EBA

09

2017

GEOSPATIAL ANALYSIS OF AID: A NEW APPROACH TO AID EVALUATION

Ann-Sofie Isaksson

Geospatial analysis of aid: A new approach to aid evaluation

Ann-Sofie Isaksson

Rapport 2017:09
Till
Expertgruppen för biståndsanalys (EBA)

Ann-Sofie Isaksson holds a PhD in Economics from the University of Gothenburg and has worked extensively on issues to do with aid, institutional development and governance in developing countries. Her current work is part of the project ”Impacts of new forms of aid: Governance challenges and distributional consequences” at the University of Gothenburg and includes geospatial analyses of aid flows, with a special focus on Chinese aid to Africa.

Acknowledgements

The author would like to thank the members of the EBA reference group – Anders Olofsgård, Andreas Kotsadam, Daniel Strandow and Joakim Molander – for useful suggestions and interesting discussions. Many thanks also to Eva Lithman from the EBA, and to Jan Pettersson and Markus Burman at the EBA secretariat for valuable input, support and smooth cooperation during the working process. Finally, a big thank you to colleagues at the University of Gothenburg for helpful comments and suggestions.

This report can be downloaded free of charge at www.eba.se

ISBN 978-91-88143-31-0
Printed by Elanders Sverige AB
Stockholm 2017

Cover design by Julia Demchenko

Table of Contents

Preface

In everyday conversations and discussions with people on development aid, they are often as interested in whether aid *reaches people living in poverty and under oppression* as in the *long term impact* on the lifes and livelihoods of the same groups. Geo-referenced aid data can potentially strengthen the ability for researchers, evaluators and students to answer *both* these questions and possibly in a cost-effective way. Recent years have seen an increased focus on results and transparency in development cooperation but also criticism from practitioners and scholars studying the "Results agenda", and "Obsessive Measurement disorder" (see EBA 2016:07). Geo-referenced aid data and Geospatial analysis and impact evaluation can make a contribution by offering a comparatively cost-effective use of available aid data where questions of aid allocations and the effects of aid can be analysed at an intermediate level (in-between the local and the national) to complement our knowledge from individual projects and programmes.

Geo-referenced aid data can be used for transparency and accountability purposes, impact evaluations of development projects and programmes, learning, improved planning and as a public good that creates improved conditions for researchers and students around the globe to study Swedish aid in diverse ways. Geo-referenced aid data is of course not a panacea but it can be one way forward trying to improve Sweden's IATI ranking (Publish What You Fund, 2016), to live up to the Swedish government's Aid Transparency guarantee and make evaluation practices in development cooperation more efficient.

The EBA has for this report commissioned Ann-Sofie Isaksson from University of Gothenburg and Örebro University to introduce and discuss this novel approach to aid analysis and evaluation. The author concludes that "not taking advantage of the opportunities geospatial analysis of aid has to offer, and the rapidly expanding geocoded data that is publicly available, would be wasteful". From a Swedish development cooperation perspective, and considering that Swedish aid flows are not yet geocoded on a wider scale, taking advantage of geospatial aid evaluation methods would require geocoding of Swedish aid data.

It is my hope that this report will find its intended audience among policymakers, evaluators and commissioners of evaluations, aid personnel working with transparency in development cooperation, statistics and aid effectiveness.

The author's work has been conducted in dialogue with a reference group chaired by Eva Lithman, member of the EBA. However, the author is solely responsible for the content of the report.

Gothenburg, October 2017

Helena Lindholm

Sammanfattning

De senaste åren har utvecklingssamarbetet blivit mer resultatinriktat, och det har uppstått en intensiv debatt om utvärdering och utvecklingspolitikens effektivitet. Syftet med den här rapporten är att presentera och diskutera en ny ansats för biståndsutvärdering, nämligen analys utifrån geokodade data. Ansatsen går ut på att använda subnationella geokodade data om utvecklingsprojekt och utvecklingsresultat för att utvärdera biståndsallokering och effekter.

Den snabbt ökande tillgången till subnationella geokodade data om biståndsinsatser och om relevanta biståndsresultat har öppnat nya möjligheter när det gäller biståndsutvärdering. Genom att kombinera geokodade biståndsdata – dvs. information om specifika utvecklingsprojekts geografiska läge – med geokodad information om utvecklingsresultat från andra datakällor, exempelvis enkätdata på individ- eller hushållsnivå, blir det möjligt att utvärdera utvecklingsprojektens subnationella fördelning och lokala effekter. Detta kan göras systematiskt och i stor skala, om man så vill för flera mottagarländer samtidigt.

Rapporten är relevant för utvärderare, akademiker och studenter på det utvecklingspolitiska området. Då den behandlar en ny ansats för biståndsutvärdering. Eftersom rapporten innehåller ett nytt tillvägagångssätt som kompletterar befintliga verktyg på området, är den särskilt användbar för organisationer som beställer oberoende utvärderingar av projekt inom utvecklingssamarbetet, både på givarsidan och på samarbetslandets sida. I rapporten diskuteras krav på data och metodfrågor, datatillgänglighet och databegränsningar, empiriska utmaningar och strategier för att hantera utmaningarna. Rapporten är också relevant som en metodologisk introduktion för utvärderare, akademiker och studenter som är intresserade av att själva analysera bistånd utifrån geokodade data.

Geokodade data vid biståndsutvärdering

Studier av biståndseffektivitet kräver en slags avvägning mellan bredd och djup. Det är erkänt svårt att bedöma effekter av det totala biståndet till ett land, samtidigt som det också ofta är problematiskt att göra generaliseringar utifrån resultat som bygger på studier av enskilda projekt. Med hjälp av subnationella analyser utifrån geokodade data är det dock möjligt att anta ett mellanliggande perspektiv. Ansatsen gör det möjligt för forskaren eller utvärderaren att systematiskt bedöma om exempelvis hälsoprojekt fördelas till de områden som har störst behov av hälsoinsatser, om projekten får en effekt på relevanta hälsovariabler i de aktuella områdena och vilka möjliga indirekta effekter som projekten kan ha på lång sikt. På så sätt kan subnationella studier baserade på geokodade data bidra till att överbrygga klyftan mellan mikro- och makroanalyser av biståndseffektivitet.

Geokodade biståndsdata kan användas både för att studera hur bistånd fördelas inom länder och för att utvärdera stödets effekter, så kallad geospatial impact evaluation (GIE). Avseende det förstnämnda så gör geokodade biståndsdata det möjligt att analysera vad som utmärker de platser där utvecklingsprojekt genomförs. Genom att jämföra platser som får bistånd med platser som inte får bistånd i fråga om exempelvis lokal infrastruktur, lokala institutioner och skillnader i invånarnas levnadsstandard kan man utvärdera om insatser når de områden där de berörda behoven är som störst.

God kunskap om den subnationella biståndsfördelningen – om var i länderna biståndsflödena hamnar – underlättar förvaltning, dialog och samordning i biståndsfrågor. Det gör det också lättare att lyfta fram eventuella ojämlikheter i biståndsfördelningen. Därmed kan man, i samarbete med biståndsgivarna, göra samarbetsländerna bättre rustade att samordna givarnas insatser och rikta biståndet till områden där det gör störst nytta. Dessutom är analysen viktig för biståndets transparens och därmed för ansvarsutkrävande. Offentliga kartor över biståndsflöden inom länder kan bidra till att hålla både givare och samarbetsländer ansvariga gentemot avsedda mottagare, då det ger medborgare en möjlighet att kontrollera att projekt genomförs på avsedd plats.

När det gäller effektutvärdering finns det tydliga empiriska utmaningar med att bedöma den kausala effekten av biståndsfinansierade utvecklingsprojekt. Eftersom biståndsmedel generellt sett inte fördelas slumpmässigt är det svårt att skilja ett projekts faktiska kausala effekter från de egenskaper som kännetecknar de individer och områden som biståndsmedlen riktas till. Randomiserade kontrollstudier kringgår detta problem genom slumpmässig exponering för insatser. Allt fler har dock ifrågasatt de randomiserade kontrollstudiernas kostnader samt möjligheten att generalisera utifrån deras resultat. Med tanke på denna kritik, samt det faktum att de allra flesta utvecklingsprojekt där data finns tillgängliga inte är baserade på randomiserade urval, är det viktigt att hitta sätt att bedöma utvecklingsprojektens kausala effekter utan tillgång till randomiserade data.

GIE (Geospatial impact evaluation) möjliggör användandet av kvasiexperimentella metoder för att utvärdera insatsers kausala effekter. Huvuddraget i dessa metoder är, enkelt uttryckt, att hitta kontrollfall som är tillräckligt lika de studerade fallen för att utgöra en lämplig kontrollgrupp. Detta är lättare när man gör jämförelser inom ett mindre område, där enskilda personer som berörs och personer som inte berörs av ett utvecklingsprojekt har liknande förutsättningar (i fråga om institutionella förhållanden, kultur, klimat osv.), än när man jämför grupper längre ifrån varandra. Ansatsen har också fördelar vad gäller generaliserbarhet, både geografiskt och över tid. Geografiska generaliseringar blir möjliga eftersom man med hjälp av GIE kan bedöma effekten av en mängd utvecklingsprojekt samtidigt, eventuellt till och med i flera olika länder. Generaliseringar över tid blir möjliga eftersom GIE ofta kan ta utgångspunkt i resultatdata som täcker långa tidsperioder som möjliggör utvärdering av utvecklingsprojektens hållbarhet och långsiktiga effekter. Slutligen är metoden förhållandevis kostnadseffektiv, eftersom den gör det möjligt för forskare och utvärderare att dra nytta av en stor mängd redan befintliga data. Som diskuteras nedan finns dock databegränsningar att ta hänsyn till.

Datatillgänglighet och databegränsningar

De senaste åren har tillgången till subnationella geokodade data om biståndsinsatser ökat kraftigt. Världsbanken publicerar numera regelmässigt latitud- och longitudkoordinater för alla projekt, och allt fler bilaterala och multilaterala biståndsorgan följer deras exempel. Dessutom har samarbetsländernas ministerier med ansvar för att förvalta inkommande biståndsflöden i allt större utsträckning börjat publicera subnationella geokodade data om utvecklingsprojekt. Geokodade biståndsdata finns också offentligt tillgängliga från organisationen AidData, som i samarbete med ovannämnda institutioner har gjort betydande insatser för att sammanställa subnationella geografiska data om utvecklingsprojekt.

Det finns dock fortfarande betydande databegränsningar att ta i beaktande. För det första är inte alla typer av utvecklingssamarbete lämpliga att analysera utifrån geokodade data. För att ett projekt ska kunna geokodas på ett meningsfullt sätt måste projektet knytas till en fysisk plats. Vissa projekt (t.ex. lokala insatser inom hälsa, utbildning eller lokal styrning) genomförs inom ett tydligt definierat geografiskt område, t.ex. i en eller flera byar, medan andra genomförs på mer övergripande nivåer, t.ex. i ett större administrativt område. Projekt som är mer geografiskt ospecifika, såsom skuldavskrivningsprogram, budgetstöd och sektorstöd, kan inte heller geokodas med någon större geografisk precision. När man drar slutsatser utifrån resultat som bygger på geokodade data från biståndsprojekt är det därför viktigt att vara medveten om de aktuella projektens egenskaper och hur dessa skiljer sig från biståndets egenskaper mer generellt. Dessutom bör luckor i geokodade biståndsdata beaktas vid analyser av allokering och effekter av geokodade projekt. Många givare och genomförandepartners specificerar inte rutinmässigt insatsers geografiska position, och även när de gör det så kan viktig information saknas. Detta begränsar urvalet av projekt som kan analyseras effektivt och gör det svårt att få en helhetsbild av samtliga utvecklingsprojekt i aktuella områden.

För att geokodade biståndsdata ska kunna analyseras på ett meningsfullt sätt måste de kombineras med geokodad information från andra datakällor, t.ex. enkätdata på individ- eller hushållsnivå eller data som bygger på fjärravläsningstekniker. Genom att använda punktkoordinater i geokodade biståndsdata kan utvecklingsprojekt kopplas till exempelvis lokala enkätrespondenter i området. Detta gör det möjligt att identifiera enkätrespondenter som bor i närheten av projektområdet och utvärdera hur dessa respondenter har det i fråga om relevanta utvecklingsvariabler jämfört med andra grupper som inte bor i direkt anslutning till de aktuella områdena.

Precis som med geokodade biståndsdata har tillgången till geokodade data av relevans för att bedöma utvecklingsresultat ökat snabbt under senare år. Tack vare omfattande geokodade surveyprojekt som Afrobarometern, Demographic and Health Survey (DHS) och Världsbankens Living Standards Measurement Study (LSMS) – som alla täcker ett stort antal länder över en längre tidsperiod – är det möjligt att studera en bred uppsättning frågor. Dessa frågor omfattar (för att nämna några exempel) fattigdom, korruption, politiskt deltagande, förtroende för institutioner, politiska attityder, konflikt, våld i hemmet, läskunnighet, fertilitet och hälsa. Dessutom finns det flera källor som tillhandahåller detaljerad geodata om exempelvis befolkningsdynamik, demografi, markanvändning, klimat och andra miljöegenskaper. Befolkningsdata hämtas främst från folkräknings- och enkätdata, och miljöindikatorer från fjärravläsningstekniker såsom satellitdata.

Trots det stora utbudet av geokodade resultatdata finns även här databegränsningar att ta hänsyn till. Framför allt är de frågor som man kan studeras med geodata, utan att samla in ytterligare data, begränsade av den information som finns tillgänglig i befintliga datamaterial. Forskare och utvärderare som utför randomiserade kontrollstudier, eller andra studier där de själva utformar enkätundersökningen, har större kontroll över vilka frågor de kan ställa. En fruktbar strategi kan här vara att kombinera geokodade data med ny insamling av uppgifter. I den här rapporten fokuserar vi dock på det stora utbudet av redan lättillgängliga geokodade resultatdata.

Så kan analyser av geokodad data användas i svenskt utvecklingssamarbete

Sverige ligger i framkant när det gäller biståndstransparens. Regeringen har inrättat en transparensgaranti som fastslår att alla offentliga handlingar och all offentlig information om svenskt utvecklingsbistånd ska göras tillgängliga på internet (se Openaid.se). Än så länge är dock inte det svenska utvecklingssamarbetet geokodat i någon större utsträckning. Detta kan komma att ändras. I Sidas årliga rapport från 2016 noteras att de datahanteringssystem som ligger till grund för OpenAid har förbättrats, och att Sidas insatser nu kan kodas utifrån deras exakta geografiska position.

Medan det finns många argument som talar för en geokodning av bistånd, finns en rad praktiska frågor, kostnader och ansvarsfördelning att ta hänsyn till. Eftersom analys utifrån geokodade data inte kan tillämpas på alla typer av utvecklingsprojekt bör man, innan eventuell geokodning inleds, undersöka dess relevans och potential för olika delar av biståndsportföljen. När man granskar den sektorsvisa fördelningen av det svenska biståndet kan man konstatera att geokodningens mervärde och genomförbarhet kan ifrågasättas för vissa sektorer (i synnerhet gäller det bistånd till flyktingar i givarlandet och den ospecificerade delen av biståndet). Om man räknar bort det bistånd som går till sektorer där möjligheten att följa pengarna eller där biståndsinsatsernas geografiska spridning är osäker så återstår dock fortfarande över halva portföljen för utvecklingssamarbete under 2016. Även om allt detta bistånd rimligtvis inte är geokodningsbart återfinns här troliga kandidater för geokodning.

För att bedöma vad det skulle kosta att geokoda tidigare och pågående projekt kan man skilja mellan geokodning på portföljnivå och detaljerad geokodning av enskilda projekt. Genom att anlita erfarna datakodare för att utvärdera projektdokumentation går det att få uppskattningar av geografiska platser för en mängd projekt i en givares biståndsportfölj till en relativt låg kostnad. För specifika projekt som givaren eller samarbetslandet är särskilt intresserade av att utvärdera kan detaljerad geokodning leda till mer precisa geografiska data som kan inbegripa projektets exakta geografiska avgränsningar, men till en högre kostnad per projekt.

Geokodade biståndsdata, som kan beskrivas som kollektiva varor, väcker slutligen frågan om vem som ska tillhandahålla resurserna och vem som ska utföra arbetet – givarna eller samarbetsländernas regeringar? Här är till synes samarbete parterna emellan det bästa alternativet. Med tanke på att utvecklingsinitiativ och deras resultat behöver lokalt ägarskap finns det goda skäl att hävda att det yttersta ansvaret bör ligga hos samarbetslandets ministerium med ansvar för att förvalta inkommande flöden. Samarbetslandet brottas dock ofta med kapacitets- och resursbegränsningar som utgör betydande hinder för geokodning. Framför allt kan samarbetsländerna behöva omfattande stöd för att hantera den inledande investeringen för att komma ikapp med geokodningen av tidigare och pågående aktiviteter i syfte att skapa data fram till idag. I ett senare skede kommer de att vara beroende av samarbete med givare för att kunna säkerställa partners kontinuerlig rapportering av geografisk information för framtida projekt.

Slutsatser

Det skulle vara slöseri att inte utnyttja de analysmöjligheter som ges av den snabbt ökande tillgången till offentligt tillgängliga geokodade data. Ur en forskares, student eller utvärderares perspektiv ger det stora utbudet av tillgängliga geokodade data om utvecklingsinitiativ och utvecklingsresultat stora möjligheter att studera systematiska mönster i biståndsfördelningen och utvecklingsprojekts lokala utvecklingseffekter på kort och lång sikt.

Om utvärderingsansatser baserade på geokodade data ska kunna användas i det svenska utvecklingssamarbetet behövs dock satsningar på geokodning. Med tanke på att Sverige ligger i framkant när det gäller givartransparens verkar detta vara ett rimligt åtagande. Ett lämpligt första steg är att gå igenom och sammanställa redan tillgängliga geokodade data för svenska biståndsflöden och införliva dessa i de svenska registren. Detta är relevant, då vissa samarbetsländers ministerier redan nu geokodar inkommande biståndsflöden och med tanke på att Sverige ofta samfinansierar projekt med andra givare som i sin tur geokodar sitt bistånd. Om det finns ett intresse av att gå vidare därefter så finns det olika alternativ, som inte på något sätt utesluter varandra.

Ett alternativ är att anlita datakodare för att göra en bred geokodning av avslutade och pågående projekt på portföljnivå. Fördelarna med detta är att det ger en överblick över var det svenska biståndet hamnar, vilket gör det möjligt att dela den sammanställda informationen med mindre resursstarka samarbetsländer som är intresserade av att offentliggöra geokodad information om inkommande bistånd.

Ett annat alternativ är att geokoda ett antal specifika insatser av särskilt intresse på ett mer detaljerat sätt. Det bidrar inte i lika hög grad till den kollektiva nyttan som att offentliggöra större mängder geokodade biståndsdata, men det kan vara ett överkomligt första steg som gör det möjligt för parterna att bilda sig en uppfattning om geokodning och vad sådana data kan användas till.

Ytterligare ett alternativ skulle kunna vara att stödja samarbetsländernas arbete med att geokoda inkommande biståndsflöden. I synnerhet kan de behöva assistans för att hantera den inledande investering som krävs för att komma ikapp med geokodningen av tidigare och pågående insatser i syfte att skapa data fram till nutid.

När man har färdigställt arbetet med att geokoda slutförda projekt finns det starka argument för att samarbetslandets regering ska leda arbetet med att utveckla det lokala systemet för geokodning av biståndsdata. Det är emellertid viktigt att påpeka att detta kräver fortsatt samarbete från utvecklingspartners sida med att kontinuerligt rapportera projekts geografiska information i den mån det är möjligt. Det är i detta sammanhang viktigt att skapa enkla rapporteringsrutiner. En viktig faktor är att undvika dubbla eller ej integrerade rapporteringssystem. Det är därmed också ett rimligt åtagande att stödja samarbetsländerna i arbetet med att utveckla system för biståndsdata i enlighet med det internationella initiativet för öppenhet i biståndet (International Aid Transparency Initiative, IATI).

Summary

Recent years have seen an increased focus on results in development cooperation, and a heated debate on the evaluation strategies and effectiveness of development policies. Against this background, the aim of this report is to introduce and discuss a novel approach to aid evaluation, namely geospatial analysis of development cooperation, utilizing subnational geocoded data on development projects and development outcomes to evaluate the allocation patterns and effects of development projects.

A rapid expansion in the availability of sub-nationally geo-referenced data on aid interventions as well as of geocoded data on relevant outcomes and covariates of aid opens for new possibilities in terms of aid evaluation. Combining geocoded aid data – that is, information on the geographic location of specific development projects – with geocoded information from other data sources, such as individual/household level survey data, makes it possible to evaluate the sub-national distribution and local effects of development projects systematically and on a wide scale, potentially across multiple recipient countries.

Introducing a new approach to aid evaluation, the report is relevant for policy evaluators, academics and students of development policy alike. Specifically, providing an overview of a novel methodological approach to include in the aid evaluation toolbox, it is useful for organisations commissioning independent evaluations of development cooperation projects, both on the donor and partner country side. Furthermore, discussing data requirements and methodological concerns, including data availability, data limitations, empirical challenges and strategies to deal with these, it is relevant as a methodological introduction for policy evaluators, academics and students interested in embarking on geospatial analysis of aid themselves.

The role of geospatial analysis in aid evaluation

Analyzing aid effectiveness there is a tradeoff between scope and depth. While estimating country-wide impacts of total aid is notoriously difficult, it is often problematic to generalize from findings based on studies of individual projects. Sub-national geospatial analysis provides an intermediate perspective. It enables the researcher or policy evaluator to systematically estimate, for instance, whether health projects are allocated to the sub-national areas in greatest need of health interventions, whether they have direct effects on relevant health outcomes in the targeted areas, as well as their potential indirect effects on other relevant development outcomes over the longer term. As such, sub-national geospatial analysis can arguably help bridge the micro-macro divide and thus fill a 'missing middle' in aid evaluation and in the academic aid effectiveness literature.

