)	(0.443)	(0.413)	(0.477)	(0.498)	(0.437)
	F	8.545	12.02	10.69	9.289	10.02	9.974	13.26
	Dynamic	1.885	0.0103	1.311*	0.143	0.415	1.180*	1.680**
		(1.313)	(0.168)	(0.671)	(0.411)	(0.771)	(0.679)	(0.675)
	F	6.376	10.46	7.542	9.947	7.020	6.541	10.47

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 ("dynamic") and without ("static") a lagged dependent variable. All variables expressed as a % of GDP. CCE IV estimation on three-year averaged data (1975-2010) using an instrument based on initial shares in donor budgets calculated over the period 1960-74, unless reported otherwise below. 991 observations from 91 countries, unless reported otherwise below. Heteroskedasticity-robust standard errors, clustered by country and corrected for degrees of freedom, in brackets. Standard errors (SE) for the long-run effects obtained via the delta method. $*$, $**$, and $***$ denote significance at 10, 5 and 1%, respectively. F shows the first-stage F-statistic.

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

Row 2 includes small countries (1148 observations from 111 countries).

Row 3 constructs the endogenous aid variable in the second stage and the initial shares in donor budgets on which the instrument is based after discarding all aid data in the years before a recipient country's independence (925 observations from 83 countries).

Row 4 replaces aid and the instrument by their first lag; sample starting with the period 1978-80 (929 observations from 91 countries).

Row 5 includes both aid and its first lag, instrumented by the current and one period lagged values of the synthetic instrument; sample starting with the period 1978-80 (868 observations from 82 countries). F is the Kleibergen-Paap Wald rk F statistic.

Row 6 replaces aid and the instrument by the unweighted average of its current and one period lagged values; sample starting with the period 1978-80 (923 observations from 91 countries).

Row 7 uses the final year values in each period for the dependent variables instead of the three-

year averages (964 observations from 89 countries).
Row 8 excludes the first period (1975-77) from estimation (924 observations from 91 countries).
Row 9 excludes the first two periods (1975-77 and 1978-80) from estimation (849 observations from 91 countries).
Row 10 excludes the first three periods (1975-77, 1978-80 and 1981-83) from estimation (771 observations from 91 countries).
Row 11 uses an instrument based on initial shares calculated over the period 1960-71 and a sample that starts with the period 1972-74 (1050 observations from 90 countries).
Row 12 uses an instrument based on initial shares calculated over the period 1960-71 and a sample that starts with the period 1975-77 (982 observations from 90 countries).
Row 13 uses an instrument based on initial shares calculated over the period 1960-71 and a sample that starts with the period 1978-80 (914 observations from 90 countries).
Row 14 uses an instrument based on initial shares calculated over the period 1960-71 and a sample that starts with the period 1981-83 (839 observations from 90 countries).
Row 15 uses an instrument based on the largest ten donors only (992 observations from 91 countries).

