Digital Threads and Regional Ties:
The Study of Global Services Trade and Regional Favoritism

INAUGURALDISSERTATION

zur

Erlangung des akademischen Grades

eines Doktors der Wirtschaftswissenschaften

der

Universität Mannheim

vorgelegt von
Carlo Moana Birkholz, M.Sc.
aus Berlin

Dekan:	Joachim Lutz
Referent:	Prof. Dr. Johannes Voget
Koreferent:	Prof. Dr. Andreas Fuchs
Tag der mündlichen Prüfung:	17. Februar 2025

Acknowledgements

To my dad, for always believing in me and giving me confidence. I miss you dearly. To my mum, for giving me love and unwavering optimism. To my brother, for teaching me about life and purpose.
To my friends: Richard, for inspiring me to be ambitious. Johannes, for your keen mind. Julius, for a shared sense of enjoyment. To Annie for sharing the road with me up to the final steps. To all of you for giving me a sense of belonging.
To my mentors: Zareh, for being my supporting wheels the first mile on the road into academia. David, for the countless hours discussing everything. I have learned the most from both of you.
To all my coauthors, thank you for working with me, teaching me, sharing your ideas, and trusting in mine.
To my supervisors: Johannes, thank you for being quick to respond when needed and instilling self-believe in my abilities as a researcher. Andi, thank you for all the feedback and comments; your guidance has always improved my work.

Carlo
December 2024, Mannheim, Germany

Contents

List of Tables

List of Figures

Preface

The global economy has undergone profound transformations over the past few decades, driven by advancements in technology, shifts in political power, and the growing interdependence of markets and regions. My dissertation, Digital Threads and Regional Ties: The Study of Global Services Trade and Regional Favoritism, uncovers factors that contribute to a persistence of regional economic disparities inspite the opportunities brought about by this transformation. It explores two themes central to this evolution - the rise of services trade, and the interplay between political power and regional resource allocation. By combining innovative data sets with causal inference methods and structural modeling, I delve into these themes to uncover the patterns and mechanisms driving regional and sectoral economic outcomes.

At its core, this dissertation is motivated by two interrelated questions. First, to what extent can the integration of global digital markets level the economic playing field for regions in developing countries? Second, how do the distribution of political power and the allocation of resources influence the geography of economic development? In addressing these questions, this work documents frictions that hinder global inclusive growth and brings evidence that points to the importance of investments in human capital and strong institutions to alleviate them. The importance of addressing regional disparities in income and opportunity is powerfully highlighted by the current political turmoil brought about by populist political agendas instrumentalizing these inequalities and the increasing migration flows caused by them. The thesis contains the following four chapters:

Chapter 1 - The Global Software Production Network[1] explores the potential of tradable services as a driver of economic growth in developing countries. It relates to a core debate in development economics on the premature deindustrialization of economies brought forward by Rodrik (2016), who, among other factors, identifies confinement to the domestic market due to a premature shift to a service economy as threatening growth. However, this argument leaves out the increasing share of the high-skilled services sector across the world and the tradable outputs produced in many of its industries such as accounting and management services or software development. The question then becomes whether developing countries can leverage lower wages compared to high-income countries, and the near-free information flow via the internet to generate exports in this sector.

Drawing on a data set of 2.55 million software projects and 2.6 million software developers across 5,400 locations, this chapter investigates this question empirically. Adopting the seminal economic geography model by Eaton and Kortum (2002) to the context of trade in software development services, it estimates the gravity equation derived from the model and identifies three key barriers limiting export opportunities in this sector:

First, there are significant productivity differences both within and across countries.

[1]Chapter 1 is based on Birkholz and Gomtsyan (2024).

While productivity levels are strongly correlated with income levels, exceptional cases exist - cities in developing countries like Bengaluru in India rank among the most productive software production locations. Yet, when comparing productivity differences between rich and poor countries, the gaps in software production are larger than in broad industry and services sectors. Moreover, locations with a higher GDP per capita exhibit a comparative advantage in the production of ideas over the production of software.

Second, despite the absence of physical trade costs, spatial frictions in software trade are nearly as large as those for trade in goods. Estimated distance elasticities range between 0.7 and 0.9, only marginally lower than conventional estimates of around 1 for trade in goods. Consistent with the notion that trade is not just hindered by transportation costs but also by information frictions that increase with distance (Allen, 2014), this result suggests that these information frictions play an even more prominent role in the trade of services, effectively offsetting the negligible transportation cost component.

Third, talent sorts in patterns that resemble a brain drain. On the individual level, more productive software developers are more likely to migrate to richer and more productive locations, even estimated among subsets of migrants from the same origin cities. At the aggregate level, locations with higher GDP per capita attract more in-migration, even after accounting for their initial stocks of human capital.

Taken together, these findings underscore the persistent barriers to equitable participation in global markets, even in sectors characterized by minimal physical trade costs. The analysis highlights the need for policy interventions that address productivity disparities, reduce information frictions, and mitigate the adverse effects of the observed talent concentration.

Chapter 2 - Favoritism and Firms: Micro Evidence and Macro Implications[2] examines the economic consequences of regional favoritism, a form of distributive politics where resources are geographically redistributed based on political connections. This chapter investigates how such favoritism influences firm outcomes and the broader economy, using detailed enterprise survey data from low- and middle-income countries. While the seminal paper by Hodler and Raschky (2014) presents first systematic empirical evidence for this phenomenon utilizing nighttime light intensity as a proxy for economic activity, the use of the firm level data allows to study the mechanisms and aggregate effects of regional favoritism.

Drawing on a difference-in-differences approach and exploiting leader transitions for identification, the chapter documents that firms located near political leaders' birthplaces experience significant increases in sales and employment during their tenure. These benefits are concentrated in the non-tradable sector and are driven by increased government demand rather than fundamental productivity improvements; neither the reported business environment of these firms improves, nor their performance in dimensions that drive

[2] Chapter 2 is based on Asatryan, Baskaran, Birkholz and Gomtsyan (2022).

firm productivity such as management practices, inputs, technology adoption or innovation improve. What is more, effects dissipate rapidly after leaders leave office. A calibrated model of resource misallocation in the spirit of Restuccia and Rogerson (2008) reveals that output losses from regional favoritism in the aggregate are however rather limited, partly because leaders' home regions are not systematically less productive.

This chapter contributes to the literature on regional favoritism by shifting the focus from aggregate proxies such as nighttime light intensity to firm-level outcomes, enabling a nuanced understanding of the productivity implications of favoritism. It also enriches the literature on misallocation by highlighting a new source of inefficiency: the spatial and sectoral distortions driven by favoritism. By doing so, it highlights the trade-offs inherent in distributive politics, where localized economic gains could come at the cost of broader misallocation of resources, with implications for policies aimed at reducing regional disparities.

Chapter 3 - Favoritism by the Governing Elite[3] extends the scope of analysis beyond primary leaders and examines regional favoritism by the governing elite as a whole. A primary contribution of this chapter is the construction of a novel global dataset georeferencing the birthplaces of federal government cabinet members in 141 countries from 1992 to 2016. This dataset, part of the Political Leaders' Affiliation Database (PLAD) project (https://www.plad.me/), represents a significant resource for future research on the regional biases of political figures and their economic implications.

Using this dataset, the chapter investigates whether ministers, like primary political leaders, engage in regional favoritism by directing resources to their home regions. The motivation for examining a broader set of political actors lies in two key considerations. First, regional favoritism is rarely the product of a single individual's choices; rather, it is an outcome of the collective dynamics of political power, shaped by competition and collaboration within and across political factions. Second, there is theoretical ambiguity regarding the scale at which politicians below the level of primary rulers engage in regional favoritism. On one hand, these politicians may face less public scrutiny than heads of state, potentially allowing them to shift resources more freely to their home regions. On the other hand, they may lack the requisite political power to direct resources at a comparable magnitude.

Employing nightlight emissions and population data derived from satellite imagery as proxies for regional economic activity, the findings indicate that ministers do indeed engage in regional favoritism, as evidenced by significant increases in nightlight intensity in their home regions following their appointments. Identification of these effects leverages the staggered timing of ministerial appointments and departures. This staggered structure with switches in treatment status also positions the study as a prime empirical application of the methodological advances brought about by the emerging literature on staggered treatment timing and treatment effect heterogeneity (Roth et al., 2023). Ana-

[3]Chapter 3 is based on Asatryan, Baskaran, Birkholz and Hufschmidt (2023).

lyzing heterogeneity in the effect size along several dimensions reveals that the magnitude of favoritism varies across ministerial posts, with the most powerful positions driving the effect. Importantly, the chapter also finds that the extent of favoritism is moderated by the strength of democratic institutions, which seem to constrain the ability of ministers to channel resources disproportionately to their regions.

By extending the study of regional favoritism to a broader set of political actors, this chapter advances our understanding of the collective role of the governing elite in shaping the spatial distribution of resources. It also deepens our understanding of the interplay between institutional constraints and distributive politics, offering policy-relevant insights into the mechanisms that drive regional economic disparities.

Chapter 4 - The Regional Economics of Mineral Resource Wealth in Africa[4] shifts focus to mineral resources which, as a significant source of government revenues, provide a critical lens through which to study redistribution mechanisms and regional favoritism. This chapter explores how the wealth generated by mineral resource extraction influences economic outcomes across different regions in African countries, considering the role of deliberate redistribution policies and unintended macroeconomic effects.

Using geocoded data on mine openings and closures and employing a difference-in-differences framework, this chapter first documents the economic booms that occur in mining regions. The effects are highly localized, with significant and persistent increases in economic activity detectable within a 30-kilometer radius of operational mines. However, the chapter's central focus lies in understanding how these localized gains ripple — or fail to ripple — throughout the rest of the country.

Non-mining regions are impacted unevenly. Politically important regions, such as capital cities and the birthplaces of national leaders, see significant economic benefits at the expense of generic non-mining regions. For instance, capital cities experience notable increases in nightlight intensity, suggesting that part of the mining revenues are preferentially channeled to the capitals. Similarly, under autocratic or corrupt regimes, leaders' birth regions disproportionately benefit, reflecting the influence of regional favoritism in shaping resource allocation. In contrast, generic non-mining regions experience reductions in luminosity.

The chapter identifies three primary mechanisms driving these patterns. First, deliberate government actions play a central role. Politically significant regions are favored, with autocratic institutions amplifying these biases. Second, macroeconomic adjustments consistent with Dutch Disease contribute to the decline of non-mining regions. Exchange rate appreciation and resource-driven shifts in national economic structures disadvantage particularly regions specialized in manufacturing. Third, non-mining regions see increased conflict incidence. This could result either from the direct use of additional resource revenues to fund conflicts or as an indirect byproduct of worsening economic conditions.

[4] Chapter 4 is based on Asatryan, Baskaran, Birkholz and Hufschmidt (2024).

Through its focus on the spatial distribution of mineral wealth, Chapter 4 reinforces the broader theme of this dissertation: the role of political power in shaping the allocation of resources and the resulting economic geography of development. It contributes to the literature on regional favoritism by illustrating how resource revenues are channeled to politically strategic regions, often to the detriment of less politically influential areas. Additionally, it extends the understanding of the resource curse by showing how the combination of political and economic dynamics affects subnational regions differently, with implications for aggregate national welfare. The findings emphasize the need for institutional reforms and policy interventions that both limit the extent of regional favoritism in the allocation of resource rents and mitigate the negative spillovers of the resource extraction.

In conclusion, this dissertation bridges the fields of economic geography, development economics, and political economy to examine how global services trade and regional favoritism shape economic disparities across regions and sectors. By integrating innovative data sources with quasi-experimental methods and structural modeling, it uncovers some of the underlying mechanisms and frictions that perpetuate the unequal distribution of resources and human capital. The thesis highlights that such disparities may not only be detrimental to the disadvantaged regions but may result in inefficient outcomes at the aggregate level. This underscores the welfare-enhancing scope of policies that can address these unevenly distributed factors, for instance through the creation and retention of human capital and the reduction in non-physical trade costs, or the strengthening of institutional checks on politicians' discretionary redistribution of resources to favored regions. These results of my thesis thus add important puzzle pieces to some of the most influential literature in economics on the roles of institutions (Acemoglu et al., 2001; North, 1990) and human capital (Becker, 1962; Benhabib and Spiegel, 1994) for economic development.

References

Acemoglu, Daron, Simon Johnson, and James A Robinson, "The colonial origins of comparative development: An empirical investigation," *American economic review*, 2001, *91* (5), 1369–1401.

Allen, Treb, "Information frictions in trade," *Econometrica*, 2014, *82* (6), 2041–2083.

Asatryan, Zareh, Thushyanthan Baskaran, Carlo Birkholz, and David Gomtsyan, "Favoritism and Firms: Micro Evidence and Macro Implications," 2022.

_ , _ , _ , and Patrick Hufschmidt, "Favoritism by the governing elite," *Ruhr Economic Papers (No. 1029)*, 2023.

_ , _ , _ , and _ , "The regional economics of mineral resource wealth in Africa," *Economica*, 2024.

Becker, Gary S, "Investment in human capital: A theoretical analysis," *Journal of political economy*, 1962, *70* (5, Part 2), 9–49.

Benhabib, Jess and Mark M Spiegel, "The role of human capital in economic development evidence from aggregate cross-country data," *Journal of Monetary economics*, 1994, *34* (2), 143–173.

Birkholz, Carlo and David Gomtsyan, "The Global Software Production Network," 2024.

Eaton, Jonathan and Samuel Kortum, "Technology, geography, and trade," *Econometrica*, 2002, *70* (5), 1741–1779.

Hodler, Roland and Paul A Raschky, "Regional favoritism," *The Quarterly Journal of Economics*, 2014, *129* (2), 995–1033.

North, Douglass C, *Institutions, institutional change and economic performance*, Cambridge university press, 1990.

Restuccia, Diego and Richard Rogerson, "Policy distortions and aggregate productivity with heterogeneous establishments," *Review of Economic dynamics*, 2008, *11* (4), 707–720.

Rodrik, Dani, "Premature deindustrialization," *Journal of economic growth*, 2016, *21*, 1–33.

Roth, Jonathan, Pedro HC Sant'Anna, Alyssa Bilinski, and John Poe, "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

The Global Software Production Network

Carlo Birkholz [*] David Gomtsyan [†]

Abstract

Can developing countries benefit from exporting opportunities in the growing sector of tradable services, given the near free information flow via the internet and wage differentials relative to developed countries? Focusing on the software development industry, we analyse data from 2.55 million software projects across 5,400 locations, and estimate an economic geography model in which locations trade tasks. The results reveal three factors limiting exports: (i) significant productivity differences within and between countries; (ii) a notable decline in trade volumes with distance; (iii) sorting patterns among software developers that are suggestive of brain drain.

Keywords: Productivity, IT, services trade, migration, sorting.
JEL: F1, L86, O15.

[*]University of Mannheim, ZEW Mannheim
[†]CREI

1.1 Introduction

Over their development path advanced economies have experienced a substantial increase in the share of the high-skilled services sector. In the US, the share of high skilled services exceeds 50% of total value added (Buera and Kaboski, 2012). Notably, many segments within this sector produce tradable output. Given that technological advances of recent decades reduced the cost of digital information flows to near zero, new exporting opportunities may arise for developing countries, where wages are lower than in developed countries. Are developing countries in a position to take advantage of these opportunities? We address this question by employing novel data that allow us to study the global software development industry, one of the fastest evolving parts of the high-skilled services sector.

Our main analysis is based on GitHub data from 2.55 million projects and 2.64 million users, and their interactions. The available data allow us to observe the locations of users at the city level, their contributions to specific projects, as well as their follower networks. We employ this information to construct flows of software code between locations from project level collaborations. Based on these flows, we propose a spatial model in the spirit of Eaton and Kortum (2002), in which software developers in different locations trade in tasks. By estimating the gravity equation derived from the model, we recover distance elasticities and productivity parameters at the city level.

