AIDDATA
A Research Lab at William & Mary

WORKING PAPER 36

March 2017

Foreign Aid and Growth at the Subnational Level

Andrea Civelli
University of Arkansas

Andrew Horowitz
University of Arkansas

Arilton Teixeira
Ibmec

Abstract

We develop a measurement strategy for the impact of foreign aid based on a regional panel vector-autoregressive model (P-VAR). We illustrate the strategy using Ugandan districts. Data for the regional units (ADM2) is assembled combining satellite sources for socio-economic activity, geo-located aid disbursements, and traditional household surveys. We find statistically significant positive and persistent effects of aid shocks on nighttime luminosity. Mapping nightlights to economic activity, the results suggest that the economic magnitude of these effects is small, but significant – with a multiplier between 2 and 3 in the medium to long-run. The P-VAR addresses endogeneity concerns associated with non-random aid assignment.

Author Information

Andrea Civelli
University of Arkansas
andrea.civelli@gmail.com
Sam M. Walton College of Business, Department of Economics, Business Building 402, Fayetteville, AR 72701

Andrew Horowitz
University of Arkansas
ahorowitz@walton.uark.edu
Sam M. Walton College of Business, Department of Economics, Business Building 402, Fayetteville, AR 72701

Arilton Teixeira
Ibmec
arilton8821@gmail.com
Ibmec, Belo Horizonte (Brazil)

Acknowledgments

This research was supported by USAID's Global Development Lab (The Lab) through a cooperative agreement with AidData (#AID-OAA-A-12-00096). The views expressed here do not necessarily reflect the views of USAID, the United States Government, or of AidData. Critical assistance was also provided by the Center for Advanced Spatial Technologies (CAST) at the University of Arkansas, USA. We thank Christopher Kilby, Jack Cothren, the participants of the ERSA-2015, AidData/William and Mary Seminar, and SEA-2016 for helpful comments. Sadek Yousuf provided excellent research assistance.

Contents

1 Introduction

The effect of official development aid (ODA) on recipient countries' economic growth has been the subject of a long-lasting and intense debate in the economic development literature. OECD countries have spent more than 3 trillion dollars on foreign aid since 1970, with the explicitly stated goals of economic development, growth, and poverty reduction in recipient countries. However, in spite of the widespread use of this policy instrument by donors, little consensus on the effect of aid on growth has been achieved.[1]

Two main challenges are faced in the measurement of ODA impact. The first is the donors' endogenous allocation of aid across recipients and the resulting difficulty in establishing causality between aid and its potential effects. Recent work by Rajan and Subramanian (2008), Deaton (2010), and Clemens et al. (2012) have argued convincingly that earlier estimates of ODA effects on recipients' growth may be subject to undermining endogeneity bias. The endogeneity issue is worsened by the difficulty of clearly disentangling donors' motivations when ODA decisions are taken. Kilby and Dreher (2010), Dreher et al. (2014), and Civelli et al. (2016) argue that the effect of aid cannot be accurately measured without considering the donor's motive.[2] Similarly, other explanations of aid, such as trade or geopolitical interests (as respectively identified by Berthélemy, 2006; Alesina and Dollar, 2000), could obfuscate the estimate of the ODA effects.

The second major issue to address in measuring the effects of ODA is the "over-aggregation" problem that could hinder the measurement of the impact of relatively small local treatments at the national level. Tierney (2011), Dreher and Lohmann (2015), and others, have argued that a subnational perspective is best for measuring aid effects since multiple sources of noise may be accumulated in the aggregation to the country level. In addition to spatial aggregation, ODA aggregation can occur in other dimensions relevant to the impact of aid on growth. Clemens et al. (2012), for instance, disaggregate aid into early and long impact varieties, finding significant growth effects only with the early impact category. Similarly, aid can be disaggregated by sector (health, education, irrigation, etc.), program, or studied at the individual project level. At the project level, specific projects have long been carefully evaluated relative to preset benchmarks. More recently, randomized field experiments (Duflo et al., 2008) have provided true control treatment analyses of projects' local effects. At this level, many projects appear effective. However, at the intermediate levels of aggregation, aid effectiveness becomes more difficult to measure, and findings become more mixed.

This gap in the apparent micro and macro effects of aid has been dubbed the micro-macro paradox (Dreher and Lohmann, 2015; Mosley, 1987). However, most of the micro-level project studies do not claim a linkage to economic growth or measurement of the full spillover effects of projects. Yet, the micro and macro effects of aid projects are linked by definition; the total impact of foreign aid upon growth must be associated with the cumulative effect of the individual projects upon growth. Consequently, a growing literature focuses on improving sub-national ODA impact measures.

[1] On this point, see for instance Clemens et al. (2012). An introduction to the ODA literature can be found in Temple (2010), Addison and Tarp (2015), and Radelet (2006).

[2] Civelli et al. (2016) show that altruistic motivations, which could help illuminate reverse causality explains only a minority of aid transactions.

This paper proposes a measurement strategy of the economic impact of foreign-aid projects in a country that simultaneously tackles the endogeneity and aggregation problems by combining a panel vector-autoregressive (P-VAR) model and multiple sources of regional data. We assess the potential of this strategy, critically discuss its limitations, and provide a demonstration using regional data from Uganda.

Utilizing a P-VAR model is particularly attractive in this context because this econometric approach provides a convenient and intuitive way to impose identifying restrictions that can address the endogeneity in the causal relation between aid and economic activity. This solution is based on the well-known recursive orthogonalization of the covariance matrix of the estimated residuals of the VAR model. This approach has been extensively used in the empirical macroeconomics literature since Sims (1980), but has seen only limited use in the development literature. The identification of the ODA effects with this methodolgy is based on a simple scheme that isolates the exogenous shocks to ODA by removing the endogenous component. This is achieved by jointly modeling ODA and economic activity as a single system and by assuming that structural innovations to ODA can affect economic activity contemporaneously, while economic shocks are assumed to have an impact on ODA disbursements only with a temporal lag. This ordering is natural in this context since the aid disbursement process in lengthy with a prior commitment phase often proceeds disbursement by more than a year.

Prior applications of the VAR model to ODA issues include Lof et al. (2015), who apply a co-integrated P-VAR to 59 countries, finding that the average long-run response of income is about $4.5 - 5$ larger than an initial increase in ODA disbursements. Similarly, Juselius et al. (2014) estimate individual countries co-integrated VAR models for 36 sub-Saharan African countries, finding a positive long-run impact of ODA flows on the macroeconomy for most countries. Gillanders (2016) estimates a P-VAR model for a set of sub-Saharan African countries and reports a positive, but small, increase in economic growth following a fairly substantial aid shock. We follow Lof et al. (2015) in the adoption of the identification strategy, but we estimate a fixed-effects P-VAR as in Gillanders (2016). All these papers focus on the dynamics of aid at national level; importantly, our paper differs from these studies in adopting a sub-national perspective.

The use of disaggregated data brings the advantage of a stronger characterization of the link between ODA and growth. However, the combination of the panel structure necessary to estimate the P-VAR and the regional disaggregation of the data introduces some unavoidable costs and some limitations. First, the P-VAR requires a sufficiently balanced panel structure, with very limited missing observations and frequently sampled time series of the endogenous variables (annual observations, at least). Second, while the cross-section dimension helps increase the number of observations, the dynamic relation between endogenous variables is derived from the time dimension of the model, which requires a sufficiently large time sample. Third, it is difficult to obtain sub-national variables at annual frequency from official statistics to use in the VAR model. Finally, the panel estimation imposes the implicit assumption of homogeneous effects across units, which implies the measurement of an average effect over the whole country.

We address these difficulties by combing information from multiple and distinct data sources. With regard to the main obstacle of data availability for P-VAR estimation in the sub-national low-income country context, we side-step the problem by utilizing nighttime luminance (nightlight) data as proxy for economic activity. ODA disbursements at the sub-national level are available from the AidData Consortium's geo-coded project mappings. We then map nightlight to more conventional economic measures, such as consumption expenditure per household, using geo-coded data from living standard measurement type surveys conducted by the Uganda Bureau of Statistics (known as the Uganda National Household Surveys - UNHS). We also rely on other geographic information system datasets to collect data about population dynamics, land use, and the surface size of the geographic units.

We choose Ugandan districts as baseline cross-section units of the analysis in an attempt to achieve a sufficiently balanced panel. We adopt the same districts geographic definition as in the AidData ODA dataset. This provides 35 districts for which both nightlights and ODA disbursements are observed frequently enough over the sample $1996-2012$. As discussed in Section 3, this subset of Uganda's 112 districts contains about 50% of Ugandan population, though only 30% of Ugandan territory. The remaining districts show extremely discontinuous luminosity series, or no light at all. These non-luminous districts together account for only 1.2% of nightlight and 11% of total aid, on average. However, in order to preserve the information coming from these non-luminous districts, we aggregate them to a synthetic district, which is added to the panel. We then standardize all observations by the district surface to maintain comparability across units. Under the assumption of homogeneity of the effects of ODA across geographic locations, our strategy of synthetic district creation to preserve the information in the non-luminous districts has the cost of reducing the benefit of disaggregation for local effect measurement for that unit.