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

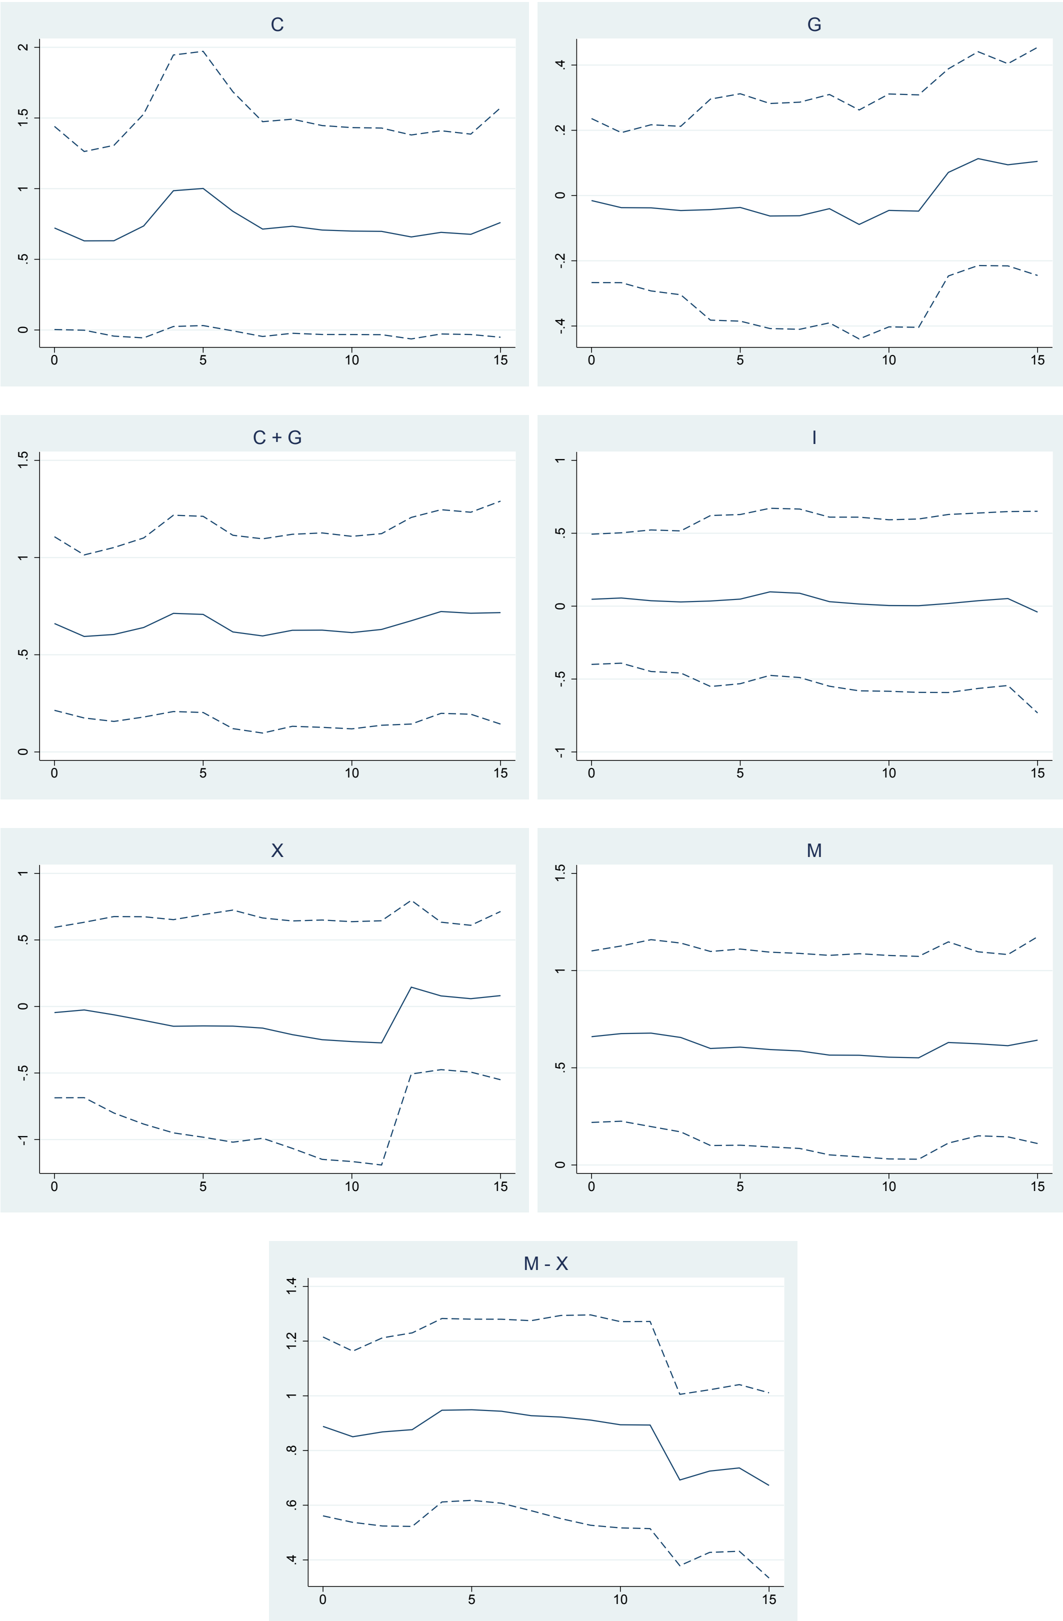

Note: graphs show how the estimated effects of aid (solid line) on household consumption (C), government consumption (G), total consumption ($C + G$), gross capital formation (I), exports (X), imports (M) and net imports ($M - X$) change when progressively dropping the countries with the largest percentage 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 2: Dropping countries with the largest relative fall in real GDP (dynamic model)

