

Foreign Aid and Economic Growth in Developing Countries: An Instrumental Variables Approach *

Arvind Magesan †

June 24, 2015

Abstract

There is little consensus on the capacity for foreign aid to cause economic growth in developing countries. This is due in large part to the fact that foreign aid recipients are selected by donors, confounding identification. This paper proposes an identification strategy that exploits exogenous variation in foreign aid receipts generated by participation in Human Rights Treaties at the UN to identify an average causal effect of aid on growth. Our approach is valid even if the effect of aid is heterogeneous across recipients for unobservable reasons. We find that an additional dollar of per capita aid causes the growth rate in a recipient country to increase by 8% over four years and 5% over a decade. The effect is explained almost entirely by an expansion of the service industry, accompanied by a large increase in household consumption, with no evidence that aid causes “Dutch disease,” as many fear it does. We conclude that aid increases growth by inducing a structural change in household demand for services.

JEL Classification: F35, F53, K33, O19, O43, P16

Keywords: Foreign Aid, Economic Growth, Services, International Organizations, Institutions, Human Rights.

*I thank Victor Aguirregabiria, Branko Boskovic, Pamela Campa, Sacha Kapoor, James Rockey, Michel Serafinelli, Atsuko Tanaka, Scott Taylor, Trevor Tombe, Jean-Francois Wen and Genet Zinabou for useful comments and suggestions that improved the paper. Miguel Olivo-Villabrille provided outstanding research assistance. Funding from SSHRC is gratefully acknowledged. All remaining errors are my own.

†Mailing Address: Department of Economics University of Calgary 2500 University Drive Calgary, Alberta T2N-1N7 Canada, Phone: 403-220-5276, e-mail: anmagesa@ucalgary.ca

1 Introduction

Foreign aid constitutes a major global flow of economic resources, and is a primary development tool of the richer countries of the world. In 2014, Official Development Assistance from the OECD alone topped \$135 billion US, at a time when many donors are themselves still yet to fully recover from the effects of the Great Recession. Yet precious little is known about the causal effect of foreign aid on the economic outcome in recipient countries that donors arguably care about most: growth. This is not for lack of effort; an active and voluminous literature aiming to assess the effectiveness of aid for economic growth has amassed over the past few decades.¹ The literature is rife with discordant results and as such, has failed to reach consensus on either the short run or long run effects of aid, nor the channel through which aid effects growth, be it positively or negatively. This lack of academic consensus also has negative consequences for the policy debate on how much aid to give and to whom. In the absence of convincing evidence on aid effectiveness, policymaking is left to the vagaries of ideological currents in donor countries.

A major obstacle to identification of a causal effect of aid on growth is a lack of powerful and plausibly exogenous variation in aid. In recent work, Clemens et al. [2012] blame the “poor quality instrumental variables that have pervaded this literature” for the lack of consensus,² arguing that the IV’s that have been employed in the literature are not obviously powerful or valid.³ Moreover, given the almost certainly large heterogeneity in the effect of aid on growth,⁴ each *valid* instrument identifies a different average causal effect under the additional assumption that the direction of response of aid to the instrument is the same for all recipients [Imbens and Angrist, 1994]. In most cases it is not obvious that the instruments that have been employed in the literature pass this

¹See Clemens et al. [2012] for a detailed discussion. Burnside and Dollar [2000] and Rajan and Subramanian [2008] are arguably the most influential papers within this literature.

²Clemens et al. [2012] take a major step toward fostering consensus on the question of aid effectiveness, by showing that the most influential studies in the field all predict a similar moderate effect of aid on short run economic growth once several key aspects of the research design are fixed across the studies. Though this finding is insightful and useful, exogenous variation in aid is absent in their identification strategy. Given the very nature of the growth-aid regression, any coefficient generated without plausibly exogenous variation in aid will be questioned on grounds of causality.

³Perhaps the most influential study of the effect of aid on growth using an IV approach is that of Rajan and Subramanian [2008]. As Clemens et al. [2012] point out, the instruments used by Rajan and Subramanian [2008] have virtually no power once population is included in the growth regression, suggesting that most of the identification power in their instruments comes from differences across recipients in population. This is worrisome because population is likely correlated with many other recipient specific variables. Moreover, the instruments employed in Rajan and Subramanian [2008] do not vary over time and are thus only useful in a cross-country regression, in which case the researcher must be concerned about permanent unobserved heterogeneity, as country fixed effects can not be differenced away.

⁴Indeed there are several reasons to expect aid to have a heterogenous effect on growth. Differences across countries in governance, technical ability, infrastructure and absorption ability are just a few.

monotonicity test, and thus identify parameters that lack clear interpretation.

Taking heterogenous effects of aid as our point of departure, we propose a powerful, plausibly exogenous and monotone instrumental variable that identifies an average causal effect - a Local Average Treatment Effect (LATE) - of aid on short run and long run growth, as well as the channel through which aid affects growth in recipient countries. In particular, we exploit exogenous variation in aid receipts per capita generated by participation in human rights treaties (HRT) at the United Nations to study whether and why aid per capita causes economic growth. The UN HRT system is the embodiment of international human rights law, and represents the norms of behavior that Western Liberal Democracies (such as the majority of OECD members) have aimed to propagate through the less developed world since the end of World War II. As such, the ratification of HRTs is often rewarded with economic aid from donors.⁵ HRT participation is a powerful instrument in our setting - it has an immediate and large impact on aid in our sample, and as we illustrate, the relationship is not spurious.⁶ We also provide ample evidence that the instrument satisfies the necessary exclusion restriction, which in this context requires that HRT participation does not cause economic growth through any channel other than aid. Additionally, conditional on observable characteristics, it is highly unlikely that countries that ratify treaties are punished with less aid for participating in HRT. To confirm this, we test for monotonicity of aid in our instrument and show that indeed, HRT participation satisfies the condition of monotonicity that is so crucial for identification in the case of heterogenous treatment effects.

In our favored specification, we find that an additional dollar per capita of aid generates an 8% larger four year growth rate and a 5% larger ten year growth rate than the country would have experienced without that dollar per capita. As the economics literature on aid effectiveness has found modest to no effects of aid in either the short run or the long run, our estimates are likely the largest estimated effects of aid to date. From a policy perspective, the results suggest that aid is a relatively cheap and useful tool for generating growth in recipient countries. To put the result in perspective, Haiti is a country of just over 10 million people. A one dollar per capita increase in aid to Haiti would amount to an almost negligible increase in the total aid budget of the OECD countries, but would have an economically significant effect on short run and long run growth.

Having established a causal relationship between aid and growth, we then use the data to

⁵That HRT participation is rewarded with economic aid is not a novel claim. The relationship between the two variables has been long understood in the political science and international law literature (see Hathaway [2007] for example) and more recently established empirically in Magesan [2013].

⁶In particular, neither FDI, trade or remittances respond to HRT participation. We also consider other falsification tests for our first stage, and show that the aid-HRT relationship is robust.

identify the channel through which per-capita aid affects growth. In doing so we directly address the question of whether aid is responsible for fostering “Dutch disease” in recipient countries, as many fear.⁷ Perhaps most prominently, Rajan and Subramanian [2011] find that “exportable industries” decline relative to others as a result of aid inflows. However the shift away from manufacturing to services (“postindustrialization”) is a well established pattern of development [Soubbotina, 2000]. Depending on the level of development, services are either a normal or luxury good, and increased household income causes a structural shift in demand for services. Thus, aid may help precipitate the shift to services by increasing household income. We first examine changes in value added in manufacturing, agriculture and services in response to aid. We find that the entire estimated increase in economic growth in response to aid is explained by an increase in per-capita value added in the services industry. We find no evidence of a change in value added in either manufacturing or agriculture. Examining the effect of aid on the expenditure components of GDP growth, we also find that aid causes a significant increase in household consumption. Taken together, the results suggest that aid causes growth by inducing a structural change in the demand for services in recipient countries. The results also suggest that a relative decline in manufacturing in response to large aid inflows should not necessarily be viewed in a negative light.

We study the causal effect of aid receipts per capita as opposed to the treatment variable typically considered in the literature, aid as a fraction of GDP. The motivation behind dividing aid by GDP in studies of aid effectiveness is to normalize aid with respect to country size. One million dollars in additional aid should have a much different effect on growth in a small country than in a large country. From an identification and interpretation perspective, however, normalizing by population is preferable for three reasons. First, changes in aid per capita are more likely to reflect changes in aid itself than changes in aid as a fraction of GDP. Changes in the latter may simply reflect positive or negative shocks to a country’s economy, which cause the denominator of the fraction to increase or decrease while aid is unchanged. This is especially true in developing countries, where GDP is quite volatile.⁸ Population is far less volatile, and changes from one year to the next in aid per capita are more likely to reflect changes in aid dollars. The second reason we prefer aid per capita is that, given the likely heterogeneity in the effect of aid on growth, it is much more amenable to a research design that relies on instrumental variables. We know that our

⁷It has been argued in the literature that aid plays a role in the decline of manufacturing in recipient countries. Aid is often targeted towards non-tradeable sectors, which causes an increase in domestic input prices (including wages) and makes tradeable sectors such as manufacturing less competitive (holding world prices fixed). Spending the higher wages can also lead to real exchange rate appreciation, further exacerbating the problem.