The use of nightlight satellite data to proxy income, both nationally and sub-nationally, has grown rapidly in the economic development literature (see Chen and Nordhaus, 2011; Henderson et al., 2012, 2011; Dreher and Lohmann, 2015). Nightlight data holds promise as a means to side-step well-known problems associated with traditional income survey data in low-income countries, such as infrequent surveys, large informal sectors, recipient data-gathering capacity constraints, and recall errors in the absence of formal income records.

The theoretical causal mechanisms of ODA to nightlight are straightforward. As discussed in the seminal nightlight papers in the economics literature (see Henderson et al., 2012, 2011; Chen and Nordhaus, 2011) nightlight, as measured by satellite, is highly correlated with income. Since the stated objective of ODA by the donors is income growth and development,[3] the theoretical causal chain from ODA to nightlight contains only one link. Of course, different types of aid will have different temporal lags and channels to growth, and different inherent effects on nightlight. For example, a bridge project that connects two areas with large potential economic synergies may have an immediate (within year) impact on income growth and nightlight. A "soft" aid project, such as one that improves the quality of primary education, might have multiple channels to growth and light. If the education project involves school construction or the hiring of new teachers, it may also generate income growth fairly quickly, assuming an output-gap exists. On the other hand, the income growth effect via the human capital formation channel will only be realized in the long-run. Finally, electrification projects could conceivably increase nightlight without a short-run impact on income. We address this concern by excluding power-supply projects in a robustness check and find similar results.

While many papers have studied the effects of aid at sub-national levels, ours is the first (to our knowledge) to utilize P-VAR estimation, nightlight data, and geo-located AidData to explore the impact of aid

[3]See the mission statement of OECD aid activities (http://www.oecd.org/about/) and note that the aid in our data set is all associated with OECD-affiliated donors.

on sub-national growth.[4] Dreher and Lohmann (2015) analyze ADM1 and ADM2 regions using data for World Bank projects for a large set of countries and an interacted instrumental variable approach, but they do not find significant causal effects of aid on growth. De and Becker (2015) use the AidData geo-coded ODA datasets and find a positive effect of health aid and water aid on the reduction of diseases incidence in Malawi using instrumental variables and propensity score matching. Similarly, Dionne et al. (2013) study the effectiveness of sector-specific aid in Malawi within a classic two-stage allocation-impact framework. Our preference instead for Uganda is mostly driven by the reliability of the AidData data for this country, which was part of the very first wave of AidData releases and has undergone several updating and cleaning revisions. However, our methodology is highly scalable, with straightforward application to many low-income countries.

We find that an initial exogenous shock to ODA is associated with a statistically significantly positive impact response of nightlight, and the positive response persists significantly for more than ten years after the shock. Mapping nightlight to economic activity, we find economic magnitudes in line with the mildly optimistic estimates usually reported at the national level. The results of our baseline specification suggest a cumulative ten-year long-run multiplier effect of ODA on per capita expenditure slightly larger than 3, defined as the ratio between change in ODA and corresponding change in economic activity. Similarly, the estimates of the short-run multiplier are around 2 for the five-year response to a temporary shock.

The remainder of the paper is organized as follows. Section 2 introduces the empirical strategy and discuss some of its limitations. Section 3 describes the data. Section 4 presents the empirical results. Section 5 concludes and discusses extensions.

2 Empirical Strategy

The empirical strategy of this paper relies on the use of a panel vector-autoregressive (P-VAR) model for the analysis of sub-national effects of foreign aid on economic growth. This approach provides a solution for two main concerns in the aid effectiveness literature. First, the VAR model allows us to solve the endogeneity of aid disbursements by imposing some restrictions on the dynamics of the model based on intuitive economic considerations. Second, analysis at the sub-national level can help address the difficulties in the measurement of aid effectiveness due to over-aggregation of aid types with different characteristics.

The P-VAR is a linear model that requires suitably balanced panel structure for estimation. Although the cross-section dimension helpfully increases the number of observations for estimation, the dynamic relation between the endogenous variables of the model can only be inferred from a sufficiently large time sample. Specifically, the low frequency of survey data (e.g., LSMSs or expenditure surveys) in many low-income countries prevents us from directly estimating a P-VAR model of sub-national economic growth using such data. Therefore, we follow a two-step procedure in order to exploit the identification advantages of the P-VAR as well as the advantages from ODA disaggregation. We first estimate a P-VAR in

[4] Among the others, examples of studies that have looked at the sub-national impact of aid projects on various outcomes are: Jablonski (2014) and Briggs (2012) for electoral outcomes, Crost et al. (2014) and Findley et al. (2011) for conflict, and Hamilton and Stankwitz (2012) for deforestation.

nightlights and ODA since nighlights data is available at annual frequency for virtually any level of geographic disaggregation. Again recall that both nightlight and ODA are endogenous in this framework. We then map the effects of ODA via lights to economic activity by adapting the predictive equation of Henderson et al. (2012)'s to our context.

The fixed-effects P-VAR(p) model is represented in equation (1):

$$Y_{i,t} = \sum_{j=1}^{p} A_j Y_{i,t-j} + u_i + e_{i,t} \tag{1}$$

where $Y_{i,t}$ is a vector of endogenous variables, i indicates the cross-sectional units of the panel and t the time dimension, p is the number of lags of the autoregressive component, u_i and $e_{i,t}$ are the vectors of panel fixed-effects and idiosyncratic errors respectively; and A_j's are coefficient matrices.

In our application, cross-section units correspond to Ugandan districts (including the synthetic district), while t is expressed in years. The endogenous vector includes the logs of the ratio of nightlight and aid disbursements to the district surface area, $light_{i,t}$ and $oda_{i,t}$ respectively

$$Y_{i,t} = \begin{bmatrix} oda_{i,t} \\ light_{i,t} \end{bmatrix}. \tag{2}$$

The choice of normalizing variables by district area follows the luminosity literature (see for example Henderson et al., 2012, 2011; Chen and Nordhaus, 2011; Dreher and Lohmann, 2015). This is standard when we want to approximate income dynamics with luminosity for two reasons. First, light growth occurs both at the extensive and intensive margin: that is, dark areas transitioning to light as well as the nightlight signal becoming more intense. Since the measurement of light by the DSMP satellites is top-coded, significant upper bound truncation in urban areas is not unusual. Where truncation occurs, light growth can only occur at the land area extensive margin. A second reason to prefer measures per land-unit area is the public goods nature of nightlight in many settings. For example, in a typical low-income country, the nightlight emissions and capacity are likely insensitive to the number of members in the household. Details regarding variable construction, data sources, and the full definition of the cross-sectional districts are provided in Section 3.

The P-VAR is estimated over the sample $1996-2012$ using the Love-Zicchino Stata package (Love and Zicchino, 2006). In our baseline specification, $p=1$ is selected, as indicated by optimal lag-selection criteria. A within estimator that removes the fixed-effects in a dynamic panel setup such as (1) introduces an additional bias in the estimation of the coefficient matrix, due to the correlation between the transformed residuals and the transformed vector of endogenous variables. As normally done in the context of dynamic panels, instrumentation with lags of Y_i is used to deal with this problem. The results we report are based on the use of the first lag of Y_i as an instrument; we also conduct a robustness check of the results using eight lags instead.[5] Robust standard errors are computed clustering by district and the model is

[5]Since the estimation procedure removes the fixed-effects using a forward orthogonal transformation, the first-lag instrumentation provides just-identifying conditions for the GMM estimation. On the contrary, using eight lags would increase the number of instruments to 32, with 28 over-identifying GMM conditions. Although this larger set of instruments passes the Hansen's test for

estimated by two-step GMM estimator.

The local disaggregation of ODA likely mitigates many of the issues that contribute to the non-random allocation of ODA, such as the strategic interplay between donor and recipient at national level or the enormous heterogeneity across recipient countries. With respect to recipient heterogeneity, the concern is attenuated as estimation is across regions of more uniform climatic, institutional, and socio-economic structure. Similarly, the capability of Ugandan sub-national governments to implement effective strategic play with multiple OECD donor countries is also likely quite limited when compared to nation-states. Nevertheless, addressing the endogeneity of ODA disbursements is also a priority at the sub-national level.

In a VAR framework, a solution to the endogeneity problem is readily available through the identification of exogenous structural innovations to ODA that can be employed to correctly measure the effects of ODA on nighlights. The VAR assumes that the reduced-form residuals of the model, $e_{i,t}$ in (1), are a linear combination of the unobservable orthogonal structural innovations of the system. These structural shocks, however, can be reconstructed from the reduced-form innovations by means of an orthogonalization of the estimated vector of residuals. Econometrically, the orthogonalization is achieved through a factorization of the estimated covariance matrix of the residuals. However, since the factorization is not unique, it is necessary to adopt some selection criteria to choose a specific one. In practice, given the linear setup of the VAR model, an easy approach to selecting an orthogonalizing scheme is to impose some restrictions on the contemporaneous relations between structural and reduced-form innovations, providing an economic justification in support of the restrictions.[6]

With two endogenous variables in the P-VAR model, only one restriction is required to fully identify the orthogonal structural shocks. The two possible cases are either that structural shocks to ODA do not affect lights on impact, or vice-versa nightlights shocks have no contemporaneous effect on ODA disbursements. We follow Lof et al. (2015) and assume that ODA shocks can affect nighlights contemporaneously, while nightlights shocks can impact ODA disbursements only with a temporal lag. The rationale for this identification strategy is that a local ODA disbursement could impact economic local activity via either the demand or supply side relatively quickly (within a year), while the response mechanism of ODA to an exogenous increase in nightlight would take more than one period to activate. This transmission mechanism is consistent with the complex process that culminates in aid allocation decisions and disbursements. Nevertheless, in Section 4 we will check the robustness of the results to the alternative identification scheme.