Geospatial analysis of development cooperation is taken to include the use of geocoded aid data both to systematically evaluate sub-national aid allocation patterns and to assess the effects of aid, i.e. geospatial impact evaluation (GIE). With respect to the former, geocoded aid data provides a tool to evaluate the sub-national distributional consequences of aid flows. With access to geocoded data on relevant covariates from before implementation of the concerned development projects started one can assess systematic variation in pre-existing characteristics of localities where development projects are implemented. To mention a few examples, this could entail comparing localities receiving and localities not receiving aid in terms of differences in local infrastructures and institutions as well as differences in the living standards of citizens residing in the areas, e.g. in terms of local poverty rates, disease burden and employment opportunities. For instance, do health interventions reach the areas within countries where health needs are the greatest?

From a development cooperation perspective, being well-informed about where aid flows go within countries comes with benefits in terms of aid management, dialogue and coordination.

Having a clear picture of the sub-national allocation of aid can help highlight potential financing gaps and inequalities in the aid distribution and thus make partner countries, in cooperation with donors, better able to coordinate donor efforts and direct aid flows to the areas where they will do most good. Furthermore, it is important in terms of transparency and thus for accountability. Publicly available mappings of aid flows within countries can help hold both donors and partner country governments accountable to their intended beneficiaries, enabling citizens to verify that projects are being implemented in their intended locations.

Turning to impact evaluation, estimating the causal impact of aid-funded development projects comes with empirical challenges originating in the fact that aid allocation is systematic as opposed to random. This implies that some individuals and sub-national areas, with certain characteristics, will be more likely to be targeted by aid than others, and consequently that it is difficult to distinguish these characteristics from the actual causal effects of development projects. Randomized controlled trials (RCTs) get around this problem by randomly assigning exposure to aid interventions. Increasingly, however, critics have raised concerns concerning the cost of RCTs and the extent to which one can generalize from their findings. Furthermore, for the great majority of past and present development projects for which there are valuable data records available, exposure has indeed not been determined through random assignment. It is thus important to find ways to assess causal impacts of development projects for the abundance of cases where randomized data on interventions is not feasible or available.

Geospatial impact evaluations have the advantage that they can control for potential confounding factors at granular geographic levels and that they are well-suited for quasi-experimental methods for causal identification. The key feature of these methods is, simply stated, to find 'control' cases that are sufficiently similar to the 'treated' cases to constitute a viable comparison group. This is easier when making comparisons within the local area, where individuals affected and not affected by a development project face similar conditions on other accounts (in terms of institutional arrangements, culture, climate, etc.), than when comparing groups further apart. Furthermore, GIEs are relatively strong in terms of generalizability – both across space and over time. In a spatial sense since GIE makes it possible to estimate the impact of a multitude of development projects, potentially across several countries, and in a temporal sense since GIEs often draw on outcome data that covers long time periods, and thus make it possible to evaluate the sustainability and long-term impacts of development projects. Finally, since the approach enables researchers and evaluators to utilize comprehensive existing data materials, it is comparatively cost-effective.

Data availability and limitations

The last few years have seen a sharp increase in the availability of sub-nationally georeferenced data on aid interventions. The World Bank now regularly publishes the latitude and longitude coordinates of all of its projects, and a growing number of bilateral and multilateral aid agencies have since followed suit. Furthermore, partner country ministries responsible for managing incoming aid flows increasingly publish sub-nationally geocoded development project data. The geocoded aid data is publicly available from AidData, who in collaboration with the above mentioned institutions have made significant efforts to map and synthesize sub-national data on development projects.

Still, there are important data restrictions to consider. First, geospatial analysis is clearly not appropriate for all types of development cooperation. To be able to geocode a project in a meaningful way it has to have a physical project site. And whereas some projects (say local interventions in terms of health, education or local governance) are implemented in a well-defined geographical area, such as one or several villages, others are realized at more aggregate levels, such as a district or greater administrative region. Furthermore, projects that are more intangible, consider e.g. debt-relief agreements, budget- and sector support, are not possible to geocode with any greater geographical precision. Hence, when generalizing from findings based on geocoded development project data, the characteristics of the covered projects and how these may differ from the characteristics of overall aid need to be kept in mind. Furthermore, gaps in the geocoded aid data should be considered when analyzing the allocation and impacts of geocoded projects. Many donors and implementing partners do not routinely map their intervention sites, and even when they do, important information may be missing. This restricts the sample of projects that can be analysed effectively and makes it difficult to get a full picture of all development projects located in the area.

To be able to analyze geocoded aid data in a meaningful way it has to be combined with geocoded information from other data sources, such as the individual/household level survey data or data based on remote sensing technologies. By using the point coordinates in the geocoded aid data, development projects can be liked to e.g. local survey respondents in the area, making it possible to identify survey respondents living near project sites, and evaluate how these respondents fare on relevant outcomes compared to other groups not living in immediate connection to the specified development project sites.

As with geocoded aid data, recent years have seen a rapid expansion in the availability of geocoded data on relevant development outcomes. Comprehensive geocoded survey projects such as the Afrobarometer, the Demographic and Health Survey (DHS) and the World Bank Living Standards Measurement Study (LSMS), which cover a great number of countries over an extended period of time, makes it possible to address a wide range of issues, including (to mention just a few examples) poverty, corruption, political participation, trust in institutions, political attitudes, conflict, health, domestic violence, literacy, fertility and disease prevalence. In addition there are numerous sources providing high resolution spatial data on e.g. population dynamics, demographics, land use, land cover, climate and other environmental features. The population data is primarily derived from census and survey data and the environmental indicators from remote sensing technologies such as satellite data.

While there is indeed a wealth of relevant geocoded data on development outcomes that can be used for geospatial analysis of aid flows, there are again data restrictions to take into account. In particular, the questions one can address with geospatial data, without engaging in further data collection, is limited by the information available in existing survey materials and data bases. Researchers and policy evaluators doing RCTs, or other studies involving designing the survey instrument yourself, have more control in terms of the questions they can ask. Although a fruitful approach can be to combine existing geospatial data with new data collection initiatives, in this report, we restrict our attention to the wealth of readily available geospatial outcome data.

The potential for geospatial analysis of Swedish development cooperation

Sweden is at the forefront in terms of donor transparency, with a government established 'Transparency guarantee' specifying that all public documents and public information on Swedish development assistance will be made available online (see Openaid.se). As of yet, however, Swedish development cooperation is not geocoded on a wider scale. This may be about to change. In Sida's 2016 annual report it is noted that improvements have been made in the data management system underlying OpenAid, and that Sida's interventions now can be marked in terms of their exact geographical position.

While the arguments for geocoding aid are many, there are practical concerns and cost considerations to take into account, as well as the issue of how to divide the responsibility for geocoding efforts. First, since geospatial analysis cannot be applied to all types of development projects, efforts to geocode development projects should be preceded by an initial screening of the relevance of and potential for geocoding different parts of the aid portfolio. Exploring the sectoral allocation of the Swedish aid portfolio, we can note that for some sectors, the value and feasibility of geocoding is questionable (in particular, aid to refugees in the donor country and the 'unallocated/unspecified' share). For other sectors, however, a considerable share of the aid activities should be likely candidates for geocoding. Deducting aid going to sectors where the ability to follow the money or the spatial differentiation of aid interventions is questionable, over half of the 2016 development cooperation portfolio still remains.

Assessing the costs of geocoding past and present projects, next, one can make a distinction between portfolio level geocoding and detailed geocoding of individual projects. Hiring experienced coders assessing project documentation, one can get estimates of geographic locations for a wide range of projects in a donor's aid portfolio at a relatively low cost. However, for specific projects that the donor or partner country may be particularly interested in evaluating, detailed geocoding could result in more precise geographic measures, potentially involving exact project location boundaries, but at a higher cost per project.

Finally, with the public good character of geocoded aid data follows the question of who should supply the resources and who should do the work – the donors or the partner country governments? Here, joint efforts are seemingly needed. Given local ownership of development initiatives as well as of results, there are good arguments for saying that the ultimate responsibility should lie with the partner country ministry in charge of managing incoming inflows. That said, though, the partner country is likely to face capacity and resource constraints that could constitute important obstacles for geocoding. In particular, they may require significant support to be able to handle the initial investment in data creation in order to catch up to the present point in time by geocoding past and ongoing activities, and at a later stage, the cooperation and compliance of development partners in the continuous reporting of geographic information of upcoming projects.

Implications

Not taking advantage of the opportunities geospatial analysis of aid has to offer, and the rapidly expanding geocoded data that is publicly available, would be wasteful. From a researcher's, student's or evaluator's viewpoint, the abundance of publicly available geocoded data on development initiatives and development outcomes provides ample opportunities for studying systematic aid allocation patterns as well as short and long term local development impacts of projects.

From a Swedish development cooperation perspective, taking advantage of geospatial aid evaluation methods requires geocoding efforts. With Sweden being a frontrunner in terms of donor transparency, this appears a reasonable undertaking. A suitable first step is to screen and compile already available geocoded data pertaining to Swedish aid flows, and incorporate it in the Swedish records. This is relevant seeing that some partner country ministries now geocode incoming aid flows and considering that Sweden often co-fund projects with other donors who geocode their aid. Being interested in proceeding from here, there are different – by no means mutually exclusive – options.

One option is to hire coders to do broad portfolio level geocoding of past and ongoing projects. This comes with advantages in terms of getting an overview of the destinations of Swedish contributions and of being able to share the compiled information with strained partner country ministries interested in publishing geocoded information on incoming aid flows.

Another option is to geocode specific projects of particular interest in a more detailed manner. While not to the same extent contributing to the public good that comprehensive publicly available geocoded aid data constitutes, it can be seen a feasible first step to get a sense of the geocoding process and what the data can be used for.

Yet another option would be to provide support to partner country initiatives to geocode incoming aid flows, in particular, for partner country ministries to be able to handle the initial investment in data creation in order to catch up to the present point in time by geocoding past and ongoing activities.

Having cleared the backlog of geocoding past projects, there are strong arguments for the partner country government taking the lead in the development of their local Aid Information

Management System. Importantly though, this will still require the cooperation and compliance of development partners in the continuous reporting of geographic information of upcoming projects. Updating the reporting routine of project managers it is important to provide simple ways to report geographic information. A key point should be to avoid dual, unaligned, reporting systems. Supporting partner countries' in developing their Aid Information Management Systems in line with the International Aid Transparency Initiative (IATI) standards is thus a sensible undertaking.

1. Introduction

Recent years have seen an increased focus on results in development cooperation, and a heated debate on the evaluation strategies and effectiveness of development policies. In a Swedish development cooperation perspective, EBAs first ever report (Olofsgård, 2014), which focused on the strengths, weaknesses and policy relevance of randomized controlled trials (RCTs), made a valuable contribution to this debate. While Olofsgård argued that RCTs are a powerful tool to evaluate the impact of aid, he also pointed to important limitations, for example in terms of generalisability and resource requirements. Similarly, in her recent EBA-report on Qualitative Comparative Analysis, Befani emphasises that rigorous quantitative research designs often are expensive, 'sometimes prohibitively so' (Befani, 2016, p.14).

The aim of this report is to introduce and discuss a novel and complementary approach to aid evaluation, namely geospatial analysis of aid flows, utilizing subnational geocoded aid, outcome, and covariate data to evaluate the allocation patterns and effects of development projects. The last few years have seen a rapid expansion in the availability of sub-nationally geo-referenced data on aid interventions as well as of geocoded data on relevant outcomes and covariates of aid. This opens for new possibilities in terms of aid evaluation. Combining geocoded aid data – that is, information on the geographic location of specific development projects – with geocoded information from other data sources, such as individual/household level survey data, makes it possible to evaluate the sub-national distribution and local effects of development projects systematically and on a wide scale, potentially across multiple recipient countries (Dreher and Lohmann, 2015; Dreher et al., 2016; Isaksson and Kotsadam, 2016 and 2017; Briggs, 2017).

Given the vast amount of resources involved, it is unsurprising that aid effectiveness remains a subject of controversy for academics and policy makers alike (Qian, 2015). Unfortunately, though, the mixed empirical evidence on aid effectiveness arguably exacerbates rather than resolves the controversy. In particular, there has traditionally been a sharp divide between the macro and micro literatures evaluating the impact of aid. Indeed, it is common to speak of a micro-macro paradox (Mosley, 1987), highlighting the tendency of scholars to be able to identify positive impacts of individual development projects but a difficulty to establish corresponding positive effects at the macro level.

Sub-national geospatial analysis provides an intermediate perspective that can help bridge the micro-macro divide and thus fill the 'missing middle' in aid evaluation and in the academic aid effectiveness literature. Specifically, rather than estimating country-wide impacts of total aid, which is notoriously difficult, or being restricted to analyzing the impact of single projects, geospatial analysis of aid enables the researcher or policy evaluator to systematically estimate, for instance, whether health projects are allocated to the sub-national areas in greatest need of health interventions, whether they have direct effects on relevant health outcomes in the targeted areas, as well as their potential indirect effects on other relevant development outcomes over the longer term.

Introducing a new approach to aid evaluation, the report is relevant for policy evaluators, academics and students of development policy alike. Specifically, providing an overview of a novel methodological approach to include in the toolbox to achieve more evidence based aid policy, the report is relevant for organisations commissioning independent evaluations of development cooperation projects, both on the donor and partner country side. In a Swedish development cooperation context, this would include Sida's evaluation function, the Ministry for Foreign Affairs (MFA), the Expert Group for Aid Studies (EBA), civil society organisations (such as The Swedish Committee for Afghanistan, SCA) and relevant partner country counterparts. Furthermore, discussing data requirements and methodological concerns, including data availability, data limitations, empirical challenges and strategies to deal with these, it is relevant as a methodological introduction for policy evaluators, academics and students interested in embarking on geospatial analysis of aid themselves. Considering its relatively broad audience, the report can be read selectively, depending on the particular interests and background of the reader. Some reading guidelines are provided below.

The next chapter discusses the role of geospatial analysis for studying aid effectiveness, including what gap it fills in relation to existing approaches and its main strengths and limitations. Chapter 3 provides a brief overview of the emerging academic literature using geocoded aid data. Chapters 4-6 focus on data and empirical considerations and are, as such, somewhat more technical. Readers who will not apply geospatial aid analysis themselves could choose to skip or browse through these chapters and still get the main message of the report. In particular, Chapter 4 discusses data considerations, including the availability of geocoded aid and outcome data, geographical matching of these data sources, and data limitations. Chapter 5 discusses empirical challenges in identifying a causal impact of development projects, and examples of strategies to deal with these within a geospatial framework. Chapter 6 presents an empirical application, focusing on the impact of aid on citizen political participation, providing some illustrative examples of how one can break the aid data down by donor, locality and sectoral focus. Chapter 7 discusses the potential for using geospatial analysis from a Swedish development cooperation perspective. Given that Swedish aid is not yet geocoded, this will involve evaluating the potential for geocoding Swedish-funded projects and how such data could be used. Chapter 8, finally, summarizes the main findings of the report and discusses implications for research, evaluation and Swedish development cooperation.

2. The role for geospatial analysis in aid evaluation

The focus of this report is the use of geospatial analysis in aid evaluation. This is taken to include the use of geocoded aid data to 1) systematically assess sub-national aid allocation patterns and 2) to assess the effects of aid, i.e. what has become known as geospatial impact evaluation (GIE).

2.1 The role of geospatial analysis to assess sub-national aid allocation patterns

The academic literature assessing aid allocation patterns has traditionally focused on country level variation in aid flows, evaluating whether aid goes to relatively poor and well-governed countries or whether aid flows are driven by economic and political ties between donor and partner countries. The geocoded aid data now available on a wider scale enables scholars to go below the country level and evaluate sub-national aid allocation patterns, addressing questions such as 'does aid reach the poorest within countries?'.

With access to data on relevant covariates from before implementation of the concerned development projects started one can assess systematic variation in pre-existing characteristics of localities where development projects are implemented (as opposed to differences that may have arisen as a result of the development projects). To mention a few examples, this could entail comparing localities receiving and localities not receiving aid in terms of differences in local infrastructures and institutions as well as differences in the living standards of citizens residing in the areas, e.g. in terms of local poverty rates, disease burden and employment opportunities. That is, factors that are relevant to consider when assessing to what extent aid is allocated to where it is most needed. For instance, do health interventions reach the areas within countries where health needs are the greatest – or specified even further, do malaria or HIV interventions reach the areas where these diseases are most prevalent – and do aid reach the very poorest areas within partner countries.

As such, geocoded aid data provides a tool to evaluate the sub-national distributional consequences of aid, and to assess to what extent the allocation pattern is needs-based or driven by other factors (e.g. based in political considerations or access to infrastructure). Section 3.1 provides a brief overview of the emerging academic literature using geocoded aid data to systematically evaluate sub-national aid allocation patterns.

From a development cooperation perspective, being well-informed about where aid flows go within countries, i.e. having a clear picture of the sub-national distribution of aid, comes with benefits in terms of aid management, dialogue, coordination and transparency (see Strandow et al., 2011). In terms of aid management, having a better understanding of the sub-national allocation of aid can highlight potential financing gaps and inequalities in the aid distribution, and thus help to ensure that aid flows to those who need it most. It can also help in the dialogue with partner countries. Sub-national data on aid allocation patterns can provide recipients with a better understanding of how and where aid should best work to their benefit, making them better able to direct aid to areas which may otherwise have been neglected. Relatedly, it should come with benefits in terms of donor coordination – having a clear picture of the aggregate distribution of aid within a country should make partner country governments, in cooperation with donors, better able to coordinate donor efforts, e.g. by letting some specialize in certain regions or by eliminating potentially wasteful project duplication. In terms of transparency, publicly available mappings of aid flows within countries will help hold both donors and recipient governments accountable to their intended beneficiaries. Through information and accountability mechanisms such as the media, audits, and research, the data can help citizens verify that projects are being implemented in their intended locations, and thus to reduce waste and to increase aid effectiveness.

2.2 The role of geospatial analysis in studying aid effectiveness

Seeking to establish a causal impact of development projects, as in studies on aid effectiveness, comes with greater empirical challenges and thus with more controversy in terms of the empirical methods used. Considering the macro and recent micro studies in the academic literature on aid effectiveness, the former tend to be criticized in terms of their ability to establish causal effects as opposed to mere correlation patterns (i.e. in terms of their internal validity), and the latter in terms of their ability to generalize from findings based on studies of individual projects (i.e. in terms of their external validity). Sub-national geospatial analysis provides an intermediate perspective that can help bridge the micro-macro divide in the aid effectiveness literature.

While useful for uncovering broad patterns, the macro literature on aid effectiveness, focusing on country level relationships between aid inflows and outcomes, is problematic for several reasons. First, it is plagued by difficulties in establishing causality. Aid is not distributed at random. Hence, receiving aid is associated with a multitude of country characteristics – known and unknown – that will tend to influence the estimates when seeking to establish the causal impact of aid (see e.g. Bräutigam and Knack, 2004). To take the most obvious example, finding a relationship between high poverty rates and high aid inflows should not be interpreted as aid making countries poor. Similarly, being interested in, say, the effect of aid on economic growth, it is of course difficult to separate the impact of aid from the effects of problems that are common in aid receiving countries. Second, the cross-country aid effectiveness literature tends to ignore the heterogeneity of aid and aggregate over aid flows that may have very different effects (Clemens et al., 2012; Bourguignon and Gunning, 2016). Furthermore, the cross-country literature is not able to account for heterogeneity within countries. While aid may have effects in targeted areas, these effects may not be sufficiently large to be measurable at the country level or they may be obscured by omitted variable bias (Dreher and Lohmann, 2015).

Indeed, many development projects are targeted at local development, meaning they should be judged against location-specific outcomes (Findley et al., 2011). Against this background, a finer lens is arguably needed when studying aid effectiveness.

At the other end of the spectrum, at the very micro level, the last decade have seen a great number of randomized controlled trials (RCTs) examining individuals' responses to specific interventions (see e.g. Kremer and Miguel, 2007; Björkman and Svensson, 2009). By randomly assigning exposure to an intervention, RCTs ensure that individuals' probability of being affected by the treatment is not correlated with its intended outcomes, thus making it possible to assess the causal impact of a project by comparing outcomes across treatment and control groups. As such, RCTs are a powerful tool to assess the causal impacts of development projects.

Increasingly, however, critics have raised concerns concerning the cost of RCTs and the extent to which one can generalize from their findings (e.g. Rodrik, 2009; Deaton, 2010). Randomized interventions require considerable planning ahead of time – since randomization is used to determine the selection of beneficiaries in a project, it must be baked into the project design from the very outset – and often involve extensive primary data collection over the treatment and control samples (BenYishay et al., 2017). This makes RCTs costly and time-consuming. Furthermore, if the investigated intervention is successful, policy makers may wish to scale it up to cover a larger share of the population. However, large scale policy is likely to be different from a small-scale experiment, e.g. because outcomes may be different when everyone is covered by the treatment rather than just a selected group of experimental subjects who are not representative of the population to be covered by the policy (Deaton, 2010). Also, while the estimated return to some projects may be substantial, that of others may be virtually zero. Low returns to individual projects can occur for a variety of reasons, including poor planning and implementation. The RCT literature provides little information on the average return to a certain type of project (Rajlakshmi and Becker, 2015). Doing so requires assessing the average treatment effects of a large number of similar interventions, and considering how costly RCTs are it is debatable whether it is viable to conduct them on the scale necessary to be able to evaluate broader donor performance in this way. Seeing that donors need to be able to learn from experience – what works in one context should help us make informed decisions of what will work in another – this is of course problematic.