According to our estimations the San Francisco Bay Area emerges as the unambiguous leader, followed by other cities located on the West Coast of the US. Among developing countries, the most productive locations are Bengaluru in India and various cities in Eastern Europe. Overall we find that there is a tight relationship between our measure and GDP per capita at the country level, and between per capita nighttime luminosity at the city level. We also find that estimated productivity differences in the software industry between the richest and poorest countries are comparable or even larger than those derived from macro data encompassing broad sectors. This means that the poorest countries are performing worse in the production of software code than in the production of goods and other services. Moreover, we construct a separate productivity measure for the generation of final software products, which presents a higher value activity than provision of coding services. We find that the comparative advantage in the generation of final software products relative to coding services increases with GDP per capita.

Despite the fact that, from a technological perspective, there are no spatial frictions to the trade in software code, our gravity equation estimates imply that distance has a negative effect on trade volumes. Specifically, our estimated distance elasticity is in the range of 0.7-0.9, which is comparable in size to the value of 1 obtained for the flow of goods within the US (Allen and Arkolakis, 2018). Our interpretation of this sizable effect is that distance affects the movement of people, and the networks in which they collaborate. The production network is shaped by collaborations formed through in-person interaction, such that online software production cannot be understood as a process that operates independently from offline location.

We then investigate the migration patterns of IT specialists within and across countries. In our data we observe the location of these software developers at different points in time. We construct a proxy for the quality of their skill set based on the centrality of the software developer in the follower network of all GitHub users, which we derive through the recursive ranking algorithm PageRank. We document that there are strong sorting patterns of migration both within and across countries based on this quality proxy. For example, we observe that IT specialists who are ranked higher in a city at time t are more likely to migrate to a more productive city (or a country with higher GDP per capita) in period $t + 1$. We further show that immigrants tend to have higher quality than the median resident in the destination. These results hold both when migrants move to places that are rated higher in terms of IT productivity than their origin location, and when they move to countries with a higher GDP per capita than their country of origin.

Taken together, our results suggest that – barring effective policy interventions – developing countries are unlikely to reap large benefits from software code exports for three reasons: First, the ability to export requires high productivity. However, our estimates show that the productivity gap in the software development sector between rich and poor countries is of a magnitude comparable or even larger to the gap in the service sector or manufacturing. Second, our estimates show that there are substantial spatial frictions which hamper trade flows. Third, the migration patterns we document indicate that developing countries experience a brain drain, which may make it harder to catch up with the technological frontier.[1]

We validate our data in several steps. First, we use two alternative approaches to measure the role of each location in the software production process. As one alternative, we construct a graph of locations in the world which are linked to each other by their observed software code flows. We again apply PageRank to recursively determine the centrality of each node (location) in the graph. As another alternative, we aggregate the individual scores we obtained from applying PageRank to the follower network at the level of locations. The results obtained according to both of these alternative approaches are closely correlated with the productivity measures obtained from the structural estimation. Second, we validate our measure for the US sub-sample by regressing it on wages of US IT specialists obtained from the American Community Survey at the location level, and for the full sample by regressing it on wages of IT specialists globally from the Stack Overflow Developer Survey at the country level. We find an economically large and statistically strong relationship. Third, we construct university rankings for the US, the UK and Germany based on individual software developers' quality scores and their reported affiliation. The list shows close resemblance with conventional rankings, such as by US News or the

[1] If migrants also facilitate the diffusion of knowledge to their home countries, then the negative effects of brain drain would be less severe. We are silent on this channel.

Academic Ranking of World Universities.

For the analysis of the questions we pose, GitHub data have important advantages over the patent data that have been widely used in the literature. First, they cover a wide range of countries with varying levels of GDP per capita, and capture an extensive membership and activity network in many developing countries, whereas the literature based on patents has focused on a small set of high income countries. Second, we observe activities at high frequency levels, while patenting is a relatively rare activity, especially at the individual level, and many inventors register only one patent during their lifetime. This makes the analysis of inventor migration complicated because economists observe inventors' locations only when they register a patent, so they need to observe the same inventor registering patents in different locations to document an event of migration.[2] Third, in the GitHub data joint participation in projects by members located in different locations is more common, which enables us to study interactions across space. Finally, software production is relatively less dependent on the investment of physical capital than other high skilled sectors, and members of teams are less confined by physical distance; they do not need to be located in laboratories with special equipment. Thus, our setting allows us to focus on the human capital and human interaction aspect of the innovation process.

There is a large literature that tries to measure productivity levels across countries (see, for instance, Klenow and Rodríguez-Clare, 1997; Hall and Jones, 1999). Methodologically we follow Waugh (2010) and use a trade model to recover productivity parameters. In contrast to the aforementioned papers we focus on one industry, but our productivity measures are at the city level rather than at the country level. Within this literature, it is worthwhile emphasizing papers that specifically focus on the level of human capital. Since software production is human capital intensive and individuals can provide their services to firms in distant locations, we believe that the human capital component in our productivity measure is large. However, it cannot be interpreted as being a measure of human capital exclusively, because other factors, such as agglomeration forces acting at the city level, are also included in our estimated productivities. Given the difficulties related to the measurement of schooling quality, researchers have used wages of migrants in destination countries to measure human capital (Clemens, 2013; Hendricks and Schoellman, 2017; Martellini et al., 2024). In this literature, researchers rely on wages to obtain measures of worker quality. However, when transitioning from one location to another, workers may face imperfect transferability of skills, discrimination, or lack of local networks. All of these factors can lead to lower estimates of migrants' true skills. Because our measure is not based on wages, it is less likely to be affected by those factors, yet still not fully void of them, or agglomeration effects, as mentioned above.

We also contribute to the literature on trade in services. The decline in communication

[2] For example, in the dataset used by Akcigit et al. (2016) 52% of inventors have only one registered patent. For this reason the authors base their analysis only on top inventors who register patents frequently.

costs has led to an increase in services trade Eckert (2019). However, a lack of data makes it difficult for researchers to measure the extent of such trade flows. Eaton and Kortum (2018), using 2010 international bilateral trade data, find a distance elasticity of 1.4 in professional services and administrative services. Other studies combine structural models with industry employment data from the US to generate trade in services without observing the actual flows (Gervais and Jensen, 2019; Eckert, 2019). Hsieh and Rossi-Hansberg (2023) study trade in non-tradeable services through the expansion of affiliates.

We also relate to other papers and emerging work using Github as a data source. Wachs et al. (2022) utilize the geolocation of software developers on Github to document the spatial distribution of software developers between and within countries. Wright et al. (2023) show that greater participation in open source development on Github at the country level leads to an increase in the number of new technology ventures in subsequent years. Wachs (2023) investigates brain drain as a consequence of conflict by following the migration of software developers on Github after the onset of the Russian invasion of Ukraine. Like us, Fackler et al. (2023) - who study collaboration in remote teams around the COVID-19 pandemic - also use Github data and estimate gravity equations at city-pair level. Their estimated distance elasticity coefficients are below 0.5, which are smaller than ours. There are, however, some key differences in our estimations related to sample selection, data construction, and the estimation specification that explain the differences in the estimated elasticity.[3]

We structure the remainder of our paper in the following way: We describe the features of GitHub data and complementary data sources in Section 1.2. In Section 1.3 we lay out our spatial equilibrium model and alternative approaches to calculate city-level productivities. We then present the results of our estimations and relate them with GDP per capita in Section 1.4. In Section 1.5 we study the migration patterns of software developers. Section 1.6 concludes.

1.2 Measuring trade in services with GitHub data

Our primary data source is a snapshot of the universe of GitHub users and their public activity on the platform in March 2021. This data is the latest available version of the GHTorent project that periodically mirrors Github's public event timeline through Github's API (Gousios, 2013). We supplement this with a snapshot of the data from June 2019 from the same source to identify changes in the reported location of users to study migration patterns.

GitHub is a service for software development and version control. It is the dominant

[3]They choose to drop locations below an arbitrary size threshold yielding only around 700 locations and appear to be using the entire sample of Github users with any location reported. As we discuss in section 1.2, we apply a number of careful data cleaning steps to the reported locations of users, to avoid introducing bias from systematic errors in the geocoding of the locations. Finally, we estimate the gravity equation with importer and exporter fixed effects at the location level, rather than their choice of country level, which is more appropriate to address multilateral resistance (Fally, 2015).

service for hosting open source software.[4] One of the main advantages of GitHub compared to other version control solutions is that it accommodates large teams of developers working independently. As a result, most widely used open source software programs have repositories on GitHub. It is also worth noting that, despite being open source, most popular programs with many users are owned by large organizations and generate revenues.[5] Some widely known names are Linux, MySQL, and Firefox. Owners of these products rely on various business models to generate revenues; the most common revenue generation model is to sell enterprise versions or additional bundles that complement the free version. Since these are sophisticated and advanced products, owners frequently hire professional software engineers for further development and updating.

Users There are a total of 45.8 million registered users in the 2021 data snapshot; these users can be uniquely identified based on their ID and user names. Registered users are mostly individuals, but can in some instances also be organizations, which are identified through a user type variable. The range of engagement and activity on the platform varies widely, as well as the completeness of the user profiles. We observe around 3.7 million users with some degree of information about their physical location. Locations are self-reported in a free text field; this information is automatically translated into a geolocation (longitude and latitude). We undertake rigorous cleaning efforts to ensure that the user input is reasonable, and that the automated geocoding is accurate. As a first step in this cleaning effort, we drop users reporting locations such as '*the internet*', '*the world*', '*anywhere*', '*remote*', '*future*', '*darknet*', '*404*', '*Earth*', '*Moon*', '*universe*', '*galaxy*', '*Milky Way*', '*Pluto*', '*Mars*', or '*space*'.[6] In a second step, we drop all users with location information that is not granular enough to map them onto cities accurately. This is crucial, as users reporting information on the country level, for example, receive the geocoordinates of the country's capital. As a third step, we manually review common user entries that represent over 1% of the observations at each location, excluding the smallest 1% of locations. This process allows us to eliminate any remaining significant errors in user allocation. We are left with a sample of 2.64 million users with cleaned locations, which is the subset of data we employ whenever our analyses rely on location information. Figure A4 in the appendix plots all unique user locations across the world. In terms of the selection of users indicating their location, we are confident that our sample reflects the active, professional users of the platform, as professional use of the platform incentivises a fully completed profile to facilitate communication and work opportunities. We provide an extended discussion of the representativeness of our sample in the Appendix section A.1.3.

[4]"What is GitHub?" The Economist, Jun 18, 2018.

[5]For example, see this commercial open-source software company index listing businesses with estimated revenues exceeding 100 million US dollars: https://notes.andymatuschak.org/Commercial_open-source_software_company_index.

[6]We manually inspect location names containing these strings to not loose valid addresses such as *Moon Vista Avenue, Las Vegas.*

For the time period up to 2019 we observe an additional aspect of the social network within GitHub, namely the followers and following of each user. The following functionality is an important feature for collaboration on the platform, as it enables users to get directed updates on other users' activities, such as changes made within shared projects or new projects started.[7] Around 3.8 million users follow at least one other user, and those who follow at least one person follow an average of 7.8 users.

Projects We observe over 189 million projects in the database, which are uniquely identified by project IDs. GitHub projects are organized into so-called repositories, which contain all of the contents of a specific project; in the following, we will use the terms "project" and "repository" interchangeably. We link users to projects via the unique project IDs. Every project has one owner, who typically holds a central role within the project, as we demonstrate in Appendix A.2, and users who – conditional on taking part in any project – belong on average to 4.5 projects. Whenever we study collaboration within projects based on geographic location, we define the projects' origin as the owner locations. Given that we do not observe locations for all users, as discussed above, these analyses rely on a subsample of 47.3 million projects for which owner location information is available. When constructing flows of code between locations in a project, we additionally require information on the locations of the contributing users. For 2.55 million projects we observe the location of the owner and the location of at least one project contributor.

Commits Commits are the primary user action to advance a project. They refer to a version of changes made to a repository's files. Changes to a project that are initially made locally are grouped and pushed to update the online version of the project. Commits typically come with a short message describing changes made, so that one can keep track of file versions. For each commit we identify the author, the committer and the project owner. The author and committer can be different users, for instance when users who are not project members suggest changes; a process explained further in the section **Forks and pull requests** below.[8]

In our analysis we construct flows of software production based on authors and owners. We clean the commits data in two main ways before constructing these flows: First, we do not consider commits where the author and owner are the same user – a construct we term self-links. Second, we alleviate potential biases stemming from bot activity by dropping users that are tagged as 'fake' by GitHub and by dropping commits that resemble the

[7]For instance, in a forum post discussing the following functionality on Github, users write "[...] when I find someone contributing to a library I use or a project that does what I need, I want to know about it immediately. [...] I follow the core developers of some of the main business critical libraries that we use (and sometimes their upstream dependency projects) so I can get a heads up on any potential breaking changes coming down the line" and "The same reasons you follow anyone on social media- to see what they're doing".

[8]Another instance can occur when multiple project members collaboratively work on the same project branch (part of the project) and only one of them commits the others changes.

automated nature of bot activity. For the latter we construct the within-project variance of the commit frequency of users with at least 25 commits, and drop them if they display a variance of zero, which means they commit in exactly steady intervals.[9]

We then define X_{ij} as the volume of code that flows from location j to location i determined by the following expression:

$$X_{ij} = \sum_{k \in K} commits_{jk} \times 1[owner_{ik} = 1], \tag{1.1}$$

where K is the set of projects, $commits_{jk}$ is the number of commits on project k by users from location j and 1[.] is the indicator function equal to 1 if the owner of project k is in location i. Intuitively this means that the volume of code flowing from location j to location i is the sum of commits from location j in projects whose owner is located in i. In Appendix A.2 we provide motivation for this approach and discuss alternatives.

Forks and pull requests Users may copy projects, in GitHub terminology "fork", and create modifications or build a different version of the parent project. There are two main rationals for doing so: First, a user may fork a project, modify it and then propose to merge the changes with the main project – an action that is referred to as creating a pull request. If accepted the changes are committed to the original project, which we record accordingly in our data as a flow of commits from the user proposing the alterations to the owner of the parent project. Second, a user may create a new independent software, which uses the original software as an input. In this case the fork represents an import of final software product.[10] While our paper focuses on trade in services that is represented by the gradual contribution of commits to the development and improvement of a software product, the trade in final software products or ideas captured by this second category of forks is an additional interesting aspect of the global software production network. For the remainder of the paper we will use the terms trade in services and trade in ideas/final software for these two dimensions of trade activity on GitHub. We empirically investigate trade in ideas in Section 1.4.4 noting however the caveat that the volume of transactions is much lower than for trade in services implying a noisier measure.

Other data In addition to the GitHub data, we use geographic information on functional urban areas (FUAs) and administrative regions, population and nighttime luminosity data from satellite images, and income data at sub-national and country level. We describe the construction of all auxiliary data we employ in the Appendix section A.1.

[9] Bots are software that run reoccurring tasks in an automated fashion.

[10] By final software product we mean a software product which can be used either by consumers or by other software developers as an input for the production of other software.

1.3 Methodology

We propose several approaches to determine the productivity of each city in the global software production process. Our main approach is based on the standard Eaton and Kortum (2002) model in which individuals in different locations produce and sell software code. This model allows us to derive a structural gravity equation and recover productivities of locations. Then, we propose two alternative reduced-form approaches for ranking cities. While each approach has its unique up- and downsides, we find that they produce consistent results.

1.3.1 A model of trade in tasks

The model is based on the standard Eaton and Kortum (2002) framework. Several papers have used this framework to impute country-specific productivity parameters (Waugh, 2010; Levchenko and Zhang, 2016). We follow the approach used in these papers to impute the level of software development productivity in specific locations. In our setting, trade takes place in software development services or tasks. We focus only on this sector and do not describe the rest of the economy. To the extent that we are interested in estimating distance elasticities and productivities for software development, the weight of software in household preferences or its contribution as an input to other sectors does not matter (see Levchenko and Zhang, 2016). The only assumption we need is that labor is the only input required to produce software code. This would not appear to be a very strong assumption because in the software development process the share of labor is likely to be higher than in most other industries. Moreover, software development tools (programs and cloud services), which are probably the next most important input, are either available as open source or highly tradeable without much variation in prices across space.