---

the validity of over-identifying conditions, the relatively small size of our panel suggests particular caution with the proliferation of instrument. For this reason, we adopt the just-identified case as baseline for our analysis. We find, however, very similar effects when we utilize larger sets of instruments in our robustness exercises. Also, notice that the removal of fixed-effects is independent of the ODA endogeneity issue and the instrumentation does not aim to solving it.

[6] This procedure is known as a Cholesky identification scheme because it exploits the Cholesky decomposition to factorize the covariance matrix of the reduced-form residuals. Let $v_{i,t}$ indicate the orthogonal structural shocks of the VAR model. Then, the reduced-form innovations can be expressed as $e_{i,t} = A_0 v_{i,t}$. Considering the $p = 1$ case for simplicity, model (1) can be written in structural form as:

$$A_0^{-1} Y_{i,t} = A_0^{-1} A_j Y_{i,t-j} + A_0^{-1} u_i + v_{i,t} \tag{3}$$

where the contemporaneous effects between the endogenous variables are now explicitly allowed for, and A_0^{-1} is the factorizing matrix that orthogonalizes the covariance matrix of the $e_{i,t}$. The Cholesky scheme imposes a set of zero-restrictions on the off-diagonal elements of A_0 (suitably re-ordered) to uniquely identify a lower-triangular factorization matrix. With a vector of two endogenous variables as here, only one off-diagonal element exists and, hence, only one restriction is necessary.

Once the structural shocks are identified, the effects of aid on luminosity can be estimated by computing the impulse response of nightlights to an initial shock to ODA. The relative magnitude of shock and response will represent the elasticity of luminosity to ODA that is at the heart of our analysis. The response functions also provide a dynamic representation of the transmission channel of transitory shocks in the short and long-run. Other tools are also used to further assess the transmission channel of ODA shocks; in particular, we analyze the long-run cumulative responses of lights to a one-time shock, the response to permanent shocks, and variance decomposition of forecast errors.

One important limitation of the VAR model is that, although the structural identification correctly reveals the dynamic interaction of the endogenous vector of variables in the system, the structural mechanism is still identified conditional on the model specification itself. Omitting other relevant endogenous variables or not controlling for relevant confounding factors might affect the dynamics of the impulse response functions and the magnitude of the identified responses. This issue has no easy solution in our context because of the difficulty of obtaining sub-national variables at annual frequency for a developing country from a reliable source.

These concerns are mitigated in part through robustness checks and the use of fixed effects. Fixed-effects in the P-VAR model eliminates the effects of district level time invariant characteristics. These can include differences in governance, cultural, climatic, and socio-economic factors that differ across district, but not over time. In an additional robustness check, we also include time-effects, which allows us to control for some common time-varying factors. The time-effects can control for business cycles or political trends at national level, for instance. As discussed in Section 4, with this specification the effects of ODA on luminosity are smaller, but still positive and significant. Finally, we re-estimate the P-VAR with ODA and nightlights standardized by population, in order to control for an important factor that can influence both light emission and ODA allocation.[7] Again, in this case, the effects of ODA are smaller but remain significant.

After the impulse response of lights to ODA is estimated with the P-VAR, it remains to connect this impact to a more standard measure of economic activity. For this we rely on the predictive stage of Henderson et al. (2012)'s approach, in which a statistical measure of economic activity (official GDP, typically) is regressed on lights.[8] We adapt this methodology in a straightforward manner to explore the link between the growth in household real expenditure at the district level and the change in luminosity for the same set of cross-sectional units used in the PVAR analysis. The predictive equation simply reads

$$x_{i,t} = \psi light_{i,t} + e_{i,t}. \tag{4}$$

Equation (4) is written in log-linear version as a two-period panel between 1999 and 2009. As in the P-VAR vector (2), $light_{i,t}$ is the log of the ratio of nightlight to the district surface area for district i in period t. We measure $x_{i,t}$ with either the log of the average household weekly consumption expenditure or, for robustness, the average household monthly expenditure in non-durable goods. Henderson et al. (2012)

[7]We construct population series at district level using the satellite source and the approximations described in Section 3.

[8]Henderson et al. (2012)'s methodology uses the estimates from this stage to ultimately infer the unobservable true growth of the underlying economic activity from an optimal combination of multiple signals correlated with it, such as GDP and lights. Since the purpose of our paper is to primarily document a transmission channel from ODA to household expenditure, we only focus on the first stage of their methodology.

measure $x_{i,t}$ at the country level using the log of national GDP instead; our estimates of $\hat{\psi}$ are very much consistent with theirs. The model is then estimated using panel fixed-effects and robust standard errors.

3 Data Description

Panel Structure – The panel used for the estimation of the P-VAR model covers the sample $1996 - 2012$, $T = 17$, and includes a cross-section of $N = 36$ sub-national regions. Of these regions, 35 correspond to administrative districts (ADM2 level units), while the last one is a synthetic region obtained from the aggregation of all the other Ugandan districts. The geographic boundaries of the districts are obtained from the world administrative divisions layer provided by ESRI-ArcGIS. We adopt the most recent definition of districts established in 2010, also shared by the AidData dataset, which consists of 112 districts. The synthetic district aggregates 77 districts with low (or no) night luminescence corresponding to about 70% of the Ugandan territory, and 48% of the population on average over this sample.

Although the introduction of the synthetic district may seem an asymmetric treatment of the data at first sight, our choice is justified by two goals. On one hand, it satisfies the need of a sufficiently balanced structure of the panel to support the P-VAR estimation. Both nightlights and ODA disbursements are regularly observed for the main 35 districts, but not for the others. Taken individually, these districts have extremely sporadic luminosity and ODA series that make them unsuitable for the VAR model. On the other hand, the aggregation corrects for these problems and allows us to preserve the information coming from this set of districts without arbitrarily dropping any observation.

Overall, the synthetic district receives a less than proportional share of total aid, on average only around 11%, and produces a very small fraction of the country nightlights, never bigger than 4%. The reason of this weaker luminosity emission can be found in the smaller urbanization share in the synthetic region. As discussed below Figure 3, there is a close correspondence between urban areas and luminosity. The districts in the synthetic region are primarily rural, with an urban share about 50 times smaller than that in the lit districts.[9] We maintain the comparability across units by standardizing the observations by the surface area of the respective district (or population in some robustness checks). Under the assumption of homogeneity of the effects of ODA across geographic locations implicitly imposed by the P-VAR approach, the use of the synthetic district weakens the benefit of the ODA geographic disaggregation strategy for this unit.

Figure 1 shows the geographic location of the cross-section units of the analysis on the Ugandan map. The districts that are part of the synthetic district are white colored in the figure; the other 35 districts are illustrated by the solid gray areas instead.

ODA Data – Data for the aid projects is drawn from AidData's Uganda Geo-coded Dataset (Release IV). This dataset maps 420 projects over the sample $1981 - 2013$ distributed across the 112 Ugandan districts. Information about the geographic precision of the disbursements for each project is reported on a scale

[9] The share of urbanized land in 2014 was $.8$% in the lit districts and $.014$% in the no-light districts. These shares are computed using the dataset provided by Pesaresi et al. (2013) - Beta release.

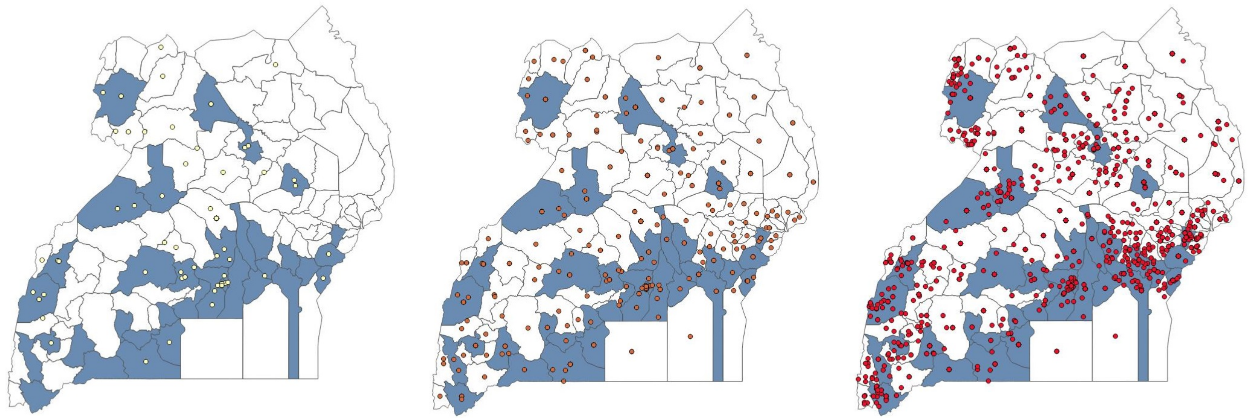

Figure 1: Distribution of the ODA project in Uganda since 1991. Data provided by AidData. Solid gray indicates the 35 districts included in the sample used in this study. White corresponds to the districts in the synthetic region.

from 1 to 8, where 1 indicates the knowledge of the exact geographic coordinates of the disbursement and 8 corresponds to projects at the central government level.[10] We enhance the aid-growth nexus at the heart of the estimation strategy by geographically filtering the projects to precision codes $1-3$, which correspond to ADM2 or lower administrative levels, in order to reduce the distance between where aid is spent and the region where it may produce positive economic effects. Figure 1 illustrates the spatial distribution of projects with precision $1-3$ over three time intervals. The number of precision 1 to 3 projects and the territory covered by them increase over time, especially in the last years of the sample.