GIE provides an intermediate perspective which rather than attempting to estimate country-wide impacts of total aid, or being restricted to analyzing the impact of single projects, utilizes subnational geocoded aid and outcome data to evaluate the sub-national distribution and local effects of development projects systematically and, possibly, on a wide scale. That is, project location-specific impacts can be estimated, but not necessarily for a single project in a single country, but potentially for a multitude of projects, across many countries. This is not to say that geospatial analysis cannot be used to analyse the impacts of individual projects. Indeed, this may be very fruitful, not the least if considering a large scale project spread across several project sites. The point is that the approach is flexible in terms of the unit and scope of analysis; the analyst can choose to focus on the local effects of a single project, or, for that matter, the local effects of all development projects in a specific sector, or of all projects from a specific donor, in one or several recipient countries.

Recent years have seen a rapid expansion in the availability of subnationally geocoded data, on development projects as well as on relevant outcomes and covariates, thus significantly improving the opportunities for GIE (AidData, 2017a). Combining geocoded development project data with geocoded individual/household level survey data or granular data based on remote sensing technologies (such as satellite data), makes it possible to address impacts across a wide variety of outcomes in different areas, including health, environment, conflict, governance, and economic development.

Geospatial impact evaluation has several attractive features (see e.g. BenYishay et al., 2017). One relates to concerns regarding the ability of non-randomized studies of aid effectiveness to infer causality. Just as any quantitative study seeking to investigate the impact of aid, studies applying geospatial analysis face empirical challenges in terms of distinguishing causal patterns from correlations. In particular, just as the distribution of aid across countries, the distribution of aid within countries is not random, implying that some individuals and sub-national areas, with certain characteristics, will be more likely to be targeted by aid than others. However, using granular geocoded data, GIEs have the advantage that they can control for potential confounding and omitted variables at fine geographic levels and that they are well-suited for quasi-experimental methods for causal identification. The key feature of these methods is, simply stated, to find 'control' cases that are sufficiently similar to the 'treated' cases to constitute a viable comparison group. This is easier when making comparisons within the local area, where individuals affected and not affected by a development project face similar conditions on other accounts (in terms of institutional arrangements, culture, climate, etc.), than when comparing groups further apart. Access to geographically precise data on aid-funded projects makes it possible to identify comparison individuals who are geographically close to the program participants, but who are unlikely to be affected by the program's presence themselves. Moreover, the possibilities for causal identification are improved by the fact that several of the outcome/covariate data sources go far back in time, making it possible to capture pre-treatment outcome measures in both treated and untreated areas.

Compared to RCTs, geospatial impact evaluation can be described as a powerful approach to assess causal impacts in a non-random setting. In particular, GIE makes it possible to rigorously and cost-effectively evaluate aid impact in cases when it is not feasible or ethical to determine which individuals participate in a program through random assignment. As noted, randomized interventions require considerable planning ahead of time and for the great majority of existing (ongoing and completed) development projects for which there are data records available, exposure has indeed not been determined through random assignment. GIE enables researchers and evaluators to utilize this valuable source of existing information and can thus be significantly cheaper and faster than RCTs. The methods can be applied retrospectively – for completed projects – or prospectively, for ongoing or future projects. And using comprehensive existing data materials, they can be implemented remotely, say from a policy evaluator's desk. The latter is not only beneficial from a cost perspective, it also means that GIE can be particularly useful to evaluators working in fragile state settings (as noted in Chapter 4 there is, for instance, geocoded aid data in countries such as Somalia, DRC, and Afghanistan).

Furthermore, whereas RCTs are often criticized in terms of their generalisability, GIEs are relatively strong in terms of the same – in both a spatial and temporal sense. First, and as noted, rather than assessing the effect of an individual project, GIE makes it possible to estimate the impact of a multitude of development projects, potentially across several countries. In addition, GIEs are often based on data covering a longer time period than RCTs (consider e.g. satellite data or data from large scale survey programmes) meaning that they can be used to measure long-run impacts. While RCTs often involve the collection of baseline data at the outset of a program, midline data during the implementation of the program, and endline data at program closure, due to the high cost of data collection they rarely evaluate impacts say five or ten years after the implementation of a project has been completed. Since GIEs often draw on outcome data that covers long time periods, they are better suited to track outcomes over the longer term, beyond project closure.

Geospatial impact evaluation is thus a valuable approach to evaluate the sustainability and broader long-term impacts of development projects. As such, the approach contrast with aid evaluation inspired by results based management (RBM) approaches, which arguably encourages practitioners to focus on short-term outputs (see e.g. Shutt, 2016). Geospatial aid data enables researchers and evaluators to study broader development outcomes, such as changes in local attitudes and behavior, over longer time horizons.

Taken together, GIE can help bridge the micro-macro divide in the aid effectiveness literature, alleviating the tradeoff between the ability to infer causality and of being able to generalize from ones findings, and make it possible to evaluate broad development impacts over the longer-term in a disaggregated and cost effective manner.

Importantly though, to incorporate geospatial analysis in the donor/partner country toolbox to achieve more evidence based development policy, it is important to be aware of its limitations as well as its opportunities. To begin with, the approach cannot be applied to all types of development projects; there are data constraints both in terms of outcome data and in terms of spatial differentiation of aid interventions. Whereas development projects implemented in specific localities, say local interventions in terms of health, education or local governance, can be evaluated using geospatial methods, projects with a national focus or without a physical project site, say general budget support or technical assistance to the central government, cannot. Strengths and limitations of the data will be discussed further in Chapter 4. Second, and as discussed above, just as any study seeking to investigate the impact of aid, studies applying geospatial analysis face empirical challenges in terms of establishing the causal impact of aid. Chapter 5 will discuss these challenges and examples of how to address them within a geospatial framework.

3. Academic applications of geospatial aid analysis

This section provides a brief overview of the emerging academic literature using geocoded aid data to systematically evaluate sub-national aid allocation patterns and impacts.

3.1 Studies of sub-national allocation patterns

As discussed in Section 2.1, while the aid allocation literature has traditionally focused on country level variation in aid flows, the geocoded aid data now available on a wider scale enables scholars to evaluate the factors driving sub-national aid allocation patterns.

Due to the good availability of geocoded aid data for World Bank and African Development Bank projects, a number of studies focus on the allocation patterns of these multilateral donors. Öhler and Nunnenkamp (2014) analyze the regional allocation of development projects financed by the WB and the AfDB in a sample of 27 recipient countries. While they find no evidence of needs based sub-national aid allocations, their findings suggest that favoritism plays a role for project location choices. Briggs (2017), similarly, use household surveys to measure the subnational distribution of a country's population by wealth quintiles and match this information to data on the location of development projects from the WB and the AfDB. In contrast with the stated preferences of the multilateral donors under study, his findings indicate that aid disproportionately flows to regions with more of the richest people and thus is not allocated effectively to alleviate extreme poverty. Nunnenkamp et al. (2012) investigate the district level allocation of World Bank aid in India. Their results suggest neither political capture nor needs based allocation of aid, but rather indicate that World Bank development projects tend to cluster in areas with favourable institutional conditions.

Powell and Findley (2012) also study the sub-national aid allocation of the World Bank and the African Development Bank, but from a donor coordination perspective. Their results point to coordination problems and inaccurate targeting of needs; often donors cluster in areas of low rather than high need. Similarly, Nunnenkamp et al. (2015) use geocoded aid data from Malawi to assess whether the country's bilateral and multilateral donors have lived up to promises to specialize and better coordinate their aid activities following the Paris Declaration of 2005. If anything, their results seem to suggest the opposite.

Other findings are slightly more positive. The results of Francken et al. (2012), who study relief aid allocation in Madagascar after the country was hit by a cyclone in 2004, suggest that compared to relief provided by the local government, assistance provided by international aid agencies was less affected by political factors, and more likely to go to poorer, but also more easily accessible communes. The results of Dionne et al. (2013) who study the sub-national allocation of education and health aid in Malawi, add some further nuance. While they find no evidence that allocation decisions are affected by need as measured by general poverty rates, their findings do indicate that health and education specific needs, measured in terms of infant survival and primary school completion rates, are strong predictors of sector-specific aid allocation.

Briggs (2014) and Jablonski (2014) both study political capture of aid in Kenya, their findings indicating that project aid is disproportionately directed to the president's political base. Dreher et al. (2016) also study political capture of aid, but focus on newly available data on Chinese development finance projects to Africa. Geocoding this data and combining it with information on African leaders' birthplaces and ethnic groups, they find that Chinese aid is disproportionately allocated to the birth regions of African leaders and (somewhat less clearly so) to areas populated by individuals who share their ethnicity. Replicating their analysis for World Bank aid, they find no evidence of corresponding regional or ethnic favoritism, suggesting that Chinese aid may be particularly easy to exploit for politicians who are engaged in patronage politics.

In sum, the reviewed studies on sub-national aid allocation patterns to some extent echo the results from the country-level aid allocation literature, pointing to insufficient targeting of poor areas and inadequate donor coordination within countries. Furthermore, several studies point to the role of convenience, or put differently – transaction costs – in the sub-national aid allocation decision, and shed light on the risk for political capture of aid. It is often suggested that since project aid tends to be relatively closely monitored, it involves comparatively little embezzlement (see e.g. Collier, 2006). As pointed out by Briggs (2014), however, these findings highlight that while the use of project aid might reduce the likelihood of outright embezzlement of funds, there is still a need to control elite capture of aid. Finally, the emerging literature on sub-national aid allocation points to donor as well as recipient country variation in allocation patterns, and highlights that sector specific needs and allocation patterns are not necessarily in line with needs measured more broadly.

3.2 Studies of aid impact

Turning to studies seeking to investigate the effects of aid, rather than what factors influence the allocation pattern of aid, here too the focus has traditionally been on aggregate country level aid flows. As discussed above, this is problematic for several reasons, one being the difficulty to isolate the causal impact of aid. The alternative has been evaluations of individual projects that instead come with challenges in terms of generalizability. Until recently, one could indeed speak of a 'missing middle' in the literature evaluating aid effectiveness. However, the geocoded aid data recently available on a wider scale now makes it possible to evaluate the local effects of development projects systematically and on a wide scale across multiple locations and recipient countries. Consequently, the last couple of years have seen an emerging literature utilizing this data to analyze aid effectiveness in a more disaggregated manner.

Some recent studies use geocoded aid data to investigate the effect of aid on local conflict. Findley et al. (2011), which was one of the first studies to introduce data on the geographic location of development projects to Africa, find that in countries engaged in armed conflicts, conflict is drawn to locations where fungible aid has been granted. Correspondingly, Wood and Sullivan (2015) examine whether humanitarian aid has negative externalities in terms of incentivizing violence and looting in a sample of African countries. Their results suggest that humanitarian aid is associated with increased rebel violence but provide less support for a relationship between aid and state violence. Strandow et al. (2016), similarly, investigate the impact of foreign aid on the intensity of violence during ongoing armed conflict, finding that funding concentration is associated with increased military fatalities, but not with civilian fatalities.

Dreher and Lohmann (2015) investigate the effect of World Bank aid on economic growth measured in terms of night time light at the sub-national regional level in 130 countries. Their findings suggest a positive correlation between aid and growth at the regional level, but no causal effects when using their preferred identification strategy (which exploits variation arising from interacting a variable that indicates whether or not a country has passed the threshold for receiving IDA's concessional aid with a recipient region's probability to receive aid).

A couple of studies examine sector specific aid impacts with encouraging results. Rajlakshmi and Becker (2015) use geocoded aid data from Malawi in combination with living standards data to assess the impact of health aid, water aid, and education aid. Their findings suggest that health aid reduces disease severity, that water aid reduces diarrhea incidence, and that education aid increases school exposure. Similarly, Odokonyero et al. (2015) use geocoded aid data to study the effects of health aid in Uganda. Their results indicate that even though health aid was not targeted to localities with the worst health conditions, it had a significant impact in reducing both disease severity and burden. However, aid effectiveness was further strengthened if resources were channeled to locations closer to communities in need. These findings point to the importance of focusing on the living standards variables actually targeted by development projects when assessing their impacts.

Some recent studies utilize the newly available data on Chinese development finance projects to Africa. Isaksson and Kotsadam (2016, 2017) match geocoded Chinese aid data with survey data for nearly 100,000 African survey respondents. Their findings indicate that Chinese development projects, unlike the aid of other examined donors, fuel local corruption and (in a smaller sample) discourage local union involvement. Similarly, Brazys et al. (2017) find that Chinese projects, unlike World Bank funded projects, are associated with higher perceptions of corruption in Tanzania.

Moreover, a couple of recent studies combine geocoded aid data with high-resolution satellite data to assess environmental impacts of development projects. Like the above mentioned studies, BenYishay et al. (2016) focus on Chinese aid, studying the effects of Chinese development projects on deforestation in the Tropical Andes, the Great Lakes region of Africa and the Mekong Delta. The authors conclude that China's development activities need not lead to widespread environmental damage when nearby ecosystems are appropriately protected, and that domestic environmental governance plays a crucial role in shaping these outcomes. Similarly, Buchanan et al. (2016) investigate the impact of World Bank development projects on sites of biodiversity significance, focusing on a range of outcomes including deforestation. They find no evidence that World Bank development projects have a negative impact on biodiversity.

Berlin et al. (2017), finally, study the impact of foreign aid, both general and gender-specific, on gender-related outcomes in Malawi and Uganda. The results from empirical estimations based on geocoded aid data matched with survey data are relatively mixed, suggesting some positive effects (e.g. in terms of reduced domestic violence), but also some negative effects (e.g. in terms of control over expenditures) of donor presence, probably reflecting the multidimensionality of the issue under study.

While estimating causal effects comes with empirical challenges (that will be discussed in Chapter 5), these studies highlight the great potential for geospatial impact evaluation provided by the newly available geocoded aid data. The substantial donor, sector and geographic variation in results point to the merits of a disaggregated approach to the analysis of aid effectiveness.

4. Data

Combining geocoded aid data – i.e. GPS coordinates for specific development projects – with geocoded information from other data sources, such as individual/household level survey data or data based on remote sensing technologies, makes it possible to uncover systematic patterns in characteristics of localities where development projects are implemented and to address aid impacts across a wide variety of outcomes. This chapter focuses on data considerations, including the availability and limitations of geocoded aid and outcome/covariate data, and how to combine these data sources.

4.1 Availability and precision of geocoded aid data

The last few years have seen a sharp increase in the availability of sub-nationally georeferenced data on aid interventions. The World Bank now regularly publishes the latitude and longitude coordinates of all of its projects, and a growing number of bilateral and multilateral aid agencies have since followed suit (BenYishay et al., 2017). Furthermore, partner country ministries responsible for managing incoming aid flows increasingly publish sub-nationally geocoded development project data. While Malawi's Ministry of Finance was the first to do so (World Bank, 2011), there are now a relatively great number of developing country governments providing sub-nationally geocoded aid data (AidData 2017c). The geocoded aid data is publicly available from AidData (2017c), who in collaboration with the above mentioned institutions have made significant efforts to map and synthesize sub-national data on development projects. Since 2010, when they released their data portal, the availability of geocoded data has expanded steadily.

In particular, for the World Bank, the African Development Bank and China the AidData portal includes geocoded aid data for a multitude of partner countries. The geocoded World Bank data covers all World Bank projects in the IBRD and IDA lending lines approved from 1995-2014 (5,684 projects spread across 61,243 locations, see AidData, 2017d). The geocoded African Development Bank data covers all AfDB activities approved in 2009-2010 (183 projects spread over nearly 2,000 locations, AidData, 2017c).The geocoded data of Chinese aid covers Chinese development finance activities in 50 African countries over the 2000-2012 period (1,952 projects across 3,545 locations, AidData, 2017c).

Furthermore, partner country ministries responsible for managing incoming aid flows publish sub-nationally geocoded development project data in Nigeria, Uganda, Senegal, Malawi, DRC, Somalia, Colombia, Honduras, Afghanistan, Bangladesh, Nepal, Iraq and Timor-Leste. The coverage of these datasets varies from country to country (see AidData, 2017c) and thus should not necessarily be seen as representative of a specific donor's involvement in that country. In Chapter 7 we briefly discuss the coverage of Swedish projects in Uganda and Afghanistan.

As noted, though, not all development projects can be geocoded effectively. Naturally, to be able to geocode a project in a meaningful way it has to have a physical project site. And whereas some projects are implemented in a limited geographical area, such as a village, others are realized at more aggregate levels, such as a district or greater administrative region. Projects that are more intangible, consider e.g. debt-relief agreements, budget or sector support (see the reasoning in Muchapondwa, et al. 2014) are not possible to geocode with any greater geographical precision.

However, instead of leaving out these projects, in order to get a full picture of aid flows, AidData follows Strandow et al. (2011) and geocodes all projects for which they have the necessary information, but classifies them into 8 categories depending on the level of precision of the geocode, ranging from 1 for most precise to 8 for least precise. Precision category 1 refers to 'an exact location, such as a populated place or a hill', category 2 to being 'near', in the 'area' of, or up to 25 km away from an exact location, category 3 refers to the second order administrative division (ADM2) in a country, such as a district, municipality, or commune, and category 4 to the first order administrative division (ADM1), such as a province, state, or governorate. Precision category 5 refers to the estimated coordinates of large features potentially spanning several ADM1, such as rivers and national parks, and 6-8 are used for projects lacking a clear sub-national destination, for instance aid flows going to the central government (for more detailed information on the precision categories and geocoding procedures, see Strandow et al., 2011, and Table A1 in the Appendix). The fact that AidData uses this approach rather than geocoding only those projects providing precise geographical information means that they are transparent in terms of data quality and leave it to the researcher or policy analyst to decide on the appropriate precision level to consider.

Since studies using geocoded aid data tend to focus on development projects with relatively precise geocodes (e.g. in precision categories 1-2, 1-3 or 1-4), and since the geographical coding precision is likely to reflect the sectoral composition of aid (Dreher and Lohmann, 2015), the implications in terms of project/sectoral coverage need to be kept in mind when generalizing from results.

As discussed above, coverage varies by donor and recipient country. Figure 1 gives the geocoding precision categories of the World Bank's projects (author's estimations based on AidData, 2017d), who geocode all their development projects. 47 percent of the project sites are attributed to a precise location (precision categories 1-2), 25 percent to the district/municipality/commune (ADM2) level (precision category 3) and 19 percent to the province/state (ADM1) level. Only for approximately 7 percent of project sites is the location unclear or only traceable to the administrative centre of the country (precision categories 6-8).

Exploring the sectoral composition of projects geocoded with different levels of precision, in the most precise categories (1 and 2) the most common sectors are, in order of magnitude, Transport (32%), Water and sanitation (17%) and Agriculture (11%). In the least precise categories (6-8), the most common sectors are instead Public Administration, Law, and Justice (33%), followed by Health and other social services (14%) and Education (9%).

Figure 1: Geocoding Precision categories of World Bank project sites

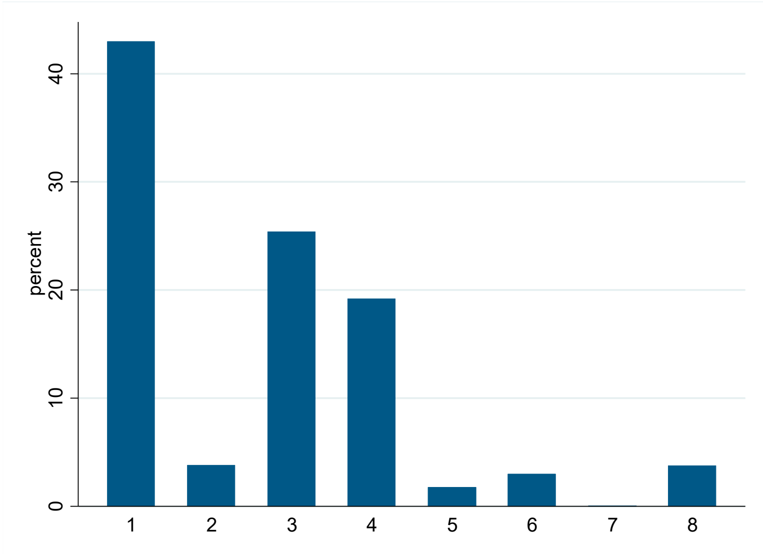

Notes: Based on AidData (2017d)

To get a sense of how comprehensive the data is, Figure 2 shows a map over all World Bank projects to Africa that are geocoded in the most precise category (precision category 1), specifically 660 projects spread across just over 4000 locations. These include ongoing as well as completed projects dating back to the mid-1990s. Hence, it is possible to assess the effects of projects that are currently in operation as well as to consider short- and long-term effects of completed projects (by utilizing outcome data that also reaches back in time).