The analytical formulation of the problem is similar to the above mentioned papers. However, given the nature of our data and the environment of open source software production, we provide somewhat different interpretations. In particular, in a conventional trade model, the unit of production is a firm located in location i that produces a differentiated good q with efficiency $z_i(q)$ by hiring labor (inputs). In our case the unit of production is an individual rather than a firm, and this individual uses his or her own labor. We assume that software developers are endowed with a fixed amount of time which they allocate to solving open source problems. In our context, the differentiated good is a specific segment of the overall code. The solutions submitted by developers require proofing and potentially additional improvements or tuning from the owner of the code, which takes owners' time. The amount of time required to improve the proposed solution is inversely proportional to the productivity $z_i(q)$ of the developer who submitted it. Additionally there is an iceberg trade cost d_{ij}. An interpretation of this cost is that the developer with productivity $z_i(q)$ may lack familiarity with a given project, as a result of which the quality of his proposed solution decreases by a factor of d_{ij}. Familiarity with a project can be built through interactions, the likelihood and intensity of which decrease with physical distance. Thus, the code owner will adopt the best solution, or equivalently the solution proposed by the developer with the highest $z_i(q)$ adjusted by the iceberg cost $(min_{i=1,...,N}\{\frac{d_{ij}}{z_i(q)}\})$.

Individual productivities are drawn from the Fréchet distribution with the cumulative distribution function $F_i(z) = e^{-T_i z^{-\theta}}$. We allow the parameter T – which governs the average of the productivity draws – to be location-specific; this is our main object of interest. We interpret it as the average level of software development productivity or skills in location i. Higher values of T_i imply higher levels of average productivity. θ captures the dispersion of productivity draws.

The final software is produced using a CES production function that aggregates a continuum of task varieties $q \in [0, 1]$ according to the following formulation

$$Q_i = \int_0^1 \left[Q_i(q)^{(\epsilon-1)/\epsilon} dq \right]^{\epsilon/(\epsilon-1)},$$

where ϵ denotes the elasticity of substitution across varieties q and $Q_i(q)$ is the amount of variety q that is used in production. Following the steps in the aforementioned literature, the fraction of software development services provided (open source problems solved) by location j in the share of total software services consumed in location i is given by the following gravity equation

$$\frac{X_{ij}}{\sum_j X_{ij}} = \frac{T_j(d_{ij})^{-\theta}}{\Phi_i},$$

where $\Phi_i = \sum_j T_j(d_{ij})^{-\theta}$ is the multilateral resistance term. Dividing X_{ij} by the analogous expression for X_{ii} and taking logs, we obtain the conventional gravity equation

$$ln\left(\frac{X_{ij}}{X_{ii}}\right) = ln\,(T_j) - ln\,(T_i) - \theta ln(d_{ij}), \tag{1.2}$$

where X_{ij} denotes the volume of the flow of goods from location j to location i, the construction of which was described in equation (1.1). Next we express the log distance cost from equation (1.2) as

$$ln(d_{ij}) = d_k + a_{ij} + b_{ij} + Lang_{ij} + im_i + \nu_{ij},$$

where d_k is the contribution to trade costs of the distance between i and j measured in miles. Other variables are an indicator if cities are in the same country (a_{ij}), an indicator if countries share a border (b_{ij}), an indicator for a common language $Lang_{ij}$ and an importer fixed effect im_i. Substituting the expression for trade costs back to the equation (1.2) we obtain

$$ln\left(\frac{X_{ij}}{X_{ii}}\right) = \underbrace{ln\,(T_j)}_{\text{Exporter FE}} \underbrace{-ln\,(T_i) - \theta im_i}_{\text{Importer FE}} \underbrace{-\theta d_k - \theta a_{ij} - \theta b_{ij} - \theta Lang_{ij}}_{\text{Bilateral observables}} -\theta \nu_{ij} \tag{1.3}$$

In equation (1.3) the first term captures exporter fixed effects, which is the main object of interest. We estimate equation (1.3) using PPML. As a result of the estimation we obtain exporter fixed effects for each location, which have the following relationship with the productivity parameter

$$exp(EFE_j) = T_j, \tag{1.4}$$

where EFE_j are the exporter fixed effects from equation (1.3). An important detail is the inclusion of the term im_i in equation (1.3). An alternative approach to that is to include a term for exporters ex_j and use importer fixed effects to recover productivities from equation (1.4). There are four reasons that motivate our choice of the former over the latter approach. First, Waugh (2010) shows that including a term ex_j in equation (1.3) implicitly assumes that unit costs of production are the same across locations. In his context, this assumption is reasonable because higher productivity locations tend to have higher wages, so both forces push in opposite directions and counterbalance each other. Given the nature of open source contributions we have assumed that developers dedicate a fixed amount of time to solving open source problems without receiving a monetary compensation, thus the counterbalancing effect that operates through wages is not present. Hence, our preferred approach is to estimate equation (1.3) with the term im_i, which implies that unit costs are lower in more productive locations because they are more efficient. Second, the specification ex_j implies that locations face different exporting costs, in addition to the gravity terms for which we control. In the case of trade in goods, this friction can be justified by the quality of infrastructure such as ports, which is typically lower in developing countries. In the case of software code, the role of these factors is arguably less important. Third, the ex_j approach requires information on the wages of software developers in cities around the world, for which precise data is not readily available. Fourth, by estimating equation (1.3) we recover a much larger number of fixed effects than with importer fixed effects. This is driven by the fact that there are more contributors (exporters) in the data than project owners (importers), which enables us to generate more variation for the identification of exporter fixed effects. Given these arguments, we prefer the use of exporter fixed effects; however, we demonstrate in Appendix section A.3 that our productivity estimates are highly correlated with a specification using importer fixed effects and imputed city-specific wages.

1.3.2 Reduced form approach

Approach 1: Page rank algorithm. We think of locations as nodes of a graph and of X_{ij}'s as the strength of the links between nodes of the graph. The position of a node in a graph depends not only on its bilateral links but also on the links of the nodes to which it is connected, and so forth. In other words, the centrality of each node is determined recursively. A widely used approach for determining the centrality of nodes is the Page Rank algorithm (Brin and Page, 1998). The scores of locations are obtained as a solution to the following equation:

$$\begin{bmatrix} Score_1 \\ Score_2 \\ \vdots \\ Score_N \end{bmatrix} = \begin{bmatrix} (1-d)/N \\ (1-d)/N \\ \vdots \\ (1-d)/N \end{bmatrix} + d \begin{bmatrix} l_{11} & l_{12} & \dots & l_{1N} \\ l_{21} & \ddots & & \dots \\ \dots & & l_{ij} & \\ l_{N1} & \dots & \dots & l_{NN} \end{bmatrix} \begin{bmatrix} Score_1 \\ Score_2 \\ \vdots \\ Score_N \end{bmatrix} \tag{1.5}$$

where d is a parameter and l_{ij} is obtained by normalizing X_{ij} ($l_{ij} = \frac{X_{ij}}{\sum_j X_{ij}}$). The normalization ensures that $\sum_{i \in N} l_{ij} = 1$. If city i has no contributor involved in any project with other cities, then $l_{ij} = 0 \ \forall j$. Links to the node itself are not counted $l_{ij} = 0$ if $i = j$. Note that the resulting matrix, which is referred to as the adjacency matrix, is not necessarily symmetric. Equation (1.5) is solved by making an initial guess ($Score_i = 1/N$) and then making iterative computations until it converges. Typically, convergence is obtained rather quickly, which also turns out to be the case in our application.

Approach 2: Follower-based ranking As we described when introducing our data, on GitHub users may follow other users. The notifications received about followed users' public activities on GitHub enable and ease interaction. At the same time, people who make important contributions, generate new ideas or manage large projects are more likely to attract followers. We employ follower information to construct a graph in which each user is a node and directional edges between nodes are based on the following and follower links of users. We then apply the same recursive ranking algorithm described above to calculate the centrality score of each user. We interpret this measure as a proxy for individual quality. Conceptually being more central in the network of followers is likely highly correlated with individuals' quality, as the more and better work you do in projects, the more likely it is for others to follow you and receive updates on your work. In order to measure productivities at the location level we aggregate individual scores. Additionally, we use individual level scores to study the pattern of positive selection into migration in Section 1.5.

1.4 Results

In this section we start by discussing our estimates of the distance elasticity for the gravity equation and present our estimates of productivity at the city level. We then relate our estimates to nightlights per capita at the city level and GDP per capita at the country level, and compare our estimated productivity gaps between rich and poor countries with macro data. We finish the section by comparing trade in tasks to trade in ideas.

1.4.1 Structural estimation results

In Table 1.1 we present the results of the gravity equation using PPML. The estimated distance elasticity is around 0.8, which is close to the absolute value of the estimates for trade in goods (Allen and Arkolakis (2018) obtain a value of 1 for the US). This large estimate implies that geography continues to play an important role in trade in tasks, even though the flow of services between locations would seem to be frictionless. Our preferred explanation for this observation is that trade flows are determined in part by offline interactions involving in-person meetings, discussing ideas and making decisions on collaborations. Online software production does not occur in a vacuum, but is shaped by offline interactions. Thus, even though new technologies and platforms such as GitHub facilitate communication, they cannot fully replace in-person interactions, but rather serve as a complement to them.

This mechanism is consistent with the idea that trade is hindered not only by transportation costs but also by information frictions which increase with distance (Allen, 2014). These information frictions are understood to be potentially large in online goods markets, particularly with a large number of market participants (Bai et al., 2022), and for trade in goods face-to-face meetings are an effective way to alleviate them (Startz, 2016). For trade in services, these frictions are likely exacerbated since product and quality details are often more difficult to define and verify, such that the information friction component in the trade costs may exceed that in goods trade.[11]

In the following columns of Table 1.1 we report the results for several additional estimations to ensure the robustness of the results. In the second column, we restrict the sample to FUAs and construct the bilateral flows by ignoring users located outside FUAs. The estimated coefficient is not affected. In the third column, we restrict the sample to US FUAs only. The estimated distance elasticity increases slightly, suggesting that there are no large differences between global and US domestic patterns. In column (4) we add a dummy variable for the same location. We expect the absolute value of the distance elasticity estimate to drop, because such pairs have 0 distance and interact with each other

[11] While from an end user perspective it is perhaps easy to verify whether a software does what it should, from a software development perspective dimensions such as the efficiency and compatibility with existing and future code need to be considered in addition to functionality.

Table 1.1: Distance elasticities for trade in tasks

	(1) X_{ij}/X_{ii}	(2) X_{ij}/X_{ii}	(3) X_{ij}/X_{ii}	(4) X_{ij}/X_{ii}	(5) $\hat{X}_{ij}/\hat{X}_{ii}$
Log distance in miles	-0.8081***	-0.8093***	-0.9129***	-0.6833***	-0.7311***
	(0.0811)	(0.0688)	(0.0834)	(0.1053)	(0.0071)
Controls	Yes	Yes	Yes	Yes	Yes
Same location dummy	No	No	No	Yes	No
Sample	FUA + Admin	FUA only	US FUA only	FUA + Admin	FUA + Admin
Observations	16,678,894	5,266,000	60,945	16,678,894	13,190,040
Pseudo R-squared	0.7067	0.7053	0.8419	0.7087	0.4920

Estimations results of equation (1.3). In columns (1), (4) and (5) the sample consists of all FUAs and Admin-2 regions. In column (2) we restrict the sample to FUAs, and in column (3) to FUAs in the United States only. In column (5) we multiply each commit by the individual quality measure of the author obtained from approach 2, in order to get a quality weighted trade flows ($\hat{X}_{ij}$). We winsorize this measure at the 99.95 level to account for extreme values produced by rare very small values in the denominator because of this multiplication. Controls include binary dummies for the same country, shared borders and shared official languages. Column (4) additionally includes a same location dummy. All specifications are estimated with PPML, and include importer and exporter fixed effects. * (**) (***) indicates significance at the 10 (5) (1) percent level.

more intensively. However, the coefficient remains sizable.[12]

One limitation of our data is that our flow variable is constructed based on counts but there might be a substantial level heterogeniety between different commits. To address this limitation, we multiply the commits made by individual j by their quality score, which we introduced in Section 1.3 under *Approach 2*. We denote the quality adjusted trade flows by $\hat{X}_{ij}$. Ideally the quality measure would be at the level of a transaction/commit, but we do not have this kind of information. Our assumption is that higher quality individuals make more valuable commits. Column (5) of Table 1.1 presents the result for this quality adjusted measure. The resulting absolute value of the distance elasticity is only slightly lower compared to the one in column (1).

1.4.2 City productivities

Productivity measures for the top 35 cities constructed according to the methodology described in Section 1.3.1 are presented in the first column of Table 1.2. It is reassuring that San Jose, which according to our FUA definition includes the entire Bay Area, appears at the top of our ranking. The positions of Portland, nicknamed Silicon Forest with its substantial technological cluster, and of Bengaluru, the IT capital of India, lend further credibility to our results. We formally validate our measure in Appendix A.4 by showing significant positive correlations with IT sector wages at the level of US FUAs and globally

[12] We also estimate the gravity equation using the June 2019 snapshot of the data. The estimated distance elasticity is very similar to the baseline with a value of -0.858. We conclude that the Covid-19 pandemic does not systematically affect the finding. We also estimate versions of the gravity equation that control for clock-hour differences (i.e., the cyclical differences in the hour of the day, where a 24-hour difference equals 0), where the estimated elasticity coefficient is -0.977. This result provides reassurance that the measured distance elasticity reflects indeed spatial frictions, rather than being confounded by temporal misalignments between more distant locations that exacerbate communication challenges.

at the country level. Additionally, we utilize information on software developers' affiliations to construct university rankings for the US, UK, and Germany and compare them with such rankings from other sources.

In columns 2 and 3 of Table 1.2 we present the results for the two reduced form approaches. One noticeable difference is that, for these approaches, the list is dominated by large cities. A key advantage of the structural model is that the results do not depend on city size. This can be seen from equation (1.3), where the outcome variable in the gravity equation is normalized by internal interactions. In the case of the recursive ranking approaches, on the other hand, it is natural that large cities receive more links; accordingly, it is not proper to interpret the scores obtained from these two methods as measures of productivity. The method based on the aggregation of individuals' scores can actually be interpreted as a proxy for total output.

Looking at some individual cities, we can see these differences. For instance, large cities with many users, such as London or Boston, rank higher in approaches 1 and 2 compared to the rank they receive through the model. Another example is Taichung, which is not a large city compared to other Asian giants but hosts Taiwan's world-beating semiconductor industry. We also find that Poughkeepsie has a relatively high rank. This is the location of IBM's headquarters. The productivity ranking by the model can deliver somewhat unexpected results as well. Specifically, we observe some locations that are not traditionally associated with the IT sector, for example, Las Palmas de Gran Canaria. Such locations might be able to selectively, due to amenities or preferential tax regimes, attract top experts, who can have a profound impact on estimated productivity.

1.4.3 Comparing software development productivity gaps with GDP per capita

In this subsection, we compare our estimated productivities with conventional measures of economic development. Since we rely on city-level data and GDP per capita data at this level of granularity do not exist, we use nighttime luminosity per capita as a proxy for income levels. One problem with nighttime luminosity is that rural or underdeveloped and sparsely populated areas may not emit any light. For this reason, we restrict the analysis to FUAs. In Table 1.3 we regress our productivity measure on nighttime luminosity per capita. In the first column, we observe a strong positive relationship between our productivity estimates and income levels, proxied by nighttime luminosity per capita, for the sample of all FUAs.