In considering the sample $1996-2012$ for P-VAR estimation, note that before 1996 only two districts received aid for precision $1-3$ projects: the capital metropolitan area Kampala and, for a couple of years, the adjacent district of Wakiso, which includes some capital suburbs. We do not include these observations in the panel for two reasons. First, having only two cross-section units for so many periods is not desirable for our empirical model.[11] Second, even though recorded with high precision codes, these are projects close to the seat of government at a time when government institutions were often the primary beneficiaries of aid. With respect to our goal of local identification of the effect of aid, these projects likely match the criteria only weakly.[12] We also exclude 2013 from the sample due to some anomalous disbursements in two districts of the Northern region of the country that suddenly increase a hundred fold. These districts are part of the synthetic unit; the total disbursements of the synthetic district became five times bigger than those of the lit districts in 2013, whereas they are usually ten times smaller. This change makes this last observation behave like an outlier.

In addition to the geo-location of individual projects, the AidData dataset includes the annual disbursement flows for each project. Total district ODA is computed as the log of the ratio between the sum of aid disbursements for precision $1-3$ projects in a district (measured in real 2005 US dollar) and the district area measured in squared kilometers. Often, the data shows multiple disbursements made under the

[10]Note that the precision scale skips the classification code 7 for unspecified reasons.

[11]Furthermore, luminosity data is available only from 1992.

[12]For instance, the description of one of the projects mentions general government and civil society as main purpose. Other two projects are for the construction of the stadium and the international airport in Wakiso.

same project i.d. in multiple locations, not necessarily in the same district or in the same year. In such cases, there is not sufficient information to assign all disbursements of a project to a single district and the records show equal aid disbursement for each of the multiple locations of a project. In these cases, we proportionally re-distribute the disbursements over the recipient districts of a project based on population size. If a district does not receive any aid in one period, we substitute the observation with a small value, $.0001$, before taking the log. This occurs in about 20% of the observations.

Finally, Figure 2 decomposes the ODA disbursements by precision code and sector of activity for the sample $1996-2012$. This is a useful illustration of the implications of our geographic disaggregation strategy for the ODA disbursements. Two observations are worth noting. First, precision codes $1-3$ account for roughly half of the ODA disbursements. The purpose of our empirical approach is to separate the possible effects of locally circumscribed projects on local economies from the impact of a large number of projects that disburse funds at higher administrative levels. As noted, aid to the national government and ministries likely have a less direct connection to local growth than a project like a bridge or road. Second, the decomposition by sector also reveals a quite different structure between regional and national disbursements. Local disbursements exhibit a higher share of projects in education and agriculture, water sanitation, infrastructure and transportation, and health. On the contrary, general budget support - the main category for precision codes 6 and 8, are typically less focused geographically. The Figure also shows that the share of projects related to energy generation and energy supply, which can inflate lights emission without an increase in economic activity, represents only about 5% of the precision $1-3$ disbursements. However, in one of our robustness checks, we repeat the analysis after removing these projects and obtain similar results.

Nightlight Data – Night-light data is obtained from the Defense Meteorological Satellite Program (DMSP). These images are annual composites processed prior to release. Since data for distinct years can be provided by different satellites, an inter-calibration is applied to harmonize nightlight values across across years-satellites. We apply the inter-calibration adjustment parameters provided by Elvidge et al. (2014). As is done for ODA, the luminosity variable in the P-VAR is then constructed as the log of the sum of the luminosity index for all the pixels within a district boundaries standardized by the district area (measured in squared kilometers). test

As noted, persistent night-light from this source is not detectable in all of Uganda's 112 districts over the DMSP sample, $1992-2012$. Figure 3 illustrates AidData project mapping overlaid to the 2010 night-light image (on the left side panel) and land coverage (on the right hand panel) for a small region around Kampala, the capital of Uganda. In the land-cover image, red is used to indicate urbanized areas, while light green corresponds to rural and agricultural land.[13] Not surprisingly, night-light signals are usually associated with greater urbanization. On the contrary, some rural regions do not produce any detectable luminosity signal for time intervals of several years. In any case, note that ODA disbursements, the red dots, are found in urban areas as well as rural areas with little detectable night-light. For this reason, the construction of the synthetic region is important to avoid the complete loss of the information coming from the darker regions.

[13]The full legend of the land-cover colors is explained in the caption of Figure 3. This image is an authors' elaboration of (NASA) Landsat 7 multi-spectral satellite data.

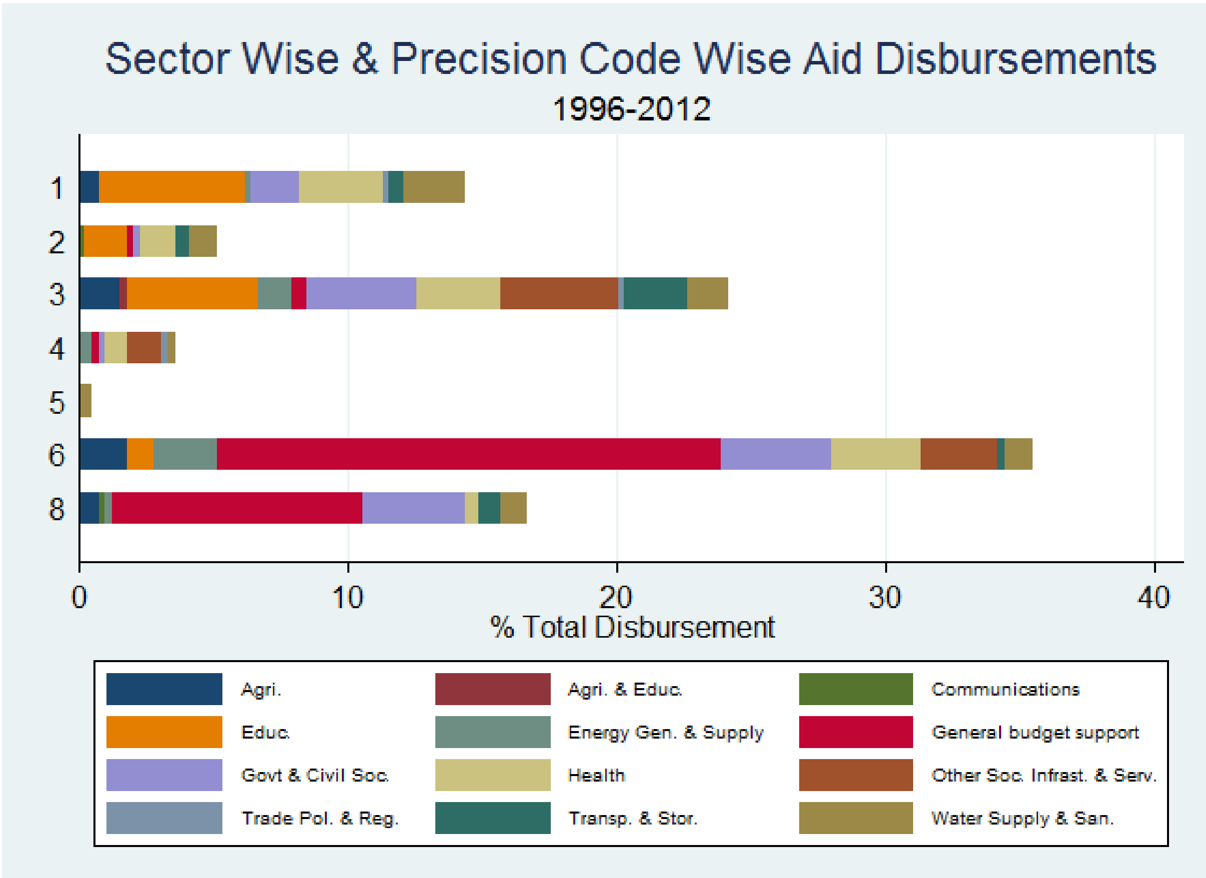

Figure 2: Disbursements classification by geo-location precision and by sector of activity.