Figure 2: World Bank projects to Africa between mid-1990's and 2014, geocoded in precision category 1

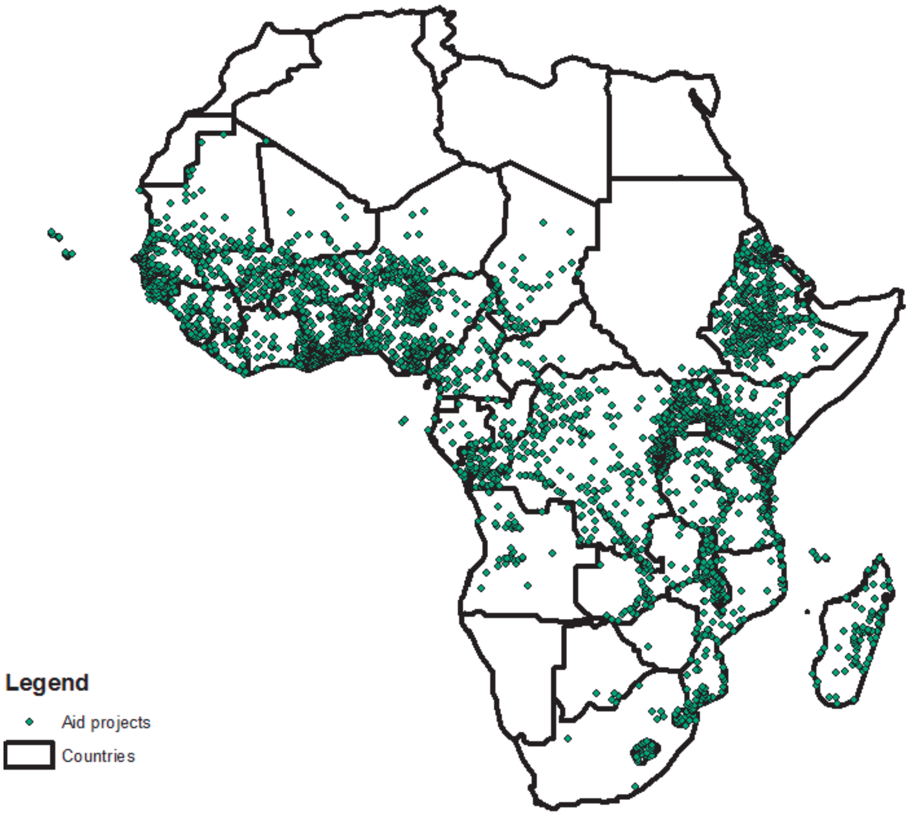

Notes: Based on AidData (2017d)

As a comparison, consider the data for Uganda, which covers all projects for the period 1978-2014 – specifically 1709 of them, from 56 different donors – in Uganda's Aid Management Platform maintained by Uganda's Ministry of Finance (AidData, 2016). Out of these, 33 percent, or 565 projects (spread over 2426 project locations) have been geocoded. Again, however, it is important to remember that these include project sites geocoded at all different precision levels, i.e. both those geocoded to exact locations and those attributed merely to, say, a central administrative unit. In the Ugandan case (Figure 3), 29 percent of the geocoded project sites are attributed to a precise location (precision categories 1-2) and 41 percent to the district/municipality/commune (ADM2) level (precision category 3). 28 percent of the project sites are geocoded with a low level of precision (precision categories 6-8). Exploring the sectoral composition of projects geocoded with different levels of precision, in the most precise categories (1 and 2) the most common sectors are, in order of magnitude, Health (31%), Education (18%), Government and civil society (17%) and Agriculture (12%). In the least precise categories (6-8), the most common sector is instead General budget support (24%), followed by Government and civil society (20%), health (12%) and agriculture (12%).[1]

Figure 3: Geocoding Precision categories of Ugandan project sites

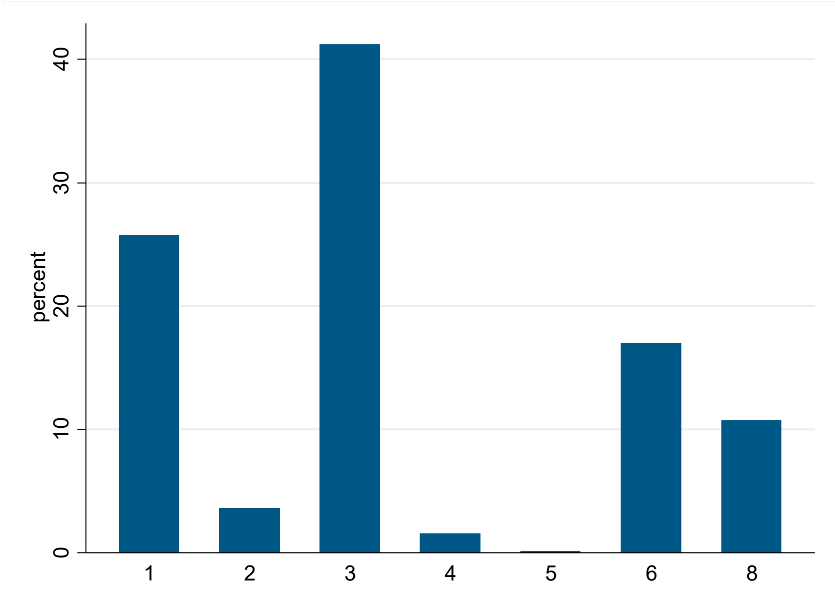

[1] Roughly the same sectoral shares hold for the projects that have not been geocoded. Here too, General budget support (20%) is the most common sector, followed by Government and civil society (20%), Agriculture (13%) and Health (12%).

Figure 4 shows a map over all Ugandan development projects geocoded in the most precise category, specifically 121 projects spread across 559 locations (some dating back to the 1990s but most with start dates during the last decade). We can note that while there is a relatively good spread of projects, there is also some clustering e.g. in the south-west and around Kampala (the blank spot in the south-east is covered by Lake Victoria).

Figure 4: 121 development projects during 1990s – 2014 in Uganda, geocoded in precision category 1

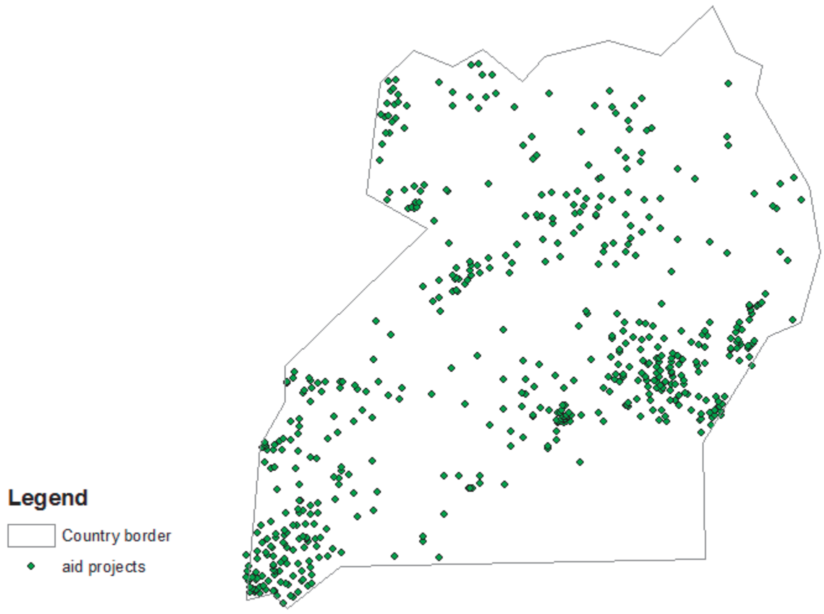

Notes: Based on AidData (2016).

The key message from the above examples is that when generalizing from findings based on geocoded development project data, the characteristics of the covered projects and how these may differ from the characteristics of overall aid need to be kept in mind. While the sectoral focus of development projects geocoded with different levels of precision overlap to some extent, they are clearly not identical.

Furthermore, gaps in the geocoded aid data should be considered when analyzing the allocation and impacts of geocoded projects. Many donors and implementing partners do not routinely map their intervention sites, and even when they do, important information may be missing, for instance with respect to the timing of implementation (BenYishay et al., 2017). Data records for analysed projects as well as other projects located in treatment and control areas is thus likely to be incomplete. The former restricts the sample of projects that can be analysed effectively. The latter is problematic considering that if seeking to assess the impact of a particular project it is useful to know (and thus to be able to control for) what other projects are located in the area. Lacking this information, the estimated effect of some particular projects may be biased by the effects of other, not geogaphically coded, projects.

Moreover, with multiple actors involved in the geocoding of aid, on the donor as well as the partner country side, double counting of projects is a potential concern. It is thus important to make sure that seemingly distinct projects identified in the same locations do, in fact, not refer to records of different donors involved in the same development project.

4.2 Availability of geocoded data on relevant outcomes and covariates

As with geocoded aid data, recent years have seen a rapid expansion in the availability of geocoded outcome and covariate data. Household survey data is increasingly geocoded to the enumeration area level (the smallest geographical sampling unit) rather than, as before, only providing geographical identifiers to the region or district level, and there are now numerous sources providing high resolution spatial data on a wide range of indicators. This section provides a brief overview of data sources (not to be interpreted as an exhaustive list) that can be used for outcome variables and covariates when examining the impact and distribution of geocoded development projects, and a discussion of the variety of questions that can be addressed by combining geocoded aid data with these data materials.

To begin with, a number of survey data sources are particularly worth mentioning. The Afrobarometer, a comprehensive multi-country survey project collecting data on political and economic attitudes and behaviour of African citizens, now allows users to apply for access to low level geographical identifiers of survey respondents (Afrobarometer, 2017).[2] The survey, which has been conducted in six waves from 1999-2015 and covers 37 African countries to date, allows for analysis of local aid outcomes and covariates relating to a multitude of issues, including corruption (see e.g. Isaksson and Kotsadam, 2016), political participation, trust in local and international institutions, and a wide range of citizen attitudes (e.g. on democracy and gender equality).

Another key data source in this context is the Demographic and Health Survey (DHS, 2017), which consists of country-representative household surveys focusing (as the name implies) on population and health indicators. This extremely comprehensive data material, containing over 300 surveys from over 90 countries, each with a sample size of between 5,000 and 30,000 households, allows for analysis of a wide array of issues relevant for assessing both the allocation and impacts of aid, including – to mention just a few examples – child health, domestic violence, literacy, family planning, maternal health, fertility, HIV prevalence and knowledge of HIV prevention. In more recent years, the DHS routinely collects geographic information in all surveyed countries and there are now geographic data (GPS coordinates at the cluster level)[3] from over 100 DHS surveys in over 45 countries available for download upon request.

---

[2] Specifically, they provide the coordinates to the surveyed Afrobarometer 'clusters', consisting of one or several geographically close villages or a neighborhood in an urban area.

[3] In order to protect the confidentiality of sample household, they randomly displace the GPS latitude/longitude positions for all surveys. In particular, the

Similarly, the World Bank Living Standards Measurement Study (LSMS) program, which provides technical assistance to national statistical offices in the design and implementation of multi-topic household surveys, now often provides geographical information at the cluster level. Apart from a large number of nationally representative cross-sectional surveys, the LSMS program contains panel data for e.g. Ethiopia, Malawi, Mali, Niger, Nigeria, Tanzania and Uganda (World Bank, 2017).

In addition there are, as noted, numerous sources providing high resolution spatial data on a wide range of indicators. In terms of population/demographics data, the Gridded Population of the World (GPW) dataset provides measures of population density across the globe based on census data (NASA Socioeconomic Data and Applications Center, 2017). Other organizations providing high resolution spatial data in this area are WorldPop (2017) and IPUMS TERRA (2017). The former combines census, survey, satellite, social media, cellphone and other spatial datasets to produce gridded maps of population distributions (population numbers per 100x100m grid squares) as well as measures of populations dynamics (migration) and characteristics (e.g. age and sex structures, births and pregnancies) in low and middle income countries. IPUMS TERRA integrates population and environmental data to provide global-scale measures of demographics, land use, land cover, climate and other environmental features. The population data is primarily derived from census and survey data and the environmental indicators from remote sensing technologies.

Several corresponding data collection initiatives provide spatial data on geographical attributes and subnational economic activity. As noted above, IPUMS TERRA (2017) also collects data on environmental characteristics, including land use and land cover as well as temperature, precipitation, and other climate-related

---

displacement is carried out so that urban clusters contain a minimum of 0 and a maximum of 2 kilometers of error, and rural clusters – where communities are more dispersed and risk of disclosure may be higher – contain a minimum of 0 and a maximum of 5 kilometers of positional error.

variables. Similarly, the US Geological Survey (USGS, 2017) collects satellite data used to assess e.g. the impact of climate and land use change. Using their tool *EarthExplorer* users can furthermore search and order geospatial data from several sources. With respect to data on subnational economic activity, the Geographically Based Economic Data project (G-Econ, 2011) has developed a GDP equivalent for grid cells (of approximately 100 x 100 km) for the years 1990, 1995, 2000 and 2005.[4] The economic data, which has been merged with demographic and geophysical data on e.g. climate, physical attributes, location indicators, population, and luminosity, is publicly available. The data on luminosity is in this context interesting in itself, since satellite data on night time light has been shown to correlate with economic activity at both the country and sub-national level why it in recent years has often been used as a proxy for economic activity at the sub-national level (see e.g. Henderson et al., 2012). Night time light data at a high-level of spatial resolution is available from the early 1990s e.g. from the Earth Observation Group of the National Oceanic and Atmospheric Administration (NOAA, 2017). From an aid evaluation perspective, data on subnational economic activity can be used both to assess impacts of local aid interventions and to get a picture of to what extent aid is allocated to the sub-national regions most in need.

Worth mentioning is also the good availability of sub-nationally geocoded conflict data, which can be used to investigate to what extent donors reach fragile areas within partner countries, as well as whether development projects have unintended consequences in terms of conflict (e.g. Strandow et al., 2016). The Uppsala

---

[4] The methodology, which varies with data availability and thus by country, is described in Nordhaus et al. (2006). For many rich countries, they rely on estimates of regional gross value added within administrative regions and then use data on population density to convert regional data to the 100 by 100 km grid cells. For many of the lowest-income countries, however, there is no regional economic data, why they instead combine population census data to estimate rural and urban populations by county with national employment and output data to estimate output per capita in agriculture and non-agricultural industries in order to get an estimate of output per capita by region.

Conflict Data Program (UCDP), based at the Department of Peace and Conflict Research at Uppsala University, has a Georeferenced Events Dataset (UCDP Georeferenced Event Dataset, 2015) covering individual events of organized violence geocoded to the level of individual villages. The most recent version of this global dataset contains 128 264 such events over the period 1989-2015 (Croicu and Sundberg, 2016). The Peace and Research Institute in Oslo (PRIO) also provides detailed geospatial data on conflict. In their PRIO-GRID (2017) dataset each grid cell contains information on armed conflicts, as well as e.g. socio-economic conditions, ethnic groups, physical attributes and climatic conditions. Furthermore, the Armed Conflict Location and Event Data Project (ACLED, 2017), which focuses on political violence and protest data for developing countries, provides information on the specific dates and locations of political violence and protest, the types of event, the groups involved, the number of fatalities, and changes in territorial control.

The above listing, which is by no means exhaustive, provides some indication of the wealth of relevant geospatial outcome and covariate data that can be used for geospatial analysis of aid flows. While this data is indeed rich, it should nevertheless be noted that the questions one can address with geospatial data, without engaging in further data collection, is limited by the information available in existing survey materials and data bases. Researchers and policy evaluators doing RCTs or other studies involving designing the survey instrument yourself – which of course comes at a significantly higher cost in terms of data collection – have more control in terms of the questions they can ask.

If the available data appears inadequate to answer the research or policy question at hand, a fruitful approach can be to combine existing geospatial data with new data collection initiatives. Existing geocoded aid data, containing information on where to find beneficiaries and non-beneficiaries of certain aid interventions, could for instance be used to guide where to implement smaller scale surveys or field experiments designed specifically to evaluate responses to these interventions. This could make the researcher or evaluator better able to assess underlying causal mechanisms that may be difficult to get at by using information in existing data sources.

In this report, however, we restrict our attention to the wealth of readily available geospatial outcome and covariate data that can be used for analyzing the local impacts and sub-national distribution patterns of development projects. In the next section we turn to how to combine these data sources with geocoded aid data.

4.3 Matching treatment and outcome data based on geographical proximity

To be able to analyze geocoded aid data in a meaningful way it has to be combined with geocoded information from other data sources, such as the individual/household level survey data or data based on remote sensing technologies described above. By using the point coordinates in the aid data, development projects can be linked to e.g. local survey respondents in the area. Doing so one can identify survey respondents living near project sites, and evaluate how these respondents fare on relevant outcomes compared to other groups not living in immediate connection to the specified development project sites.

Commonly used programs to work with spatial data include ArcGIS and QGIS. However, since these programs take time to understand and to learn to use effectively, extracting information from, and combining, spatial datasets can be challenging for non-expert users. Indeed, in the past the computational costs associated with obtaining relevant spatial data (e.g. the intensity of nighttime lights) at meaningful units of analysis (e.g. village administrative boundaries) has prevented spatial analyses in a wide range of fields (Goodman et al., 2017). A number of recent initiatives, e.g. from PRIO grid and IPUMS TERRA, have significantly simplified access to spatial data for non-experts.

Specifically worth mentioning in this context is AidData's recently created tool 'geoquery' (AidData, 2017e) that allows users to extract and combine spatial data from a multitude of different sources within 'geo framework' – their open source, spatial data management platform – without having training in GIS. Using the online menus, users get to choose country and geographical unit of analysis (e.g. ADM1, ADM2, ADM3) and select indicators from 22 geocoded aid and outcome datasets. These include the geocoded datasets on development projects in the AidData portal, as well as geocoded outcome/covariate data on environment and population (e.g. satellite- and census-based data on population density, vegetation density, land cover, temperature, precipitation, slope and elevation), conflicts, economic development (e.g. subnational measures of GDP and nighttime lights) and access to infrastructure (e.g. markets, population centers, and roads). Data on the requested indicators are then merged together at the geographical unit of analysis that the user specified. Finally, the user is emailed a copy of the customized dataset in csv-format, i.e. in a format that is compatible with multiple other programs that the user might prefer to work in (Excel, STATA, R, etc.).[5]

To illustrate, let us consider an example. Let us say we are interested in the sub-national allocation of World Bank health aid in Kenya. Specifically, is World Bank health aid allocated to regions of the country in particular need of health interventions? Through the online menus at geoquery we first get to choose what country to focus on, here Kenya, and then to select the preferred unit of analysis. For Kenya you can choose to get data at the provincial (ADM1), district (ADM2) or sub-district (ADM3) level. To keep things simple, let us consider the provincial level. Next we get a list of datasets that are available for this unit of analysis. For Kenya, there is province level data from 20 different datasets. First of all, we select World Bank aid data. Via the menus we then add some filters to narrow down our search. Specifically, we choose to focus on total disbursements to the health sector since the turn of the millennium. Next, we are interested in an indicator capturing health needs at the province level. It turns out there is a relevant

[5] For more detailed information on AidData's geo framework an online tools, see Goodman et al. (2017).

indicator on child mortality that we could use. In order to capture health needs in the beginning of the studied period – i.e. pre-existing health needs that the World Bank may base their health aid allocation decisions on, rather than potential outcomes of World Bank health projects – we focus on child mortality in the year 2000. After submitting our request we get an email with the requested data in csv-format.

Figure 5: Province level child mortality in 2000 and World Bank health aid since 2000 in Kenya

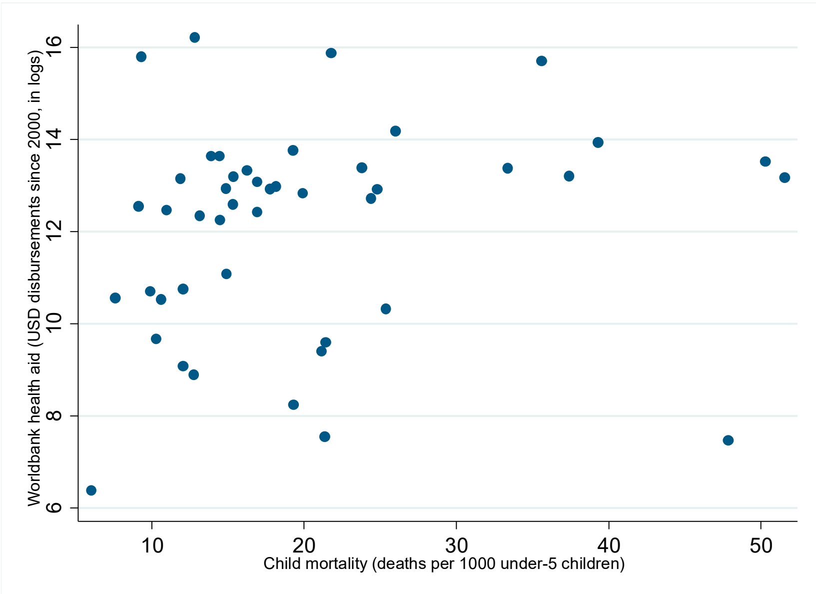

Importing the data to STATA (or for that matter, to Excel) we get a small dataset with one observation for each of the 47 Kenyan provinces (for illustration, a smaller version of the dataset, focusing on the key indicators only, is presented in Table A2). Based on this we can examine the relationship between province level child mortality and subsequent World Bank health aid. Judging from the scatter plot in Figure 5, this relationship is, while seemingly positive, by no means clear-cut. Confirming this, a simple regression between province level child mortality in 2000 and subsequent World Bank health aid (in logs) reveals no statistically significant relationship. The results of this little exercise should obviously not be taken to indicate that the sub-national allocation of World Bank health aid in Kenya is not guided by sub-national variation in health needs. Such conclusions would of course require a much more thorough investigation. However, the example serves to illustrate that extracting and combining spatial aid and outcome/covariate data is simple when using online tools like geoquery.

For somewhat more experienced users, programs such as STATA offer commands that make it possible to match datasets based on geographical proximity (latitude and longitude coordinates) without using GIS. One such command is 'nearstat' in STATA, which calculates distance based variables based on coordinates in the data. Combining geocoded survey data and geocoded aid data it could for instance be used to identify, and calculate the distance to, the development project located nearest a survey cluster, or to identify (and calculate distances to) all development projects located within a specified distance of a survey cluster.

5. Identifying causal impacts: empirical challenges and strategies

As noted, geospatial analysis of geocoded aid and outcome/ covariate is useful both to uncover systematic patterns in characteristics of localities where development projects are implemented and to explore aid impacts. The former is relatively straight forward. With access to covariate data from before implementation of the concerned development projects started one can assess systematic variation in pre-existing characteristics of localities in question. Establishing a causal impact of development projects, as opposed to systematic differences in the characteristics of localities where development projects are implemented, is not equally straight forward. This chapter will discuss the empirical challenges involved, and strategies to deal with these within a GIE framework.

5.1 Empirical challenges

The challenges originate in the fact that the sub-national allocation of aid, just as the distribution of aid across countries, is non-random. This implies that some individuals and sub-national areas, with certain characteristics, will be more likely to be targeted by aid than others. Indeed, as described above, these systematic patterns are, in themselves, interesting to investigate using geospatial aid and covariate data. Nonetheless, they are problematic when seeking to estimate the causal impact of aid interventions. For instance, just as finding a relationship between high poverty rates and high aid inflows at the country level should not be interpreted as aid making countries poor, finding that people living close to health projects tend to be less healthy should not be interpreted as health aid making people sick.