⁸The coefficient of variation on GDP is greater than 1 in our sample, making it a high variance variable.

instrument monotonically shifts aid, but there is no obvious relationship between HRT participation and GDP, making the relationship between the instrument and the ratio of aid to GDP unclear. This is not the case with aid per capita, as again, population is highly persistent and not likely correlated with participation in human rights treaties. Finally, we also prefer our measure because it is easier to make clear statements about aid policy based on a causal parameter identified off variation in aid itself as opposed to variation in aid as a fraction of GDP.

While our estimate is “local,” we argue that we are identifying a causal parameter. To this end, we systematically exhaust the possible ways in which our identification conditions may be violated. While it is difficult to imagine channels through which ratification of human rights treaties dealing with issues such as torture and statutory limits on war crimes may affect growth, HRT participation is a *choice*, and it is potentially not randomly assigned even conditional on country fixed effects, time effects and a host of control variables. We thus address the possibility of selection on the instrument as a function of unobservable time varying characteristics using two separate methods. We first use the method suggested by Altonji et al. [2005] who show that one can identify selection on *observable* time varying characteristics and use this to quantify the amount of selection on *unobservables* that would be required to nullify the coefficient estimate of interest. We show that if anything, we are *underestimating* the effect of aid on growth. This is because the selection on observables is minimal and *negative*, meaning that the selection on unobservables would have to be implausibly large and work in the opposite direction of the selection on observables in order for the population coefficient to be zero. As a second test for selection, we consider a “trimmed” sample, where observations with apparently non-random HRT participation are excluded from the sample.⁹ We find that in spite of the inevitable loss of statistical power that accompanies considering a smaller sample, our estimated effects are even larger and highly statistically significant. This is consistent with our finding using Altonji et al. [2005], as it suggests that if anything we are underestimating the effect of aid.

We also check for other channels through which HRT participation may indirectly affect economic growth. In particular, we consider the possibility that HRT participation occurs at the same time as other more economically meaningful reforms which directly cause growth. We show that even conditioning on economic reforms the reduced form relationship between HRT participation and economic growth remains. Another channel we consider is unobservable region-specific time trends.

⁹This was, to the best of our knowledge, first proposed in Lim et al. [2015]. The idea is that HRT participation appears to be selected on observables for individuals whose covariates predict participation extremely well, and is thus likely to be selected on unobservables too.

Regions grow together and often reform at the same time [Buera et al., 2011]. We show that our results are robust to the inclusion of these trends. As a more general test of the exclusion restriction, we also consider a “placebo” test. We show that in a subsample where HRT participation does not affect aid, there is no reduced form relationship between economic growth and HRT participation, consistent with aid being the only channel through which HRT participation affects growth.

While we interpret our findings as causal, we reiterate that the causal effects we identify are “local” in the usual sense. We are identifying the effect of aid on growth for the subpopulation of countries whose aid receipts respond to HRT participation - the *compliers* in the nomenclature of Imbens and Angrist [1994]. There are countries who receive the same amount of aid regardless of their participation (*always takers*), and there are countries who will not receive aid regardless of their participation (*never takers*). We can not make any claims about the effect of aid on growth in these countries. But as long as there are no *defiers* - countries who are punished with *less* aid for participation - the existence of whom we find highly implausible, we are identifying a very useful causal parameter. Moreover, given that we identify a large positive effect for the set of countries who participate in HRT at the UN, our findings suggest an important role for international agreements such as Human Rights treaties: a way of effectively allocating foreign aid.

2 Data and Background

Much of the data in the study are standard and publicly available, and detail on data sources is relegated to the appendix. We focus our attention here on the institutional detail that underlies our first stage relationship. In particular, we discuss the construction of our instrument - HRT participation at the UN- and illustrate that the variable contains the variation necessary for identification, and then discuss and illustrate the power of the first stage. We then argue and provide empirical evidence that our instrument satisfies the necessary monotonicity condition. Finally, we discuss the variables used in the study and present some basic summary statistics.

2.1 The UN HRT System and the Treaty Participation Measure

The United Nations HRT system is a byproduct of the horrors of two World Wars. On December 10, 1948 the United Nations General Assembly adopted the Universal Declaration of Human Rights (UDHR), which has been translated into more than 300 languages and dialects.¹⁰ Since the inception

¹⁰The UDHR was recognized by the Guinness Book of World Records for the UDHR being the world’s most translated document. See <http://www.ohchr.org/EN/UDHR/Pages/WorldRecord.aspx>, last retrieved Nov 28, 2014.

of the UDHR, which is viewed as a general declaration and definition of the rights and freedoms that all human beings are entitled to, a large number of treaties dealing with a variety of specific human rights issues have come into existence. The treaties are designed by UN member countries, and then open to signature (non-binding) and ratification (binding) In constructing our measures of treaty participation, we consider all UN treaties with a human rights element.¹¹ While other international institutions such as the ILO also house treaties that deal with human rights issues, we choose to limit ourselves to UN treaties to minimize heterogeneity, and because we believe that UN treaties are more likely to satisfy our exclusion restriction than say ILO treaties that deal with labor issues. Our primary measure of HRT participation of a country at a given time is simply the cumulative sum of ratified treaties. Formally, indexing countries by i , the year by t , and treaties (in chronological order) by r , our measure is given by:

$$T_{it} = \sum_r h_{irt} \tag{1}$$

where h_{irt} is a binary variable taking the value 1 if country i has ratified treaty r at any year up to t and 0 otherwise. As global human rights standards change over time, one could be concerned about using an absolute measure such as this, but our regressions contain time fixed effects and region specific trends, and we are not worried about this issue. But as a check, we ensure that our results are robust to a relative measure of treaty participation, the difference between a country's own cumulative sum of ratifications and the average of other countries' cumulative ratifications.¹²

For the purposes of identification, we would like to check that rates of HRT participation can shift significantly across countries in the same region at the same time, and are not just due to level differences over time across the countries. As our regressions will control for time and in some cases region-time effects, in figures 1 - 3 we plot the *relative* HRT participation of countries over time for sets of countries in three different regions, South America, Sub-Saharan Africa and Central Asia. As we see, in all three cases the relative rankings of the countries in terms of participation changes over time. Not all differences in HRT participation across countries in the same region at the same time can be explained by permanent differences, key to our identification approach. We now discuss our identification strategy in much more detail.

¹¹ For a comprehensive and detailed study of human rights treaties see Elliot [2011]. The list of treaties we consider are listed in the appendix.

¹²Buera et al. [2011] use a similar distance weighted measure in their study of country adoption of market-oriented policy in response to learning about other countries' experience. See Magesan [2013] for an application to the decision to ratify HRT.

Figure 1 South America

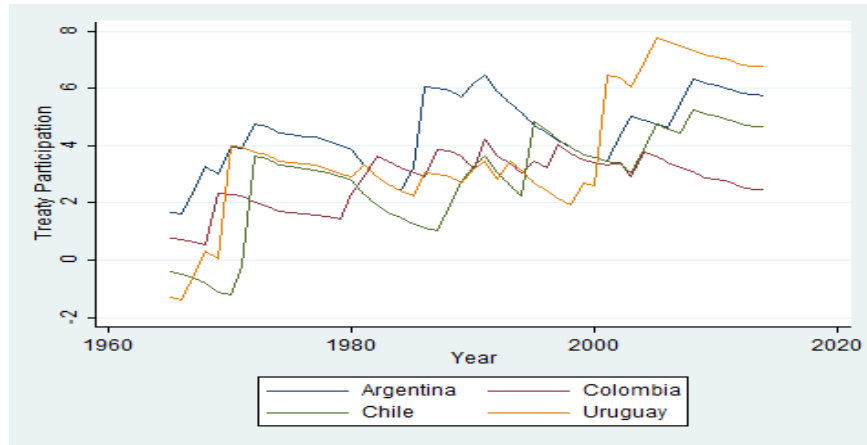


Figure 2 Sub Saharan Africa





Figure 3 Central Asia



2.2 The First Stage

We exploit the relationship between participation in UN Human Rights Treaties and foreign aid receipts to identify the effect of aid on growth. For many years the ratification of HRT posed a puzzle to researchers in political science and international relations. Unlike treaties that delineate rules of the game in other spheres of international interaction such as trade and the environment, human rights treaties don't solve a collective action problem, and countries don't mutually benefit from each other's participation. Human rights treaties simply constrain the behavior of governments with respect to their own populations, which does not confer a direct benefit on others. Even if leaders of countries did care about how other leaders treat their own populations, they do not need to ratify the treaty in order to reap the benefits of another's ratification. To resolve the puzzle, the observed ratification of HRTs has been rationalized as a process where material benefit, in particular foreign aid, is conferred upon ratifiers as a reward for ratification in exchange for the sovereignty costs that are a consequence of ratification (see Hathaway [2007] for a discussion). Donor countries, in particular, wealthy liberal democracies (most members of the OECD), reward participation in HRT either implicitly or explicitly with economic aid, because widespread ratification implies a propagation of norms of behavior and consolidation of international institutions that donors value.¹³