Next, we compare our productivity measure with GDP per capita data from the WDI. As was mentioned above, we need to aggregate our productivity measures at the country level. We use three alternative approaches. First, we calculate the average productivity in the top 5% of locations within each country. Second, we use population shares of each location within each country and construct population weighted aggregate productivity at the country level. Third, we use the GitHub user shares of each location within each

Table 1.2: Ranking of the top 35 cities across the world

Rank	Model	Approach 1	Approach 2
1	San Jose	San Jose	San Jose
2	Prague	New York	New York
3	Bengaluru	Seattle	London
4	Las Palmas de Gran Canaria	Boston	Beijing
5	Los Angeles	London	Seattle
6	Nuremberg	Washington D.C.	Shanghai
7	Portland (Oregon)	Los Angeles	Portland (Oregon)
8	Ottawa	Paris	Boston
9	New York	Beijing	Los Angeles
10	Seattle	Tokyo	Tokyo
11	Detroit	Atlanta	Berlin
12	Taichung	Chicago	Paris
13	Krasnoyarsk	Portland (Oregon)	Guangzhou
14	Toronto	Berlin	Toronto
15	Berlin	Denver	Austin
16	Ho Chi Minh City	Austin	Hangzhou
17	Sydney	Shanghai	Chicago
18	Tokyo	Toronto	Denver
19	Cape Town	Amsterdam	Washington D.C.
20	Cambridge	Bengaluru	Melbourne
21	Arrecife	Seoul	Pittsburgh
22	London	Philadelphia	Stockholm
23	Dallas	Tijuana	Moscow
24	São Paulo Nanjing	Guangzhou	Sydney
25	Krakow	Vancouver	Vancouver
26	Boston	Zurich	Bengaluru
27	Oslo	São Paulo	Montreal
28	Vancouver	Stockholm	Amsterdam
29	Moscow	Montreal	São Paulo
30	Beijing	Sydney	Atlanta
31	Dutchess County US (Poughkeepsie)	Cambridge	Philadelphia
32	Austin	Moscow	Madrid
33	Melbourne	Delhi [New Delhi]	Barcelona
34	Nanjing	Melbourne	Munich
35	Tijuana	Hangzhou	Seoul

This table displays the top 35 locations ranked by the three different methodologies described in Section 1.3.

Table 1.3: Correlations between IT productivity and nighttime luminosity per capita and GDP per capita globally

	(1) Log productivity	(2) Log productivity	(3) Log productivity	(4) Log productivity
Log nightlights per capita	0.5248*** (0.0634)			
Log GDP per capita		0.8448*** (0.1162)	0.8367*** (0.1228)	0.9028*** (0.1259)
Sample	FUA	Country level	Country level	Country level
Aggregation method		Average of top 5%	Population weighted	User weighted
Observations	2,639	121	121	121
R-squared	0.0239	0.3252	0.3145	0.3251
F	68.45	52.86	46.45	51.40

The dependent variables are log productivity estimated from the model. For the country level regressions producicitivities are aggregated using three different approaches: first, by averaging productivity in top 5% locations (column 2); second, by applying population weights in each location (column 3); third, by applying GitHub user weights in each location (column 4). For the country level regressions, we restrict the sample to those countries with multiple locations to reduce the influence of outliers, however the results are robust to using all countries. Standard errors are robust. * (**) (***) indicates significance at the 10 (5) (1) percent level.

country and construct user weighted aggregate productivity at the country level. The results, presented in columns (2)–(4) of Table 1.3, show that there is a strong positive relationship between GDP per capita and all three productivity measures.

Having established a positive relationship between our estimated productivity measure and various measures of income, we also want to assess whether gaps in software development productivity are different from gaps in GDP per capita between high and low income countries. To this end, we calculate the difference in average log GDP per capita of countries in the top and bottom GDP per capita deciles. We fix the set of these countries in both groups and also calculate the difference between the average log of productivity. The difference in GDP per capita is 4.61 log points (see Table 1.4). The equivalent figures are 4.27 log points for within-country population-weighted productivity, 4.15 log points for the average productivity of top 5% locations, and 4.64 log points for GitHub user-weighted productivity. According to all three approaches, the productivity differences are very close to each other and also to the differences in GDP per capita. However, we know from the macro development literature that the agricultural sector is a major contributor to per capita GDP differences between rich and poor countries (Gollin et al., 2002). Productivity differences in other sectors are smaller. Thus, we want to compare our estimated productivity gaps with non-agricultural sectors. We use data from the WDI on value added and employment in the industry and services sectors and construct productivity gaps for the same set of countries that we classified as belonging to the top and bottom deciles

Table 1.4: Productivity gaps between rich and poor countries

Variables	Productivity gap
GDP per capita	4.61
Industry VA per worker	3.71
Services VA per worker	3.73
IT productivity, top 5%	4.15
IT productivity, population weighted	4.27
IT productivity, user weighted	4.64

This table present log productivity differences between top and bottom 10% of countries sorted by GDP per capita. The sample is restricted to those countries with multiple locations to reduce the influence of outliers, however the results are robust to using all countries. Productivity gaps are calculated as $log(\bar{X}_{top10}) - log(\bar{X}_{bot10})$, where $\bar{X}$ is the average of the variable shown in the rows of this table in top or bottom income group. Data for GDP per capita, sectoral value added and employment were obtained from WDI. IT producitivities are aggregated at the country level by using three approaches: first, by averaging productivity in top 5% locations; second, by applying population weights in each location; third, by applying GitHub user weights in each location.

based on GDP per capita. The productivity gap for industry is 3.71 and for services 3.73, which are smaller than our estimated IT productivity gaps. This means that in terms of productivity in the software development sector, poor countries perform slightly worse than they do in other non-agricultural sectors.

1.4.4 Trade in ideas

In Section 1.2 under **Forks and pull requests** we discussed that our data allow us to study trade of ideas and final software utilizing forks. In this case the analogue of equation (1.1), which formalized the construction of the flow of code, is given by:

$$\tilde{X}_{ij} = \sum_{k \in K} fork_{ik} \times 1[owner_{jk} = 1], \tag{1.6}$$

where $\tilde{X}_{ij}$ is the flow of final software from city j to city i, $fork_{ik}$ is the number of forks on project k by other projects with owners from city j and $1[.]$ is the indicator function equal to 1 if the owner of project k is located in city j. Note that for the construction of this measure we use the second category of forks we described in the data section only, as those capture the dimension of trade in ideas.

Going back to the model described in Section 1.3.1, we now assume that the unit of production is a project owner located in city i who produces a differentiated software q. On the demand side other project owners decide from which project to fork. We follow the same steps as above to estimate a gravity equation and obtain measures of productivity of final software generation. Column 1 of Table 1.5 present the results of the distance elasticity for trade in ideas/final software. The estimated coefficient is smaller compared to trade in software code, which suggests that ideas flow more freely in space, yet not fully void of frictions.

Table 1.5: Trade in ideas

	(1)	(2)	(3)	(4)
	$\tilde{X}_{ij}/\tilde{X}_{ii}$	Comparative advantage in ideas over services		
Log distance in miles	-0.4376***			
	(0.0072)			
Log GDP per capita		0.8082***	0.3396***	0.1280
		(0.1751)	(0.1241)	(0.1048)
Controls	Yes	No	No	No
Sample	FUA + Admin	Country level	Country level	Country level
Aggregation method		Average of top 5%	Population weighted	User weighted
Observations	11,922,149	119	119	119
R-squared	0.6629	0.1363	0.0611	0.0139
F		21.30	7.492	1.493

In column (1), the dependent variable is log productivity for trade in ideas estimated from the model. Controls include binary dummies for the same country, shared borders and shared official languages. In columns (2) - (4), the dependent variable is the ratio of productivities for trade in ideas and trade in services aggregated to the country level using three different approaches: first, by averaging the ratio in top 5% locations (column 2); second, by applying population weights in each location (column 3); third, by applying GitHub user weights in each location (column 4). For the country level regressions, we restrict the sample to those countries with multiple locations to reduce the influence of outliers. The results are robust to using all countries, and the estimate in column (4) becomes statistically significant. Standard errors are robust. * (**) (***) indicates significance at the 10 (5) (1) percent level.

Final product ownership generates more value than coding, which is why developers individually and software production locations collectively strive to move up in the value chain and provide successful final software products (Arora et al., 2001). To better understand the positions of different geographic locations in the value chain, we construct a measure of comparative advantage for idea production versus provision of coding services. We back out productivities in idea production equivalently to the approach for software development services that were presented in Table 1.2, however using the flow data based on equation (1.6). Then we construct the ratio of productivity in ideas over productivity in services at the location level and aggregate it to the country level following the previous three aggregation approaches. We regress the resulting ratio on GDP per capita. The results of this exercise are presented in columns (2)-(4) of Table 1.5. For all three aggregation approaches we observe a positive relationship between GDP per capita and comparative advantage in idea production, while in two cases the estimated coefficients are statistically significant. These results suggest that higher income countries have a comparative advantage in idea production compared to coding services.

1.5 Migration and sorting

In this section, we turn to the migration of human capital across and within countries. We are particularly interested in determining whether there is quality-based selection into locations. To assess this, we construct an individual-level migration variable, which requires that we observe individuals in both our 2019 and 2021 snapshot of the data and that they report their location in both years.[13] The resulting sample comprises about 1.56 million users, of whom about 98,000 migrate, 38,000 between countries and 60,000 within countries. At the country level, the largest gross outflows of migrants are from the US, India, the UK, Canada and Brazil, while countries with largest gross inflows are the US, the UK, Germany, Canada and the Netherlands. Figure A7 illustrates some of the largest bilateral migration flows.

We combine this information about migration decisions with the individual-level quality scores that were constructed as an intermediate step to assemble the city ranking according to *Approach 2*. We regress a dummy that indicates whether an individual migrated or not on this measure. The results are presented in panel A, columns (1) – (3) of Table 1.6. We observe a positive and statistically highly significant coefficient that is robust to different fixed effect structures – the most rigorous of which includes destination country and origin city fixed effects. In this case, migrants from the same city of differing quality lend the identifying variation. In panel B we assign individuals to quartiles based on their score and estimate the same specifications by using indicator variables for each quartile. We observe that the estimated coefficients increase monotonically in all specification. In columns (4) and (5) of the table, we study differences in within and across country migration in relation to our measure. The observed effects are similar for both types.

To address the question of quality-based sorting, we construct an indicator for upward and downward migration. The indicator is equal to 1 if the destination city of the migrant is ranked higher than the origin city based on the estimated productivities from the model. The results for the continuous score and quartiles dummies are presented in panels A and B of Table 1.7, respectively. In both panel A and B, we observe that the coefficient on upward migration is larger than that on downward migration. The fact that the coefficient on downward migration is positive is not unexpected, because we know from the literature that higher-skilled individuals are more mobile (Borjas et al., 1992). In columns (3) and (4) of Table 1.7, we restrict the sample to migrants only, thereby eliminating potential confounding effects arising from selection into migration. This restriction also addresses concerns that reporting a change in location may be correlated with the quality measure of software developers. In these specifications, the estimated coefficients for upward and downward migration have opposite signs. The results of this table indicate that (i) higher quality software developers are more likely to migrate in general; (ii) among migrants, those of higher quality are more likely to migrate to better locations and those of lower quality to worse locations.

While we demonstrated that our measure of a location's productivity is well correlated with income levels, it might be the case that individuals choose to migrate to a lower quality location with higher income levels. To investigate this, we regress our individual

[13] We apply the same data cleaning efforts to the 2019 snapshot of the data that we described in Section 1.2 for the 2021 snapshot of the data.

Table 1.6: Individual quality and likelihood to migrate

	(1) Migrated	(2) Migrated	(3) Migrated	(4) Migrated within country	(5) Migrated across country
Panel A:					
Log individual score	0.1902***	0.1639***	0.1898***	0.1902***	0.1838***
	(0.0091)	(0.0081)	(0.0052)	(0.0052)	(0.0123)
Observations	939,034	938,552	933,943	921,550	909,621
Pseudo R2	0.0175	0.0630	0.108	0.106	0.222
Panel B:					
2nd quartile	0.6303***	0.5971***	0.6201***	0.6804***	0.5001***
	(0.0224)	(0.0252)	(0.0188)	(0.0160)	(0.0404)
3rd quartile	0.9101***	0.8504***	0.8814***	0.9439***	0.7497***
	(0.0160)	(0.0215)	(0.0218)	(0.0184)	(0.0446)
4th quartile	1.2919***	1.1739***	1.1991***	1.2919***	1.0106***
	(0.0166)	(0.0278)	(0.0279)	(0.0219)	(0.0635)
Observations	1,566,353	1,565,559	1,558,279	1,539,900	1,519,561
Pseudo R2	0.0439	0.0902	0.133	0.123	0.244
Origin country FE	Yes	Yes	No	No	No
Destination country FE	No	Yes	Yes	No	Yes
Origin city FE	No	No	Yes	Yes	Yes
Number migrants	97,438	97,438	97,438	60,122	37,316

In columns (1) - (3) the dependent variable is an indicator variable that is equal to one if an individual's location changed comparing the 2019 and 2021 snapshots of the GitHub database. In column (4) we consider location changes within the same country only, and in column (5) changes to locations in another country only. The individual quality score is based on the centrality of the individual in the follower network. Panel A presents results for the log of this individual score, whereas in panel B we construct dummies for the quality score quartile an individual belongs to. All specifications are estimated by PPML. The fixed effects employed in each regression are marked in the table. Standard errors are clustered at the level of origin cities. * (**) (***) indicates significance at the 10 (5) (1) percent level.

level quality scores on a dummy indicating an upward or downward migration based on the origin and destination countries' relative GDP per capita. The results are presented in Table 1.8 and are similar to the ones based on locations' productivities. We observe that individuals with higher quality scores are more likely to migrate in both directions, but the coefficient on upward migration is higher. In columns (3) and (4) we again restrict the sample to cross-country migrants to remove systematic differences between migrants and non-migrants, as well as within-country migrants and cross-country migrants. The results show that among migrants, the higher-skilled ones are more likely to move up.

1.5.1 Migrants in their destinations

Next we assess migrants' relative quality compared to the quality of residents in their destination location before migrating. To this end we construct a dummy variable that

Table 1.7: Directional migration of individuals based on individual quality

	(1) Up migration	(2) Down migration	(3) Up migration	(4) Down migration
Panel A:				
Log individual score	0.2124***	0.1515***	0.0307***	-0.0343***
	(0.0064)	(0.0081)	(0.0034)	(0.0070)
Observations	872,287	878,591	69,184	66,393
Pseudo R2	0.186	0.128	0.0907	0.127
Panel B:				
2nd quartile	0.6368***	0.5832***	0.0104	-0.0276**
	(0.0214)	(0.0284)	(0.0104)	(0.0119)
3rd quartile	0.9155***	0.8246***	0.0558***	-0.0787***
	(0.0217)	(0.0356)	(0.0091)	(0.0107)
4th quartile	1.2668***	1.0687***	0.0954***	-0.1364***
	(0.0288)	(0.0452)	(0.0101)	(0.0148)
Observations	1,465,610	1,467,499	85,657	82,480
Pseudo R2	0.202	0.147	0.0927	0.131
Destination country FE	Yes	Yes	Yes	Yes
Origin city FE	Yes	Yes	Yes	Yes
Sample	All	All	Migrants	Migrants
Number migrants	52,256	37,763	52,256	37,763

The dependent variable *up migration* (*down migration*) is an indicator variable that is equal to one if an individual migrated to a location more (less) productive than their previous location. In columns (3) and (4) we restrict the sample to migrants only. The individual quality score is based on the centrality of the individual in the follower network. Panel A presents results for the log of this individual score, whereas in panel B we construct dummies for the quality score quartile an individual belongs to. All specifications are estimated by PPML. The fixed effects employed in each regression are marked in the table. Standard errors are clustered at the level of origin cities. * (**) (***) indicates significance at the 10 (5) (1) percent level.

indicates whether an individual is above or below the median quality of GitHub users in their destination city. In panel A column (1) of Table 1.9 we regress the migration dummy on this measure, employing destination city fixed effects. By design the outcome has a sample mean close to 0.5, such that a positive coefficient in this regression indicates that migrants are on average better than the median user in their destination. Vice versa, a negative coefficient would suggest the opposite. The estimated effect implies that an average migrant is better than the median of users in 74% of cases in our sample.[14] In columns (2) and (3) we decompose migration into upward and downward migration based on locations' productivities as in Table 1.7. The results show that on average this finding holds even in the case of an upward migration move. Naturally, the estimated coefficient is larger for downward migration moves, as the median quality of software developers is lower

[14]We transform the semi-elasticity of 0.3937 according to the following formula: $(100 * (exp(\beta) - 1))$. Multiplying the baseline likelihood of 0.5 with the resulting 48.245% yields around 24% higher likelihood of being above the median quality in the destination.