In effect, the ultimate purpose of aid could be said to create a fundamental identification problem; if aid is allocated so as to help people in particular need, it will be difficult to separate the impact of aid from the effects of all the problems that are common in aid receiving localities. The systematic, i.e. non-random, distribution of aid creates important selection problems implying that areas receiving aid – and the citizens residing in these localities – are likely to differ systematically from other areas and the people living there. Furthermore, given that donors and partner country governments do not allocate donor funded projects at random, the same factors which make one donor-funded project be located in a particular area may result in other projects being located there as well. With clustering of aid in certain sub-national regions it may be difficult to separate the impact of one donor-funded project from that of another.

To the extent that the systematic variation in the characteristics of localities where development projects are implemented is known, and there is data to capture it, the analyst will be able to control for it in a regression framework in order to avoid biased estimates of the causal effect of development projects. However, the systematic variation is likely to go beyond the factors the analyst is aware of and able to control for. To take a concrete example, since not all aid is geocoded, the analyst is likely to be able to control for the co-location of some projects but not all. If unaccounted for, unobserved systematic variation will bias estimates intended to capture the causal effect of development projects.

As an example, consider the study of Isaksson and Kotsadam (2016), which investigates the impact of Chinese development projects on local corruption in Africa. While the authors observe that people living close to active Chinese project sites are more likely to report experiences with corruption, they emphasize that this – on its own – should not be interpreted as indicating a causal effect of Chinese aid on local corruption. The reason is that local corruption levels, and other factors correlated with corruption (such as population density, economic activity and infrastructure access) are likely to influence Chinese project location decisions. For instance, a donor not prepared to use bribes may be less inclined to implement a project in a highly corrupt area. An alternative, more cynical, position is that a donor prepared to engage in corrupt activity themselves is more likely to set up development projects (and possible connected business ties) in particularly corrupt locations. While the exact nature of the relationship is difficult to ascertain, it does not seem plausible to assume that there is *no* relationship between Chinese project localization and the pre-existing institutional characteristics of project sites. We return to how the authors deal with this issue below.

Summing up, estimating the causal impact of aid comes with empirical challenges originating in the fact that, for the great majority of cases, the allocation of aid is non-random. Randomly assigning exposure to aid interventions, as is done in RCTs, is a neat way to get around the described selection problems. However, for the great majority of existing (ongoing and completed) development projects for which there are valuable data records available, exposure has indeed not been determined through random assignment. It is thus important to find ways to assess causal impacts of development projects for the abundance of cases where randomized data on interventions is not feasible or available.

5.2 Empirical strategies to address these challenges

With access to detailed geocoded data stretching over an extended time period, the analyst can utilize both spatial and temporal variation in aid and outcome data to assess the impact of aid interventions. This entails comparing localities affected and not affected by development projects over time – before and after development project implementation – and controlling for potential confounding and omitted variables at fine geographic levels. As such, GIEs are well-suited for quasi-experimental methods for causal identification, i.e. finding 'control' cases that are sufficiently similar to the 'treated' cases to constitute a viable comparison group. What strategies can be used depend on the specific question asked and the availability of aid and outcome data with spatial and temporal variation.

5.2.1 Difference-in-differences: pre- and post-intervention changes in the outcomes

With access to outcome data containing temporal variation, it is possible to compare the pre- and post-intervention changes in the outcome of interest for a treatment group relative to a control group in a difference-in-differences (DiD) setup (see e.g. Angrist and Pischke, 2009). A relevant 'treatment' in our case could be the implementation of a development project, reaching some areas but not others. As the name implies, the estimation strategy entails comparing the average change in the outcome variable over time (before and after project implementation) for a treatment group exposed to the development project and a control group not exposed to the development project.[6] This makes it possible to capture the effect for the observations in the treatment group in the period after treatment, controlling for common time trends, i.e. factors that would cause changes in the outcome even in the absence of a treatment, as well as possible pre-existing differences between treatment and control group.

A precondition for the approach to hold, however, is the 'parallel trends assumption' saying that the treatment and control

---

[6] In a regression framework, the DiD setup for outcome *Y* of individual (or the relevant unit of analysis) *i* in time period *t* could be illustrated as follows: $Y_{it} = \beta_0 + \beta_1(after)_{it} + \beta_2(treated)_i + \beta_3(after \times treated)_{it} + \alpha\mathbf{X}_{it} + e_{it}$. Here, a time period dummy indicating that the observation took place after treatment (let us call it *after*) captures potential common time trends, and a dummy variable for belonging to the treatment group (let us call it *treated*) captures possible differences in *Y* between the treatment and control groups that existed prior to the treatment. $\mathbf{X}$ is a vector of possible control variables and *e* is the regression error term. The coefficient of interest, β_3, refers to an interaction term between *after* and *treated*, and captures the effect for the observations in the treatment group in the period after treatment, controlling for common time trends and pre-existing differences between the groups. Specifically, β_3 – or the difference-in-differences estimate – captures the difference in the outcome change between the two groups that arises as a result of the treatment: $\beta_3 = (Y_{after} - Y_{before}\ for\ treatment\ group) - (Y_{after} - Y_{before}\ for\ control\ group)$

group would exhibit the same trend in terms of the concerned outcome in the absence of the treatment. If, for instance, the implementation of the development project comes as part of a wider reform package in the exposed areas, this could mean that the trends between treatment and control group would have diverged even in the absence of the development project, thus making it problematic to attribute the change solely to the treatment. Hence, while controlling for common time trends and pre-existing differences between groups is an important step ahead, the comparability of treatment and control groups in terms of potential group-specific time-trends still needs careful consideration.

With access to panel data, that is repeated observations for the same individuals (or household, or whatever the unit of analysis might be), the difference-in-difference approach becomes even stronger, allowing the analyst to control for unobserved individual variation (i.e. variation not captured by other control variables in the regression framework) and thus to minimize estimation bias that may arise from endogeneity in the treatment variable.

Consider the study by Odokonyero et al. (2015) on the impacts of health aid in Uganda. As described by the authors, the empirical evidence on the effectiveness of health aid is inconclusive. In line with the micro-macro paradox, a number of cross-country studies fail to find an association between aid and improvements in various health indicators (e.g. Williamson, 2008; Wilson, 2011), whereas some influential micro-level studies show that certain health interventions effectively reduce a number of diseases common in low-income countries (e.g. Kremer and Miguel, 2004, on deworming medication). Taking the intermediate perspective focused on in this report, i.e. utilizing subnational geocoded aid and outcome data to evaluate the local effects of health projects, the findings of Odokonyero et al. (2015) are optimistic. Using a DiD-approach combining household panel data with geographically-referenced health aid data allows them to control for common time trends and pre-existing differences between groups as well as unobserved household variation that may bias the estimate if there is endogeneity in the treatment variable. Such endogeneity could result e.g. from certain households being more likely to be targeted both by the health projects and other interventions taking place in parallel. Their findings indicate that health aid reduced both disease severity and burden, especially when aid is channeled to locations in particular need.

Importantly, however, even in cases where the ideal data is not available, e.g. individual level panel data or local measures of the outcome of interest from before and after an aid intervention, rich geospatial data enables flexibility in terms of estimation strategies. Returning to the paper on Chinese aid and local corruption (Isaksson and Kotsadam, 2016), discussed above, using four Afrobarometer survey waves the authors have access to temporal variation in the corruption experiences of African citizens. However, since the survey does not have a panel structure, they are not able to follow the same localities over time and evaluate variation in corruption occurring around a project site before and after project implementation. Simply comparing the corruption experiences in areas that are close to and far away from project sites is likely to involve selection bias since Chinese project location decisions probably are affected by pre-existing characteristics of project sites. Hence, the authors instead compare sites where a Chinese development project was actually under implementation at the time of the survey and sites where a project will be opened but where implementation had not yet been initiated at the time the Afrobarometer covered that particular area. This gives a difference-in-differences type of estimate [7] intended to minimize the possible bias from unobserved characteristics that may influence selection into being a Chinese project site.

[7] Specifically, the authors compare the difference between post-treatment individuals (with an ongoing Chinese project within 50 km) and control individuals (with no Chinese project – ongoing or future – within 50 km) with the difference between pre-treatment individuals (with a future Chinese project within 50 km) and control individuals within the same country/region and year (due to country/region and year fixed effects).

5.2.2 Matching strategies

Another popular approach is to use matching strategies, i.e. strategic subsampling of observational data to identify treatment and control cases that are very similar except for the presence or absence of the intervention (see e.g. Angrist and Pischke, 2009). First of all, we should note that matching, just as standard regression frameworks, relies on controlling for observable factors, and hence that it is no magic bullet in terms of causal identification. However, it shifts focus from controlling for factors that may influence the outcome variable to factors mattering for treatment assignment. This could be compelling seeing that we may have more information about the factors determining treatment, since treatment assignment tends be the product of human institutions and government regulations, than about what factors affect the outcome variable of interest. That is, we may be better informed about which factors determine if a development project is implemented in a particular area than about what factors impact, say, different living standards outcomes in the areas under study, and thus be better able to control for factors mattering for the former than for the latter.

A commonly used method in this context is propensity score matching. The propensity score, as defined by Rosenbaum and Rubin (1983), is the conditional probability of an observation being assigned to a particular treatment given a set of observable characteristics. It is often referred to as a balancing score – conditional on the propensity score, the distribution of measured baseline covariates should be similar between treated and untreated subjects. In an RCT the true propensity score, i.e. the probability of being in the treatment group, is known. In fact, it is defined by the analyst in the study design. In a non-randomized intervention, on the other hand, this probability is not known, but can be estimated in a regression of treatment status on observed baseline characteristics. The estimated propensity score is the predicted probability of treatment derived from the regression model. Estimation based on propensity score matching thus takes place in two steps: first the probability of treatment is estimated using e.g. logit or probit, then estimates of the effect of treatment are computed by matching on the estimated score from the first step. There are different propensity score methods for removing the effects of confounding factors when estimating the effect of treatment in the second stage, one being one-to-one matching in which pairs of treated and untreated subjects with similar values on the propensity score are formed.[8] Once a matched sample has been constructed, the treatment effect can be estimated by directly comparing outcomes between treated and untreated subjects in the matched sample.

An advantage of propensity score matching is that it makes the (potential lack of) common support, or overlap in terms of covariates, between treatment and control group explicit (Austin, 2011). In fact, common support is a key assumption underlying the model, specifying that for each value of the covariates used in the matching, there must be a positive probability of being both treated and untreated. This assumption ensures that there is sufficient overlap in the characteristics of treated and untreated subjects to find adequate matches. For instance, if the treated subjects include only men, the match would also have to be a man, meaning that women would be excluded from the analysis and the results would not pertain to them. Hence, when using propensity score matching, one is explicitly comparing outcomes between treated and untreated subjects who have a similar distribution of observed baseline covariates. If there are substantial differences in baseline covariates between treated and untreated subjects, this will be clear from the small number of matched subjects and the analyst is faced with a choice between restricting the analysis to the minority of treated and untreated subjects who have similar covariate patterns, or to stop the investigation with reference to that treated and untreated subjects are so different that a meaningful comparison of outcomes between the two groups is not plausible. In a standard regression framework, the degree of overlap between the covariates of the two groups is not as explicit,

[8] For more information about this approach and other propensity score methods (stratification on the propensity score, inverse probability of treatment weighting (IPTW) using the propensity score, and covariate adjustment using the propensity score), see e.g. Austin (2011).

and the not-so-careful analyst may proceed with the analysis without being aware that the regression model is extrapolating from one relatively distinct population to another.

Geospatial aid studies using propensity score matching include Rajlakshmi and Becker's (2015) study of health aid in Malawi, Wood and Sullivan's (2015) study of whether humanitarian aid incentivizes violence, the paper by Strandow et al. (2016) on the impact of foreign aid on the intensity of violence during armed conflict, and the paper by Berlin et al. (2017) on the impact of foreign aid on gender-related outcomes (see Chapter 3).

6. Empirical illustration

This chapter will provide some examples of how one can break geocoded aid data down by donor, locality and sectoral focus and analyze the impact of development projects' on different outcomes, running some illustrative regressions in the process. Specifically, we will focus on the impact of aid on citizen participation.

Many Western donors have a clearly expressed ambition to promote good governance and citizen engagement in recipient countries. For Swedish aid, strengthening democracy is a key objective and promoting citizen political participation a clearly expressed ambition (Regeringskansliet, 2016). Similarly, the World Bank aid branch has 'Governance and Institutions' as one of its central pillars, and emphasizes the goal to promote citizen participation to hold governments and private sector partners accountable (IDA, 2016). Lacking geocoded data on Swedish aid, let us consider the role of World Bank aid for promoting citizen participation.

6.1 World Bank aid and citizen engagement

I geographically match geocoded data on the World Bank's development projects to Africa over the period 1995-2014 (AidData, 2017d) to 98,449 respondents from 4 Afrobarometer survey waves (geocoded by Knutsen et al., 2016) in 29 African countries over the period 2002-2013 (Afrobarometer, 2017b)[9]. In line with the discussion in Chapter 4, the point coordinates in the aid data are used to link development projects to local survey respondents in the Afrobarometer.

---

[9] Namely Benin, Botswana, Burkina Faso, Burundi, Cameroon, Cape Verde, Cote D'Ivoire, Ghana, Guinea, Kenya, Lesotho, Liberia, Madagascar, Malawi, Mali, Mauritius, Morocco, Mozambique, Namibia, Niger, Nigeria, Senegal, Sierra Leone, South Africa, Tanzania, Togo, Uganda, Zambia and Zimbabwe.

Specifically, the coordinates of the surveyed Afrobarometer clusters (consisting of one or several geographically close villages or a neighborhood in an urban area) are used to match individuals to project sites for which there are precise point coordinates (precision category 1 and 2) in STATA. Measuring the distance from the cluster center points to the aid project sites makes it possible to identify the clusters located within a cut-off distance of at least one project site.

To begin with, let us consider the most aggregate level, including all World Bank aid and not making a distinction between projects going to different sectors. This is interesting since the mere presence of an influential donor in the local area could, irrespective of the specific project content, presumably affect citizen norms and behavior, and considering the possibility that projects with very different focal areas may still contain participatory elements. That is, even if promoting citizen participation is not the main objective of a project, it may well require beneficiaries to actively participate and cooperate, thereby potentially affecting participatory norms and behavior. Importantly though, the aggregation implies that a null result should not be interpreted as a failure of World Bank development projects explicitly targeting citizen participation. Including all World Bank development projects irrespective of sectoral focus, but restricting the sample to include only projects with precise geocodes (precision category 1 and 2) we are left with 688 projects spread across 6663 project locations.

To measure citizen participation, I use two Afrobarometer questions, asking 1) whether the respondent is a member of any voluntary organization or community group, and 2) whether the respondent, during the past year, has attended a community meeting. The dependent variables used in the estimations are two dummy variables taking the value one if the respondent answered these respective questions in the affirmative, and zero otherwise.

Following Isaksson and Kotsadam (2016, 2017) I use a spatial-temporal estimation strategy distinguishing between sites where a World Bank project is actually under implementation and sites where a project will be opened but where implementation had not yet been initiated at the time of the survey.[10] The main explanatory variables will thus focus on living near either an ongoing or future World Bank project site. While the fact that the Afrobarometer does not have a panel structure hinders us from following specific localities over time, before and after the arrival of a World Bank project, with this estimation strategy we can still compare areas selected as project sites before and during project implementation, thus making use of the time variation in the data.

Being interested in whether World Bank projects leave a footprint on local institutions – specifically, whether they promote citizen participation – we need to make an assumption about the geographical reach of this mark. As discussed in Knutsen et al. (2016), the appropriate cut-off distance from a project – within which an individual will be considered treated – is a trade-off between noise and size of the treatment group. With a too small cut-off distance, we get a small sample of individuals linked to ongoing and future project sites. On the other hand, a too large cut-off distance would include too many untreated individuals into the treatment group, leading to the estimated coefficient being biased towards zero due to measurement error in the explanatory variable. Following Isaksson and Kotsadam (2016, 2017), I use a 50 km cut-off (from the project coordinates to the coordinates of the cluster center point) in the main specification, but also present results using alternative cut-offs (25 and 75 km).

Assuming that citizen participation is affected within a 50 cut-off distance of a project, the estimation strategy thus includes three groups of individuals, namely those 1) within 50 km of at least one site with an ongoing World Bank project, 2) within 50 km of a site where a World Bank project will start (but where implementation was yet to start at the time the Afrobarometer visited that particular location) but not close to any ongoing projects, and 3)

[10] Originally used in the papers by Knutsen et al. (2016) and Kotsadam and Tolonen (2016) on the effects of mining.

more than 50 km from any World Bank project site, and the linear probability (OLS)[11] regression will take the form:

$$Participation_{ivt} = \beta_1 \cdot project\_ongoing + \beta_2 \cdot project\_future + \alpha_s + \delta_t + \gamma \cdot \boldsymbol{X}_i + \varepsilon_{ivt}$$

That is, the political participation of an individual *i* in cluster *v* (an Afrobarometer cluster consists of one or several geographically close villages or a neighborhood in an urban area) at year *t* is regressed on a dummy variable *project_ongoing* capturing whether the individual lives within 50 kilometres of an ongoing World Bank project, and a dummy *project_future* for living close to a site where a World Bank project is planned but not yet implemented at the time of the survey (the excluded reference group being those not living close to any World Bank project site). To control for variation in average citizen participation rates across time and space, the regressions include sub-national region fixed effects (α_s) and year fixed effects (δ_t). To control for individual variation in participation, a vector $(\mathbf{X}_i)$ of individual-level controls from the Afrobarometer are included.[12] The standard errors are clustered at the geographical clusters (enumeration area, town or neighborhood).

As described in detail in Isaksson and Kotsadam (2016, 2017), and as implied from the discussion in Chapter 5, interpreting the coefficient on project_ongoing (β_1) in isolation as capturing an effect of World Bank projects on citizen participation would necessitate that the location of World Bank projects is not correlated with pre-existing rates of citizen engagement. Considering that the World Bank project location decisions most likely are influenced by the pre-existing characteristics of project sites, such as the need for aid interventions or the infrastructure available to support aid interventions, this assumption does not appear reasonable. However, including project_future allows us to compare sites with ongoing projects to other areas selected as locations for World Bank projects, but where the project were yet

[11] Instead calculating marginal effects after probit regressions does not change the interpretation of any results.

[12] Age, age squared, gender and urban/rural residence.

to be initiated at the time of the survey. That is, we can compare areas before a project has been implemented with areas where a project is currently under implementation, and not only areas close to and far away from project sites. Below the estimated regression coefficient I therefore present test results for the difference between project_ongoing and project_future (i.e. $\beta_1 - \beta_2$), giving us a difference-in-difference type of measure that should help control for unobservable characteristics that may influence selection into being a World Bank project site.

The results are presented in Table 1. Looking at the coefficients on project_ongoing, we can note that living within 50 kilometres of a site where a World Bank project is currently being implemented is indeed associated with a greater probability of both being a member of a voluntary organization or community group and of having attended a community meeting during the past year. With respect to the former, the point estimate is sizeable; compared to people not living close to any World Bank project site, individuals with an ongoing project in their vicinity are approximately 8 percentage points more likely to be a member of a community group or voluntary organization. For community meeting attendance the equivalent difference is around 3 percentage points. Taken at face value, this thus seems to suggest that World Bank projects could help promote local citizen participation.

Table 1: World Bank aid citizen participation

VARIABLES	(1) Member of organisation	(2) Attend community meetings
project_ongoing	0.083***	0.029**
	(0.014)	(0.012)
project_future	0.031**	0.050***
	(0.015)	(0.015)
Individual controls	YES	YES
Year FE	YES	YES
Region FE	YES	YES
Diff-in-diff ongoing-future	0.0524	-0.0209
F test: ongoing-future=0	10.91	2.090
p value of F test	0.000960	0.148
Observations	76,593	76,593
R-squared	0.097	0.157

Robust standard errors (clustered by the survey clusters) in parentheses; *** p<0.01, ** p<0.05, * p<0.1

For reasons discussed above, however, the coefficient on *project_ongoing* should not be interpreted in isolation. Considering the parameter on *project_future* we can note that this too is positive in both estimations. This tells us that the World Bank tends to locate their projects in areas with higher rates of citizen participation to begin with. There are several potential reasons for this, one being that these areas offer well-functioning infrastructure to support aid interventions, and hence that project location decisions are motivated by convenience rather than (or at least as well as) need. An important caveat, however, is that we consider overall World Bank aid here; the picture might well be different if we focused on aid more specifically target to promoting citizen participation. Anyway, if this selection effect is not taken into account, the above estimation will over-estimate the effect of World Bank aid on citizen participation. Indeed, for community meeting attendance, the positive parameter on *project_future* appears larger than that on *project_ongoing,* suggesting that World Bank projects, if anything, might have a negative effect on attendance. To see if the differences between *project_ongoing* and *project_future* are statistically significant, however, we need to consider the test results presented in the lower part of the table.

As it turns out, with respect to membership in a community group or voluntary organization the difference between project_ongoing and project_future ($\beta_1 - \beta_2$) is positive and statistically significant, suggesting that compared to people living close to a future World Bank project site, individuals with an ongoing project in their vicinity are approximately 5 percentage points more likely to be a member. The negative difference between project_ongoing and project_future observed for community meeting attendance is, however, not statistically significant at conventional levels. Using alternative geographic cut-offs than the admittedly agnostic 50 kilometres around development project sites, specifically 25 and 75 kilometres (see Table A3), does not change the interpretation of results.

Hence, in line with the World Bank's stated efforts to promote civil society development and community participation, these findings offer some suggestive evidence that their projects stimulate membership in community groups and voluntary organizations. Notably, though, this effect would have been overstated had we not accounted for selection into being a World Bank project site.