Ultimately the data will reveal the strength of the relationship between HRT participation and foreign aid. The relationship was first empirically tested in Magesan [2013] who found an economically large and robust effect of treaty participation on a country's subsequent foreign aid receipts. In figure 4 we present a plot of aid per capita and treaty participation, both purged of country and time effects. The figure reveals a strong positive relationship.¹⁴


2.2.1 Monotonicity

Given the likely heterogeneity across countries in the effect of aid on growth, identification of an average treatment effect requires monotonicity of the aid variable with respect to the instrument [Imbens and Angrist, 1994]. It is difficult to imagine a case where countries are punished with less aid, all else constant, for participating in HRT at the UN. In other words, we cannot imagine a scenario that would lead to there being defiers in our sample. We argue, however, that previous

¹³ Even if individual donors don't *directly* condition their own bilateral aid decisions on recipient HRT participation, many donors take the lead on aid decisions from International Governmental Organizations (IGO's), and there is evidence that IGO's do condition decisions on human rights commitments (see Oestreich [2007] for a discussion).

¹⁴We checked for non-linearities in the first stage relationship. In particular we included a squared HRT participation in the regression. The coefficient on the non linear term is not significant at conventional levels.

Figure 4 The First Stage



studies employing instrumental variables do not obviously satisfy this property due both to the nature of the aid variable as well as the instrument. By and large, the treatment variable considered in the existing literature is aid as a fraction of GDP. Having GDP in the denominator introduces an additional complication as one must argue that the fraction as a unit is monotone in the instrument, instead of just having to argue that aid is monotone.

The case is difficult to make for many of the instruments that have been used. Some studies have used lagged aid as an instrument for current aid.¹⁵ Yet it is not obvious whether a high level of aid receipts in a previous year should cause an increase or decrease in current aid. While in some cases aid is persistent, donors may find it in their interest to reduce aid to countries that have been well-funded to those previously less well-funded as a way of spreading wealth. The fact that past aid may affect current GDP further complicates the relationship between instrument and treatment. Other studies (i.e., Burnside and Dollar [2000]) use population of the recipient (and interactions of population with measures of policy) and arms imports as a fraction of total imports among other variables representing the strategic importance of a recipient. The issue here is that there are many motives for giving aid in addition to strategic value, in particular, recipient need and behavior. Thus, while countries with a smaller population may be easier to manipulate for strategic purposes, they also have a greater need. Similarly, though arms imports may represent strategic value, it

¹⁵See for example Hansen and Tarp [2001]

may also represent bad behavior. Similar arguments apply to Rajan and Subramanian [2008] who, in the context of a cross country regression, use a vector of common characteristics between donor and recipient (for example history of a colonial relationship), as well as the interaction between this vector and the ratio of donor to recipient population to measure strategic value and influence that the donor can exert over the donor. While Rajan and Subramanian [2008] argue that the smaller the recipient is the easier it is to influence and thus the more strategically useful is aid, one could counter that larger recipients require more aid than small recipients for any non strategic purpose, and the net effect of population on aid is not obvious.¹⁶ Moreover, as is well known since Acemoglu et al. [2001], a country’s colonial history affects GDP through other (non-monotone) channels, further obfuscating the relationship between the instrument and treatment.¹⁷

Given the likely non-monotonicity in the instruments employed in past studies, it is perhaps not so surprising that the different studies yield different results. If the effect of aid on growth is indeed heterogenous, the resulting estimate is a combination of the average effects for *compliers* and *defiers*, where the set of countries in these groups are instrument specific. It is difficult to imagine in the context of our instrument that there are defiers in the population of aid recipients, as participation in HRT is encouraged by the UN, and not likely to be viewed as “bad behavior” by the type of aid donors in our sample - the OECD countries. As such we believe our instrument satisfies the necessary monotonicity condition. Given that our treatment is continuous, we can informally test whether our instrument is monotone. As Angrist and Imbens [1995] point out, if the (variable) treatment is monotone in the instrument, then the CDF of the treatment when the instrument is “switched on” should lie under the CDF of the treatment when the instrument is “switched off.” That is, the distribution of the treatment when the instrument is switched on should first-order stochastically dominate the distribution of the treatment when the instrument is switched off. Here our instrument is not binary, so we compare the CDFs of aid (purged of controls, time effects and country fixed effects) for above and below median HRT participation and for above 75th percentile and below 25th percentile HRT participation in figures 5 and 6. Both figures clearly suggest that the distribution of aid for ratifiers first order stochastically dominates the distribution of aid for non-ratifiers.

¹⁶One could even make the case that larger recipients require more aid for strategic reasons as well.

¹⁷A recent study by Galiani et al. [2014] employs a potentially powerful and monotone instrument - crossing of the IDA’s threshold for aid eligibility - to study the causal effect of aid on growth. However only 35 countries cross the threshold during their sample, and they do so (by definition) at a very specific point in their economic development, and external validity with respect to other countries or even other points of development is not obvious.

Figure 5 CDFs of Aid for Above and Below Median HRT Participation





Figure 6 CDFs of Aid for Top and Bottom 25 percent HRT Participation



2.3 Summary Statistics

In Tables I and II we present some basic summary statistics. Table I presents statistics on our outcomes of interest - growth in GDP and growth in the variables that comprise GDP. All growth variables are in per capita terms. Following the standards in the literature, we consider two horizons for growth, four year and ten year, and in each case we consider both overlapping and non-overlapping windows of data.¹⁸Not surprisingly given the countries that comprise the sample, all growth variables display high variance - the standard deviation is many times larger than the mean in most cases. Also of note is the relatively slow growth of value added in agriculture. This is consistent with conventional wisdom on development, which suggests that development is characterized initially by a transition from agriculture to manufacturing, and then a subsequent transition to services.

In Table II we present the summary statistics for our causal variables of interest, our instrument and controls and other variables we use to check robustness of our results. The foreign aid flows we consider are disbursements of Overseas Development Assistance (ODA) from the Development Assistance Committee (DAC) of the OECD. This measure is standard and commonly used in the literature. However as Clemens et al. [2012] point out, a primary reason why many studies do not find a significant effect of aid on growth could be due to the fact that much of the aid that comprises total flows is not geared towards growth, at least in the short run. Aid for vaccines, education, or humanitarian reasons, for example, should not be expected to affect growth in the near future. Aid for immediate infrastructure building, by contrast, should. As such, Clemens et al. [2012] divide aid into three mutually exclusive categories: “early impact” aid, “late impact” aid and humanitarian aid, and suggest that only “early impact” aid should be considered when studying the effects of aid on growth (see Clemens et al. [2012] for a detailed description of how “early impact” aid is calculated). As such, we consider all our main specifications using Clemens et al. [2012] definition of “early impact” aid. The average country in our sample receives almost 33 dollars per capita of aid, 12 of which are defined as “early impact” aid.

While HRT participation likely has no direct effect on GDP growth, we must be concerned with the possibility of selection (non-random assignment) and alternative channels (besides aid) through which HRT participation may affect growth in the short run. The next set of variables comprise our

¹⁸Using non-overlapping windows of data means the researcher must “throw away” data, foregoing statistical power, but reduces the mechanical autocorrelation that is induced by overlapping data in a growth context. We present estimates using both overlapping and non overlapping windows for robustness.