Table 1.8: Migration to higher and lower income locations based on individual quality

	(1) Migration to > GDP per capita	(2) Migration to < GDP per capita	(3) Migration to > GDP per capita	(4) Migration to < GDP per capita
Panel A:				
Individual quality	0.3021***	0.1936***	0.0196***	-0.0248***
	(0.0111)	(0.0116)	(0.0040)	(0.0070)
Observations	839,292	807,682	27,416	25,410
Pseudo R2	0.125	0.125	0.141	0.226
Panel B:				
2nd quartile	0.5330***	0.6941***	-0.0086	0.0049
	(0.0306)	(0.0379)	(0.0108)	(0.0153)
3nd quartile	0.8936***	0.9535***	0.0078	-0.0150
	(0.0272)	(0.0368)	(0.0090)	(0.0139)
4nd quartile	1.3681***	1.2778***	0.0344***	-0.0584***
	(0.0268)	(0.0490)	(0.0089)	(0.0150)
Observations	1,393,561	1,345,274	33,800	31,156
Pseudo R2	0.140	0.138	0.142	0.230
Origin city FE	Yes	Yes	Yes	Yes
Sample	All	All	Cross-country migrants	Cross-country migrants
Number migrants	22,913	14,403	22,913	14,403

The dependent variable in columns (1) and (3) ((2) and (4)) is an indicator variable that is equal to one if an individual migrated to a country with higher (lower) GDP per capita than their previous location. In columns (3) and (4) we restrict the sample to cross-country migrants only. The individual quality score is based on the centrality of the individual in the follower network. Panel A presents results for the log of this individual score, whereas in panel B we construct dummies for the quality score quartile an individual belongs to. All specifications are estimated by PPML. The fixed effects employed in each regression are marked in the table. Standard errors are clustered at the level of origin cities. * (**) (***) indicates significance at the 10 (5) (1) percent level.

in these cases. In columns (4) and (5) we replicate the specification but for upward and downward migration defined by GDP per capita differences as in Table 1.8. The general patterns and estimated coefficients turn out to be very similar to the productivity based results.

In panel B of Table 1.9 we investigate how migration decisions affect the migrants' individual position in the quality score distribution. We calculate the change in quality score quartile based on the distribution of quality scores in origin and destination location in 2019, that is prior to migration taking place. We regress the change in quartile on the different migration dummies we have employed in panel A. The results are consistent with the evidence we compiled so far. Migrants move on average down the quality score distribution, which is driven by moves to more productive and higher income locations.

Table 1.9: Migrants comparative quality in the destinations

	(1) Above median score in destination	(2) Above median score in destination	(3) Above median score in destination	(4) Above median score in destination	(5) Above median score in destination
Panel A:					
Migrated	0.3937***				
	(0.0091)				
Up migration (productivity)		0.3469***			
		(0.0079)			
Down migration (productivity)			0.4332***		
			(0.0134)		
Up migration (GDP per capita)				0.3284***	
				(0.0151)	
Down migration (GDP per capita)					0.3851***
					(0.0155)
Observations	1,560,104	1,553,869	1,553,869	1,560,104	1,560,104
Pseudo R2	0.0050	0.0033	0.0034	0.0025	0.0025
	(1) Δ quartile individual score	(2) Δ quartile individual score	(3) Δ quartile individual score	(4) Δ quartile individual score	(5) Δ quartile individual score
Panel B:					
Migrated	-0.0496***				
	(0.0125)				
Up migration (productivity)		-0.1224***			
		(0.0121)			
Down migration (productivity)			0.0561***		
			(0.0192)		
Up migration (GDP per capita)				-0.1449***	
				(0.0201)	
Down migration (GDP per capita)					0.0039
					(0.0168)
Observations	1,566,039	1,553,926	1,553,926	1,566,039	1,566,039
R-squared	0.4388	0.1012	0.0714	0.4438	0.4346
Destination city FE	Yes	Yes	Yes	Yes	Yes
Number migrants	97,438	52,256	37,763	22,913	14,403

The dependent variable in panel A is an indicator variable that is equal to one if an individual has a higher quality score than the average user in the destination location. In panel B the dependent variable is the difference of individuals' quality score quartiles between their location in 2019 and their location in 2021, calculated according to the distribution of quality scores in 2019 in both locations. Explanatory variables are: *Migration* - a dummy for migration; *Up migration* a dummy if migration takes place to a location with higher productivity or to a country with higher GDP per capita; *Down migration* a dummy if migration takes place to a location with lower productivity or a country with lower GDP per capita. The individual quality score is based on the centrality of the individual in the follower network. All specifications in panel A are estimated by PPML, in panel B by OLS. The fixed effects employed in each regression are marked in the table. Standard errors are clustered at the level of origin cities. * (**) (***) indicates significance at the 10 (5) (1) percent level.

Moves to less productive places see the migrant on average move up the quality score distribution.

Table 1.10: Migration flows at the country level

	(1) Net migration	(2) Out-migration	(3) In-migration
Panel A:			
Log GDP per capita	0.0213*	0.0128**	0.0323**
	(0.0115)	(0.0055)	(0.0129)
Observations	146	146	146
R-squared	0.0177	0.0269	0.0442
Panel B:			
Log GDP per capita	0.0327***	-0.0042	0.0250***
	(0.0075)	(0.0060)	(0.0082)
Observations	108	108	108
R-squared	0.1053	0.0037	0.1028

In Panel A we require countries to have more than 20 GitHub users. For the outcomes net migration and in-migration 13 countries, and for out-migration 14 countries do not meet this condition. In Panel B we restrict the sample to countries with more than 150 users. Standard errors are robust. * (**) (***) indicates significance at the 10 (5) (1) percent level.

1.5.2 Aggregate flows of migration

In the previous subsection we documented strong sorting patterns using individual level migration decisions. These patterns imply that locations and countries with an initially low stock of individuals with high quality are loosing their best experts. In the literature this phenomenon is referred to as brain drain. In this subsection we investigate whether the migration pattern at the individual level has tractable implications at the aggregate level. To this end, we construct three measures: net migration flows, gross inflows and gross outflows.

We aggregate the individual quality scores at the country level in 2019 to calculate the initial stock of human capital. We then construct our measure of gross inflow, as the sum of scores of individuals who migrated to a country in 2021. Equivalently, we calculate the measure of gross outflow as the sum of scores of migrants leaving the country. We divide both the inflow and the outflow measure by the initial stock of human capital we calculated for 2019, to express them in relative terms. Net migration is constructed as the ratio of the stock of human capital in 2021, over the initial stock in 2019. In Table 1.10 we regress these measures on GDP per capita. To reduce the noise in this regression, we drop countries that have less than 20 users in 2019 in panel A. In panel B we increase the threshold to at least 150 users.

The results show that countries with higher GDP per capita experience positive net migration. This appears to be driven by larger inflows, indicated by the positive coefficients in both panels in the third column, which are larger than the coefficients for outflows in the second column. The small positive coefficient for out-migration becomes insignificant for the specification in panel B. We think, however, that the tentatively positive coefficient on out-migration makes intuitively sense, indicating that there is stronger movement in both directions in higher income countries. This resembles a setting in which software developers from high-income countries might migrate to other high income countries, and software developers from low-income countries tend to migrate strictly upwards. The results confirm our conjecture based on the individual level regressions that wealthier countries are attracting talent, while poorer countries are losing talent.

1.6 Conclusions

In this paper we bring new empirical evidence to the debate on the role high-skilled tradable services play in economies around the world, and for the development process of low-income countries.

We study the software development industry, specifically the large and commercially important sector of open source development, by utilizing detailed data at the level of individual software developer. Our main contribution is the estimation of productivity levels in 5,400 locations around the world. Our results show that there are large differences in productivity levels within and across countries, which are indicative of human capital differences across space. We find that the productivity gaps between the richest and poorest countries in software development are somewhat larger than for the broadly defined manufacturing and services sectors. Developing countries are seemingly not able to leverage the fact that transportation costs are near zero to generate exports, likely because of information frictions that are captured in the sizable distance elasticities we measure. Moreover, we find evidence of "brain drain" – that is, a sorting pattern in which the best software developers from less developed countries or cities with low levels of productivity move to more productive locations. This exacerbates existing differences.

These findings present a rather bleak picture for low-income countries. Nevertheless, there are some locations in developing countries, such as Bengaluru, which have very high productivity levels and are ranked among the global leaders. Understanding the evolution of the ICT sector in these places can provide valuable lessons for other locations in developing countries on how to boost productivity in this sector.

There are a number of important questions that require further attention. Follow-up research should, for example, investigate the role of agglomeration effects in the software development sector. Another important question pertains to the potential knowledge spillovers from emigrating software developers back to their origin locations, and whether these spillovers might offset human capital losses from brain drain over the long term. The challenges in tackling these questions involve utilizing a solid identification strategy based on plausibly exogenous shocks, and, in this connection, the need for a longer time horizon. Despite the fact that GitHub has existed as a platform since 2008, the user base was comparatively small in the early periods, such that the utilization of a longer time horizon comes with the trade-off of a much smaller sample size. We believe that it will be possible to answer these questions credibly as more data become available to researchers.

References

Akcigit, Ufuk, Salomé Baslandze, and Stefanie Stantcheva, "Taxation and the International Mobility of Inventors," *American Economic Review*, October 2016, *106* (10), 2930–81.

Allen, Treb, "Information Frictions in Trade," *Econometrica*, 2014, *82* (6), 2041–2083.

_ and Costas Arkolakis, "Modern Spatial Economics: A Primer," in Lili Yan Ing and Miaojie Yu, eds., *World Trade Evolution*, Routledge, 2018.

Arora, Ashish, V.S. Arunachalam, Jai Asundi, and Ronald Fernandes, "The Indian software services industry," *Research Policy*, 2001, *30* (8), 1267–1287.

Bai, Jie, Maggie X Chen, Jin Liu, Xiaosheng Mu, and Daniel Yi Xu, "Stand Out from the Millions: Market Congestion and Information Friction on Global E-Commerce Platforms," 2022.

Bias, Barbara and Song Ma, "The Education-Innovation Gap," Working Paper 2023.

Borjas, George J., Stephen G. Bronars, and Stephen J. Trejo, "Self-selection and internal migration in the United States," *Journal of Urban Economics*, 1992, *32* (2), 159–185.

Brin, Sergey and Lawrence Page, "The anatomy of a large-scale hypertextual Web search engine," *Computer Networks and ISDN Systems*, 1998, *30* (1), 107–117. Proceedings of the Seventh International World Wide Web Conference.

Buera, Francisco J. and Joseph P. Kaboski, "The Rise of the Service Economy," *American Economic Review*, May 2012, *102* (6), 2540–69.

Clemens, Michael A., "Why Do Programmers Earn More in Houston Than Hyderabad? Evidence from Randomized Processing of US Visas," *American Economic Review*, May 2013, *103* (3), 198–202.

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

_ and _ , "Trade in Goods and Trade in Services," in Lili Yan Ing and Miaojie Yu, eds., *World Trade Evolution*, Routledge, 2018.

Eckert, Fabian, "Growing Apart: Tradable Services and the Fragmentation of the U.S. Economy," Working Paper 2019.

Fackler, Thomas, Michael Hofmann, and Nadzeya Laurentsyeva, "Defying Gravity: What Drives Productivity in Remote Teams?," Technical Report, Discussion Paper 2023.

Fally, Thibault, "Structural gravity and fixed effects," *Journal of International Economics*, 2015, *97* (1), 76–85.

Freire, Sergio, Kytt MacManus, Martino Pesaresi, Erin Doxsey-Whitfield, and Jane Mills, "Development of new open and free multi-temporal global population grids at 250 m resolution," *Population*, 2016, *250.*

Garicano, Luis, "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy*, 2000, *108* (5), 874–904.

Gervais, Antoine and J. Bradford Jensen, "The tradability of services: Geographic concentration and trade costs," *Journal of International Economics*, 2019, *118*, 331–350.

Goldbeck, Moritz, "Bit by bit: colocation and the death of distance in software developer networks," Technical Report, Discussion Paper 2023.

Gollin, Douglas, Stephen Parente, and Richard Rogerson, "The Role of Agriculture in Development," *American Economic Review*, May 2002, *92* (2), 160–164.

Gousios, Georgios, “The GHTorent dataset and tool suite,” in “2013 10th Working Conference on Mining Software Repositories (MSR)” IEEE 2013, pp. 233–236.

Hall, Robert E. and Charles I. Jones, “Why do Some Countries Produce So Much More Output Per Worker than Others?*,” *The Quarterly Journal of Economics*, 02 1999, *114* (1), 83–116.

Hendricks, Lutz and Todd Schoellman, “Human Capital and Development Accounting: New Evidence from Wage Gains at Migration*,” *The Quarterly Journal of Economics*, 12 2017, *133* (2), 665–700.

Hsieh, Chang-Tai and Esteban Rossi-Hansberg, “The Industrial Revolution in Services,” *Journal of Political Economy Macroeconomics*, 2023, *1* (1), 3–42.

Klenow, Peter J. and Andrés Rodríguez-Clare, “The Neoclassical Revival in Growth Economics: Has It Gone Too Far?,” *NBER Macroeconomics Annual*, 1997, *12*, 73–103.

Levchenko, Andrei A. and Jing Zhang, “The evolution of comparative advantage: Measurement and welfare implications,” *Journal of Monetary Economics*, 2016, *78*, 96–111.

Martellini, Paolo, Todd Schoellman, and Jason Sockin, “The Global Distribution of College Graduate Quality,” *Journal of Political Economy*, 2024, *0* (0), 000–000.

Moreno-Monroy, Ana I, Marcello Schiavina, and Paolo Veneri, “Metropolitan areas in the world. Delineation and population trends,” *Journal of Urban Economics*, 2021, *125*, 103242.

Ruggles, Steven, Sarah Flood, Ronald Goeken, Megan Schouweiler, and Matthew Sobek, “Integrated Public Use Microdata Series: Version 12.0 [dataset],” Technical Report, Minneapolis, MN: IPUMS 2022.

Startz, Meredith, “The value of face-to-face: Search and contracting problems in Nigerian trade,” *Available at SSRN 3096685*, 2016.

Wachs, Johannes, “Digital traces of brain drain: developers during the Russian invasion of Ukraine,” *EPJ Data Science*, 2023, *12* (1), 14.

_ , Mariusz Nitecki, William Schueller, and Axel Polleres, “The geography of open source software: evidence from github,” *Technological Forecasting and Social Change*, 2022, *176*, 121478.

Waugh, Michael E., “International Trade and Income Differences,” *American Economic Review*, December 2010, *100* (5), 2093–2124.

Wright, Nataliya Langburd, Frank Nagle, and Shane Greenstein, “Open source software and global entrepreneurship,” *Research Policy*, 2023, *52* (9), 104846.

A Appendix

A.1 Additional data description

A.1.1 Spatial data

We employ a number of supplementary data sources, which we combine with our main data by spatial proximity.

Locations We use shape files from the Global Human Settlements Functional Urban Areas dataset, which identifies metropolitan areas and their surrounding commuting zones around the world. The methodology of creating these functional urban areas (FUAs) is laid out in Moreno-Monroy et al. (2021).[15] We map GitHub users based on their geocoordinates to the FUAs. To capture less densely populated areas as well, we then group together users that fall outside the borders of FUAs and assign them to the admin-2 region they are located in. Shapefiles for administrative borders come from the Database of Global Administrative Areas (GADM). In the remaining paper we use the terms locations and cities interchangeably. We drop locations with less than 10 unique users to avoid calculating very noisy aggregate measures at the location level. The top 20 locations in terms of the number of users are displayed in Table A3. We arrive at a final sample of 5,424 locations in 179 countries. We map all our other data sources into these geographic areas; Figure A5 provides a visual example of this approach for nighttime luminosity, GitHub users and FUAs.

Population We extract population numbers for the locations we consider from the Global Human Settlements population grid, which is a spatial raster that depicts the distribution of the residential population. We utilize the grid at a resolution of 1 kilometer; each cell has a value for the predicted number of people living in that area. The construction of the raster is explained in Freire et al. (2016). We overlay that raster with the FUA and admin-2 borders shape files to extract the sum of population at our level of observation.

Nightlights We obtain nighttime luminosity by overlaying a spatial raster of nighttime luminosity provided by the Earth Observation Group with our FUA and admin-2 border shape files. We utilize the V2.1 annual version of VIIRS to extract the average sum of nocturnal light omitted at the location level. This version of nighttime data has the advantage that it is not top coded, making cross-country comparisons of cities with potentially strongly diverging luminosity levels more precise.