In a very limited number of cases (7 periods in 4 districts) there is no luminosity in one year, even though these periods are preceded and followed by very clear nightlight signals. Since the DMSP satellites are meant to record the luminosity emission from stable human-based sources of light, we believe it is more plausible to interpret the lack of observations in these years as missing observations rather than actual zero values. Therefore, we substitute these observations with an *spline* piecewise polynomial interpolation.[14] Overall, 7 of the 35 individual districts have continuous detectable nightlight signals and received aid every year since 1996; the large majority of them receive ODA every year after 2000. Also the synthetic district exhibits positive ODA and nightlight signal for the entire sample.

Household Surveys – We construct the measures of economic activity at the regional level necessary to estimate the predictive stage of Henderson et al. (2012) using the Uganda National Household Surveys (UNHS) administrated in 1999 and 2009.[15] We present estimates based on both the district average household weekly consumption expenditure and the average household monthly expenditure in non-durable goods. Household income data is also available from the UNHS, but given the well-known problems with income data in such contexts (e.g., high levels of informal activity, resistance to disclosing income, and recall error in the absence of written records), we believe that change in expenditure would be a better

[14] The choice of using these substitutions bears a small impact on the final results. We find nearly identical effects when we add a small value to the zero observations instead of applying the interpolation.
[15] These surveys are similar in design to the World Bank Living Standard Measurement Surveys (LSMSs).

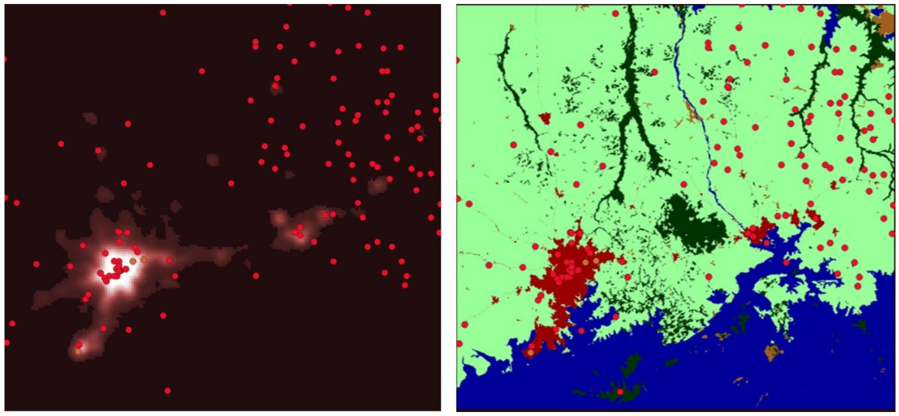

Figure 3: Comparison between land coverage and night luminosity. On the left side, the night lights are reported for on an area of about 30,000 Km^2 around the capital Kampala, for 2010. The red dots indicate ODA projects locations. The intensity of lights is represented in a black to white red dots. On the right side, land coverage for the same area and the same year is shown. The color legend of the different uses is: red for urban coverage, dark green for forests, light green for agriculture and pasture.

indicator of economic growth. Moreover, income and expenditure should be highly correlated in this poor environment with relatively little saving. Since aid disbursements may impact household durable expenditures as well, our impact estimates could be interpreted as conservative with respect to total expenditure growth, and by extension, to income growth.

Following Henderson et al. (2012), we then use the logs of the average expenditure in real 2005 US dollar terms by district as the dependent variable in the fixed-effects long-run model (4), and the log of lights per squared kilometer described above for the independent variable. Paralleling the panel of the main P-VAR estimation, we use the 35 districts and the synthetic district as cross-sectional units in the estimation of (4). The UNHS is designed to be representative at national and macro-region level; however, geographic coordinates are available for each surveyed household. We use this information to assign the households to the correct districts based on the 2010 administrative definition of districts. We then compute the district expenditure as the weighted average of the expenditures reported by the households within a district, using the survey multipliers to construct district rescaled weighting schemes. [16]

Population Series – We also construct district level population series to normalize ODA and night-lights (rather than districts areas) as a robustness checks. The geographic distribution of population at relatively high resolution (about 100 squared meters) is obtained from the Socioeconomic Data and Application Center (SEDAC) at $5-$year frequency for the period $1995-2005$. These are raster images of world population, harmonized with national and sub-national administrative population counts and United Nation

[16]The surveys basically adopt a stratified two-stage sampling design, in which Enumeration Areas (EA) are first sampled with probability proportional to their population relative to the national aggregate, and then households are randomly sampled within each Enumeration Area. On average, we had at least 60-80 observations per district, although a few of the 35 districts are not sampled in either year of the surveys. Since the EA are sub-units of the districts, the sampling procedure does not necessarily cover every district, neither it can guarantee the fully coverage of a district included in the survey. For each district in our panel, we identify all the EA belonging to that district and use the (national) multipliers of the EA to construct the relative within-district weights for the households in each EA. Since the multipliers designed to reflect the representativeness of the EA at national level, this approach is only an approximation, in particular when a small portion of the the EA of a district are sampled. However, the approach is quite satisfactory for the synthetic region and it is arguably more effective than simply treating the households sampled within a district as $i.i.d.$ observations, for example. As a robustness check, we estimate (4) also under this second scenario finding still very significant, but 30% smaller, estimates of ψ.

country statistics, which can be aggregated at the desired geographic unit consistently over time.[17]

The low frequency of observations requires a further manipulation to construct annual time series: we interpolate the 5-year data with an *spline* piecewise polynomial. Though the interpolation returns annual observations, this remains a quite noisy measure of population change, with only a mechanically predicted variability over time. However, since population dynamics are relatively slow and predictable, we can utilize it with some confidence to standardize the other variables of the model, but not directly as a control variable.

4 Empirical Results

As discussed above, our methodology consists of two linked estimation stages. In the first stage we estimate the responses of nightlights to an ODA shock in the P-VAR. In the second stage we map the responses of luminosity to changes in local expenditure in the spirit of Henderson et al. (2012).

4.1 Impulse Response Functions

We first discuss the results for our preferred baseline specification of the model as expressed in equation (1), providing a set of robustness checks at the end of this Section. It is worth noting again that the baseline model is a fixed-effects P-VAR(1) in the logs of ODA and nightlights normalized by districts' areas (as in Henderson et al., 2012). The first observation is that the model has two stable roots largely inside the unit circle, the largest is $.89$, and it strongly satisfies the invertibility condition. This means the non-stationarity of the variables is not an issue with our data and a model in levels is fully satisfactory.

Figure 4 illustrates the impulse response functions of the model for a 10 year horizon to a one standard deviation temporary shock to ODA disbursements, along with their 95% confidence intervals (in gray) computed by Monte Carlo simulation. On the left-hand side of the Figure, we can see the disbursement shock is large and persistent. Similarly, reported on the right-hand side, the response of luminosity is positive, significant, and persistent over the ten year horizon after the shock. Even though the response of night-lights is smaller and decays over the longer horizon, it generates a non-trivial impact on luminosity. The impact shock to log-aid is around $.65$ and translates into an increase of about $.13$ units of log-luminosity five years from the impulse. The peak occurs two years from the impulse at $.15$. These responses correspond to a $5-$year elasticity of about 19.5%.

Figure 5, focuses on the long-run effect of the temporary ODA shock, illustrating the cumulative response of luminosity. The cumulative impact on night-lights is large and statistically strongly significant; numerically, this cumulative response is 1.31 at ten years and 1.59 at fifteen. We can compute the cumulative elasticity of night-light with respect to ODA by comparing this cumulative response to the cumulative change in ODA after the shock, which is 3.96 and 4.64 at ten and fifteen years after the shock. We find that

[17] The raster images are part of the two Gridded Population of the World series provided by SEDAC, the GPWv3 and the GPWFE (CIESIN-CIAT, 2005; CIESIN-FAO-CIAT, 2005).

	Horizon: years									
	1	2	3	4	5	6	7	8	9	10
Light shock	1.00	0.99	0.98	0.96	0.95	0.94	0.93	0.92	0.92	0.91
ODA shock	0.00	0.01	0.02	0.04	0.05	0.06	0.07	0.08	0.08	0.09

Table 1: Variance decomposition of lights under identification ordering: oda, light.

the long-run elasticity settles at 33% and 34% respectively for the ten and fifteen years horizons.

Another way to assess the the long-run effects of ODA on lights is to compute the response of lights to a permanent shift in ODA, which is obtained using the estimates of the coefficients in the coefficient matrix A_1 and of the factorization matrix of the covariance matrix of the estimated reduced-form residuals. The nightlights equation in the structural form of model (1), under our baseline structural identification, can be re-written as:

$$light_{i,t} = \tilde{\alpha} light_{i,t-1} + \tilde{\gamma} oda_{i,t} + \tilde{\beta} oda_{i,t-1} + \tilde{u}_{2,i} + v_{2,i,t} \quad (5)$$

where the coefficients are $\tilde{\alpha}$, $\tilde{\gamma}$, and $\tilde{\beta}$ are combinations of the coefficients in A_1 and the elements of the factorization matrix, $\tilde{u}_{2,i}$ is the fixed-effect of the structural form, and $v_{2,i,t}$ is the structural shock for the light equation (see Footnote 6). The long-run elasticity to a permanent shift in ODA is computed as the ratio $\frac{\tilde{\gamma}+\tilde{\beta}}{1-\tilde{\alpha}}$ and it is estimated to be 35.6%.[18]

As a last result, we look at the variance decomposition of the forecast error of nighlights at different time horizons to assess the relative contribution of ODA and luminosity shocks to the total variance of lights under the same baseline orthogonalization of the shocks structure; this is reported in Table 1. As the time horizon increases, the ODA shock explains a growing share of the volatility of lights: 5% at five years and 9% at ten years. The majority of the variance of observed luminosity is determined by its own shocks, but the aid feedback can cause non-negligible fluctuations as well.