Table I Summary Statistics: Outcomes

	Obs	Mean	Std. Dev.	Min	Max
Four Year Growth (Overlapping Sample)					
GDP	3484	.070	.155	-.959	1.740
Household Consumption	2280	.064	.160	-1.297	.730
Government Consumption	2299	.078	.241	-1.198	1.528
Physical Investment	2417	.104	.372	-1.551	2.865
Exports	3252	.277	.446	-2.802	3.251
Imports	3249	.271	.426	-1.608	2.019
Manufacturing Value Added	2378	.089	.235	-1.311	2.509
Services Value Added	2700	.083	.192	-1.399	2.400
Agriculture Value Added	2725	.020	.148	-1.090	.761
Four Year Growth (Non - Overlapping Sample)					
GDP	896	.070	.155	-.793	1.398
Household Consumption	557	.065	.164	-1.118	.730
Government Consumption	562	.075	.236	-.967	1.240
Physical Investment	592	.109	.377	-1.511	2.865
Exports	795	.284	.457	-1.340	3.020
Imports	794	.279	.441	-1.107	1.800
Manufacturing Value Added	575	.095	.261	-1.052	2.509
Services Value Added	654	.083	.215	-1.399	2.400
Agriculture Value Added	659	.017	.159	-.880	.680
Ten Year Growth (Overlapping Sample)					
GDP	2974	.168	.305	-.944	3.003
Household Consumption	1887	.151	.262	-1.172	1.062
Government Consumption	1906	.178	.395	-1.375	2.179
Physical Investment	2006	.225	.550	-1.908	3.444
Exports	2748	.661	.702	-3.145	3.923
Imports	2745	.649	.653	-1.528	3.070
Manufacturing Value Added	1934	.203	.406	-1.634	2.039
Services Value Added	2232	.187	.337	-1.838	2.900
Agriculture Value Added	2257	.042	.221	-1.045	1.401
Ten Year Growth (Non - Overlapping Sample)					
GDP	340	.171	.294	-.641	2.802
Household Consumption	170	.123	.262	-.667	1.017
Government Consumption	172	.146	.392	-1.030	1.281
Physical Investment	184	.177	.603	-1.525	3.444
Exports	246	.556	.699	-1.243	3.278
Imports	245	.505	.646	-.921	2.141
Manufacturing Value Added	169	.160	.414	-1.634	1.060
Services Value Added	201	.150	.347	-1.219	2.152
Agriculture Value Added	203	.011	.214	-.678	.690

Table II Summary Statistics: Aid, Treaties, Controls, and Alternative Channels

	Obs	Mean	Std. Dev.	Min	Max
Aid per capita	3655	32.937	36.547	-41.427	345.407
Early Impact Aid per capita	3110	11.909	18.378	-35.485	205.391
HRT participation	3739	7.082	4.412	0	20
Controls					
Polity	3693	-.405	6.620	-10	10
Civil Freedoms	3169	4.432	1.460	1	7
UNSC Membership	3739	.070	.255	0	1
Sanctions	3739	.341	.474	0	1
Market Orientation	3548	.387	.487	0	1
Possible Alternative Channels					
Merchandise Exports Per Capita	3512	360.751	986.538	.515	24163.21
Merchandise Imports Per Capita	3510	345.397	618.313	2.062	8379.936
Foreign Direct Investment Per Capita	2791	38.485	119.988	-1206.447	2417.4
Remittance Inflows Per Capita	2300	46.432	92.01	.001	736.637

controls - variables upon which we would expect countries to possibly select into HRT participation. As we do not want to introduce endogenous variables into our specification, we restrict ourselves to variables that measure institutions we'd expect to be highly correlated with HRT participation and growth performance. Ideally, once we condition on these variables (and country and time fixed effects) HRT is as good as randomly assigned. In section 3.3 we examine this assumption in detail using [Altonji et al., 2005].

As we can see in the table, the average country is slightly non-democratic as defined by the commonly used Polity measure of Marshall and Jaggers [2002], and is slightly “not free” as defined by the measure of civil liberties of Gastil [1991]. This is not so surprising, as the typical aid recipient is poorer than the world average and income is negatively correlated with institutional quality. The average country is a member of the (non permanent) security council with 7% probability, faces some type of sanction with a 34% probability, and has a market oriented economy with 39% probability.

The rest of the variables in the table comprise possible alternative channels through which HRT participation may affect growth. We should be concerned in particular that HRT affects other economic inflows. For example participation in HRT could bring trade benefits or FDI from developed world, who interprets ratification as a signal that the country is “open for business.” Or perhaps countries that ratify HRT more frequently have better institutions for remittances from international migrants. As it is not difficult to imagine that these inflows affect growth, we should be concerned that the effect of aid that we estimate is in fact the effect of some other flow that enters the country at the same time as the aid does. We also test for this possibility below.

3 Human Rights Treaty Participation, Aid, and Economic Growth

In this section we uncover the relationship between treaty induced aid - the aid that is disbursed as a result of treaty participation - and economic development in recipient countries. Our causal regression of interest is given by:

$$g_{it+1} = \alpha_i a_{it} + \mathbf{x}'_{it} \beta^g + e_{it} \quad (2)$$

Where g_{it+1} is economic growth between periods t and $t + 1$, a_{it} is aid per capita received in period t , \mathbf{x}_{it} is a vector of observable variables that are potentially correlated with HRT participation and explain economic growth, and e_{it} represents factors unobservable to the econometrician that explain growth. We allow for country and time specific unobservable differences in growth:

$$e_{it} = \omega_i^g + \delta_t^g + \tilde{e}_{it} \quad (3)$$

Note the i subscript on the coefficient α_i . We allow for heterogeneous treatment effects, and aim to identify an average causal effect of aid on growth, which we denote as α . We expect that $cov(a_{it}, \tilde{e}_{it}) \neq 0$, and as such simple OLS regression results in inconsistent estimates of α . There are many reasons why this may be the case. For example, needy countries are more likely to receive aid and are less prone to grow and thus $cov(a_{it}, \tilde{e}_{it}) < 0$. On the other hand, countries that use aid effectively are more likely to be rewarded by donors and are more prone to growth, and so $cov(a_{it}, \tilde{e}_{it}) > 0$. To identify an average causal effect of aid on growth we exploit exogenous variation in aid receipts generated by participation in HRT at the UN. We model the relationship between HRT participation, T_{it} , and aid - aid supply - as:

$$a_{it} = \theta T_{it} + \mathbf{x}'_{it} \beta^a + u_{it} \quad (4)$$

where:

$$u_{it} = \omega_i^a + \delta_t^a + \tilde{u}_{it} \quad (5)$$

We provided evidence above that $\theta > 0$ - HRT participation causes an increase in foreign aid per capita for the average country. As we expect that there are heterogeneous treatment effects of aid on growth, a key further assumption is that the effect of HRT participation on aid receipts is (weakly) positive for all recipients. This is the monotonicity assumption that underlies identification of Local Average Treatment Effects [Angrist and Imbens, 1995]. In our setting, monotonicity means that either countries get more aid when they ratify HRT (they are *compliers*) or their aid receipts don't depend on ratification (they are *never takers* or *always takers*). We must not have countries who, conditional on covariates, are *punished* for participating in HRT (*defiers*). We provided evidence in favour of this assumption above.

Finally, to be valid the instrument must satisfy the usual condition:

$$\text{cov}(T_{it-1}, e_{it}) = 0 \tag{6}$$

While it is difficult to think of a way that HRT participation might *directly* affect economic growth, at least in the short run, it is possible that there is either an *indirect* link between HRT participation and economic growth through some channel other than aid, or that HRT participation is not randomly assigned with respect to growth. We carefully examine these possibilities in a series of robustness checks in the next section. For the moment, however, we try to control for any channel through which HRT participation may affect growth outside of its effect through the aid channel in \mathbf{x}_{it} . For the purposes of identification we are only concerned with observable variables that are correlated with HRT participation, such as institutional quality measures, so the observable characteristics we use have been identified in the political science literature as primary determinants of HRT participation.¹⁹ We would like to avoid including any type of control variable that is jointly determined with economic growth, and for this reason we keep our specification as parsimonious as possible.

If countries with better institutions tend to ratify more treaties and have better economic performance, we would expect to overestimate α in the absence of institutional controls. Similarly with human rights practices - if countries with better human rights practices tend to ratify more treaties and have better economic performance, we would expect to overestimate α in the absence of controls for human rights practices. As such we include measures of political institutions and human rights practices in \mathbf{x}_{it} . Another possible channel we must consider is a (non-permanent) place on

¹⁹See i.e., Hathaway [2007]

the United Nations Security Council. Kuziemko and Werker [2006] find that countries receive more US aid when they are rotated onto the Council. It is possible that non permanent members are more likely to participate in UN institutions like HRT, and that other growth enhancing benefits are bestowed upon non-permanent members as well. Even in the event that there are no such additional benefits, it should be noted that countries are elected to serve as non-permanent members, and the reasons for which they are elected to serve could be related to expected future growth. We include membership on the Security Council in \mathbf{x}_{it} . Finally, countries that face sanctions likely receive less foreign aid and are also “bad” global citizens. We thus include a binary indicator of whether a country is under some type of sanction in a given year. While the inclusion of these controls does make our assumption in equation 6 more plausible, there are other possible reasons the exclusion restriction does not hold. We return to this after we present and discuss the main results.