[15] For some countries alternative definitions of urban areas are available – for example, the Metropolitan Statistical Areas or Commuting Zones for the US – but such maps are not available for all countries and approaches may differ across countries.

A.1.2 Income data

We are interested in relating the differences we measure in human capital across space to income differences. We do so at the level of FUAs for the United States, and globally at the country level.

American Community Survey (ACS) We use the ACS data provided by Ruggles et al. (2022) to construct wages at the level of Public Use Microdata Areas (PUMAs), which are the smallest identifiable geographic unit in that dataset. They are non-overlapping statistical areas containing no fewer than 100,000 people each. Given that FUAs do not exactly align with PUMAs, we intersect them, and re-weight the average wages thus obtained. We calculate the weights as follows:

$$Weight_{p,F} = \frac{Share\ intersected\ area_{p,F} * Population_p}{Population_{P,F}}, \tag{A1}$$

where the index p depicts the individual PUMA, F the FUA it is intersecting with, and P, F all PUMAs intersecting with the same FUA. Figure A6 in the Appendix visualizes the intersection of PUMAs and FUAs.

We use occupational information to identify individuals who are employed in software-related occupations. We identify 14 such occupations, which are listed in Table A2. We have also extended the list by including a broader list of occupations that may require software development skills, such as economist and physicist. This extended list yielded similar results. However, we believe a stricter definition is more appropriate because the fraction of economists engaged in software development is unlikely to be high and this is not their main activity.

Software developer wages We are not aware of any global administrative database on the earnings of software developers. For this reason we utilize data from a survey conducted by *Stack Overflow*, which is a question-and-answer website for programmers and has over 20 million registered users. Every year *Stack Overflow* conducts a survey among its users on various issues related to their professional activity including their salaries. We use the *2023 Developer Survey* since it has broader coverage compared to previous years. Ninety thousand developers from 87 countries responded to the survey. We drop survey responses from users who stated something other than being a software developer by profession or programmer as part of their work, in order to focus on the earnings of IT professionals. Of this sub-sample the number of respondents with non-missing wage income responses ranges from 16409 in the US to 12 in Senegal, Kuwait and Bahrain. The country with the median number of observations has 135 respondents. We winzorise the wages at the 99% level to reduce the impact of outliers, in particular in the small sample countries. Clearly, this survey comes with limitations but we believe that a comparison of our estimated productivity measure with wages from a survey from a different source is a useful exercise that can potentially support the validity of our estimates.

WDI We obtain GDP per capita in constant 2015 US dollars for the years 2019 and 2021 at the country level. We merge this information to our remaining data by 3-letter country codes. From this source we also obtain data for value added per worker for the industry and services sectors.

A.1.3 Representativeness

In the following paragraphs, we provide a more detailed discussion of the representativeness of our sample, given that we are able to map only a sub-sample of users accurately into locations. We refer to information provided in Section 1.2, which introduces the users and commits data, along with the individual quality scores generated through *Approach 2* outlined in Section 1.3.2.

We require the information of users location to attribute commits, which form the basis of the trade flows we construct, to locations. Our dataset comprises 218,848,238 commits from users whose locations were accurately identified following our data cleaning procedures. Additionally, we identify 380,053,481 commits from users without location information. While this constitutes a share of 36.5%, it is noteworthy that users with location information are far more active; They average 82.6 commits compared to 12.1 commits for users lacking location details. To address the potential skew in commit volume caused by less meaningful commits from users with incomplete profiles, we compute a quality-adjusted share by weighting each commit with the respective user's individual quality score. Consequently, when adjusting for quality scores, we are able to attribute 67.4% of the commit volume to specific locations. Notably, our gravity estimations using raw commit counts and quality adjusted commits deliver similar results (see columns (1) and (5) of Table 1.1). The fact that there is a large difference in the covered share of commit volume between both approaches, yet the gravity estimation results being close to each other suggests that it is unlikely that there are systematic patterns in terms of not reporting location information.

Table A1: Share of local connections by team size

Team size	Observations	Local share
2-5	269,053	0.598
6-20	152,971	0.492
21-100	80,064	0.406
>100	83,041	0.158

This table shows the average share of local connections across projects of a given size team. A connection is an undirected link between two users.

Table A2: IT occupations

Code	Description
1005	Computer and information research scientists
1006	Computer systems analysts
1007	Information security analysts
1010	Computer programmers
1021	Software developers
1022	Software quality assurance analysts and testers
1031	Web developers
1032	Web and digital interface designers
1050	Computer support specialists
1065	Database administrators and architects
1105	Network and computer systems administrators
1106	Computer network architects
1108	Computer occupations, all other
1240	Other mathematical science occupations

This table presents the list of occupations in the ACS, which we classify as IT-related. The first column displays occupation codes according to variable *occ*.

Table A3: City user counts

	Location	User count		Location	User count
1	San Jose	101,242	11	Toronto	33,329
2	New York	79,778	12	Guangzhou	32,560
3	London	64,576	13	São Paulo	32,339
4	Bengaluru	62,438	14	Moscos	32,066
5	Beijing	60,909	15	Tokyo	30,909
6	Seattle	46,213	16	Boston	29,773
7	Los Angeles	42,568	17	Chicago	28,983
8	Shanghai	39,951	18	Berlin	23,813
9	Delhi [New Delhi]	38,054	19	Pube	23,221
10	Paris	34,714	20	Seoul	22,137

A.2 The organization of teams

In this section, we study the structure of production teams. Our primary reason for doing so is to understand how to define the flows of software code between locations. However, this touches upon a much broader aspect in the theory of the firm and there is a large literature studying the hierarchies in organizations (Garicano, 2000).

Production teams can be organized in different ways. At one extreme, the production process may be organized in the shape of a star, such that every worker or production unit delivers its output to the central unit. Alternatively, production may be organized as a chain in which each unit delivers its output to the next. Production can also be organized as a fully connected graph in which each individual interacts with everyone else.

We utilize our data to shed light on the structure of software production teams. We construct linkages between individuals based on the follower network within a project. Then, we test whether the owner of the project stands out among others. To that end, we estimate the following specification:

$$y_{ij} = \alpha + \beta_1 Owner_j + \beta_2 Owner_i + \epsilon_{ij}, \tag{A2}$$

where y_{ij} is a dummy if individual i follows individual j , *Owner* is a dummy if the person is the owner of the project and ϵ_{ij} is the error term. If the team is organized as a chain or if everyone interacts with everyone within the network, then the owner should not have a special status and the coefficient $\beta_1 = 0$.

We present the results of our estimations in Table A4. Estimations are conducted for all projects that have more than two participants. In the first column, the only explanatory variable is whether user j is the owner. The estimated coefficient indicates that owners are much more likely to be followed by other project members. Project owners are thus the central figures in projects, and other team members want to be informed about their contributions as well as the issues and pull requests they open (for instance, specifically those labeled "help wanted" or "good first issue"). In terms of the organizational form of production, this resembles the star mentioned above.

In the second column we include the $Owner_i$ control and find that the estimated coefficient is also sizable. However, the larger coefficient of $Owner_j$ that is statistically significantly different from $Owner_i$ shows that the owner is more likely to be followed than follow others. The average for y_{ij} is 0.015. This indicates that within an average team there are few interactions between a randomly selected pair. By contrast, owners play a central role and maintain bilateral interactions with other contributors.

In the following columns we add an indicator variable if a pair of members are located in the same country and city. The estimated coefficients on our variable of interest decrease somewhat but they are still large and statistically significant. In column (5) we report results for the sample of teams with 100 participants or more. The comparison with the

Table A4: The structure of collaboration in software production teams

	(1) i follows j	(2) i follows j	(3) i follows j	(4) i follows j	(5) i follows j	(6) Share of follows
Owner$_j$	2.0161***	2.1468***	1.4894***	1.3300***	1.2989***	0.9352***
	(0.0014)	(0.0015)	(0.0028)	(0.0036)	(0.0141)	(0.0018)
Owner$_i$		1.9697***	1.2169***	1.0627***	-7.2051***	
		(0.0016)	(0.0032)	(0.0041)	(0.9721)	
Same country			0.9506***	0.6787***	0.4621***	
			(0.0018)	(0.0027)	(0.0040)	
Same location				0.4514***	0.2389***	
				(0.0026)	(0.0047)	
Team size	> 2	> 2	> 2	> 2	> 100	> 2
Mean	0.015	0.015	0.030	0.031	0.015	0.161
Observations	244,177,260	244,177,260	47,869,198	30,712,310	24,947,588	3,419,080
Pseudo R^2	0.0303	0.0548	0.0517	0.0502	0.0106	0.0323

Columns (1)-(5) present the estimation results of equation (A2), where the dependent variables are dummies taking a value of 1 if contributor i follows contributor j. Column (6) presents the results of a regression where the dependent variable is the share of follower links of individual i among all following links in a given project. All specifications are estimated with PPML. In column (5) the sample is restricted to projects with more than 100 contributors. * (**) (***) indicates significance at the 10 (5) (1) percent level.

results in column (4) reveals that in large projects the role of the owner is as central as in smaller projects. In larger projects the owner is much less likely to follow others, which given the larger team size seems to be intuitive. The distinction between large and small projects is important because in our data such projects contribute disproportionately more to non-local links. More specifically, in teams with 2 to 5 members, links to local members account for 60% of all links, while in teams with more than 100 members such links account for only 15% (see Table A1). In the last column of Table A4, we regress the share of follower links on the owner dummy. Again we obtain a very large and precisely estimated positive coefficient.

In Figure A1, we provide further evidence that within teams a few individuals attract disproportionately more connections than all others. In this figure the blue line shows the correspondence between the share of followers and the share of projects by the top individual. More specifically, the figure shows that in almost one-quarter of projects the top individual gets 100% of all follower links. If we interpret the following as a proxy for interactions, this suggests that in a quarter of projects there are no horizontal interactions between other members. Moving further along this line we see that in over 40% of projects the leading individual gets 50% of all links.[16] The other lines under the blue one show the same relationships for individuals ranked from second to fifth in terms of the follower share received. The figure considers projects involving more than five members. Raising this

[16]We should emphasize that when the leading individual follows others, this also generates a follower link. That implies that even for follower shares below 100% there does not have to be horizontal interaction between project members that are not the leading individual.

threshold, the distance between the top individual and the subsequent members becomes larger.

Figure A1: The hierarchy of following structures in project teams

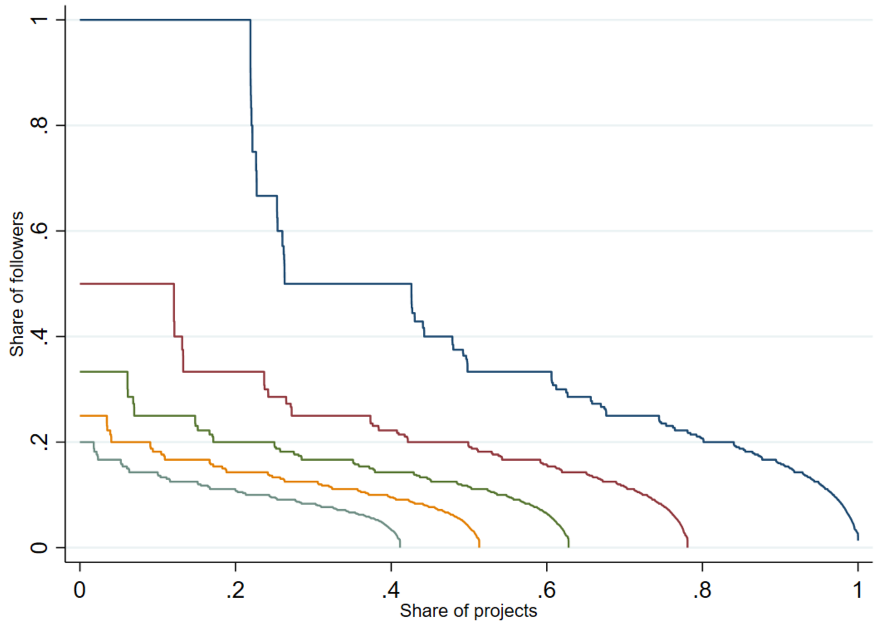

The figure plots the cumulative distribution of the share of followers within projects held by the top 5 team members. The line at the top corresponds to the individual with the highest follow share; the lines below show the follow share of the 2nd, 3rd, 4th and 5th most followed individual.

When constructing trade flows, a key decision that we need to make is whether code generated by a person in a given city flows to all other locations from which the project has members, or whether it flows to the city of the owner. Our results presented in Table A4 and Figure A1 provide strong support for the latter approach. Assuming that the code flows to all other cities will vastly exaggerate trade flows because, as suggested by our analysis, many team members do not interact with each other and work independently. To make this more intuitive, we can consider the following example from commodities trade. Imagine that a Chinese phone assembly plant imports separate components from Japan and South Korea. All three countries are thus part of the same supply chain, but the trade volumes generated by this production process do not directly affect bilateral trade between South Korea and Japan, even if all three production units are part of the same multinational company.[17]

[17] In a parallel paper, Goldbeck (2023) is interested in estimating the distance elasticity within the US. The author assumes that every member of a project interacts with every other member in a symmetric way. It is not surprising that, under this assumption, the author obtains a zero distance elasticity. Additionally, the author uses dummy variables at city-pair level, which ignores the intensity of the collaboration both at individual level and how many individuals collaborate between a city-pair. Our approach takes care of the intensive margin.

A.3 Location productivity measures from importer fixed effects

Here we describe a specification in which we recover city-specific productivities from importer fixed effects. In this case we no longer assume that software developers supply labor at a constant marginal disutility. Instead following the conventional model we assume that software developers (workers) supply labor at wage w_i at importing location i. Then the productivity at the city level is given by:

$$T_i = \left(\frac{FE_i}{FE_{SJ}}\right)\left(\frac{w_i}{w_{SJ}}\right)^{\theta} \tag{A3}$$

where subscript SJ denotes San Jose, which we use for normalization. Following Waugh (2010) we set $\theta = 0.18$. Because IT specialists' wage data is not globally available at the location level, we construct an approximation utilizing both the ACS and Stack Overflow survey data. To this end, we regress population numbers on average hourly wages for US cities from the ACS data to establish a relation between city size and software developers' average wages. We then estimate country level average hourly wages of software developers by dividing the Stack Overflow country level average yearly compensation of software developers by the average number of hours worked by these IT specialists also from the ACS data, implying that the number of hours worked are uniform across countries. Further assuming that the city-size and wage relationship is constant across countries, we calculate the location level wages as:

$$w_i = \beta_{ACS} * pop_i + w_c \tag{A4}$$

where w_c is the country level wage component from the Stack Overflow survey data, β_{ACS} the coefficient from the wage and city size regression and pop_i the population size of city i.

Figure A2 presents a scatter plot of our productivity estimates based on exporter fixed effects against the one based on importer fixed effects with wages. There is a tight fit between both measures with a correlation coefficient of 0.9. Note that the sample is restricted to locations for which both an importer and exporter fixed effect can be derived and to countries for which we have data from the *Stack Overflow* survey.