We conclude this Section by conducting a set of six robustness checks for the dependence of the response of night-lights on the specification of the model and treatment of data. Figure 6 illustrates the responses of night-lights to the ODA shock for these six cases. The first check is with the alternative identification scheme of the structural shocks, reported in panel (a), which constrains to zero the response of lights on impact. Panel (a) shows that the response is somewhat smaller, but preserves significance and shape. The second check, in panel (b), excludes aid disbursements from projects that are related to energy generation and power supply network enhancement. These projects could increase lights emission without a direct effect on the real economic activity. The Figure shows that the results are virtually identical.

In panel (c), we use eight lags for the instrumentation of the forward orthogonal transformation of $Y_{i,t-1}$, necessary to correct the bias introduced in a dynamic panel model by the removal of the fixed effects. The Hansen's J-statistic confirms the validity of these instruments at a comfortable level of confidence

[18]Point estimates of the P-VAR coefficients and of the Cholesky decomposition matrix are reported in Appendix A.

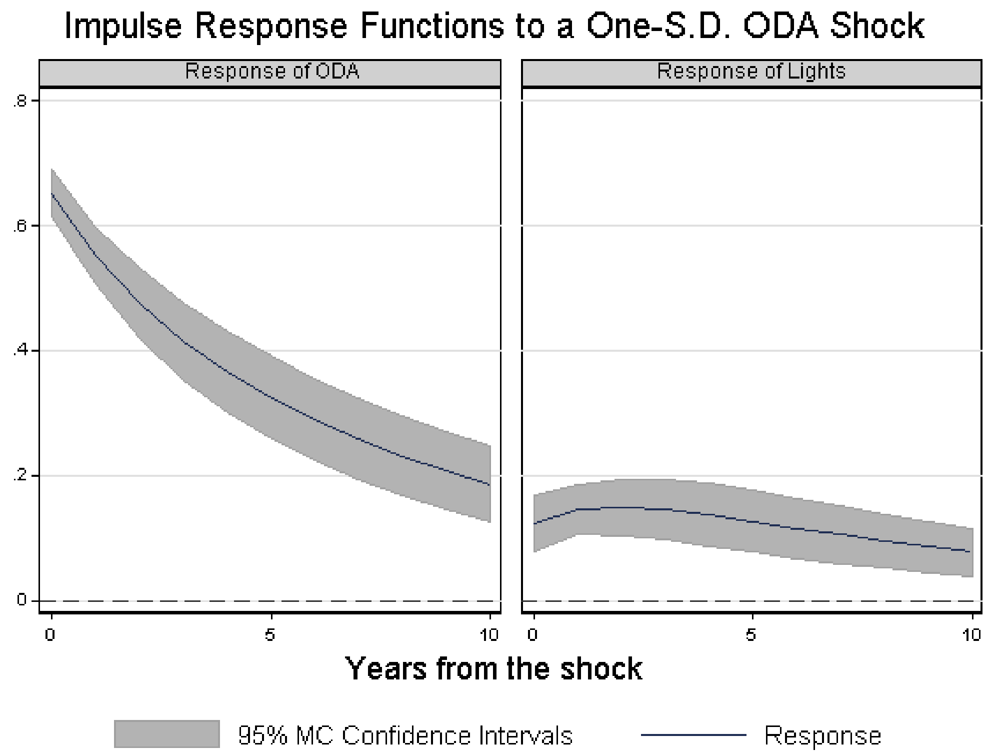

Figure 4: Response functions to a one standard deviation shock to disbursements. Identification ordering: light, oda. Years from the shock on the x-axis.

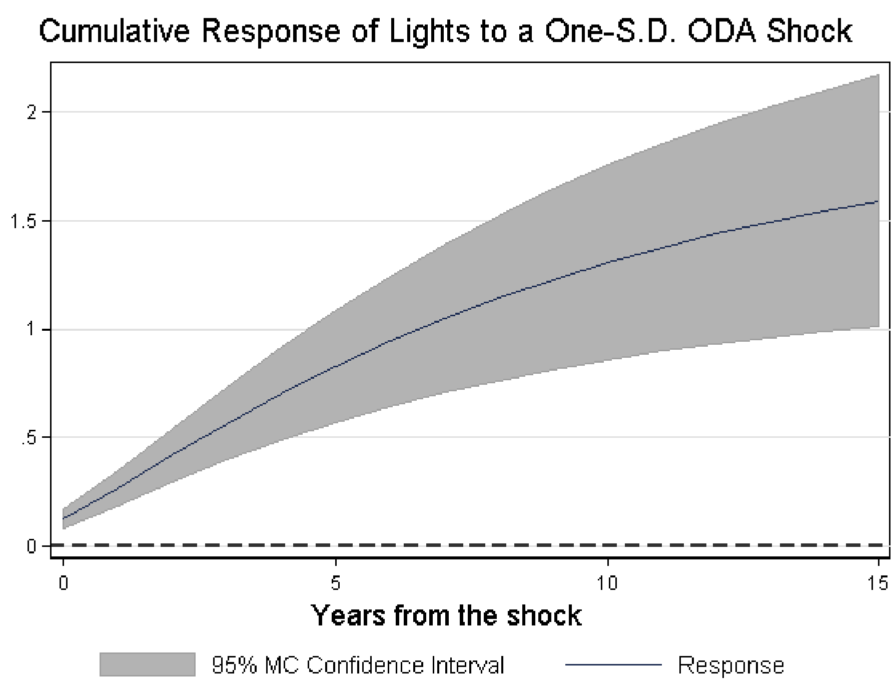

Figure 5: Cumulative response of lights to a one standard deviation shock to disbursements. Identification ordering: light, oda. Years from the shock on the x-axis.