3.1 Results

In table III we present the estimates of the effects of total OECD aid on 4-year and 10-year economic growth. The results point to a strong positive effect of aid per capita on economic growth. For comparison’s sake we include OLS-FE estimates with and without our control variables. The first important thing to note is that when we add controls the estimates of α change only slightly and move further away from zero, suggesting that if anything, we are underestimating the effect of aid on growth. In columns 3-4 we present the instrumental variable estimates using all the data. The first stage results tell us that, unlike many previous studies using instrumental variables methods to uncover the effect of aid and growth, we should not be concerned about weak instruments. As reported in Magesan [2013], net ODA responds very strongly to treaty participation. While the estimated effect of aid on four year growth (top panel) is always positive, comparing the OLS-FE estimates of α with the IV estimates, we see that the IV estimates are several times larger in magnitude than the corresponding OLS-FE estimates. This is consistent with a downward bias in the OLS-FE estimates, which we may expect for two reasons. First, we anticipate that countries that expect large growth rates in the immediate future would receive less aid as they are less in need, while countries expecting negative growth shocks would make good candidates for foreign aid. Second, in the event that aid is measured with error, the OLS-FE estimates would suffer from attenuation bias. In the fifth and final column we replicate column 4 except that we use non-overlapping windows of data, which is more standard in the literature. While this implies a substantial loss of data (and statistical power), it avoids the potentially serious autocorrelation

Table III Treaty Participation, Aid, and Growth

4yr Growth	OLS FE(1)	OLS FE(2)	IV(1)	IV(2)	IV(3) (No overlap)
Aid per capita	.001*** (.0001)	.001** (0.0002)	0.005** (0.002)	0.007** (0.003)	0.0055** (0.0028)
First-stage:			Aid per capita	Aid per capita	Aid per capita
HRT participation			2.215*** (.662)	2.110*** (.664)	2.69*** (.784)
Reduced Form:			4 yr growth	4 yr growth	4 yr growth
HRT participation			.011** (.004)	.015*** (.005)	.015** (.007)
Countries	85	84	85	84	84
Observations	3484	2830	3484	2830	659
<hr/>					
10yr Growth					
Aid per capita	.002*** (.0005)	.002*** (0.0007)	0.007** (0.003)	0.01*** (0.004)	0.009*** (0.003)
First-stage:			Aid per capita	Aid per capita	Aid per capita
HRT participation			2.587*** (.696)	2.694*** (.715)	3.661*** (1.222)
Reduced Form:			10 yr growth	10 yr growth	10 yr growth
HRT participation			.018** (.008)	.026** (.01)	.032*** (.010)
Countries	85	84	85	84	84
Observations	2974	2332	2974	2332	292
Country FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes	Yes

Notes: Dependent variable is difference in logs of GDP. Standard errors are clustered at the country level. * Indicates statistical significance at the 10% level, ** indicates statistical significance at the 5% level, *** indicates statistical significance at the 1% level.

problem present in overlapping windows of the outcome variable. The results in this fifth column essentially echo those of the fourth column. Importantly the loss of data does not affect the power in the first stage.

The result in the top panel of column 5, which is our favoured specification, implies that a one dollar per capita increase in aid receipts implies an additional 0.55 percentage point increase in 4-year growth above what the country would expect. As the average four year growth rate in the whole sample is 6.7 percent, an increase of 0.55 is economically significant. It represents an 8% higher growth rate. Looking at 10-year growth in the bottom panel, we see a similar story. The IV estimates are several times larger than the OLS-FE estimates, and the result in column 5 implies that a dollar per capita increase in foreign aid causes a country to experience a 0.9 percentage point higher growth rate than it would have in the absence of the aid increase. As the average 10-year growth rate is 17.4 percent, the effect of an increase in one dollar per capita of aid is about a 5 percent larger growth rate over ten years than the average country would experience.

As we discussed above, our treatment variable (aid per capita) differs from the one typically used in the literature (aid as a fraction of GDP), and so directly comparing our estimates with those in the literature is not straightforward. Clearly our estimates are larger than those in Rajan and Subramanian [2008], who find little to no effect of aid on growth. Clemens et al. [2012] make the important point that ODA contains many types of aid, some of which should be expected to have an effect on growth within a few years (“early impact” aid), and some which should not. They find that an OLS fixed effects specification predicts that early impact aid has a modest effect on growth across a broad range of studies, many of which predicted that (total) aid had no effect on growth. Note that our IV estimates suggest that even *total* aid has economically significant effects on growth in the short and long run. To best compare our results with their study, we divided ODA into the three categories and found that the IV estimate of the effect of HRT participation on aid receipts is economically and statistically larger for early impact aid as opposed to late impact aid, and that it has no effect on humanitarian aid, and that when we replace Net ODA with early impact aid, our estimated effect of aid on growth is significantly larger.²⁰ The findings are in line with Clemens et al. [2012], but the IV estimates suggest a much larger effect than the OLS-FE estimates that they obtain.

In sum, the estimated Local Average Treatment Effect of aid per capita on short run and long run growth is economically and statistically significant. Conditional on our identification

²⁰These results are available on request.

assumption $cov(T_{it-1}, \tilde{e}_{it}) = 0$ holding, the results point to an effect of aid larger than what has been estimated previously in the literature. While we believe we have controlled for the most likely alternative channels through which treaty participation may affect growth, as discussed above, we still have a way to go in establishing causality. We turn to this issue in the next subsection.

3.2 Robustness: Establishing Causality

In this subsection we systematically examine and rule out possible reasons that the assumptions under which our estimate of α has a causal interpretation may be violated. The reasons can be divided into three broad groups: 1) Selection on unobservables (non random assignment of the instrument) 2) The instrument explains growth through a channel other than aid, 3) Spurious first stage. We address each set of concerns in turn.

3.3 Selection on Unobservables

It is possible that countries prone to growth select into treaty participation. High quality institutions have been long linked to positive growth outcomes (see for example Acemoglu et al. [2001] for causal evidence on this issue) and countries with good institutions are perhaps more likely to be “good global citizens” and participate in HRTs. That is, there is potentially positive selection on unobservable time varying country specific characteristics, in which case we are over estimating the effect of aid on growth.

As Altonji et al. [2005] pointed out, selection on *observable* variables can be useful for quantifying how much selection on unobservable variables one would need to be estimating a non zero coefficient when the true population coefficient is in fact zero. In the case of an instrumental variables estimator, under the assumption that a first stage exists, this amounts to comparing the estimated reduced form effect of the instrument on the outcome in a regression with and without observable control variables. In Table IV we compare the reduced form estimates of the effect of treaty participation on 4-year and 10-year economic growth) with and without the full set of controls.

The table suggests that there is little change in the estimated reduced form relationship between HRT participation and economic growth when we add observable variables. In fact, in our favoured specification with non overlapping 4-year growth, the selection on unobservables would have to be 47 times as large as the selection on observables and work in the *opposite* direction in order for the population coefficient to be zero. This is highly implausible. In other words, if anything, we are underestimating the true effect of aid on growth.

Table IV Selection on Observables I

	4yr Growth	4yr Growth	4yr Growth	4yr Growth
HRT Participation	.0146*** (.0053)	.0147** (.0045)	.0143** (.0067)	.0146** (.0069)
Countries	84	84	84	84
Observations	2830	2830	659	659
Controls	No	Yes	No	Yes
	10yr Growth	10yr Growth	10yr Growth	10yr Growth
HRT Participation	.026*** (.0095)	.026** (.0098)	.032*** (.010)	.032*** (.010)
Countries	84	84	84	84
Observations	2332	2332	659	659
Controls	No	Yes	No	Yes
Overlap	Yes	Yes	No	No

Notes: Standard errors are clustered at the country level. * Indicates statistical significance at the 10% level, ** indicates statistical significance at the 5% level, *** indicates statistical significance at the 1% level.

We can also use a method of sample trimming to check for selection on treatment.²¹ The idea here is to trim from the sample observations where the observable controls are “too good” at predicting HRT participation or non participation. We would expect that if, conditional on some value of the observables, the ratification or non ratification of treaties is highly certain (non random), that unobservables also predict ratification for these units. Of course if the unobservables predict growth as well then we have an identification problem. So we try to identify and purge these cases from the data as follows. Let d_{it} indicate whether country i 's HRT participation is greater than the sample median participation at time t . We estimate $P(d_{it} = 1|\mathbf{x}_{it})$ where \mathbf{x}_{it} is the full vector of control variables. We then drop all observations where the controls \mathbf{x}_{it} are such that $P(d_{it} = 1|\mathbf{x}_{it}) > 0.75$ or $P(d_{it} = 1|\mathbf{x}_{it}) < 0.25$ - these are units especially prone or not prone to HRT ratification for observable -and thus likely unobservable - reasons.²² We then consider our main specifications with the trimmed sample in Table V.

While trimming the sample inevitably results in a loss of statistical power, we note that our coefficient estimates remain statistically significant, and if anything are larger than the estimated effects presented in Table III In particular, we can say that the reduced form effect of HRT participation on economic growth is stronger in a subsample of the data where selection on unobservables

²¹This idea is proposed in Lim et al. [2015].

²²There is no rule for how much trimming one should do. We ensure, however, that in our trimmed sample the covariance between the part of growth predicted by the controls and HRT participation is economically and statistically insignificant, meaning that within this sample, if the observables are a good measure of the unobservables, then there is likely no selection on unobservables.