Still, one could argue that these results are surprisingly strong seeing that we consider overall World Bank aid, irrespective of focus, rather than aid aimed specifically at strengthening democracy or promoting citizen political participation. Unfortunately, the broad sectoral classifications in the World Bank data contain no category focusing on, say, strengthening democracy or promoting citizen engagement. The closest we get is arguably the sector 'Public Administration, Law, and Justice', but as the name implies, this is more focused on strengthening formal regulations than popular participation. There is, however, an indicator focusing on the specific 'theme' of the intervention, with one sub-category being 'Participation and civic engagement'. In a next step, let us therefore focus on World Bank aid going to this relatively narrow group of projects. Due to the limited number of projects listed in this category,[13] we have to restrict the analysis to Kenya and Lesotho, which are the only countries with a decent number of respondents connected to ongoing and future project sites.

The results for this restricted sample, which are presented in Table 2, suggest an effect of projects focused on participation and civic engagement on community meeting attendance but not on membership in voluntary organizations or community groups. In particular, compared to people living close to a future World Bank participation project site, individuals with an ongoing participation project in their vicinity are approximately 12 percentage points more likely to have attended a community meeting during the past year. Hence, while we naturally cannot draw any far-reaching conclusions based on this small sample and quick break-down[14] of World Bank aid, the results offer some suggestive evidence that World Bank projects focused on participation and civic engagement do have some success in stimulating the same.

---

[13] There are 118 project sites within 50 kilometres of an Afrobarometer cluster, 17 percent of the total Afrobarometer sample has an ongoing project within 50 kilometres and only 1 percent has a future project within this distance.

[14] Further efforts to classify World Bank aid project contents – for example by going through project titles and project descriptions in the data – would most likely give a more accurate classification.

Table 2: World Bank projects focused on 'Participation and civic engagement' and participatory outcomes in Kenya and Lesotho

VARIABLES	(1) Member of other organisation	(2) Attend community meetings
project_ongoing	0.000	0.066***
	(0.026)	(0.025)
project_future	0.079	-0.052*
	(0.056)	(0.030)
Individual controls	YES	YES
Year FE	YES	YES
Region FE	YES	YES
Diff-in-diff ongoing-future	-0.0791	0.118
F test: ongoing-future=0	1.654	11.92
p value of F test	0.200	0.000676
Observations	5,185	5,182
R-squared	0.118	0.132

Robust standard errors (clustered by the survey clusters) in parentheses; *** p<0.01, ** p<0.05, * p<0.1

6.2 Chinese aid and citizen engagement: a comparison

In the 2014 'Aid Policy Framework' the Swedish government noted that "New donors do not always start out from democratic values or respect for human rights" and that this presents the international community "with new challenges and makes it even more important to emphasise human rights and democracy in international aid" (Regeringskansliet, 2014, p.19). Largest among the 'new donors' is China, and indeed their aid policy is quite different from that of many Western donors. In particular, the fact that China claims to follow a policy of non-interference in the domestic affairs of recipients (see e.g. Bräutigam, 2009; Tan-Mullins et al., 2010) arguably implies that they are unlikely to use their funds to promote civil society development and citizen participation. The principle, which is clearly spelled out in official Chinese documents (see e.g. State Council, 2014), contrasts with that of Western donors, who often tie their aid to economic and political reforms in recipient countries, and whose aid visions, as described above, often tend to focus on improvements in democracy, human rights and governance. It is thus interesting to compare findings on the effects of World Bank aid on local citizen engagement to results for Chinese aid.

The data on Chinese development projects is obtained from geo-referenced project-level data of AidData's Chinese Official Finance to Africa dataset (AidData, 2017c), introduced by Strange et al. (2015) and geocoded by Dreher et al. (2016). Since the Chinese government, unlike the OECD-DAC donors, does not release official project-level financial information about its foreign aid activities, this data is based on an open-source media based data collection technique (AidData's Tracking Underreported Financial Flows (TUFF) methodology, described in detail in Strange et al., 2013 and 2015), synthesizing and standardizing a large amount of information on Chinese development finance to African countries.[15]

Running equivalent estimations for Chinese development projects, the results, presented in Table 3, are indeed different. In particular, the estimations do not suggest a statistically significant difference between people living close to ongoing and future Chinese project sites, neither in terms of their probability to be members of community groups and other voluntary organizations nor in terms of their inclination to attend community meetings.

[15] For a more detailed description of the Chinese aid data, see Isaksson and Kotsadam, 2016 and 2017.

Table 3: Chinese aid and citizen participation

VARIABLES	(1) Member of organisation	(2) Attend community meetings
project_ongoing	0.023*	0.008
	(0.012)	(0.011)
project_future	0.026**	0.014
	(0.012)	(0.010)
Individual controls	YES	YES
Year FE	YES	YES
Region FE	YES	YES
Diff-in-diff ongoing-future	-0.00327	-0.00529
F test: ongoing-future=0	0.0569	0.205
p value of F test	0.812	0.651
Observations	65,646	65,646
R-squared	0.103	0.167

Robust standard errors (clustered by the survey clusters) in parentheses; *** p<0.01, ** p<0.05, * p<0.1

In sum, the results of these empirical examples illustrate the importance of exploring donor heterogeneity in aid impacts. In particular, they provide some suggestive evidence that World Bank projects, in line with the Bank's stated objective to promote civil society development and community participation, indeed promote citizen engagement. The equivalent results for China, who rather stress their policy of non-interference in the domestic affairs of recipients, suggest no such effect. Furthermore, the empirical exercise highlights how flexible the analysis of geospatial aid data can be in terms of scope versus depth. Prioritizing scope rather than depth, the analyst can focus on overall aid – i.e. all development projects irrespective of sector – from one or several donors. Being interested in aid with a specific focus, one can instead break the analysis down by e.g. sectoral or stated objectives of development projects, or for that matter, choose to follow the roll out of a specific development project. Finally, the empirical examples shed light on the need to account for selection effects when seeking to estimate causal effects of development projects, the World Bank results in particular indicating that the effect of living close to an ongoing project would have been overstated had we not accounted for selection into being a World Bank project site.

7. Potential for geospatial analysis of Swedish aid

This chapter will discuss the possibilities for geospatial analysis of Swedish development cooperation. Given that Swedish development cooperation is not yet geocoded on a wider scale, this will involve evaluating the potential for geocoding Swedish contributions and how the resulting data could be used. After a brief discussion of Swedish aid transparency initiatives and interest in geocoding, we turn to a general discussion of the costs, benefits and practical concerns involved when geocoding development projects. Next, we consider the structure of Swedish development cooperation and its implications for the value and feasibility of geocoding, and the coverage of Swedish contributions in existing geocoded datasets. Finally, we discuss ownership considerations and donor vs. partner country responsibility for geocoding initiatives.

7.1 Swedish aid transparency initiatives and interest in geocoding

Sweden has a long tradition of promoting government transparency. In fact, it was the first country in the world to enact a bill enforcing the principle of public access to information ("Offentlighetsprincipen") in 1776, making it a requirement for all authorities to publish documents unless legislation specified their classified status (Clare et al., 2016). Likewise, Sweden is at the forefront in terms of donor transparency. With a government established 'Transparency guarantee' (Government Offices of Sweden, 2010) specifying that all public documents and public information on Swedish development assistance will be made available online on 'OpenAid', a web-based information service about Swedish aid based on the International Aid Transparency Initative (IATI) data standard (OpenAid, 2017), it got a top-ten position in the 2016 Aid Transparency Index (Publish What You Fund, 2016).

As of yet, however, OpenAid does not provide geographical information about Swedish contributions to development projects. In particular, whereas it is explicitly stated that the publicly available information must explain "when, to whom and why money has been made available, and what results have been achieved" (Government Offices of Sweden, 2010), no mention is made of specifying where the money have gone within countries and OpenAid provides no sub-national mapping of projects.

While, to date, Swedish development cooperation is not geocoded on a routine basis, Sida has indeed expressed interest in geocoding initiatives. In 2011 Sweden, together with the World Bank and a handful of bilateral donors announced their endorsement of the Open Aid Partnership "to visualize their development assistance through innovative technologies, particularly through web-based open collaborative maps", with the main goal to synchronize the mapping tools among donors, building on the World Bank Mapping for Results (M4R) initiative (World Bank, 2012). Since then, however, they have seemingly taken different paths. As discussed in Chapter 4, The World Bank, now regularly publishes the latitude and longitude coordinates of all of their projects (available at AidData, 2017d) and display their interventions in detailed maps publicly available online (World Bank, 2017b).

Yet again, this may be about to change, in Sida's 2016 annual report it is noted that improvements have been made in the data management system underlying OpenAid. A specific example mentioned is that Sida's interventions now can be marked in terms of their 'exact geographical position' (Sida, 2017, p. 134). This is confirmed by Carl Elmstam (e-mail correspondence 2017-06-29), responsible for Openaid and transparency at Sida, who clarifies that OpenAid has prepared for geocoding in their data management system, in line with the IATI standard, but not yet finalized how to organize the reporting. He further notes that it is not yet decided how to deal with activities started before the tools to geocode were made available.

The IATI, set up to help provide information about aid flows to enable developing countries to plan and manage those resources effectively and citizens in developing as well as donor countries to hold their governments accountable for the use of those resources, encourages the publication of geocoded data. In their 2016 annual report it is noted that "Although a majority of publishers are publishing valid data for the core elements, we are keen to see an increase in those using the 'value added' elements of the Standard […] for example on geolocation and results" (International Aid Transparency Initiative, 2016, p. 15). They go on to point out that the number of IATI publishers providing 'at least some valid data' on the geographic coordinates of projects has gone up from 14 publishers doing so in 2014 to 47 doing so in 2016 (see Figure 7 in International Aid Transparency Initiative, 2016). Similarly, in a report on the impact of publicly available aid data, focusing on the Swedish OpenAid initiative, a key takeaway of the authors is that geocoded data could further improve the usefulness of the standardized data sets provided (Clare et al., 2016).

7.2 Geocoding aid: benefits, costs and considerations

Geocoding aid comes with a number of benefits. These relate to the opportunities for using geocoded data in evaluation of development cooperation, discussed in this report. A key argument has been that geospatial analysis of sub-national aid flows provides an important intermediate perspective, which rather than attempting to estimate country-wide impacts of total aid, or being restricted to analyze the impact of single randomized interventions, utilizes subnational geocoded aid and outcome data, often available over long time-periods, to systematically evaluate the subnational distribution and local effects of development projects. As such, the approach makes it possible to compare aid allocation patterns and development outcomes across sectors and localities in a disaggregated fashion.

Unlike monitoring inspired by results based management (RBM) approaches, which arguably encourages practitioners to think about development in terms of short-term easily measurable results like the number of people reached by a programme (see e.g.

Shutt, 2016), geospatial aid data enables researchers and evaluators to study broader development outcomes, such as changes in local attitudes and behavior, over longer time horizons. Moreover, seeing that the data is publicly available and stimulates a wealth of international research on the possible transformative impacts of development projects, it should arguably ease rather than add to the workload of donor and partner government evaluation functions. Increased data availability is likely to imply that there will be more valuable studies from external sources to draw on.

Relatedly, and as discussed in Section 2.1, geocoding aid also comes with benefits in terms of aid management, dialogue, coordination and transparency (Strandow et al., 2011). Describing the impact of geocoding their aid portfolio, Malawi's minister of finance and development planning indicated that "Being able to see in a map all the donor funded activities in Malawi has transformed the way we think about development and positively helped our own planning effort," (Hon. Ken Lipenga, as cited in World Bank, 2012).

As these benefits suggest, geocoding aid has an important public good dimension – the more geocoded aid data that is publicly available, the more useful the data will be. In terms of the above benefits, if the aid flows of the majority of donors are geocoded and publicly available it will be easier to get a clear picture of the aggregate distribution of aid within a country and thus for development partners, on both the donor and recipient side, to coordinate efforts so as to avoid financing gaps and wasteful project duplication and make sure that aid flows to those who need it most. The same goes for its function in terms of transparency and accountability. Publicly available mappings of aid flows will help citizens hold both donors and recipient governments accountable. The fact that geocoded aid data, unlike much of the data analyzed in RCTs or other studies of individual projects, tends to be publicly available and easily accessible online in principle means that anyone – a student at a local university, an investigative journalist or an independent evaluator – could conduct or replicate a geospatial analysis of aid flows. Mapping aid flows should thus not only be seen as something that helps the donor or partner country's own aid management and evaluation efforts, but also as a contribution to a public good.

That said, there are of course practical concerns as well as cost considerations to take into account. First of all, and as noted, geospatial analysis cannot be applied to all types of development projects. Effective use of geospatial aid data requires spatial differentiation of development programmes. Whereas development projects implemented in specific localities, say local interventions in terms of health, education or local governance, can be evaluated using geospatial methods, programmes with a national focus or without a physical project site, say general budget support or technical assistance to the central government, cannot. Hence, efforts to geocode development cooperation should be preceded by an initial screening of the relevance of and potential for geocoding different parts of the aid portfolio.

Assessing the costs of geocoding past and present projects, it is reasonable to make a distinction between portfolio level geocoding and detailed geocoding of individual projects. According to Bradley Parks, executive director at AidData (interview, 2017-05-19), one can get point based estimates of geographic locations for a wide range of projects in a donor's aid portfolio at a relatively low cost. However, for specific projects that the donor or partner country may be particularly interested in evaluating, detailed geocoding could result in more precise geographic measures, potentially involving exact project location boundaries, but at a higher cost per project.

For portfolio geocoding, AidData uses a double blind technique, whereby two trained coders independently go through aid project documentation in search for geographic identifiers, assigning uniform latitude and longitude coordinates, precision codes, and standardized place names to each geographic feature (for details about the methodology used, see Strandow et al., 2011). To determine which codes a project will be assigned, the two code rounds are checked against each other. If they are in agreement, these become the codes released for public use. If they do not agree, the project is moved into an "arbitration round" where a geocoding project manager reconciles the codes and assigns a master set of geocodes for all of the locations described in the project documentation. Due to this double blind technique and possible arbitration round, AidData estimates that geocoding one project requires 2.5 labour hours – one from each coder plus an additional 0.5 hour (on average) if it goes to arbitration. With the coders being paid around 10 USD per hour this amounts to an estimated geocoding cost of approximately 25 USD per geocoded project in the portfolio (Bradley Parks, interview 2017-05-19).

Another approach is to geocode specific projects and programmes that may be particularly suitable for geospatial analysis and that the donor or partner country may be especially interested in evaluating in a more detailed manner. While not to the same extent contributing to the public good that comprehensive publicly available geocoded aid data constitutes, it can be seen a feasible first step to get a sense of the geocoding process and what the data can be used for. Detailed geocoding of specific projects can result in much more precise geographic measures, but comes at a significantly higher cost per project (at least if using the services of AidData). Even if an aid project has not georeferenced its intervention sites during implementation, it is often possible to do so retrospectively and remotely. To mention an example, in order to conduct a GIE of the Government of Liberia's spatial development corridor strategy, which requires foreign investors in the natural resource sector to contribute to public goods, AidData assembled a geospatial database of all natural resource sector FDI in Liberia between 2004 and 2013 (see Bunte et al., 2017). By digitizing hard copies of maps from ministries and constructing geographical polygons that correspond to the specific tracts of land granted to concessionaires based on the field survey instructions in the concession contracts, they were able to calculate at a high-level of spatial resolution (1km x 1km grid cells) whether a particular location had been 'treated' with an FDI project. It is estimated that it took approximately 12 months and $100,000 to construct this dataset (BenYishay et al., 2017).

7.3 The potential for geocoding Swedish aid

In line with the above discussion, geocoding Swedish development cooperation– either on a detailed, but selective, project basis or on a more large-scale portfolio basis – would have to entail an initial screening of the geocoding potential of Swedish contributions to development projects. While a comprehensive screening is beyond the scope of this report, we can consider the structure of Swedish development cooperation and some illustrative cases in order to get a sense of how to approach the issue.

Utilizing the data available at OpenAid (2017), we can gain some valuable insights with respect the structure and content of the Swedish portfolio for development cooperation. In 2016 Swedish aid amounted to a total of 4.5 billion USD, distributed across 4396 aid activities. Figure 6 displays the sectoral allocation of the aid portfolio.

For some of these sectors, the value and feasibility of geocoding is questionable. The biggest post – aid to refugees in the donor country – involves money deducted from the aid budget to cover expenditures for the reception of refugees from ODA-eligible countries during their stay in Sweden. Being used to cover expenses in Sweden rather than on site in a partner country this part of the aid portfolio is arguably not a reasonable candidate for geocoding. For the second biggest post, 'unallocated/unspecified' aid, the feasibility of geocoding is also questionable. A large portion of this aid (830 out of 983 million USD) is multilateral core support, i.e. non-earmarked support to multilateral organizations such as the EU, the UN and the World Bank. Being non-earmarked, these flows are of course difficult to follow and thus to geocode.

Figure 6: Main sectors of Swedish aid in 2016

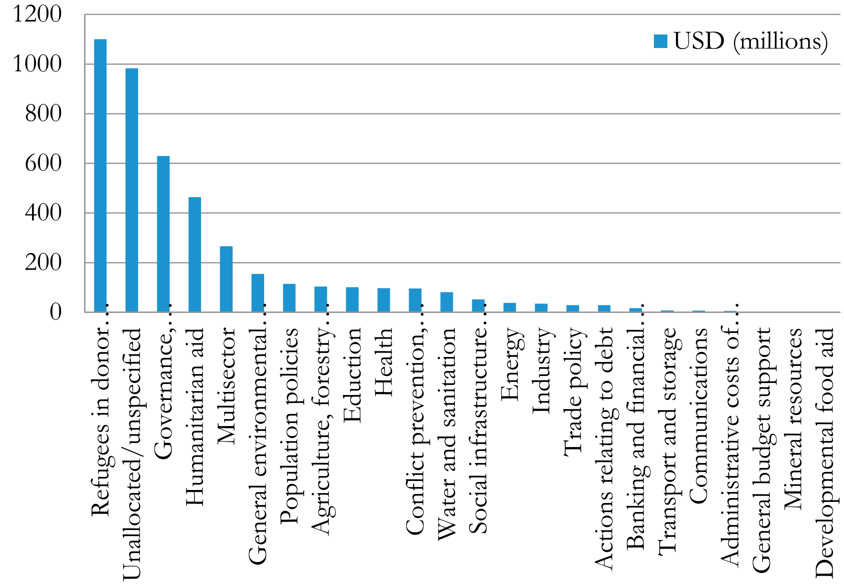

Source: Based on data from OpenAid (2017)

However, for the important sectors that follow – i.e. governance activities, humanitarian aid, multi-sector, environmental and population initiatives, and activities in the agricultural sector, education, health and conflict prevention – a large share of the aid activities are likely to be relevant candidates for geocoding. To get a sense of the magnitude, if we deduct from the USD 4.5 billion aid portfolio aid going to sectors where the ability to follow the money or the spatial differentiation of aid interventions is questionable,[16] over half (52%) of the aid budget still remains. While all of these flows are presumably not possible to trace to precise geographic locations, a reasonable share should be feasible

[16] In particular, the 1.1 billion going to refugees in donor countries, the 983 million of unallocated/unspecified aid, the 29.3 million to trade policy, the 29.1 million to actions relating to debt, the 16.6 million to banking and financial services, the 6.4 million to administrative costs of donors and the 392 thousand to general budget support.

to trace to, say, the district level (recall the discussion of the different geocoding precision categories in Section 4.1).[17] Indeed, in some cases the geographical data may already be available, seeing that some partner country ministries now geocode incoming aid flows and considering that Sweden often co-fund projects with other donors (such as the World Bank) who geocode their aid.

Figure 7: Five main sectors of Swedish aid to Tanzania in 2016 (in millions of USD)

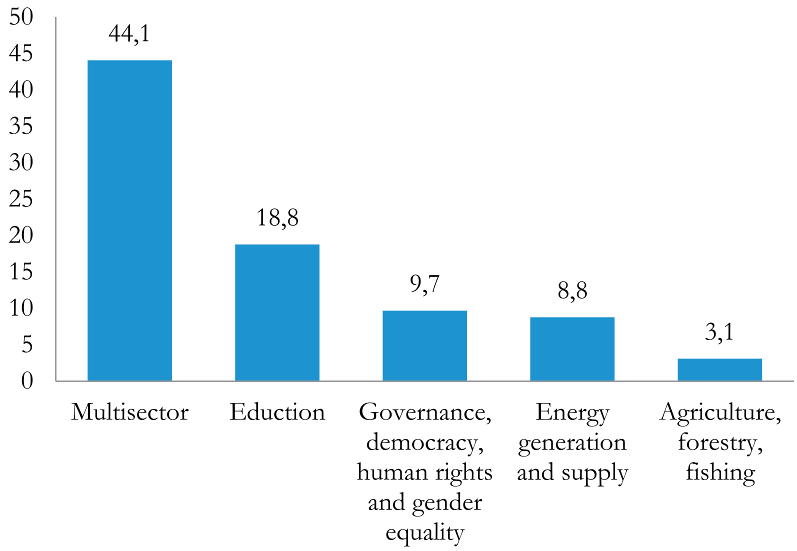

Source: Based on data from OpenAid (2017)

Consider the case of Tanzania, which has consistently been among the largest recipients (in 2015 it was the largest recipient) of Swedish aid over the recent years (OpenAid, 2017). In 2016 Sida's aid portfolio to Tanzania amounted to USD 89 million distributed over 102 activities. In terms of sectoral allocation, the clearly biggest post was Multisector aid (USD 44 million) followed by

[17] Furthermore, in order to get a full picture of aid flows to a particular country, even programs with no clearly defined and precise geographical location should be geocoded, although with a specified lower level of precision.

Education, Governance, Energy and Agriculture (see Figure 7). While at first sight, multisector aid may seem diffuse and difficult to track, actually considering the specific projects involved modifies this picture. Specifically, the 44 million USD in multisector aid to Tanzania includes 24 activities, out of which the by far biggest project is the Productive Social Safety Net (PSSN) program to which Sweden contributed nearly 36 million USD.