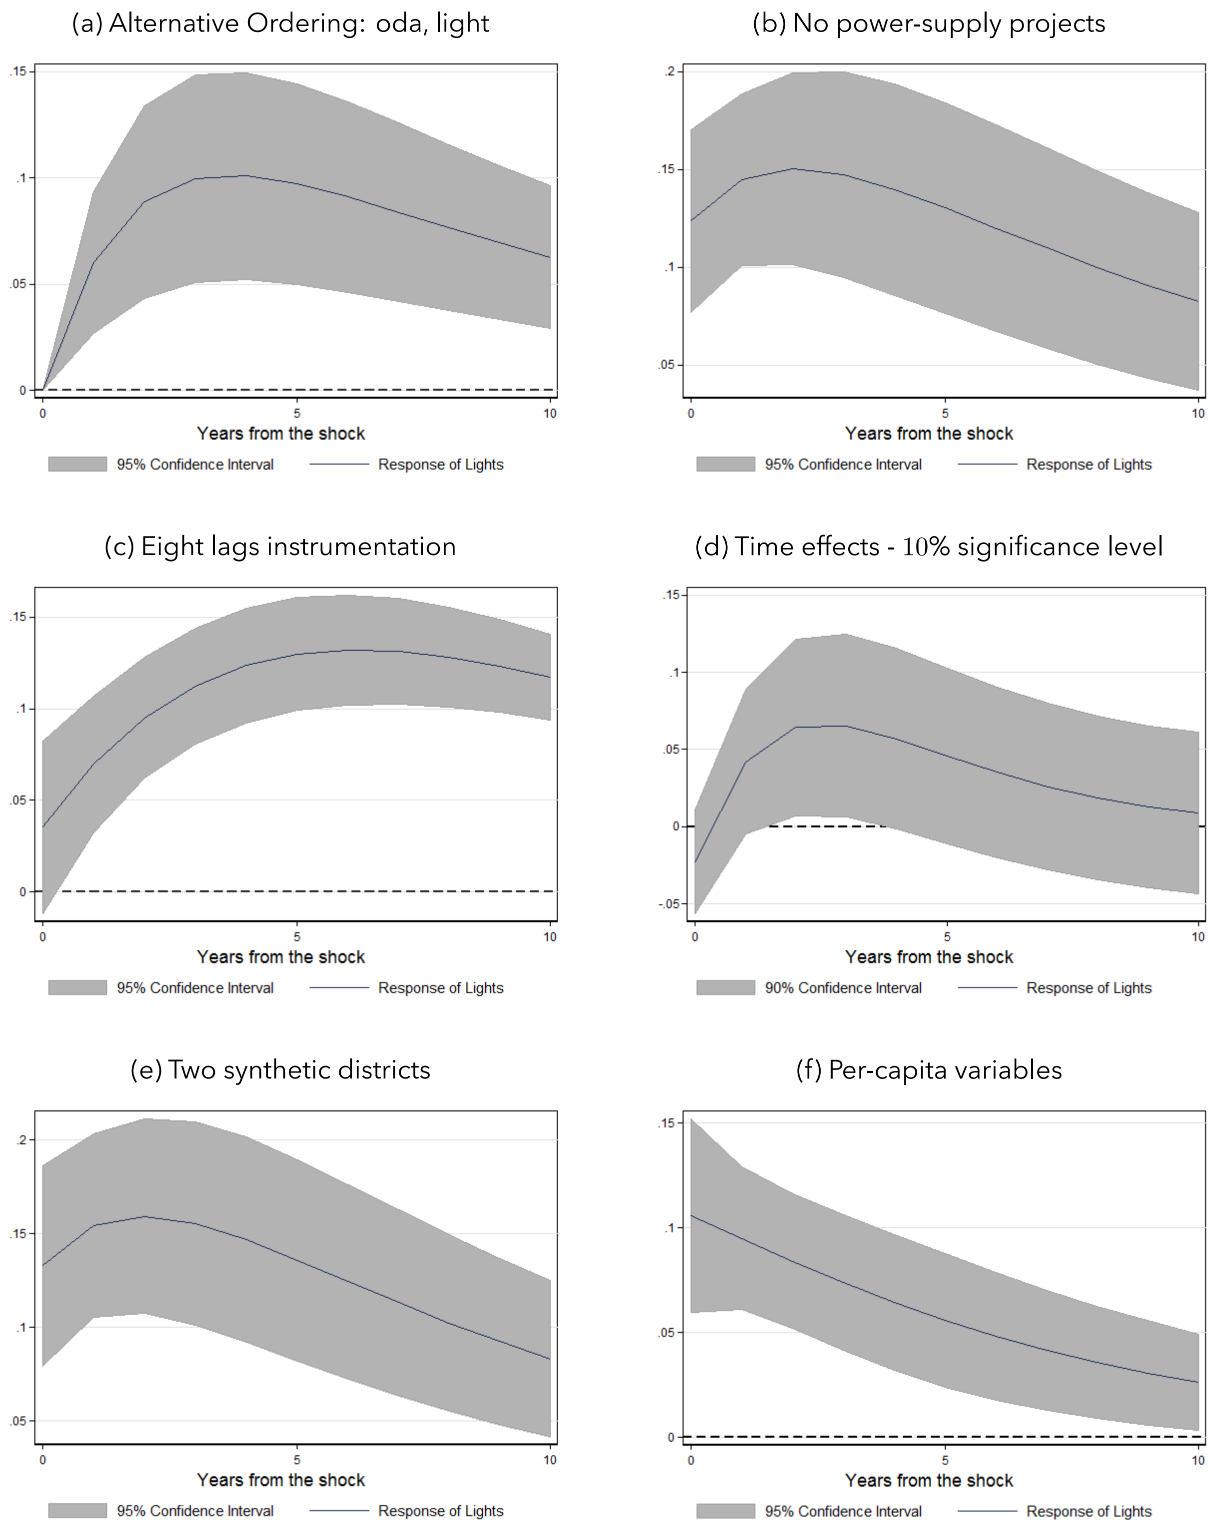

Figure 6: Robustness checks – Responses of nightlights to a one standard deviation shock to ODA. Years from the shock on the x-axis.

(the null hypothesis of joint validity of the instruments is not rejected with a p-value of $.16$). The response becomes more persistent, with a peak at 6 years from the shock, but the overall response preserves a hump-shaped trajectory and exhibits a similar maximum value of $.13$ units of log-luminosity. The fourth robustness check introduces time effects in the panel specification and is illustrated by Panel (d). The use of time effects controls for some time-varying factors common across units that can affect the interaction between ODA and night luminosity, such as national business cycle, political cycles, or government fiscal cycles. However, these effects do not capture, regional transfers from the central government to the peripheries, or other local policies. The effects of the ODA shocks are reduced to one third of those in the baseline specification and the impact response becomes slightly negative, but no longer significant. Importantly, however, the hump-shaped dynamics of the response function is preserved. As a consequence of the weaker effects, the statistical significance of the response decreases as well. The figure shows that the response is significant at 10% between 1 and 4 years.

In the fifth robustness check, we explore the role of the synthetic district. The model in Panel (e) divides the synthetic district in two regions, separating North-Eastern districts from the Southern ones. The model includes 37 cross-sectional units then, and the results are very similar again. Finally, Panel (f) considers variables normalized by district population rather than district area. The response remains significantly positive and persistent, with the same impact response around $.1$. The key difference using this data treatment is that the response is now monotonically decreasing, while it is characterized by a clear hump over the medium horizon in the main results. Our preference for data normalized by district area is based on the evidence by Henderson et al. (2012), who show that night-lights per area is a strong predictor of economic activity. Lights normalized by district surface are also more tightly connected to household expenditure than lights per capita in our sub-national context, and are less distorted by the upper-bound light truncation. This truncation will limit the sensitivity of light to population growth in dense urban areas. Normalization by area, on the other hand, will capture light growth at the land extensive margin.

4.2 Link to Economic Activity

We now turn to the second stage of the analysis that maps the ODA shocks to a traditional measure of economic activity. As discussed previously, given the type of data available from the household surveys, we believe household expenditure is the best measure of economic activity at the district level in this context. Many alternative measures of economic impact and time horizons are possible so these estimates might be considered baseline exercises. Table 2 below reports the ten-year elasticities of two average household expenditure measures to district luminosity (these are the $\hat{\psi}$, in equation 4).

Column (a) is the elasticity of household weekly consumption expenditure to lights and Column (b) is the elasticity of average household monthly expenditure in non-durable goods to lights. We find highly significant estimates in both cases, with a larger elasticity for the expenditure in non-durable goods (39.5% compared to 23.1% for the weekly consumption expenditure). The magnitude of both the effects is also strongly consistent with the $.32$ estimated by Henderson et al. (2012) for the analogous regression at country level.[19]

[19] Recall, they utilize country-level GDP rather than district-level household expenditure.

	weekly expend.	monthly expend.
	(a)	(b)
log(lights/area)	.231	.395
	(.068)***	(.099)***
F.E.	Y	Y
Obs.	70	70
B/W R^2	.32	.57

Table 2: Estimates of the ten-year elasticity of household expenditure to district luminosity from model (4). Column (a): dependent variable is the log of the average real household weekly consumption expenditure; Column (b): dependent variable is the log of the real average household monthly expenditure in non-durable goods. Standard errors indicated in parenthesis, with significance levels of respectively 1%, 5%, and 10% indicated by *** , ** , and *. Survey years: 1999 and 2009.

Using the Table 2 $\hat{\psi}$ below estimates and the baseline P-VAR specification, we can estimate the cumulative effects of an aid shock on household expenditure. The ten-year cumulative elasticity between ODA disbursements and night-lights found in Section 4.1 was 33%. This indicates that a district-level cumulative increase in aid disbursements of 1% over ten years would generate an increase of about 13 basis points in the district average cumulative household expenditure in non-durable goods ten years after the shock. This increase is 7.6 basis points if we measure economic activity through average household consumption expenditure instead. Similarly, a temporary shock to ODA of 1% would cause a response in the average household non-durable expenditure of 7.7 basis points five years from the shock, and 4.5 points in the average household consumption expenditure.

As an alternative perspective on the economic magnitude of these effects, we can convert them into average dollar per-capita terms. For example, in 2009 for which we have both household expenditure and ODA disbursement data, the average 1% increase in per-capita real dollars ODA disbursements across the Ugandan districts corresponds to 4.5 cents per-capita. We find that a temporary positive shock to aid of this magnitude returns an increase of 9.3 cents in per-capita expenditure in non-durable goods and of 8.4 cents in per-capita consumption expenditure over five years.[20] This is equivalent to a multiplier very close to 2. Naturally, the effects of ODA strengthen in the long-run if we consider cumulative estimates. A cumulative increase of per-capita ODA equivalent to 18 cents over ten years produces a cumulative increase in non-durable expenditure of 62.3 cents and an increase in consumption expenditure of 56.3 cents. The multiplier is around $3.2 - 3.4$ in the long-run then. It is worth noting that these results are in line with those reported by Lof et al. (2015), who find an average multiplier effects around 4.5 for a panel of 59 countries.

[20] The 1% increase in ODA in per-capita terms is computed from the corresponding increase in land-unit terms, which is the unit of measure used to compute the compounded elasticities from the empirical exercise. Since districts' areas are constant, the increase in ODA would fully come from a change in disbursements. The increase in ODA is then divided by the 2009 population size of each district to obtain the per-capita value. For the dollar responses of the expenditure measures, we start from the district average households expenditures in the same year. The real average household expenditure in non-durable goods is $\$605$; since the average household size is around 5 people, the per-capita expenditure in non-durable goods is equal to $\$121$. The corresponding values for the average annual consumption expenditure are $\$935$ and $\$187$ respectively. We finally apply the compounded elasticities to these averages.

5 Conclusion

The low-income country recipients of US and other OECD donor of foreign aid contain over 4 billion people, the majority of the global population. Yet the large literature attempting to measure aid impact at the country level using traditional data sources and estimation techniques has produced no consensus on the effects of aid on growth. This research shifts the analysis of aid impact from the country to the sub-national level, combines traditional data sources with remote sensing (satellite) data, and employs an estimation technique that accounts for the endogenous allocation of aid across the sub-national (district) units. We overcome a previous constraint to the use of subnational luminosity data when a subnational administrative unit does not generate a significant night-light signal by creating a synthetic district.