Table V Selection on Observables II

	4yr Growth	4yr Growth (No Overlap)	10yr Growth	10yr Growth (No Overlap)
Aid per capita	0.011** (0.005)	.008* (.004)	.013* (.007)	.013* (.007)
First Stage				
HRT Participation	1.820 ** (0.776)	2.073** (0.877)	2.019** (.849)	2.988 (1.891)
Reduced Form	4yr Growth .019*** (.007)	4yr Growth (No Overlap) .016** (.007)	10yr Growth .027* (.014)	10yr Growth (No Overlap) .039** (.016)
Countries	79	79	78	75
Observations	1900	493	1462	181
Controls	Yes	Yes	Yes	Yes

Notes: Standard errors are clustered at the country level. * Indicates statistical significance at the 10% level, ** indicates statistical significance at the 5% level, *** indicates statistical significance at the 1% level.

is less likely. This is consistent with our findings above - if there is selection on unobservables it is such that we are underestimating the effect of foreign aid on growth.

3.4 Alternative Channels for the Instrument

We now turn our attention to ruling out alternative channels through which HRT participation may cause economic growth. To this end, it is useful to re-write the unobservable given by equation 3 to potentially include lagged HRT participation:

$$e_{it} = \lambda T_{it-1} + \omega_i^g + \delta_t^g + \tilde{e}_{it} \quad (7)$$

Combining equations 2 with 4 and 7 we have the following reduced form relationship between economic growth and HRT participation:

$$g_{it} = (\alpha\gamma + \lambda)T_{it-1} + \mathbf{x}'_{it-1}(\beta + \alpha\theta) + e_{it}^* \quad (8)$$

Clearly for identification of α given knowledge of γ from the first stage, we must have that $\lambda = 0$. This is simply another way to view the exclusion restriction. λ will not be zero, and thus $cov(T_{it-1}, e_{it}) \neq 0$, if HRT participation systematically affects growth through some other channel. We can not formally test the exclusion restriction, as any estimate of the reduced form relationship between growth and HRT participation in equation 8 will include the aid channel $\alpha\gamma$. However we

systematically test and rule out alternative channels through which HRT participation may affect growth. We first examine specific observable channels that HRT participation may operate through to cause growth. Of course we can not observe every possible channel through which HRT may cause growth, so we then consider a more general placebo test to determine whether treaty participation still affects growth in a subpopulation where the aid channel does not exist.²³

The first specific channel we consider is reform. Specifically, we consider the possibility that countries ratify treaties at the same time as they undertake other reforms, and these reforms cause growth. This seems unlikely, as most economic and political reforms are binary (adopt or not) while our treaty participation variable is an intensity measure which can continue to increase or decrease once the binary reform has been undertaken. Additionally, our controls of political and human rights institutions should capture any major political reforms, which are the types of reform most likely to be correlated with HRT participation. However we are especially concerned that economic reform such as market liberalization, which potentially causes growth, occurs simultaneously with HRT participation. The measure of adoption of market-oriented policy that we use is that of Sachs and Warner [1995], extended by Wacziarg and Welch [2008].²⁴ The measure is commonly used in the literature,²⁵ and is a binary variable that takes the value of 1 if several conditions of market openness are met.²⁶

In Table VI we present estimates of our favored specifications, controlling for the adoption of market oriented policies. Table VI illustrates that when we include the adoption of market oriented policy as a control the estimated relationship between treaty participation, aid, and growth remains qualitatively unchanged and highly statistically significant. Note however that the estimated effect of aid on growth is smaller than in our baseline estimates. We conclude that HRT participation does not mask other reforms which themselves cause growth.

Another possibility we consider is that regional trends explain changes in growth as well as participation in international initiatives such as UN human rights treaties. We thus include region specific trends in our favored specification, and present the estimates in Table VII. As we can see, the estimates are robust to the inclusion of region-specific trends, and the estimated effects are

²³This type of test was first proposed in Angrist [1990].

²⁴We extend the variable to the end of our sample assuming that the value in 2008 is the same as in subsequent years, which is likely the case for adopters in particular, as the decision is permanent in most cases.

²⁵A recent example is Buera et al. [2011] who study whether countries learn from each others' experience with market-oriented policies before deciding to adopt the policies themselves.

²⁶In particular, the following five criteria must be satisfied: The average tariff rate on imports is below 40%, Non-tariff barriers cover less than 40% of imports, The country is not a "socialist economy," The state does not hold a monopoly of the major exports, The black market premium is below 20%

Table VI Does Treaty Participation Just Proxy for Other Reform?

	4yr Growth	4yr Growth	10yr Growth	10yr Growth
Aid Per Capita	.004** (.002)	.004** (.002)	.006** (.003)	.006 *** (.002)
Market Oriented Policy	.000 (.020)	-.001 (.024)	-.036 (.042)	-.036 (.046)
First-stage: HRT participation	Aid per capita 2.401*** (.584)	Aid per capita 2.799*** (.615)	Aid per capita 2.934*** (.632)	
Market Oriented Policy	3.272 (3.754)	3.387 (5.353)	4.514 (3.401)	
Country FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Overlap	Yes	No	Yes	No
Countries	78	78	78	78
Observations	2671	702	2209	278

Notes: Standard errors are clustered at the country level. * Indicates statistical significance at the 10% level, ** indicates statistical significance at the 5% level, *** indicates statistical significance at the 1% level.

economically larger than our baseline estimates in table III.

We finally consider a placebo test as in Angrist [1990]. If there exists a subpopulation where we know that $\alpha\gamma = 0$, by estimating equation 8 on data from this subpopulation we will be able to test the hypothesis that $\lambda = 0$, because any statistically significant relationship between growth and HRT participation in this subpopulation implies that HRT participation must affect growth through some other channel. Two important facts help identify a useful subpopulation for our purposes. First, many countries experience a period of not receiving aid or receiving small amounts of aid before becoming full time recipients. Clearly if countries are receiving little to no aid, increases in HRT can not be associated with increased aid and any changes in growth we observe in response to changes in HRT must operate through a channel other than aid. Second, human rights only became a driving factor in the machine of international relations in the 1970s [Moyn, 2010], and we should expect that human rights did not factor into aid allocation decisions until after 1970.²⁷ A search on Google N Gram for historical mentions of the phrase “foreign aid and human rights” between 1948

²⁷ As Eric Posner writes of the US case in *The Guardian*, US presidents starting from Jimmy Carter (US President from 1977-1981), regardless of party affiliation, used the term “human rights” far more often than the Presidents who preceded him. See <http://www.theguardian.com/news/2014/dec/04/-sp-case-against-human-rights>, last retrieved June 4, 2015.

Table VII Region specific trends

	4yr Growth	4yr Growth	10yr Growth	10yr Growth
Aid Per Capita	.008** (.004)	.006* (.003)	.010** (.005)	.010** (.005)
First-stage: HRT participation	Aid per capita	Aid per capita	Aid per capita	
	1.618** (.641)	1.888*** (.664)	2.047*** (.720)	2.818*** (1.362)
Country FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Overlap	-	Yes	No	Yes
Region specific trends	Yes	Yes	Yes	Yes
Countries	84	84	84	84
Observations	2830	741	2332	292

Notes: Standard errors are clustered at the country level. * Indicates statistical significance at the 10% level, ** indicates statistical significance at the 5% level, *** indicates statistical significance at the 1% level.

(the year the UDHR was written) and 2008 (the last year of available data) supports this. Figure 7 illustrates a large spike in the use of this phrase from near nothing in the early 1970s. Ultimately, we can statistically test whether $\gamma = 0$ prior to 1970, and if it is the case, we have a useful placebo test.

In Table VIII we present estimates of γ and λ for this subpopulation in the case of four year growth (we do not have sufficient data to consider this placebo test for 10-year growth). We can not reject the hypothesis that $\gamma = 0$, and we thus have a valid test for $\lambda = 0$.²⁸ In both columns the estimate of λ is far from statistical significance at conventional levels or of the opposite sign from α . We take this table to be evidence suggesting that either $\lambda = 0$ or $\lambda < 0$. In the latter case, our estimate of the effect of aid on growth is biased downwards toward zero, and we identify a lower bound on the true effect.

3.5 Spurious First Stage

One other possibility we would like to rule out is that our first stage estimates are picking up a spurious relationship between foreign aid and HRT participation. Flows of aid to a country are potentially related with other economic inflows such as trade, foreign investment and remittances.

²⁸We can not include fixed effects in the regression with non overlapping windows, as we have fewer than two observations per country and in the period of the placebo test much of the variation in HRT participation is purely cross sectional.

Figure 7 Mentions of Foreign Aid and Human Rights, 1948-2008




Table VIII Placebo Test

	4yr Growth	4yr Growth
Treaty Participation	-.022* (.011)	.040 (.034)
First-stage: HRT participation	Aid per capita .933 (6.438)	Aid per capita -1.086 (3.273)
Country FE	Yes	No
Year FE	Yes	Yes
Controls	Yes	Yes
Overlap	Yes	No
Countries	56	48
Observations	330	88

Notes: Standard errors are clustered at the country level. * Indicates statistical significance at the 10% level, ** indicates statistical significance at the 5% level, *** indicates statistical significance at the 1% level.