Figure A2: Correlation of productivity measure derived from importer and exporter fixed effects

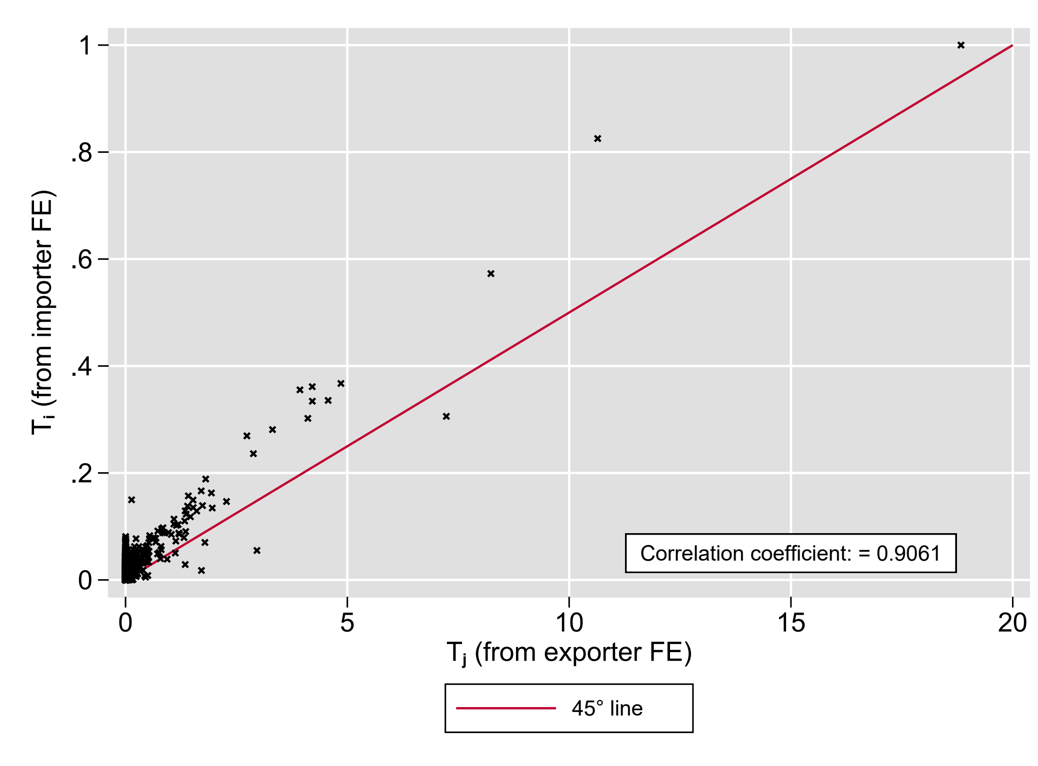

The figure shows a scatter plot of productivity parameters, where the y-axis marks the values derived from importer fixed effects and the x-axis the values from exporter fixed effects. The correlation coefficient between both values is 0.91.

A.4 Validation

We take two steps to validate our estimated measures. First, we compare our productivity measure with wages. Second, we use our data and construct university rankings and compare them with such rankings from other sources.

Using wages to proxy productivity In the absence of direct measures of productivity, one solution is to use the wages of software developers, which are closely related to productivity, especially in an industry where the share of labor is high.

We begin by restricting our sample to the US and regress our productivity measure on the wages of IT specialists in US cities. Our wage data come from the ACS, as described in Section 1.2. The results are displayed in panel (a) of Figure A3. We observe that both variables move together, also indicated by a significant correlation coefficient of 3.02. In panel (b) of Figure A3 we explore the relationship between our measure and the wages of software developers around the world. The wage data are constructed from a survey conducted by *Stack Overflow*. The data are at the country level, so we need to aggregate our productivity measures as well. To this end, we use the share of GitHub users of each location within each country and construct user weighted aggregate productivity at the country level. We restrict the sample to countries with multiple locations to reduce the influence of outliers, however the results are robust to using all countries. For this specification we also observe a positive relationship between our aggregated productivity measure and wages of software developers across countries. Clearly, the survey data have limitations, but both results together lend credibility to our estimated productivity measure. The advantage of the survey is that it covers many countries around the world, while the advantage of the US data is that they come from an official source and are less likely to suffer from selection bias.

Comparing university rankings We take advantage of information on the reported affiliations of users. Using this information we construct a ranking of universities. This approach is similar to *Approach 2*. However, instead of aggregating individual scores at the city level, we aggregate individual scores at the university level. More specifically, we identify university affiliated users for the US, the UK and Germany, and sum their individual scores for the identified institutions. Table A5 below lists the top 35 universities that emerge from this approach.

This exercise bears some similarities to the recent paper by Martellini et al. (2024), who use data from the website Glassdoor to construct university rankings. We should emphasize that our ranking is field-specific and includes computer science, mathematics, engineering and some other technical fields whose representatives are intensively involved in computer programming. Also, the ranking does not directly measure the quality of university graduates because individuals with a university affiliation can be faculty members,

Figure A3: Estimated productivities and IT-sector wages

(a) US FUAs productivity and IT-related professions' wages

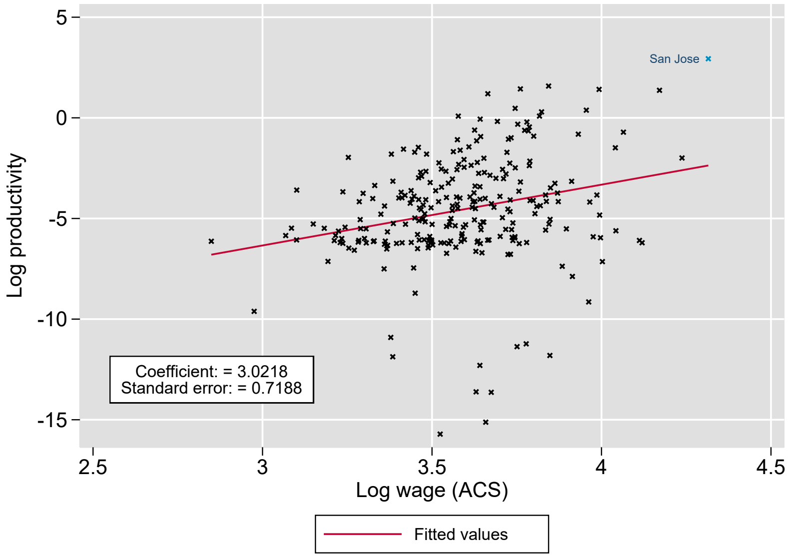

(b) User weighted productivity and IT wages country level

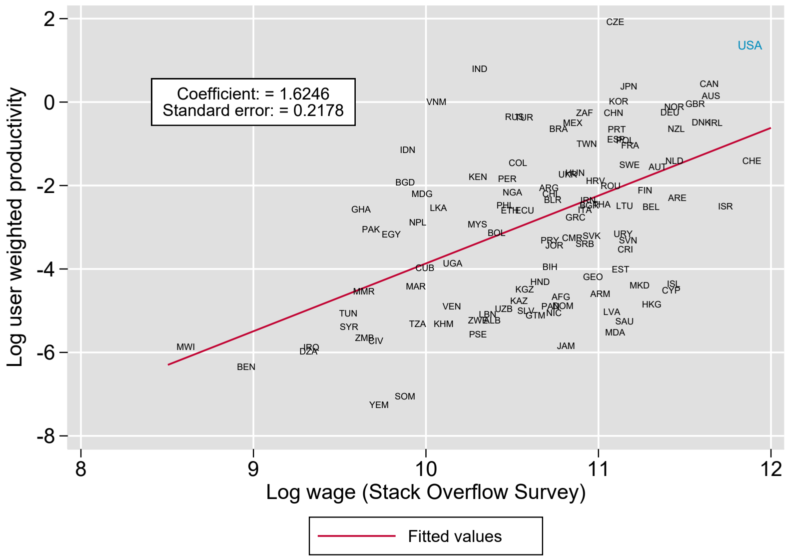

Panel (a) plots the relationship between log productivity estimated from the model and wages of IT specialists, constructed from the ACS, across FUAs in the US. Panel (b) plots the relationship between log productivity aggregated at the country level by applying user weights across locations within each country and wages of IT specialists from the 2023 Stack Overflow Developer Survey.

people working at university labs and students. Even if it only includes faculty members, it is still a valuable measure because it captures the knowledge and contributions of faculty to frontier software projects, which is an important input to the educational process. Importantly, these software projects have real life applications and commercial uses, so

Table A5: RANKING OF THE TOP 35 UNIVERSITIES IN THE US, THE UK AND GERMANY

Rank	**University**	**Rank**	**University**
1	MIT	19	Northeastern University
2	University of California, Berkeley	20	University of Saarland
3	Carnegie Mellon University	21	Columbia University
4	University of California, Los Angeles	22	University of California, San Diego
5	Stanford University	23	University of Duesseldorf
6	University of Oxford	24	University of Applied Sciences Munich
7	Vanderbilt University	25	Arizona State University
8	Technical University Berlin	26	Harvard University
9	University of Wisconsin-Madison	27	Brown University
10	Johns Hopkins University	28	Purdue University
11	University of Edinburgh	29	California Institute of Technology (Caltech)
12	University of Washington	30	University of California, Davis
13	Cornell University	31	Technical University Munich
14	Brigham Young University	32	University of Cambridge
15	University of Colorado Boulder	33	University of Hawaii
16	University of Arizona	34	University of Essen
17	New York University	35	University of Michigan
18	Washington University in St. Louis		

our measure does not capture some abstract theoretical knowledge.[18] Compared with the results of Martellini et al. (2024) our ranking is highly correlated with conventional rankings, such as the US News Best Colleges Ranking or the Academic Ranking of World Universities.[19] The fact that the university ranking produced from our data is so closely related to rankings produced by independent sources lends further credibility to our results and indicates that it is unlikely that our data suffers from systematic selection issues.

[18]From this point of view our exercise is also related to Bias and Ma (2023) who construct a distance measure between university course syllabi and academic articles to measure the "education-innovation gap".

[19]See http://www.shanghairanking.com/rankings/gras/2021/RS0210 for the 2021 ranking of universities regarding Computer Science and Engineering.

A.5 Additional figures

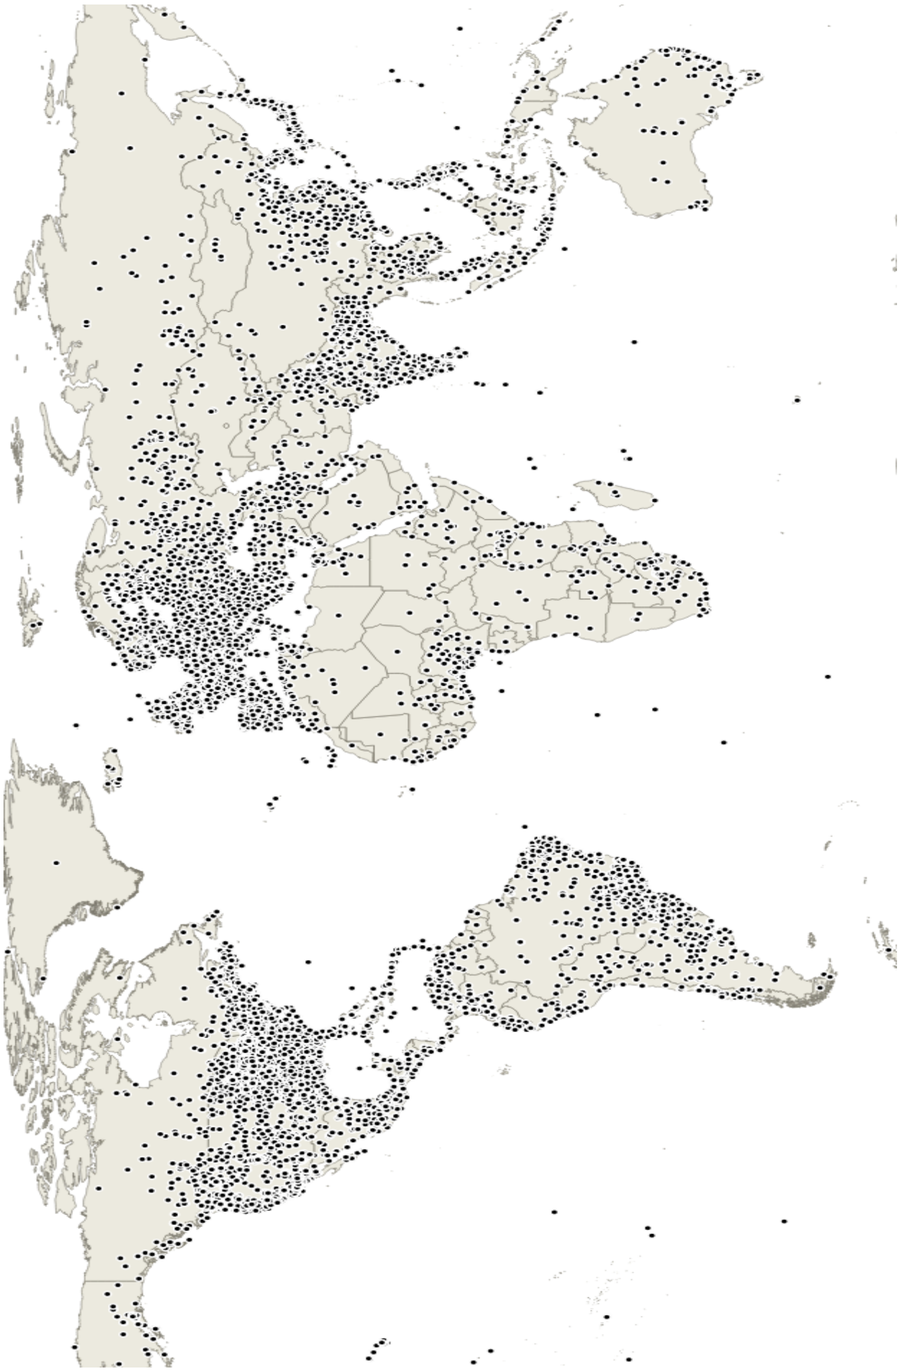

Figure A4: Visualization of GitHub users' locations across the world

Figure A5: Example of sample construction - nightlights (white shading), Functional Urban Areas (blue shading), GitHub users (red dots)

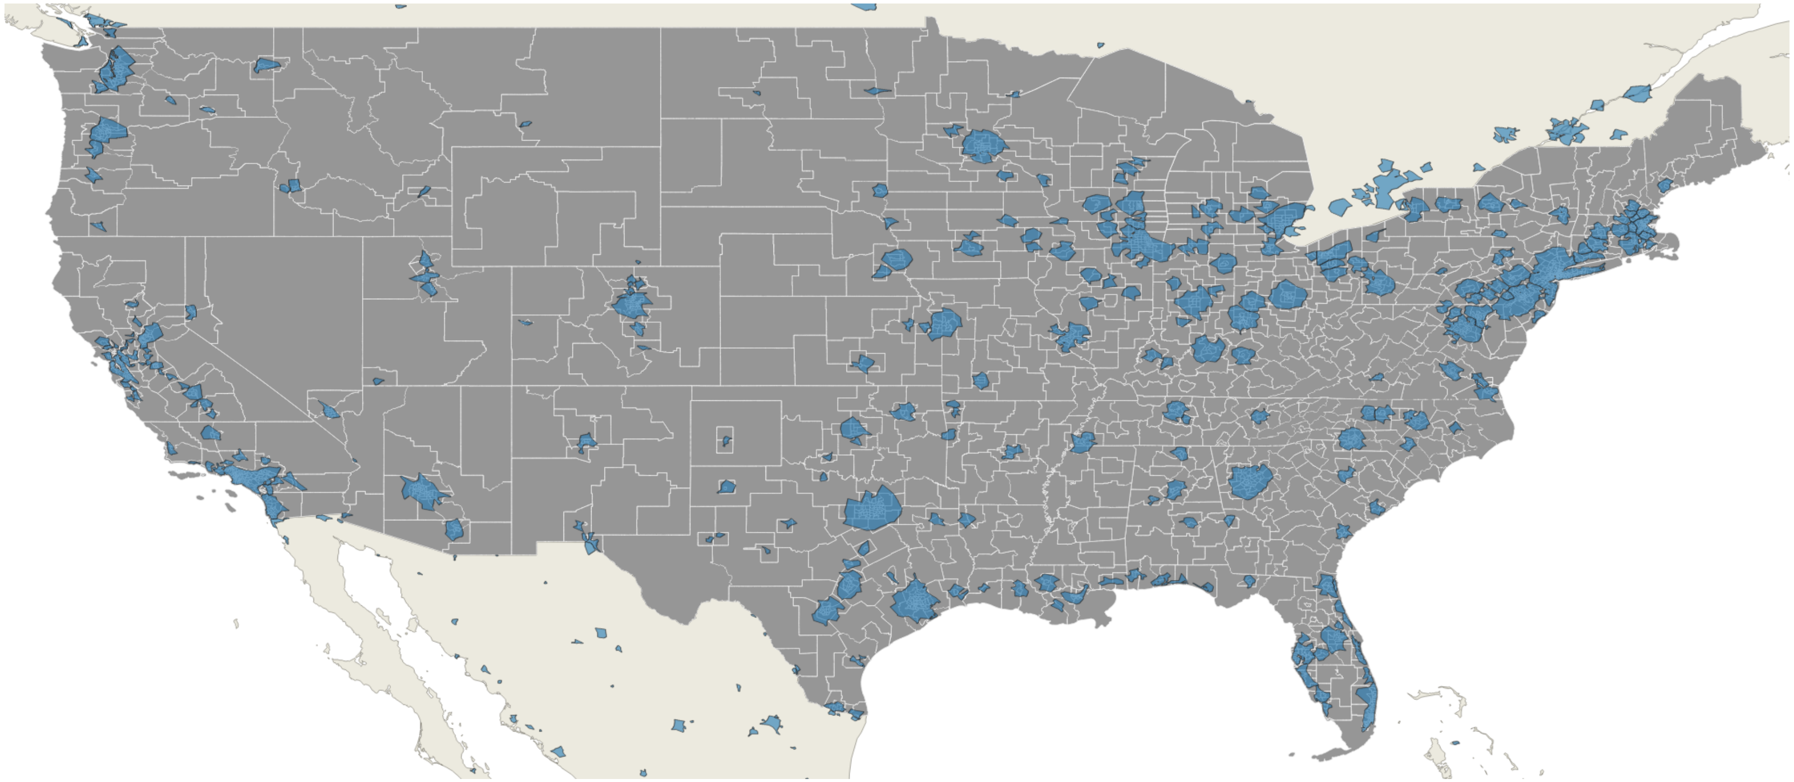

Figure A6: Visualization of the intersection of PUMAs and FUAs.