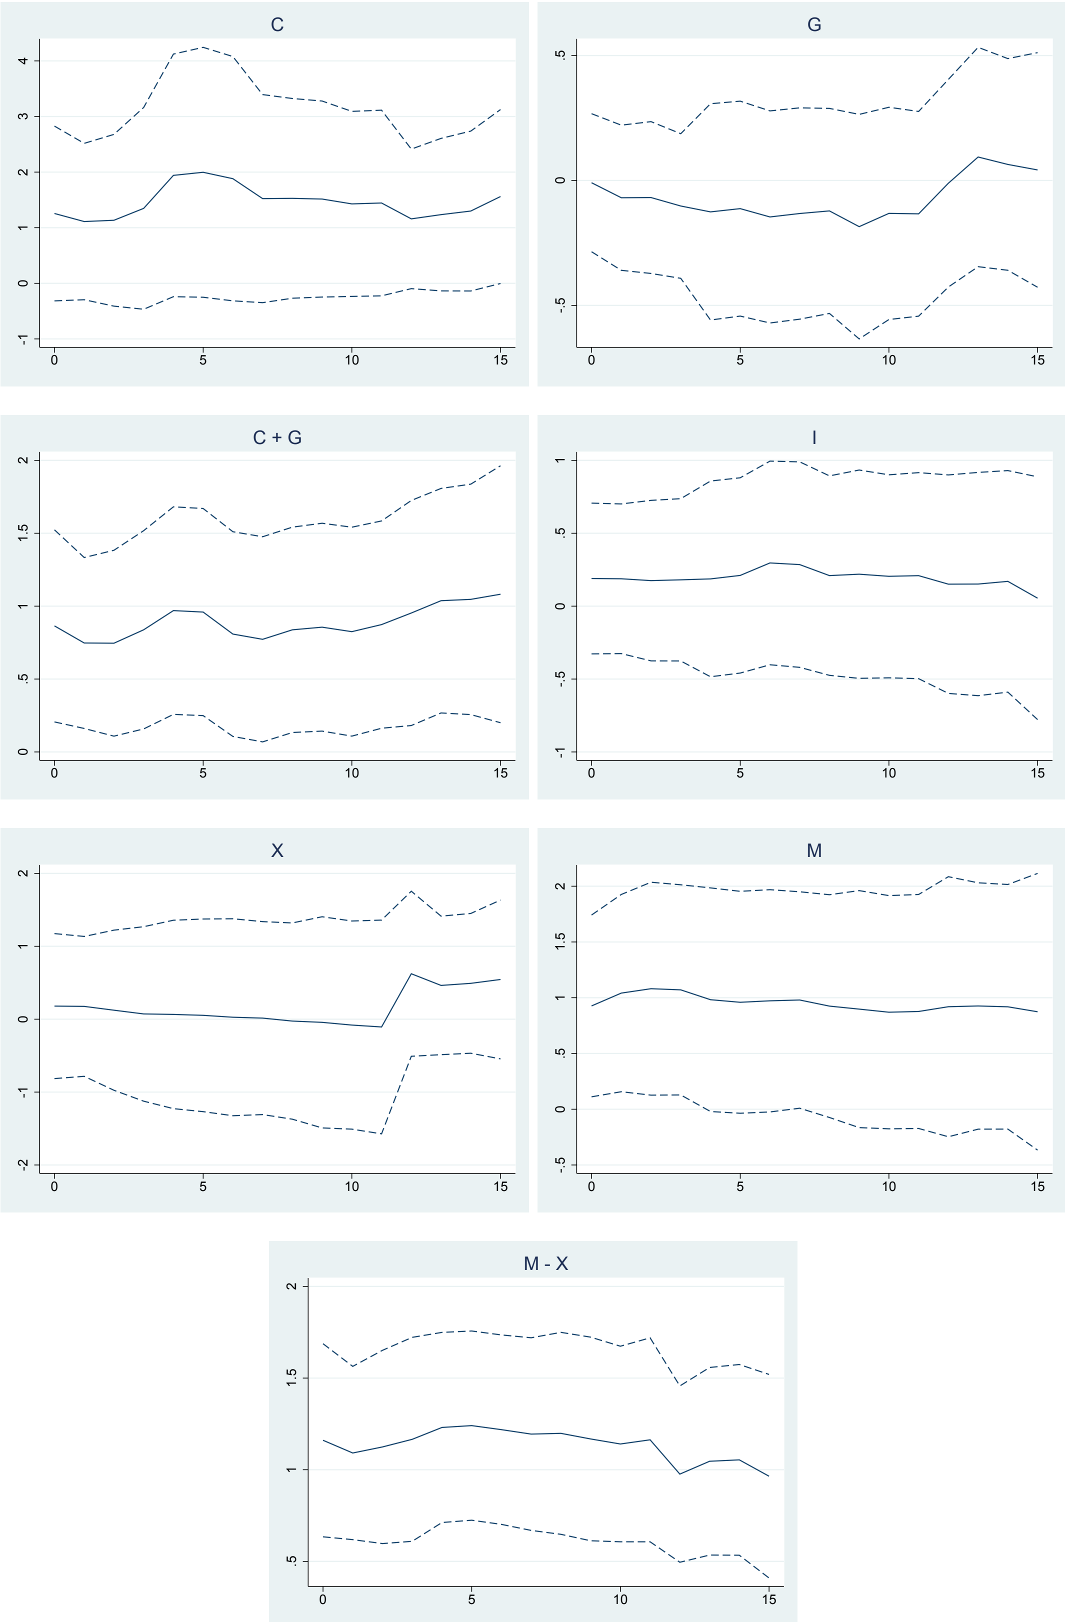

Note: see note Figure 1. These graphs show the long-run effects of aid (solid line) with 90% confidence intervals (dashed lines) when progressively dropping the countries with the largest percentage declines in real GDP, based on CCE IV estimation of a model that includes a lagged dependent variable.