For example, if HRT participation credibly signals that a country is “open for business,” we would likely observe increased trade and investment in the country along with more aid. We would like to be sure that our first stage estimates are not simply a reflection of this. In Table IX we examine whether HRT participation causes an increase in other economic flows.²⁹ The estimates suggest that HRT participation does not cause an increase in other economic variables that may explain growth.

We are also concerned with the possibility that the positive relationship we estimate in the first

²⁹There are massive outliers and extremely large variance in the case of trade volume and foreign direct investment, so we use the natural logarithm of trade and FDI to reduce the influence of these properties of the variables on the estimates.

Table IX HRT participation and other economic inflows

Economic Flow	Remittance Per Capita	Trade Per Capita	FDI Per Capita
HRT participation	-.500 (2.678)	.005 (.014)	-.002 (.004)
Country FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Countries	83	83	83
Observations	2264	2919	2484

Notes: Standard errors are clustered at the country level. * Indicates statistical significance at the 10% level, ** indicates statistical significance at the 5% level, *** indicates statistical significance at the 1% level.

Table X Does future HRT participation explain past aid?

	Aid PC 4yr lag	Aid PC 4yr lag	Aid PC 10yr lag	Aid PC 10yr lag
R_{it+1}	-.334 (.640)	-.388 (.678)	-.198 (.778)	.790 (.963)
T_{it}	1.950** (.796)	2.226*** (.760)	2.506** (1.161)	4.355*** (1.257)
Country FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Overlap	Yes	No	Yes	No
Countries	84	84	84	84
Observations	2830	741	2332	292

Notes: Standard errors are clustered at the country level. * Indicates statistical significance at the 10% level, ** indicates statistical significance at the 5% level, *** indicates statistical significance at the 1% level.

stage is actually driven by a reverse relationship, that perhaps *future* HRT participation explains *past* aid receipts, or that perhaps there is a common trend in the two variables explained by some confounding factor. As such, we would like to consider an equation of the following form:

$$a_{it} = \gamma_0^* R_{it+1} + \gamma_1^* T_{it} + \mathbf{z}'_{it} \theta^* + u_{it}^* \quad (9)$$

where R_{it+1} is the number of treaties ratified in period $t + 1$, that is, $R_{it+1} = T_{it+1} - T_{it}$.

If $\gamma_0^* = 0$ we can reject the hypothesis that future ratifications explain current aid receipts. In Table X we present the estimates of γ_0^*, γ_1^* .

In columns 1 and 2 we consider the effect of HRT participation in the subsequent four years from aid receipt, while in columns 3 and 4 we consider HRT participation in the subsequent 10-years. As we can see, the number of treaties ratified in the future never has a statistically significant effect on past aid, and the estimated coefficient is either negative or economically insignificant in most cases. Importantly, once we condition on future ratifications, past ratifications continue to be economically and statistically significant.

The recent work of Clemens et al. [2012] offers a further way to check for a spurious first stage relationship. They suggest a method to classify aid as having either an “early impact” on growth, a “late impact” or being humanitarian in nature and thus not expected to have an effect on growth.

Table XI Which type of aid responds to HRT participation?

	Early Impact Aid PC	Late Impact Aid PC	Humanitarian Aid PC
HRT participation	.706** (.274)	.525* (.293)	-.022 (.103)
Country FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Countries	84	84	84
Observations	3075	3075	3075

Notes: Standard errors are clustered at the country level. * Indicates statistical significance at the 10% level, ** indicates statistical significance at the 5% level, *** indicates statistical significance at the 1% level.

If any of these types of aid should have an impact on growth in a window of four years, early impact aid should. Thus, if our first stage mechanism is correct, given we are finding that the aid induced by treaty participation has a positive effect on growth it should be the case that early impact aid is most sensitive to HRT participation. Table XI illustrates that the effect of HRT participation is more economically and statistically significant in the case of early impact aid than late impact aid. In the case of humanitarian aid there is no effect. This is further evidence in support of our causal chain from HRT participation to aid to growth.

Finally we consider the possibility that our results are explained by mean reverting behavior in aid and growth. It is possible that periods of low growth are followed by periods of high growth, and similarly with aid, and that these series' line up in such a way that we incorrectly infer a positive effect of aid on growth. In particular, it could be that countries that have low growth receive more aid, and then because growth reverts to its mean in the subsequent period we incorrectly infer that aid caused this reversion. Note that in order for this to be an issue in our setting, ratifications of HRT must also be timed in a specific way - they must occur in low growth/high aid periods preceding high growth periods. Our evidence in support of the exclusion restriction above suggests this is not the case, but we still consider the possibility that mean reversion explains our results, by studying whether current aid receipts are explained by past growth in table XII. We can see that the mean reverting pattern does not exist, as there is no statistically significant relationship between past growth and current aid.

Table XII Mean Reversion

	Aid Per Capita	Aid Per Capita	Aid Per Capita	Aid Per Capita
Lagged 4-year Growth	-0.858 (8.310)	-2.698 (8.505)	- -	- -
Lagged 10-year Growth	- -	- -	1.742 (8.709)	-6.676 (6.691)
Country FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Overlap	Yes	No	Yes	No
Countries	84	84	82	82
Observations	2650	634	1965	255

Notes: Standard errors are clustered at the country level. * Indicates statistical significance at the 10% level, ** indicates statistical significance at the 5% level, *** indicates statistical significance at the 1% level.

4 Mechanism

Having established the causal link between aid and growth, we now move to study the mechanism through which aid affects economic growth. In particular, we try to identify what is driving the estimate in Table III of the effect of aid on growth over 4-years, where we do not use overlapping windows of data.³⁰ In doing so we also study the possibility that aid causes “Dutch disease” in recipient countries. Aid may cause “Dutch disease” because it is often targeted towards non-tradeable sectors, which causes an increase in domestic input prices (including wages) and makes tradeable sectors such as manufacturing less competitive (holding world prices fixed). Spending the higher wages can also lead to real exchange rate appreciation, further exacerbating the problem. In table XIII we first look at aid’s effects on value added in the three major sectors of the economy in recipient countries: agriculture, manufacturing and services.

We see that aid causes an increase in value added in the services sector.³¹ Note that the coefficient is identical to the estimate in the final column of Table III. That is, the effect of aid on growth that we identify in our main regression results is almost completely explained by increased value added in the services. Moreover there is no statistically significant effect on value added in either manufacturing or agriculture. If anything, the effect is positive. These results suggest that aid precipitates the shift to services, a well established pattern of development [Soubbotina, 2000]. A primary mechanism through which growing economies shift from a manufacturing to service

³⁰There is not sufficient data to consider 10-year growth.

³¹The World Bank classifies “services” as including: trade, transport, communications, government, financial and business services and personal, social and community services [Soubbotina, 2000].

Table XIII Aid's effect on Growth in Value Added

	Agriculture	Manufacturing	Services
4-year growth			
Aid per capita	.003 (.002)	.0041 (.0032)	.0055* (.0031)
First Stage			
HRT Participation	2.77*** (.831)	3.220*** (.900)	2.770*** (.833)
Countries	585	524	584
Observations	79	75	79
Controls	YES	YES	YES

Notes: Standard errors are clustered at the country level. * Indicates statistical significance at the 10% level, ** indicates statistical significance at the 5% level, *** indicates statistical significance at the 1% level.

Table XIV Aid's effect on Expenditure

	Household Consumption	Govt Spending	Physical Investment	Exports	Imports
Aid per capita	.004** (.002)	.002 (.002)	-.003 (.005)	.003 (.004)	-.025 (.019)
First Stage					
HRT Participation	3.709*** (.923)	3.469*** (.860)	3.714*** (.889)	2.627*** (.888)	2.627*** (.888)
Countries	466	468	490	707	707
Observations	66	66	69	86	86
Controls	YES	YES	YES	YES	

Notes: Standard errors are clustered at the country level. * Indicates statistical significance at the 10% level, ** indicates statistical significance at the 5% level, *** indicates statistical significance at the 1% level.

based economy is structural changes in demand that accompanies higher incomes. Services are either a normal or luxury good depending on the service and the stage of development, and so as household incomes increase we observe growth in services relative to other industries. We can thus use data on household consumption to verify whether this is the channel through which services expand in response to aid. In Table XIV we present estimates of the effect of aid on household consumption, government consumption, physical investment, exports and imports.

While table XIII suggests that aid expedites the shift to services in growing economies, table XIV confirms the channel through which the shift to services occurs. In particular, table XIV is consistent with aid causing households to spend more on services. Our results also suggest that studies that find evidence of “Dutch disease” (for example Rajan and Subramanian [2011]) should not necessarily be taken in a negative light. Aid induced “Dutch disease,” even if it occurs, does

not necessarily imply a loss in welfare in aid receiving countries as other industries may benefit from the increased aid in a way that more than compensates for the loss in the negatively affected industries.³² It is worth noting that aid has no effect on exports, more evidence against the “Dutch Disease” hypothesis. The estimate of the effect of aid on government spending suggests as well that the effect on services can not be explained purely by aid dollars being used by the government to fund services for aid workers. In contrast to ?, who find that aid causes economic growth among the 35 countries who have graduated from the IDA’s aid lending program primarily through physical investment, we find that treaty induced aid has no effect on physical investment. Of course, our sample of aid receiving countries is much larger and it should not be surprising that we find a different effect.