Our regional estimation strategy entails two-stages. In stage one we use a panel vector-autoregressive model to generate the impulse response of luminosity to aid shocks. We find a robust, statistically significant, and persistent response to the shock – though it is of modest magnitude. The second stage uses a traditional regression approach to generate coefficients estimates to map the impulse response to traditional local economic variables. Connecting the two stages indicates that the shock impact on expenditure is small, but non-negligible. We find the multiplier effect of ODA on household expenditure ranges between 2 and 3 in the medium to long-run horizon.

Our approach is highly scalable across location, sector, and outcomes, and it holds promise as a flexible tool for policy analysis. The most immediate opportunity is application of our methodology to the other countries with geo-coded AidData (over fifteen countries at this writing). Examples of scalability beyond location include using this approach to measure local effects of alternative "treatments" (to official foreign aid) that can be tracked over time and for which the auto-regressive methodology would be suitable. For example, a straightforward application would be to measure the local impacts of aid disbursements from large private foundations and NGO. Examples of scalability beyond economic growth could include the impact of a treatment upon investment decisions, health conditions, governance, and the environment. Additionally, the impact of specific categories of aid that we would not expect to have strong light-generating consequences could also be captured by satellite signals other than night-lights. For instance, agricultural land-use-change associated with irrigation or farmers' education projects could be measured using infrared and near-infrared satellite data. This approach is the subject of ongoing research.

References

Addison, T. and F. Tarp (2015): "Aid Policy and the Macroeconomic Management of Aid," *World Development*, 69, 1–5.

Alesina, A. and D. Dollar (2000): "Who gives foreign aid to whom and why?" *Journal of Economic Growth*, 5, 33–63.

Berthélemy, J.-C. (2006): "Bilateral Donorsnterest vs. Recipientsevelopment Motives in Aid Allocation: Do All Donors Behave the Same?" *Review of Development Economics*, 10, 179–194.

Briggs, R. C. (2012): "Electrifying the base? Aid and incumbent advantage in Ghana," *The Journal of Modern African Studies*, 50, 603–624.

Center for International Earth Science Information Network, Columbia University (CIESIN-CIAT) (2005): "Gridded Population of the World Version 3: Population Grids (GPWv3)," Palisades, NY: Socioeconomic Data and Applications Center (SEDAC), Columbia University, downloaded on 4-11-2015.

Center for International Earth Science Information Network, Columbia University; United Nations Food and Agriculture Programme; and Centro Internacional de Agricultura Tropical (CIESIN-FAO-CIAT) (2005): "Gridded Population of the World: Future Estimates (GPWFE)," Palisades, NY: Socioeconomic Data and Applications Center (SEDAC), Columbia University, downloaded on 4-11-2015.

Chen, X. and W. D. Nordhaus (2011): "Luminosity Data as a Proxy for Economic Statistics," *Procedings of the National Academy of Sciences*, 108, 8589.

Civelli, A., A. W. Horowitz, and A. Teixeira (2016): "Is Foreign Aid Motivated by Altruism or Self Interest? A Theoretical Model and Empirical Test," *BE Journal of Economic Policy and Analaysis - Forthcoming*.

Clemens, M. A., S. Radelet, R. R. Bhavnani, and S. Bazzi (2012): "Counting Chickens when they Hatch: Timing and the Effects of Aid on Growth," *The Economic Journal*, 122, 590–617.

Crost, B., J. Felter, and P. Johnston (2014): "Aid under Fire: Development Projects and Civil Conflict," *American Economic Review*, 104, 1833–56.

De, R. and C. Becker (2015): "The Foreign Aid Effectiveness Debate: Evidence from Malawi," AidData Working Paper 15/6.

Deaton, A. (2010): "Instruments, randomization, and learning about development," *Journal of Economic Literature*, 48, 424–455.

Dionne, K. Y., E. Kramon, and T. Roberts (2013): "Aid Effectiveness and Allocation: Evidence from Malawi," Paper Presentation, Conference search Frontiers in Foreign Aidt Princeton University.

Dreher, A., V. Z. Eichenauer, and K. Gehring (2014): "Geopolitics, Aid and Growth," Cepr discussion paper no. dp9904.

Dreher, A. and S. Lohmann (2015): "Aid and Growth at the Regional Level," IMF Working Papers 15/196, International Monetary Fund.

Duflo, E., R. Glennerster, and M. Kremer (2008): "Using randomization in development economics research: A toolkit," in *Handbook of development economics 4*, ed. by T. P. Schultz and J. Strauss, Amsterdan: North-Holland, 3895–3962.

Elvidge, C. D., F.-C. Hsu, K. E. Baugh, and T. Ghosh (2014): "National Trends in Satellite-Observed Lighting: 1992-2012," in *Global Urban Monitoring and Assessment Through Earth Observation*, ed. by Q. Weng, London: CRC Press.

Findley, M. G., J. Powell, D. Strandow, and J. Tanner (2011): "The Localized Geography of Foreign Aid: A New Dataset and Application to Violent Armed Conflict," *World Development*, 39, 199509.

Gillanders, R. (2016): "The Effects of Foreign Aid in Sub-Saharan Africa," *The Economic and Social Review*, 47, 339–360.

Hamilton, S. E. and C. Stankwitz (2012): "Examining the relationship between international aid and mangrove deforestation in coastal Ecuador from 1970 to 2006," *Journal of Land Use Science*, 7, 177–202.

Henderson, J., A. Storeygard, and D. Weil (2011): "A bright idea for measuring economic growth," *American Economic Review*, 101, 1949.

Henderson, J. V., A. Storeygard, and D. N. Weil (2012): "Measuring Economic Growth from Outer Space," *American Economic Review*, 102, 994–1028.

Jablonski, R. (2014): "How Aid Targets Votes: The Effect of Electoral Strategies on the Distribution of Foreign Aid," *World Politics*, 66, 293–330.

Juselius, K., N. F. Moller, and F. Tarp (2014): "The Long-Run Impact of Foreign Aid in 36 African Countries: Insights from Multivariate Time Series Analysis," *Oxford Bulletin of Economics and Statistics*, 76, 1534.

Kilby, C. and A. Dreher (2010): "The impact of aid on growth revisited: Do donor motives matter?" *Economics Letters*, 107, 338–340.

Lof, M., T. J. Mekasha, and F. Tarp (2015): "Aid and Income: Another Time-series Perspective," *World Development*, 69, 19–30.

Love, I. and L. Zicchino (2006): "Financial development and dynamic investment behavior: Evidence from panel VAR," *The Quarterly Review of Economics and Finance*, 46, 190–210.

Mosley, P. (1987): *Foreign Aid, its Defense and Reform*, Lexington, KY: University Press of Kentucky.

Pesaresi, M., G. Huadong, X. Blaes, D. Ehrlich, S. Ferri, L. Gueguen, M. Halkia, M. K. T. Kemper, L. Lu, M. A. Marin-Herrera, G. K. Ouzounis, M. Scavazzon, and P. S. V. S. L. Zanchetta (2013): "A Global Human Settlement Layer From Optical HR/VHR RS Data: Concept and First Results," *Selected Topics in Applied*

Earth Observations and Remote Sensing, IEEE Journal of, 6, 2101–2131.

RADELET, S. (2006): "A primer on foreign aid," Working paper, Center for Global Development.

RAJAN, R. AND A. SUBRAMANIAN (2008): "Aid and Growth: What does the Cross-Country Evidence Really Show?" *Review of Economics and Statistics*, 90, 643–665.

SIMS, C. A. (1980): "Macroeconomics and Realty," *Econometrica*, 48, 1–48.

TEMPLE, J. R. (2010): "Aid and conditionality," in *Handbook of development economics*, ed. by D. Rodrik and M. Rosenzweig, Elsevier, chap. 67, 4415–4523.

TIERNEY, M. J. (2011): "More dollars than sense: refining our knowledge of development finance using AidData," *World Development*, 39, 1891–1906.

Appendix

A Point Estimates from Baseline P-VAR

	$oda_{i,t}$ (a)	$light_{i,t}$ (b)
$oda_{i,t-1}$	.801	.095
	(.039)***	(.025)***
$light_{i,t-1}$	.230	.671
	(.103)**	(.057)***
F.E.	Y	
Obs.	540	
N. of Panels	36	
Instruments	L1	

Table A1: Estimates of the fixed-effects P-VAR(1) model in equation (1). Cross-section units are the Ugandan districts (including the synthetic district); t is expressed in years. The endogenous variables are the logs of the ratio of nightlight to the district surface area, $light_{i,t}$, and the ratio of aid disbursements to the district surface area, $oda_{i,t}$. Column (a): ODA equation; Column (b): nightlight equation. Standard errors indicated in parenthesis, with significance levels of respectively 1%, 5%, and 10% indicated by *** , ** , and *. Sample years: 1996 and 2012.

	oda	$light$
oda	.652	0
$light$	.124	.527

Table A2: Estimates of the Cholesky factorization matrix. Variables order: light, oda.