PSSN is described as a cash transfer project that is nationwide in scope and directed to approximately 1.1 million households identified as extremely poor. While OpenAid provides useful information about the project, including the implementing agency (Tanzania Social Action Fund, TASAF) and other donors involved (e.g. the World Bank and DFID), it does not provide any geographical information.

Figure 8: Geographical roll out of PSSN in Tanzania 2013-2015

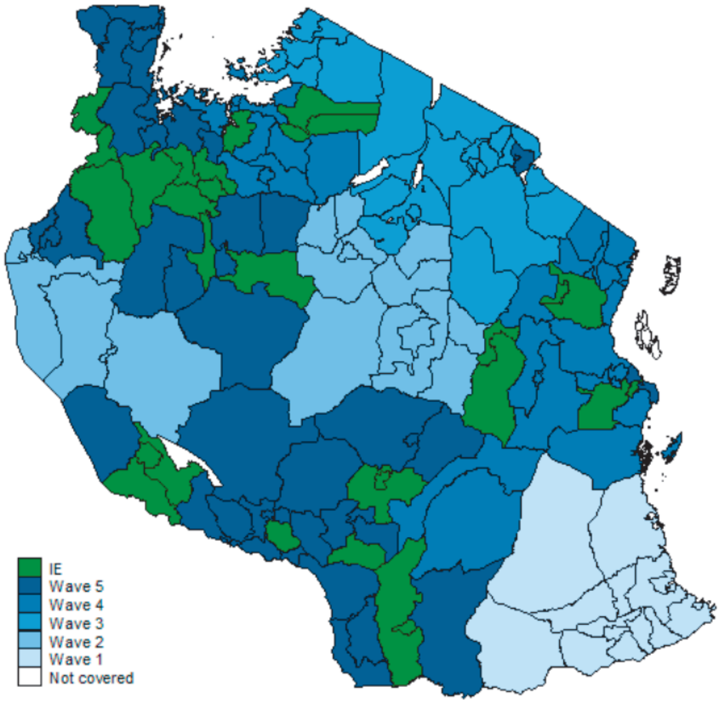

Source: World Bank (2016)

However, a quick online search makes clear that such information is available. In particular, documentation from the World Bank and the Tanzania National Bureau of Statistics reveals that the project was rolled out in five waves during 2013-2015 (see Figure 8), and includes statistics on the 161 project areas covered as well as the number of villages/communities covered in each project area (World Bank, 2016).

Geocoding this project, the coders would, in line with the discussion in the previous section, dig into the geographical information that could be obtained from the official project documentation, and possibly, depending on the requested level of detail and what information the implementing agency is willing to disclose, obtain more detailed geographic information from project boundary files. Considering the World Bank's involvement, the information should in this case be relatively easy to obtain.

Interestingly, the PSSN project is the subject of an impact evaluation, where 16 project areas were randomly selected for evaluation. If the project were to be geocoded, the analyst could utilize the temporal and spatial variation in project coverage given by the roll out of the project and compare findings from the impact evaluation of selected areas to findings from a GIE evaluation of the scaled up project in its entirety.

Let us consider another Tanzanian example, this time from the education sector, i.e. the second biggest sector in terms of the amount of Swedish aid received in 2016. The 18.8 million USD that went to the Tanzanian education sector in 2016 was split across five different projects. The by far biggest of these is the 'Big Results Now in Education' project, where the Swedish contribution amounted to nearly 17 million USD. The overall objective of this program, as described on OpenAid, is to improve education quality and learning outcomes in primary and lower secondary school through improved transparency, incentives, teacher conditions and provision of support where needed. Specific activities mentioned include performance based school incentive grants, non-financial performance incentives for teachers, school improvement toolkits and teacher training programs. The need to identify schools that are lagging behind is stressed.

Judging from this description, which suggests support taking place at the school level, it should clearly be possible to geographically trace the intervention. While OpenAid does not provide any geographical information about the project, a quick online search gives for instance a listing of all individual schools in Tanzania, ranked in terms of a number of school performance indicators, and a district level map of school performance, both compiled by the National examinations council of Tanzania (2017) for the Big results now project. Hence, there seems to be great potential for identifying the geographic location of schools covered (and schools not covered) by the project and thus for evaluating the project (e.g. in terms of intended school performance outcomes) using geospatial methods. Indeed, considering the magnitude of the program, it is quite likely that the geographic information already exist in the project documentation.

In some partner countries geographic information on development initiatives where Sweden is involved may in fact be readily available on a somewhat wider scale. As noted in Section 4.1, partner country ministries responsible for managing incoming aid flows now publish sub-nationally geocoded development project data in a number of countries. The coverage of these datasets varies from country to country (see AidData, 2017c), though, and hence they should not be seen as providing a full representation of a specific donor's involvement. In a report on Malawi's geocoding experiences, for instance, it is noted that a few donors are missing entirely (as an example they mention DANIDA) either because they do not report through Malawi's Aid Management Platform or because their activities are not well captured in reporting systems (Weaver et al., 2014).

To get a picture of the coverage of Swedish contributions to these countries, let us consider the case of Uganda. In the Ugandan geocoded aid data, described in Section 4.1, Sweden shows up as part of a few selected interventions, either grouped together with other donors or specifying Sweden's individual contribution to a development intervention. In particular, focusing on projects that are geocoded with some precision – at least to the district level – Sweden is involved in 10 projects (spread across 54 project sites), including for instance support to the work of the Naguru Teenage Health Centre on reproductive health. While this project can also be identified on OpenAid, where they specify that it has involved disbursements of around 3 million USD (or approximately 20 million SEK) over the period 2011-2014, they do not disclose the geographic information available from AidData. For some projects it is thus simply a question of reconciling the information available from the partner country ministry (via AidData) with what information is displayed on OpenAid.

Next, let us consider the geocoded data for Afghanistan, which covers the period since 2001 and contains information on 1580 projects spread across 7403 project sites (AidData, 2017c). Here, Sweden shows up in a small number of joint programs (over the entire period, 7 programs, spread across 45 sites, to be specific). Considering that Afghanistan was the largest recipient of Swedish aid in 2016, with 54 listed activities only in that particular year (OpenAid, 2017), this is clearly not representative of the Swedish involvement in the country.

Having a closer look at the distribution of Swedish funds to Afghanistan, the largest post in 2016 (approximately USD 20.8 million) was Core Support to the Swedish Committee for Afghanistan (OpenAid, 2017). SCA is not mentioned in the geocoded data for Afghanistan. However, given their reported focus on education, healthcare, community governance and rural livelihoods, many of their interventions should be geographically traceable. With respect to their education projects, for instance, they report that in 2015 68,654 children were enrolled in 515 SCA supported schools (Swedish Committee for Afghanistan, 2017). The specificity of this information suggests that it should be possible to geographically identify these schools. In their reporting of results (Svenska Afghanistankommittén, 2016), they indicate that while they can follow developments in the number of students, patients etc., changes in other outcomes, such as attitudes and human rights, require access to other sources of information over a longer time horizon. Here, being able to geographically connect data on their interventions with survey data (from for instance the Demographic and Health Survey) on relevant outcomes provides a promising alternative given data collection constraints due to security considerations in the country.

The second largest post in Sweden's development cooperation with Afghanistan in 2016 was a 16.6 million USD contribution to the Afghanistan Reconstruction Trust Fund (OpenAid, 2017). While the Swedish contribution to this fund is not specified in the geocoded data for Afghanistan, the projects of the fund itself is part of the geocoded data, and could thus be referenced in the project description on OpenAid.

Hence, while geocoded data from partner country ministries at this stage does not provide a full representation of Sweden's involvement in the countries in question, they clearly provide useful information that should be utilized. For some development cooperation initiatives, such as the Swedish support of the Naguru Teenage Health Centre in Uganda, specified in the geocoded dataset provided by AidData, one option is to simply incorporate the geographical information from AidData in the project description provided on OpenAid. For others, such as the Swedish contribution to the Afghanistan Reconstruction Trust Fund, where the geocoded data includes project level geocodes for the concerned Fund but does not specify Sweden's particular involvement, it seems a good idea to at least provide a reference to the geocoded dataset on OpenAid. A reasonable first step is thus to screen and compile the data that is available from selected partner country ministries (via AidData) so as to get an overview of existing geographic information of Swedish-funded projects on OpenAid.

Importantly though, the Tanzanian examples above point to the potential for geocoding development initiatives to which Sweden has contributed even when data from the partner country ministry is not available (which indeed applies to the great majority of aid flows). As noted, here too, the geographical data may in some cases already be available, seeing that Sweden often co-fund development programs with other donors who geocode their aid. In other cases, coders could go through project documentation, either for detailed geocoding of individual projects or for more comprehensive portfolio level geocoding, in line with the description in Section 7.2.

7.4 Geocoding initiatives: for whom, by whom?

The above examples demonstrate that multiple actors are involved in the geocoding of development initiatives. This is unsurprising. With the public good character of geocoded aid data follows the question of who should provide it. Basically, who should supply the resources and who should do the work – the donors or the partner country governments? Joint efforts are seemingly needed.

There are good arguments for saying that the ultimate responsibility should lie with the partner country ministry responsible for managing incoming inflows. Given local ownership of development initiatives, with projects and programs being initiated by local stakeholders and often being funded by multiple donors, the partner country government should arguably be in the best position to provide an overview of ongoing initiatives and implementing agencies. Likewise, it is reasonable to argue for local ownership of development results; local actors should be the ultimate stakeholders when assessing outcomes of development initiatives (which furthermore are often co-funded by the partner country government). Furthermore, to get a full picture of the sub-national aid allocation pattern in a country, and thus be able to reap the benefits of the data in terms of aid management, it is important to capture a large share of all development projects in the country, as opposed to most activities of a few selected donors.

Then again, donors too benefit from geocoded aid data, both in terms of aid management and for evaluation purposes. In particular, information on what works better and worse when implementing development projects is, if deemed relevant in the partner country, also clearly relevant from a donor perspective. In addition, the partner country is likely to face capacity and resource constraints that could constitute important obstacles for geocoding.

Against this background, there seems to be a need for collaborative efforts. In Malawi, for instance, which in 2011 became the first partner country to make sub-nationally geocoded development project data on incoming aid flows publicly available, the geocoding initiative was not driven by the government alone, but led by a team at the University of Texas in collaboration with AidData, Development Gateway and the Government of Malawi (Weaver et al., 2014).

Weaver et al. (2014) describe some key lessons from the geocoding experiences of Malawi on this account. They emphasize the importance of government ownership in all aspect of data management, from data creation to data dissemination and analysis. However, they also note that there is likely a substantial up-front investment in data creation in order to catch up to the present point in time by geocoding past and ongoing activities in the local Aid Information Management System (AIMS), and that this work may best be done by outside groups who have the resources and technical capacity (e.g. AidData or Development Gateway). Yet again, they stress that such a third-party solution is not sustainable, and that the government and its development partners must be prepared to take over the data creation process immediately upon completion of this up-front investment in order to avoid a large backlog of new projects in need of geocoding. The authors further suggest that the government should take a lead role in requesting the compliance of its development partners. Since donor offices are often subjected to data and reporting requests from a variety of actors, there is likely to be some resistance to providing the required geographical information. The authors find that this resistance is softened when a data request is accompanied by a letter from a high-ranking official in the relevant ministry. In a similar vein, the authors stress that it is critical to be sensitive to the existing workload of donor country office staff and provide simple ways to report information, noting for instance that AidData has developed a location collection template and recommends that donor staff simply share this excel template with each project manager as a survey that will take them only a few minutes to complete.

Hence, while there are seemingly strong arguments for the partner country government taking the lead in the development of the local Aid Information Management System (AIMS), they may require significant support to be able to handle the initial investment in data creation (i.e. to do the country level equivalent of portfolio geocoding, in line with the description in Section 7.2), and at a later stage, the cooperation and compliance of development partners in the continuous reporting of geographic information of upcoming projects. This shared responsibility does of course not only apply to the reporting of geographic information on aid funded projects. As described in a briefing note on the Aid Information Management System of the Government of the Republic of South Sudan (Republic of South Sudan, 2014, pp. 1-2), whereas the day-to-day operation of the AIMS is undertaken by the Department for Aid Coordination in the Ministry of Finance and Economic Planning, the data input to the AIMS is the responsibility of all development partners. It is specified that "Each donor agency is responsible for entering data relating to their own bilateral projects" and that "For pooled funds, the ultimate responsibility for data input lies with the lead donor of the fund".

A key point in this context is to avoid dual, unaligned, reporting systems. With standardized reporting frameworks it should be possible to reconcile data from different sources, e.g. for Sweden to incorporate geographic information on projects obtained from the partner country ministry in their own project description provided on OpenAid, and for partner country ministries to draw on information routinely collected by donors involved. Supporting partner countries' in developing their Aid Information Management Systems in line with IATI standards is thus a sensible undertaking. As noted, AidData provides such services (see AidData, 2017b). Furthermore, in 2016 Development Gateway launched the 'Open Aid Geocoder', which is a geocoding tool developed in partnership with the World Bank's Open Aid Partnership in line with IATI reporting standards (Development Gateway, 2016).

8. Concluding remarks

A rapid expansion in the availability of sub-nationally geocoded data on aid interventions as well as of geocoded data on relevant development outcomes opens for new possibilities in terms of aid evaluation. This report introduced and discussed a novel approach to aid evaluation utilizing these new sources of information, namely geospatial analysis of development cooperation.

The approach combines geocoded aid data – that is, information on the geographic location of specific development projects – with geocoded information from other data sources, such as individual/household level survey data, to evaluate the sub-national distribution and local effects of development projects systematically and on a wide scale, potentially across multiple recipient countries. That is, rather than estimating country-wide impacts of total aid, which is notoriously difficult, or being restricted to analyzing the impact of single projects, geospatial analysis of aid enables the researcher or policy evaluator to systematically estimate, for instance, whether health projects are allocated to the sub-national areas in greatest need of health interventions, whether they have direct effects on relevant health outcomes in the targeted areas, as well as their potential indirect effects on other relevant development outcomes over the longer term.

The report has highlighted several attractive features of geospatial aid evaluation. Having a clear picture of the sub-national allocation of development projects can help highlight potential financing gaps and inequalities in the aid distribution and thus make partner countries, in cooperation with donors, better able to coordinate donor efforts and direct aid flows to the areas where they will do most good. Furthermore, publicly available mappings of aid flows within countries can help hold both donors and partner country governments accountable to their intended beneficiaries. In terms of impact evaluation, geospatial methods have the advantage that they can control for potential confounding factors at fine geographic levels and that they are well-suited for quasi-experimental methods for causal identification. Furthermore, GIEs are relatively strong in terms of generalizability – both across space and over time. In a spatial sense since GIE makes it possible to estimate the impact of a multitude of development projects, potentially across several countries, and in a temporal sense since GIEs often draw on outcome data that covers long time periods, and thus make it possible to evaluate the sustainability and long-term impacts of development projects. Finally, since the approach enables researchers and evaluators to utilize comprehensive existing data materials, it is comparatively cost-effective.

The key limitation of geospatial analysis of aid is that it is not appropriate for studying all types of development cooperation. To be able to geocode a project in a meaningful way it has to have a physical project site (or several). And whereas some projects, say local interventions in terms of health, education or local governance, are implemented in a well-defined geographical area, others are realized at more aggregate levels, such as a district or greater administrative region, or do not have a clear project site.

That said, though, not taking advantage of the opportunities geospatial analysis of aid has to offer, and the rapidly expanding geocoded data that is publicly available, would be wasteful. From a researcher's, student's or evaluator's viewpoint, the abundance of publicly available geocoded data on development initiatives and development outcomes provides ample opportunities for studying systematic aid allocation patterns as well as short and long term local development impacts of projects.

From a Swedish development cooperation perspective, considering that Swedish aid flows are not yet geocoded on a wider scale, taking advantage of geospatial aid evaluation methods would require geocoding efforts. With Sweden being a frontrunner in terms of donor transparency, this appears a natural step to take. Increased publication of geocoded data is encouraged by the International Aid Transparency Initiative, and was in a recent report emphasized as something that would increase the usefulness and impact of Sweden's publicly available aid data. And indeed, there are now reports to suggest that OpenAid will soon be able to mark Sida's interventions in terms of their geographical position.

A reasonable first step is to screen and compile already available geocoded data pertaining to Swedish aid flows, and incorporate it in the Swedish records. This is relevant seeing that some partner country ministries now geocode incoming aid flows and considering that Sweden often co-fund projects with other donors who geocode their aid. Being interested in proceeding from here, there are different – by no means mutually exclusive – options.

One option is to hire coders to do broad portfolio level geocoding of past (how far back in time to go will of course have to be decided) and ongoing projects. This comes with advantages in terms of getting an overview of the destinations of Swedish contributions and of being able to share the compiled information with strained partner country ministries interested in publishing geocoded information on incoming aid flows. However, since all development cooperation projects cannot be geocoded effectively, such efforts should be preceded by an initial screening of the relevance of and potential for geocoding different parts of the aid portfolio.

Another option is to geocode specific projects of particular interest in a more detailed manner. While not to the same extent contributing to the public good that comprehensive publicly available geocoded aid data constitutes, it can be seen a feasible first step to get a sense of the geocoding process and what the data can be used for.

Yet another option (again, they are not mutually exclusive) would be to provide support to partner country initiatives to geocode incoming aid flows, in line with what the concerned ministries have done in e.g. Malawi and Uganda. Given local ownership of development initiatives as well as of results, there are good arguments for saying that the ultimate responsibility for geocoding should lie with the partner country ministry in charge of managing incoming inflows. That said, though, the partner country is likely to face capacity and resource constraints that could constitute important obstacles for geocoding. In particular, they may require significant support to be able to handle the initial investment in data creation in order to catch up to the present point in time by geocoding past and ongoing activities.

Having cleared the backlog of geocoding past projects, there are seemingly strong arguments for the partner country government taking the lead in the development of their local Aid Information Management System. Importantly though, this will still require the cooperation and compliance of development partners in the continuous reporting of geographic information of upcoming projects. Updating the reporting routine of project managers it is important to provide simple ways to report geographic information. A key point should be to avoid dual, unaligned, reporting systems. Supporting partner countries' in developing their Aid Information Management Systems in line with the International Aid Transparency Initiative (IATI) standards is thus a sensible undertaking.

References

Afrobarometer (2017) "Data usage and access policy", available at: http://afrobarometer.org/data/data-use-policy (accessed 2017-04-28).

Afrobarometer (2017b) "Data", available at: http://afrobarometer.org/data/ (accessed 2017-05-21).

AidData (2017a) "AidData & USAID Geospatial Impact Evaluation Opportunities", available at: http://aiddata.org/aiddata-usaid-geospatial-impact-evaluation-opportunities

AidData (2017b) "Geocoding", available at: http://aiddata.org/geocoding (accessed: 2017-06-30).

AidData (2017c) "Subnational Geospatial Research Datasets" available at: http://aiddata.org/subnationalgeospatial-research-datasets.

AidData (2017d) "WorldBank_GeocodedResearchRelease_Level1_v1.4.2 geocoded dataset", Williamsburg, VA and Washington, DC: AidData. Available at: http://aiddata.org/research-datasets (accessed 2017-05-02).

Aiddata (2017e) "Geoquery", available at: http://geo.aiddata.org/query/ (accessed 2017-05-24).

AidData (2016) "UgandaAIMS_GeocodedResearchRelease_Level1_v1.4.1 geocoded dataset", Williamsburg, VA and Washington, DC: AidData. Available at: http://aiddata.org/research-datasets (accessed 2017-05-02).

Angrist, J.D. and J-S. Pischke (2009) Mostly Harmless Econometrics: An Empiricist's Companion, Princeton, New Jersey: Princeton University Press.

Armed Conflict Location and Event Database (2017) available at http://www.acleddata.com/ (accessed 2017-05-10).

Austin, P.C. (2011) "An Introduction to Propensity Score Methods for Reducing the Effects of Confounding in Observational Studies", Multivariate Behavioral Research, 46(3), pp. 399–424.

Befani, B. (2016) "Pathways to change: Evaluating development interventions with qualitative comparative analysis (QCA)", Rapport 2016:5, Expertgruppen för Biståndsanalys (EBA), Stockholm.

BenYishay, A., Parks, B., Runfola, D. and R. Trichler (2016) "Forest Cover Impacts of Chinese Development Projects in Ecologically Sensitive

Areas", AidData Working Paper 32, December 2016, Williamsburg, VA: AidData.

BenYishay, A., Runfola, D., Trichler, R., Dolan, C., Goodman, S., Parks, B., Tanner, J., Heuser, S., Batra, G. and A. Anand (2017) "A Primer on Geospatial Impact Evaluation Methods, Tools, and Applications", mimeo, Williamsburg, VA: AidData.

Berlin, M. P., Bonnier, E. and A. Olofsgård (2017) "The donor footprint and gender gaps", mimeo, SITE Stockholm School of Economics.

Björkman, M. and J. Svensson (2009) "Power to the people: evidence from a randomized field experiment on community-based monitoring in Uganda", Quarterly Journal of Economics. 124, pp. 735-769.

Bourguignon, F. and J. W. Gunning (2016) "Foreign aid and governance: A survey", mimeo.

Brazys, S., Elkink, J. A. and G. Kelly (2017) "Bad neighbors? How co-located Chinese and World Bank development projects impact local corruption in Tanzania", The Review of International Organizations, 12(2), pp 227–253.

Bräutigam, D. (2009) The Dragon's Gift: The Real Story of China in Africa, Oxford, UK: Oxford University Press.

Briggs, R. C. (2014) "Aiding and Abetting: Project Aid and Ethnic Politics in Kenya" World Development, vol. 64, pp. 194-205.

Briggs, R. C (2017) "Does Foreign Aid Target the Poorest?", International Organization, vol. 71, Winter 2017, pp. 187–206.