We note in passing that the estimates in XIV contradict a common criticism of foreign aid, that it often doesn’t reach those it is intended to help, and is instead consumed or “wasted” by recipient governments before it can reach households. We find that aid has an economically large and statistically significant effect on household consumption. In fact, growth in household consumption explains most of the GDP growth in recipient countries. This suggests that aid, or at least treaty-induced aid, does indeed reach households in developing countries.

5 Conclusion

In this paper we exploit exogenous variation in foreign aid receipts generated by participation in Human Rights Treaties (HRT) at the UN to identify the causal effect of aid on economic growth. As the effect of aid on growth is heterogenous across countries, our instrument is particularly useful in that it satisfies the monotonicity condition for identification of a Local Average Treatment Effect. We find that aid has a significant positive effect on short run and long run growth, and that the estimate is robust to tests of selection and violation of the exclusion restriction. The estimated effect of aid on growth is explained completely by an increase in value added in the services, likely driven by an increase in consumer demand. This suggests that studies that find that aid precipitates a decline in exporting industries (“Dutch disease”) should perhaps be viewed in a different light, as this decline does not necessarily represent a loss in welfare. Our results should be taken with caution. Our estimate is a LATE, meaning that it is specific to the instrument, in

³²One possibility is that the effect of aid on services (and ultimately growth) is transitory. For example aid may be used to set up services to cater to aid workers. Given our results for 10-year growth, however, the effect on services is not likely transitory.

this case participation in HRT. In principle, given another valid instrument one would identify a different effect. Given the lack of credible instrumentation strategies in the literature, our results make an important contribution to the understanding of aid’s effects on economic development in spite of the “local” nature of the causal effect we identify. Additionally our results suggest an important role for international agreements such as HRT, as they serve to help effectively allocate aid to countries who can use it.

References

- Daron Acemoglu, Simon Johnson, and James A. Robinson. The colonial origins of comparative development: An empirical investigation. *The American Economic Review*, 91(5):1369–1401, 2001.
- Joseph G. Altonji, Todd E. Elder, and Christopher R. Taber. Selection on observed and unobserved variables: Assessing the effectiveness of catholic schools. *Journal of Political Economy*, 113(1):151–184, 2005.
- Joshua D. Angrist. Lifetime earnings and the vietnam era draft lottery: evidence from social security administrative records. *The American Economic Review*, pages 313–336, 1990.
- Joshua D. Angrist and Guido W. Imbens. Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *Journal of the American statistical Association*, 90(430):431–422, 1995.
- Francisco J. Buera, Alexander Monge-Naranjo, and Giorgio E. Primiceri. Discrete choice with social interactions. *Econometrica*, 79(1):1–45, 2011.
- Craig Burnside and David Dollar. Aid, policies, and growth. *American Economic Review*, pages 847–868, 2000.
- Michael A. Clemens, Steven Radelet, Rikhil R. Bhavani, and Samuel Bazzi. Counting chickens when they hatch: Timing and the effects of aid on growth. *The Economic Journal*, 122:590–617, June 2012.
- Michael A. Elliot. The institutional expansion of human rights, 18632003: A comprehensive dataset of international instruments. *Journal of Peace Research*, 48(4):537–546, 2011.

- Sebastian Galiani, Stephen Knack, Lixin Colin Xu, and Ben Zou. The effect of aid on growth: Evidence from a quasi-experiment. *World Bank Policy Research Working Paper 6865*, (2), May 2014.
- Raymond D. Gastil. *Freedom in the world: Political rights and civil liberties*. Freedom House, 1991.
- Henrik Hansen and Finn Tarp. Aid and growth regressions. *The Journal of Development Economics*, 64(2):547–570, 2001.
- Oona A. Hathaway. Why do countries commit to human rights treaties? *The Quarterly Journal of Economics*, 51(4):588–621, 2007.
- Guido Imbens and Joshua Angrist. Identification and estimation of local average treatment effects. *Econometrica: Journal of the Econometric Society*, 62(2):467–475, 1994.
- Ilyana Kuziemko and Eric Werker. How much is a seat on the security council worth? foreign aid and bribery at the united nations. *Journal of Political Economy*, 114(5):905–930, 2006.
- Claire Lim, James Snyder, and David Stromberg. The judge, the politician, and the press. *American Economic Journal: Applied Economics*, 2015.
- Arvind Magesan. Human rights treaty ratification of aid receiving countries. *World Development*, 45:175–188, September 2013.
- Monty G. Marshall and Keith Jagers. Polity iv project: Political regime characteristics and transitions, 1800-2002. 2002.
- T. Clifton Morgan, Navin Bapat, and Yoshi Kobayashi. The threat and imposition of sanctions: Updating the ties dataset. *Conflict Management and Peace Science*, 31(5):541–558, 2014.
- Samuel Moyn. *Last Utopia: Human Rights in History*. Harvard University Press, 2010.
- Joel E. Oestreich. *Power and principle: human rights programming in international organizations*. Georgetown University Press, 2007.
- Raghuram G. Rajan and Arvind Subramanian. Aid and growth: What does the cross-country evidence really show? *The Review of Economics and Statistics*, pages 643–665, November 2008.
- Raghuram G. Rajan and Arvind Subramanian. Aid, dutch disease and manufacturing growth. *The Journal of Development Economics*, (94):106–118, 2011.

J.D. Sachs and A. Warner. Economic reform and the process of global integration. *Brookings Papers on Economic Activity, 1995*, 1:1–118, 1995.

Tatyana P. Soubbotina. *Beyond Economic Growth: Meeting the Challenges of Global Development*. The World Bank, Washington, D.C., 2000.

R. Wacziarg and K.H. Welch. Trade liberalization and growth: New evidence. *World Bank Economic Review*, 22(2):187–231, 2008.

Table XV A1: Data Sources

Variable	Source
Security council membership	UN
Treaty participation	UN Treaty Database
Foreign Aid Receipts	DAC of the OECD
Democracy	Polity IV Project [Marshall and Jaggers, 2002]
Human Rights	Freedom House: Civil Freedoms [Gastil, 1991]
Market Orientation/Liberalization	Sachs and Warner [1995] and Wacziarg and Welch [2008]
Sanctions	Threat and Imposition of Sanctions (TIES) 4.0 [Morgan et al., 2014]
GDP and population	World Bank
Exports	WDI: Merchandise exports (current US)
Imports	WDI: Merchandise imports (current US)
Physical Investment	WDI: Gross capital formation (constant 2000 US)
Foreign Direct Investment	WDI: Foreign direct investment, net inflows (BoP, current US)
Government Consumption	WDI: General government final consumption expenditure (constant 2000 US)
Household Consumption	WDI: Household final consumption expenditure (constant 2000 US)
Agriculture Value Added	WDI: Agriculture, value added (constant 2000 US)
Manufacturing Value Added	WDI: Manufacturing, value added (constant 2000 US)
Services Value Added	WDI: Services, value added (constant 2000 US)

Notes: All growth variables are constructed using differences in logs of the variable levels.
WDI denotes the World Bank's World Development Indicators.

6 Appendix

6.1 Data Sources

Table A2: Treaties

Treaty	Date Opened
1. Convention on the Prevention and Punishment of the Crime of Genocide	9/12/1948
2. International Convention on the Elimination of All Forms of Racial Discrimination	7/3/1966
3. International Covenant on Economic, Social and Cultural Rights	19/12/1966
4. International Covenant on Civil and Political Rights	19/12/1966
5. Convention on the Non-applicability of Statutory Limitations to War Crimes and Crimes Against Humanity	26/11/1968
6. Convention on the Elimination of All Forms of Discrimination against Women	1/3/1980
7. Convention against Torture and Other Cruel, Inhuman or Degrading Treatment or Punishment	10/12/1984
8. Convention on the Rights of child	20/11/1989
9. Second Optional Protocol to the International Covenant on Civil and Political Rights, aiming at the abolition of the death penalty	15/12/1989
10. International Convention on the Protection of the Rights of All Migrant Workers and Members of their Families	18/12/1990
11. Convention relating to the Status of Refugees	28/7/1951
12. Protocol relating to the Status of Refugees	22/04/1954
13. Convention on the Reduction of Statelessness	30/08/1961
14. Convention Relating to the Status of Stateless Persons	06/06/1960
15. Convention on Consent to Marriage, Minimum Age for Marriage and Registration of Marriages	07/11/1962
16. Slavery Convention	25/09/1926
17. Protocol Amending the Slavery Convention	25/10/1926
18. Supplementary Convention on the Abolition of Slavery	07/09/1956
19. Protocol to Prevent, Suppress and Punish Trafficking in Persons	12/12/2000
20. International Convention Against the Recruitment, Use, Financing and Training of Mercenaries	04/12/1989