Bräutigam, D. and S. Knack (2004) "Foreign Aid, Institutions, and Governance in Sub-Saharan Africa ", Economic Development and Cultural Change, 52(2), pp. 255-285.

Buchanan, G. M., Parks, B. C., Donald, P. F., O'Donnell, B. F., Runfola, D. Swaddle, J. P., Tracewski, L. and S.H.M. Butchart (2016) "The Impacts of World Bank Development Projects on Sites of High Biodiversity Importance", AidData Working Paper 20, February 2016, Williamsburg, VA: AidData.

Bunte, J. B., Desai, H., Gbala, K., Parks, B., and D. M. Runfola (2017) "Natural Resource Sector FDI and Growth in Post-Conflict Settings: Subnational Evidence from Liberia", AidData Working Paper 34, February 2017, Williamsburg, VA: AidData.

Clare, A., Verhulst, S. and A. Young (2016) "Openaid in Sweden: Enhanced Transparency and Accountability in Development Cooperation", Open Data's impact, January 2016, available at: http://odimpact.org/static/files/case-study-sweden.pdf (accessed: 2017-06-29).

Clemens, M. A., Radelet, S., Bhavani, R. R. and S. Bazzi (2012) "Counting Chickens when they Hatch: Timing and the Effects of Aid on Growth", Economic Journal, 122(561), pp. 590–617.

Croicu, M. and R. Sundberg (2016) "UCDP GED Codebook version 5.0", Department of Peace and Conflict Research, Uppsala University.

Deaton, A. (2010) "Instruments, Randomization, and Learning about Development", Journal of Economic Literature, 48(2), pp. 424-455.

Demographic and Health Survey (2017) "Data" available at: http://dhsprogram.com/Data/

Development Gateway (2016) "Development Gateway launches the Open Aid Geocoder", March 15 2016, available at: http://www.developmentgateway.org/blog/development-gateway-launches-open-aid-geocoder (accessed 2017-06-30).

Dionne, K. Y., Kramon, E. and T. Roberts (2013) "Aid Effectiveness and Allocation: Evidence from Malawi", mimeo.

Dreher, A., Fuchs, A., Hodler, R., Parks, B. C., Raschky, P. A., and M. J. Tierney (2016) "Aid on demand: African leaders and the geography of China's foreign assistance", mimeo.

Dreher, A. and S. Lohmann (2015) "Aid and growth at the regional level", Oxford Review of Economic Policy, 31(3-4), 420-446.

Findley, M.G., Powell, J., Strandow, D. and J. Tanner (2011) "The Localized Geography of Foreign Aid: A New Dataset and Application to Violent Armed Conflict", World Development, 39(11), pp. 1995-2009.

Francken, N., Minten, B. and J. Swinnen (2012) "The Political Economy of Relief Aid Allocation: Evidence from Madagascar", World Development, 40(3), pp. 486-500.

Geographically based Economic data (2011) "Geographically based Economic data (G-Econ)", Yale University, Last revised May 2011, available at http://gecon.yale.edu/ (accessed 2017-05-10).

Goodman, S., BenYishay, A. and Runfola, D. (2017) "Overview of the geo Framework", AidData, Available online at: http://geo.aiddata.org/.

Government Offices of Sweden (2010) "A transparency guarantee in Swedish development assistance", Diarienummer UD10.052, June 11, 2010, available at: http://www.regeringen.se/informationsmaterial/2010/06/ud10.052/ (accessed 2017-06-1).

Henderson, J. V., Storeygard, A. and D. N. Weil (2012) "Measuring Economic Growth from Outer Space" American Economic Review, 102(2), pp. 994-1028.

International Aid Transparency Initiative (2016) "Annual report 2016", available at: http://www.aidtransparency.net/wp-content/uploads/2017/03/IATI-Annual-Report-2016-EN.pdf (accessed: 2017-06-29).

IPUMS TERRA (2017) "About IPUMS TERRA", available at https://www.terrapop.org/about-us (accessed 2017-05-10).

Isaksson, A-S and A. Kotsadam (2016) "Chinese aid and local corruption", Working Papers in Economics, no 667, University of Gothenburg.

Isaksson, A-S and A. Kotsadam (2016) "Chinese financial flows and union involvement", mimeo, University of Gothenburg.

Jablonski, R. S. (2014) "How Aid Targets Votes: The Impact of Electoral Incentives on Foreign Aid Distribution", World Politics, 66(2), pp. 293-330.

Kotsadam, A. and A. Tolonen (2016) "African Mining, Gender, and Local Employment", World Development, Vol. 83, pp. 325–339.

Knutsen, C. H., Kotsadam, A., Hammersmark Olsen, and T. Wig (2016) "Mining and local corruption in Africa", forthcoming in American Journal of Political Science (for an older version see Memo 09/2015-v1, University of Oslo).

Kremer, M. and E. Miguel (2007) "The illusion of sustainability", Quarterly Journal of Economics, 122, pp. 1007-1065.

Kremer, M. and E. Miguel (2004) "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities", Econometrica, 72(1), pp. 159–217.

Odokonyero, T., Ijjo, A., Marty, R., Muhumuza, T. and G. O. Moses (2015) "Sub-national Perspectives on Aid Effectiveness: Impact of Aid on Health Outcomes in Uganda", AidData Working Paper 18, December 2015, Williamsburg, VA: AidData.

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

Muchapondwa, E., Nielson, D., Parks, B., Strange, A.M. and M.J. Tierney (2014) "Ground-truthing' Chinese development finance in Africa: Field evidence from South Africa and Uganda", WIDER Working Paper, No. 2014/031.

NASA Socioeconomic Data and Applications Center (2017) "Gridded Population of the World", Available at: http://sedac.ciesin.columbia.edu/data/collection/gpw-v3 (accessed 2017-05-08).

National examinations council of Tanzania (2017) "Big results now", available at: http://www.necta.go.tz/brn (accessed: 2017-06-02).

National Oceanic and Atmospheric Administration (2017) "Earth Observation Group (EOG)", available at https://ngdc.noaa.gov/eog/index.html (accessed 2017-05-10).

Nordhaus, W., Azam, Q., Corderi, D., Hood, K., Victor, N.M., Mohammed, M., Miltner, A. and J. Weiss (2006) "The G-Econ Database on Gridded Output: Methods and Data", Yale University, May 12 2006.

Nunnenkamp, P., Sotirova, A. and R. Thiele (2015) "Do Aid Donors Specialize and Coordinate within Recipient Countries? The Case of Malawi", Aid Data Working Paper 10, June 2015.

Nunnenkamp, P., Öhler, H. and M. Sosa Andrés (2012) "Need, merit, and politics in multilateral aid allocation: A district-level analysis of World Bank projects in India", Kiel Working Paper no. 1783, Kiel Institute for the World Economy, July 2012.

Olofsgård, A. (2014) "Randomized controlled trials: Strengths, weaknesses and policy relevance", Rapport 2014: 1, Expertgruppen för Biståndsanalys (EBA), Stockholm.

Peace and Research Institute Oslo (2017) "About the PRIO-GRID 2.0", available at http://grid.prio.org/#/about (accessed 2017-05-10).

Powell, J. and M. G. Findley (2012) "The Swarm Principle? A Sub-national Spatial Analysis of Donor Coordination in Sub-Saharan Africa", mimeo, may 17 2012.

Publish What You Fund (2016) "Aid transparency index 2016" available at: http://ati.publishwhatyoufund.org/ (accessed 2017-06-01).

Rajlakshmi, D. and C. Becker (2015) "The Foreign Aid Effectiveness Debate: Evidence from Malawi", Aid Data Working Paper 6, March 2015.

Regeringskansliet (2014) "Aid policy framework– the direction of Swedish aid", Government Communication 2013/14:131, 13 March 2014, Stockholm.

Regeringskansliet (2016) "Policyramverk för svenskt utvecklingssamarbete och humanitärt bistånd", Regeringens Skrivelse 2016/17:60, 14 December 2016, Stockholm.

Republic of South Sudan (2014) "The Aid Information Management System (AIMS)", Briefing Note 1.5, available at: http://grss-mof.org/wp-content/uploads/2014/03/BN-1.5-AIMS-04-02-12.pdf (accessed 2017-06-30).

Roodman, D. (2007). The anarchy of numbers: aid, development, and cross-country empirics. The World Bank Economic Review, 21(2), 255-277.

Rodrik, D. (2009) "The New Development Economics: We Shall Experiment, but How Shall We Learn?", In What Works in Development? Thinking Big and Thinking Small, ed. Jessica Cohen and William Easterly, pp. 24–47. Washington, D.C.: Brookings Institution Press.

Rosenbaum, P. R. and D. B. Rubin (1983) "The Central Role of the Propensity Score in Observational Studies for Causal Effects", Biometrika, Vol. 70(1), pp. 41-55.

Shutt, C. (2016) "Towards an alternative development management paradigm?", Rapport 2016:07, Expertgruppen för biståndsanalys (EBA), Stockholm.

Sida (2017) "Sidas årsredovisning 2016", Artikelnummer: Sida 62023sv, available at: http://www.sida.se/Svenska/Publikationer-och-bilder/publikationer/2017/februari/sidas-arsredovisning-2016/ (accessed 2017-06-01).

State Council (2014) "White Paper on China's Foreign Aid", Information Office of the State Council, The People's Republic of China, July 2014, Beijing, available at: http://english.gov.cn/archive/white_paper/2014/08/23/content_281474982986592.htm (accessed: 2017-05-22).

Strandow, D., Findley, M., Nielson, D. and J. Powell (2011) "The UCDP and AidData codebook on Geo-referencing Foreign Aid. Version 1.1", Uppsala Conflict Data Program, Paper no. 4, Uppsala University.

Strandow, D., Findley, M. G. and J. K. Young (2016) "Foreign Aid and the Intensity of Violent Armed Conflict", AidData Working Paper 24, May 2016, Williamsburg, VA: AidData.

Strange, A. M., Parks, B. C., Tierney, M. J., Fuchs, A., Dreher, A. and V. Ramachandran. 2013. "China's Development Finance to Africa: A Media-Based Approach to Data Collection" CGD Working Paper 323. Washington, DC: Center for Global Development.

Strange, A. M., Parks, B., Tierney, M. J., Fuchs, A. and A. Dreher, (2015), "Tracking under-reported financial flows: China's development finance and the aid-conflict nexus revisited", Journal of Conflict Resolution, pp. 1-29.

Svenska Afghanistankommittén (2016) "Effektrapport Svenska Afghanistankommittén avseende 2015", Effektrapport Svenska Afghanistankommittén fastställd 2016-06-23.

Swedish Committee for Afghanistan (2017) "Our work", available at: https://swedishcommittee.org/our-work (accessed 2017-06-22).

Tan-Mullins, M., Mohan, G. and M. Power (2010) "Redefining 'aid' in the China-Africa context", Development and Change, 41(5), pp. 857-881.

Qian, N. (2015). Making Progress on Foreign Aid. Annual Review of Economics, 7(1), 277-308.

UCDP Georeferenced Event Dataset (2015) "UCDP Georeferenced Event Dataset (GED) Global version 5.0 (2015)", available at http://ucdp.uu.se/downloads/ (accessed 2017-05-10).

US Geological Survey (2017) available at https://www.usgs.gov/ (accessed 2017-05-10).

Weaver, C. Davenport, S. Baker, J., Findley, M., Peratsakis, C. and J. Powell (2014) "Malawi's Open Aid Map", Washington DC: The World Bank.

Williamson, C.R. (2008) "Foreign aid and human development: The impact of foreign aid to the health sector", Southern Economic Journal, 75(1), pp. 188-207.

Wilson, S.E. (2011) "Chasing Success: Health Sector Aid and Mortality", World Development, 39 (11), pp. 2032-2043.

Wood, R. M. and C. Sullivan (2015) "Doing Harm by Doing Good? The Negative Externalities of Humanitarian Aid Provision during Civil Conflict", AidData Working Paper 11, July 2015, Williamsburg, VA: AidData.

World Bank (2017) "Living Standard Measurements Study (LSMS)", Available at http://econ.worldbank.org/WBSITE/EXTERNAL/EXTDEC/EXTRESEARCH/EXTLSMS/0,,menuPK:3359053~page

PK:64168427~piPK:64168435~theSitePK:3358997,00.html(accessed 2017-05-08).

World Bank (2017b) "Global reach", available at: http://maps.worldbank.org/p2e/mcmap/index.html?org=ibrd (accessed 2017-06-01).

World Bank (2016) "Evaluating Tanzania's productive social safety net: targeting performance, beneficiary profile, and other baseline findings". Washington, D.C.: World Bank Group.

World Bank (2012) "Open Aid Partnership: World Bank and Partners call for Open Collaborative Maps to improve development aid impact", Press Release No:2012/170/WBI, available at:

http://web.worldbank.org/WBSITE/EXTERNAL/NEWS/0,,contentMDK:23054525~pagePK:34370~piPK:34424~theSitePK:4607,00.html (accessed 2017-05-30)

World Bank (2011) "Aid Partnership: World Bank and Partners Call for Open Collaborative Maps to Improve Development Aid Impact". Press release, 25 November 2011, available at: http://www.worldbank.org/en/news/press-release/2011/11/25/open-aid-partnership-world-bank-partners-call-open-collaborative-maps-improve-development-aid-impact (accessed 2017-02-27).

WorldPop (2017) "About WorldPop" available at http://www.worldpop.org.uk/ (accessed 2017-05-10).

Öhler, H. and P. Nunnenkamp (2014) "Needs-Based Targeting or Favoritism? The Regional Allocation of Multilateral Aid within Recipient Countries", Kyklos, 67(3), pp. 420-446

Appendix

Table A1: Geocoding precision categories from the UCDP/AidData codebook (Strandow et al., 2011)

1. The coordinates corresponds to an exact location, such as a populated place or a hill. The code is also used for points that join a location which is a line (such as a road or railroad). Lines are not coded only the points and areas that connect lines.
2. The location is mentioned in the source as being "near", in the "area" of, or up to 25 km away from an exact location. The coordinates refer to that adjacent, exact, location.
3. The location is, or is analogous to, a second order administrative division (ADM2), such as a district, municipality, or commune.
4. The location is, or is analogous to, a first order administrative division (ADM1), such as a province, state, or governorate.
5. The location can only be related to estimated coordinates, such as when a location lies between populated places; along rivers, roads, and borders; more than 25 km away from a specific location; or when sources refer to parts of a country greater than ADM1 such as a National Park which spans across several provinces
6. The location can only be related to an independent political entity, meaning the pair of coordinates that represent a country.
7. Unclear. The country coordinates are entered to reflect that sub-country information is unavailable.
8. The location is estimated to be a seat of an administrative division (local capital) or the national capital. If aid goes to Luanda, for example, without further specification on the location, and there is an ADM1 and a capital called Luanda, then code the coordinates of the capital with precision 8. If it is not spelled out that aid goes to the capital; but if it is clear that it goes to a government ministry or to government financial institutions; and if those institutions are most likely located in the capital; then the coordinates of the capital are coded with precision 8. (However, if it can be verified that the recipient institution is located in the capital then the coordinates of the capital with precision 1 are used.)

Table A2: Data on province level World Bank health aid and child mortality obtained from geoquery

ID	Province name	World Bank health aid (USD disbursements since 2000)	Child mortality (deaths per 1000 under-5 children)
1	Baringo	838188	14,46
2	Bomet	0	14,99
3	Bungoma	1897	21,34
4	Busia	645583	33,39
5	Elgeyo-Marakwet	833604	13,89
6	Embu	7216339	9,33
7	Garissa	333089	24,42
8	Homa Bay	738468	50,27
9	Isiolo	293160	15,33
10	Kajiado	7783233	21,79
11	Kakamega	543144	37,41
12	Kericho	3804	19,31
13	Kiambu	15883	10,30
14	Kilifi	1449386	26,01
15	Kirinyaga	258576	10,97
16	Kisii	30428	25,38
17	Kisumu	1130920	39,31
18	Kitui	479170	16,93
19	Kwale	952620	19,28
20	Laikipia	38736	7,61
21	Lamu	371747	19,93
22	Machakos	283018	9,13
23	Makueni	412181	17,76
24	Mandera	535020	15,39
25	Marsabit	610656	16,25
26	Meru	46601	12,04
27	Migori	1754	47,83
28	Mombasa	10968861	12,83
29	Murang'a	7265	12,77
30	Nairobi	247903	16,93

31	Nakuru	8798	12,05
32	Nandi	404477	24,78
33	Narok	44483	9,91
34	Nyamira	12053	21,17
35	Nyandarua	0	11,07
36	Nyeri	589	6,03
37	Samburu	516367	11,89
38	Siaya	527772	51,53
39	Taita Taveta	416263	14,89
40	Tana River	650079	23,79
41	Tharaka-Nithi	228751	13,17
42	Trans Nzoia	37458	10,63
43	Turkana	14658	21,41
44	Uasin Gishu	64810	14,90
45	Vihiga	6604662	35,58
46	Wajir	434208	18,15
47	West Pokot	208644	14,50

Table A3: World Bank aid and citizen participation, using alternative geographic cut-offs

Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1

VARIABLES	(1) Member of organisation	(2) Member of organisation	(3) Attend community meetings	(4) Attend community meetings
ongoing25	0.069***		0.015	
	(0.014)		(0.012)	
future25	0.040**		0.033*	
	(0.018)		(0.017)	
ongoing75		0.075***		0.028**
		(0.014)		(0.013)
future75		0.010		0.046***
		(0.014)		(0.014)
Individual controls	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
Region FE	YES	YES	YES	YES
Diff-in-diff ongoing-future	0.0290	0.0649	-0.0175	-0.0177
F test: ongoing-future=0	2.392	19.36	1.227	1.677
p value	0.122	1.10e-05	0.268	0.195
Observations	68,667	81,102	68,667	81,102
R-squared	0.096	0.099	0.158	0.158

Previous EBA-reports

2017:08 *Research capacity in the new global development agenda*, Måns Fellesson

2017:07 *Research Aid Revisited – a historically grounded analysis of future prospects and policy options*, David Nilsson, Sverker Sörlin

2017:06 *Confronting the Contradiction – An exploration into the dual purpose of accountability and learning in aid evaluation*, Hilde Reinertsen, Kristian Bjørkdahl, Desmond McNeill

2017:05 *Local peacebuilding – challenges and opportunities*, Joakim Öjendal, Hanna Leonardsson, Martin Lundqvist

2017:04 *Enprocentmålet – en kritisk essä*, Lars Anell

2017:03 *Animal health in development – it's role for poverty reduction and human welfare*, Jonathan Rushton, Arvid Uggla, Ulf Magnusson

2017:02 *Do Anti-Discrimination Measures Reduce Poverty Among Marginalised Social Groups?*, Rachel Marcus, Anna Mdee, Ella Page

2017:01 *Making Waves: Implications of the irregular migration and refugee situation on Official Development Assistance spending and practices in Europe*, Anna Knoll, Andrew Sherriff

2016:11 *Revitalising the policy for global development*, Per Molander

2016:10 *Swedish Development Cooperation with Tanzania – Has It Helped the Poor?*, Mark McGillivray, David Carpenter, Oliver Morrissey, Julie Thaarup

2016:09 *Exploring Donorship – Internal Factors in Swedish Aid to Uganda*, Stein-Erik Kruse

2016:08 *Sustaining a development policy: results and responsibility for the Swedish policy for global development*, Måns Fellesson, Lisa Román

2016:07, *Towards an Alternative Development Management Paradigm?*, Cathy Shutt

2016:06 *Vem beslutar om svenska biståndsmedel?*, *En översikt*, Expertgruppen för biståndsanalys

2016:05 *Pathways to change: Evaluating development interventions with Qualitative Comparative Analysis (QCA)*, Barbara Befani

2016:04 *Swedish responsibility and the United Nations Sustainable Development Goals*, Magdalena Bexell, Kristina Jönsson

2016:03 *Capturing complexity and context: evaluating aid to education*, Joel Samoff, Jane Leer, Michelle Reddy

2016:02 *Education in developing countries what policies and programmes affect learning and time in school?*, Amy Damon, Paul Glewwe, Suzanne Wisniewski, Bixuan Sun

2016:01 *Support to regional cooperation and integration in Africa – what works and why?*, Fredrik Söderbaum, Therese Brolin

2015:09 *In search of double dividends from climate change interventions evidence from forest conservation and household energy transitions*, G. Köhlin, S.K. Pattanayak, E. Sills, E. Mattsson, M. Ostwald, A. Salas, D. Ternald

2015:08 *Business and human rights in development cooperation – has Sweden incorporated the UN guiding principles?*, Rasmus Klocker Larsen, Sandra Atler

2015:07 *Making development work: the quality of government approach*, Bo Rothstein and Marcus Tannenberg

2015:06 *Now open for business: joint development initiatives between the private and public sectors in development cooperation*, Sara Johansson de Silva, Ari Kokko and Hanna Norberg

2015:05 *Has Sweden injected realism into public financial management reforms in partner countries?*, Matt Andrews

2015:04 *Youth, entrepreneurship and development*, Kjetil Bjorvatn

2015:03 *Concentration difficulties? An analysis of Swedish aid proliferation*, Rune Jansen Hagen

2015:02 *Utvärdering av svenskt bistånd – en kartläggning*, Expertgruppen för biståndsanalys

2015:01 *Rethinking Civil Society and Support for Democracy*, Richard Youngs

2014:05 *Svenskt statligt internationellt bistånd i Sverige: en översikt*, Expertgruppen för biståndsanalys

2014:04 *The African Development Bank: ready to face the challenges of a changing Africa?*, Christopher Humphrey

2014:03 *International party assistance – what do we know about the effects?* , Lars Svåsand

2014:02 *Sweden´s development assistance for health – policy options to support the global health 2035 goals*, Gavin Yamey, Helen Saxenian, Robert Hecht, Jesper Sundewall and Dean Jamison

2014:01 *Randomized controlled trials: strengths, weaknesses and policy relevance*, Anders Olofsgård