Figure A7: Bilateral migration flows

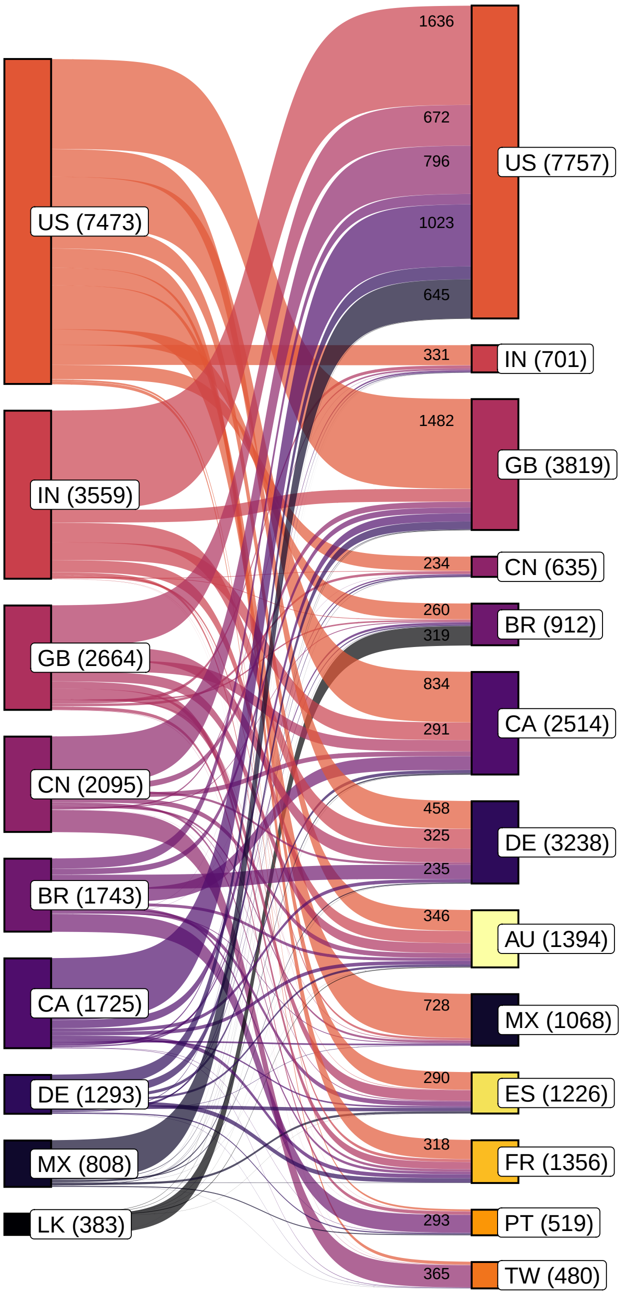

The figure presents bilateral migration flows between origin countries on the left side and destination countries on the right side. We selected all countries that send at least one flow of 200 or more migrants. For the largest individual flows the numbers in black represent the size of the flow. The numbers in brackets behind the country codes signal the total amount of migrants send or received by a country.

Favoritism and Firms: Micro Evidence and Macro Implications *

Zareh Asatryan † Thushyanthan Baskaran ‡ Carlo Birkholz §
David Gomtsyan ¶

Abstract

We study the economic implications of regional favoritism, a form of distributive politics that redistributes resources geographically within countries. Using enterprise surveys from low- and middle-income countries, we document that firms located close to leaders' birthplaces grow substantially in sales and employment after leaders assume office. Firms in favored areas also experience increases in sales per worker, wages, and measured total factor productivity. These effects are short-lived, and operate through rising government demand in the non-tradable sector. We calibrate a simple structural model of resource misallocation in a two-sector and two-region economy on our estimates. This exercise implies that, despite large firm-level effects, output losses caused by favoritism are small because leaders do not tend to redistribute funds towards less productive regions.

Keywords: Regional favoritism, firm performance, enterprise surveys, resource misallocation.
JEL: D22, D72, O43, R11.

*We thank Francesco Amodio, Xavier D'Haultfoeuille, Georg Duernecker, Andreas Fuchs, Leonardo Giuffrida, Patrick Hufschmidt and Konrad Stahl for feedback. We also thank Nikoloz Chkheidze for excellent research assistance, and Joshua Wimpey from the World Bank Group for sharing and helping us use the enterprise data. We acknowledge financial support from the German Research Foundation (DFG) within the Project "Regional Favoritism and Development" (Grant no. 423358188 / BA 496716-1).

†ZEW Mannheim
‡Ruhr University Bochum
§University of Mannheim, ZEW Mannheim
¶CREI

2.1 Introduction

Regional favoritism - that is, the geographical redistribution of resources within countries based on preferential political treatment - is a large phenomenon observed in many parts of the world (Hodler and Raschky, 2014). Many authors have claimed that such cases of distributive politics, which are especially prevalent in lower income and less democratic countries (Golden and Min, 2013), lead to distortionary economic policies that contribute to sustaining or even widening the income gap between high and low income countries.

However, the magnitude of the effect and the mechanisms through which distributive policies between regions generate aggregate inefficiencies have not been rigorously explored. The ultimate answer depends on the productivity levels of the recipient regions compared with the rest of the economy and the efficiency of firms that benefit the most within the recipient regions. If leaders divert too many resources to beneficiaries in their home regions, which happen to be unproductive, for example because the distributive policy is steered by private rent seeking motives or by the sole desire to hand out benefits to gain political support, then favoritism can lead to sizeable efficiency losses on aggregate. On the other hand, if the redistribution is such that it benefits productive firms, for example because leaders are able to target them given their informational advantages about home regions, then favoritism will not substantially diminish welfare.

To understand the aggregate implication of regional favoritism, we employ cross-sectional survey data from up to around 150,000 enterprises in 120 low and middle income countries, and utilize transitions of national political leaders which provide us with identifying variation in up to 33 countries. With these data and identification at hand, we document the firm-level effects of regional favoritism, trace the channels leading to these effects, and calibrate a simple structural model of resource misallocation in a two-sector and two-region economy to estimate the aggregate effects of favoritism.

Our first contribution is to document the existence of strong regional favoritism in firm outcomes using a difference-in-differences approach. Firms located around the birthplaces of political leaders are larger in terms of their sales and number of employees than firms located in other regions during the leaders' term in office. Exploiting information on the exact geo-location of firms, we show that these effects of favoritism are strongest in very close proximity to the leaders' birthplaces, and that the effects diminish by distance. In our baseline specification, we employ a country size and shape-specific distance measure to define treatment and find that treated firms have 14% higher sales and 8% more employees compared to control firms. For an average firm, these effects translate into $1.1 million higher sales and 6 additional employees. We show that these results are robust to several alternative definitions of the treatment area. Our placebo analysis does not find evidence for the existence of pretrends in firm outcomes, suggesting that the causality likely runs from leader changes to firm outcomes. Although we note that the cross-sectional nature of our data is not ideal for testing for such dynamic effects, the very local effects in close proximity to leaders' birthplaces that we identify make it unlikely that causality goes in the opposite direction, that is, local growing firms affecting leaders' choice. A further robustness exercise uses propensity score weights from random forest classification to balance out differences in many observable characteristics between treated and control firms in our cross-sectional data and confirms our baseline findings.

Second, we exploit the richness of our enterprise survey data and study the mechanisms that lead to these outcomes. We find that firms located in favored regions are not only larger in size but also produce more output per worker, pay higher wages, and have higher total factor productivity compared to other firms. Prima facie, this evidence suggests that regional favoritism may be considered as an efficiency enhancing policy. However, our further results indicate that the effects are driven by the non-tradable sector only, partly fueled by direct government transfers, and that they are temporary, fading away almost immediately after leaders leave office. This evidence goes in contrast to the hypothesis that favoritism induces general productivity improvements, since these should lead to more balanced growth in both the tradable and non-tradable sectors, as well as extend to the longer term (van der Ploeg, 2011). Additionally, we do not find evidence that any of the important correlates of productivity – exports, management practices, quality of inputs, or research & development activities – improve in firms located in favored regions, nor that the general business and regulatory environment – measured by firms' perceptions on business constraints – improves among these firms.

Overall, these results are consistent with the interpretation that leaders divert public resources to their home regions, thereby generating higher demand for output produced by firms operating in the non-tradable sector. This redistribution comes at the cost of other regions and is thus indicative of misallocation of resources.

As a final step, we set up a simple misallocation model in the spirit of Restuccia and Rogerson (2008). We use the model to quantify the aggregate implications of regional favoritism. We consider an economy with two regions and two sectors, where firms face wedges driven by favoritism. In our setting, redistribution between regions increases the level of income in the leader's region and thus demand. Since demand for non-tradable goods can be satisfied only by local production, factors of production reallocate towards the non-tradable sector in the leader's region and towards the tradable sector in the other region. This higher concentration of labor in the two sectors decreases the marginal productivity of firms and results in aggregate losses. We calibrate the model to match the moments that we estimate empirically. Our counterfactual exercise shows that in a country with spatial wedges driven by favoritism, output is 0.07% lower compared to a distortion-free economy. One of the reasons why output losses are small is that on average leaders' regions do not tend to be less productive than the rest of the economy.

Our paper is related to two strands of literature. First, we contribute to the evolving literature on regional favoritism. Miquel et al. (2007) were one of the first to develop a theoretical framework for favoritism and Hodler and Raschky (2014) were one of the first to document evidence for it. In particular, they use satellite data from across the globe and find higher intensity of nighttime light in the birthplaces of the countries' political leaders compared to other regions within countries. A closely related literature documents similar favoritism effects in political leaders' ethnic homelands.[1] Several papers extend the work on ethno-regional favoritism to specific sets of policies.[2] Our contribution is to study the effects of favoritism on firms, which allows to better understand the productivity implications of such distributional polices.

Second, our paper relates to the literature on how the misallocation of factors of production leads to differences in aggregate total factor productivity. This literature goes back to Restuccia and Rogerson (2008); Hsieh and Klenow (2009, 2010), and is surveyed by Hopenhayn (2014); Restuccia and Rogerson (2017); Martinez-Bravo and Wantchekon (2021). In this context, several studies have used enterprise survey data to estimate aggregate output losses caused by various institutional frictions (Ranasinghe, 2017; Besley and Mueller, 2018). Our contribution is to highlight a new source of misallocation that is driven by regional favoritism, which is caused by the endogenous concentration of production factors in tradable and non-tradable sectors in each region. Several related papers study efficiency losses caused by policy distortions in spatial contexts. Brandt et al. (2013) study China's economy in a model with multiple provinces, and private and state-owned firms. Desmet and Rossi-Hansberg (2013) introduce labor wedges to a model with cities to assess efficiency losses in the US and China. Fajgelbaum et al. (2018) use an economic geography model to estimate welfare losses caused by heterogeneity in tax systems across US states.

The remainder of the paper is structured as follows: Section 2.2 presents the data and our identification approach. Section 2.3 discusses our baseline empirical results as well as the robustness tests. In Section 2.4 we develop the mechanisms that drive our main findings. Section 2.5 sets up the quantitative model and calibrates it to arrive at aggregate implications. Section 2.6 concludes.

2.2 Empirical design

2.2.1 Data

Firms Our firm-level data are a repeated cross-section drawn from the World Bank Enterprise Surveys. The surveys have been conducted since 2006, and they span over

[1] De Luca et al. (2018); Dickens (2018) observe higher nighttime light intensity in political leaders' ethnic homelands, and Franck and Rainer (2012); Kramon and Posner (2016); Amodio et al. (2019); Asatryan et al. (2021b) find evidence for improved human capital outcomes among individuals belonging to either the same ethnicity, or coming from the same region as those holding political power.

[2] These policies include road building in Kenyan districts (Burgess et al., 2015) and Sub-Saharan Africa more broadly (Bandyopadhyay and Green, 2019), infrastructure projects in Vietnam (Do et al., 2017), school construction in Benin (André et al., 2018), enforcement of audits (Chu et al., 2021) and taxes (Chen et al., 2019) in China, mining activities in Africa (Asatryan et al., 2021a), and the allocation of foreign aid in Africa (Dreher et al., 2019; Anaxagorou et al., 2020), among others.

140 countries, of which 98 countries have been surveyed more than once. Among these countries, the survey is typically repeated in two to five year intervals, leading to an average of 2.5 survey waves per country. Firms are drawn by stratified random sampling, with stratification performed based on firm size, geographic location within the country, and sector of activity.[3] The surveys cover non-micro formal firms in the non-agricultural private sector. Thus, by design, they exclude firms that are fully government owned, are informal, have less than five employees, or are classified as agricultural firms. In general, our data will be representative of the manufacturing and service sectors, but not for the above mentioned sectors or firms.

The enterprise surveys contain information on general firm characteristics such as their age, ownership structure, and sector, as well as indicators of their performance in terms of sales, employment, and input factors. In addition, firms are asked about their management practices, relations to the government, crime and corruption, and the business environment. These latter aspects allow us to study the channels of how favoritism operates in greater detail.

For the main part of our empirical analysis, we consider the sub-sample of surveys carried out since 2009, as they provide us with the geocoded location of firms.[4] In additional specifications we use the general sample, where we can identify the location of firms on a regional level. We give priority to the smaller sub-sample of geocoded data to achieve greater precision, and to perform detailed spatial analysis, while we rely on the latter sample to test the robustness of our baseline findings on a larger sample.

Political leaders To identify political leaders in power we use the Archigos database of political leaders (version 4.1). The database includes information on the start and end date of the primary effective leader's time in power. Archigos data are available up to 2015 and we manually extend these data by including leaders from 2016 to 2020. We then utilize a plug-in that automatically parses a leader's birthplace to Google Maps' API, and retrieves the latitude and longitude of the city or town. We manually validate no matches or faulty matches that can arise due to cities sharing the same names, special characters in city names, or other reasons. We exclude any leader with less than one year of tenure.

We merge these data on leaders with the enterprise data by country. In the geocoded sub-sample we can calculate the distance of every firm to each leader's birthplace in the sample period. In the larger sample with regions as the spatial dimension, we generate a dummy indicating whether a firm is within a leader's birth region. In total, we have 250 leaders from 120 countries. Figure 2.1 plots the leaders' birthplaces and firms in a map. Since our empirical strategy builds on leader transitions, our identifying variation comes from a much smaller sample than the 250 leaders. First, as discussed above, the enterprise surveys have only been carried out 2.5 times within each country on average. Second,

[3]Further information on the sampling and stratification procedure can be found at https://www.enterprisesurveys.org/en/methodology.

[4]For data privacy reasons the latitudes and longitudes are precise within 0.5 to 2 kilometers.

Figure 2.1: Birthplaces of leaders and locations of firms in the sample

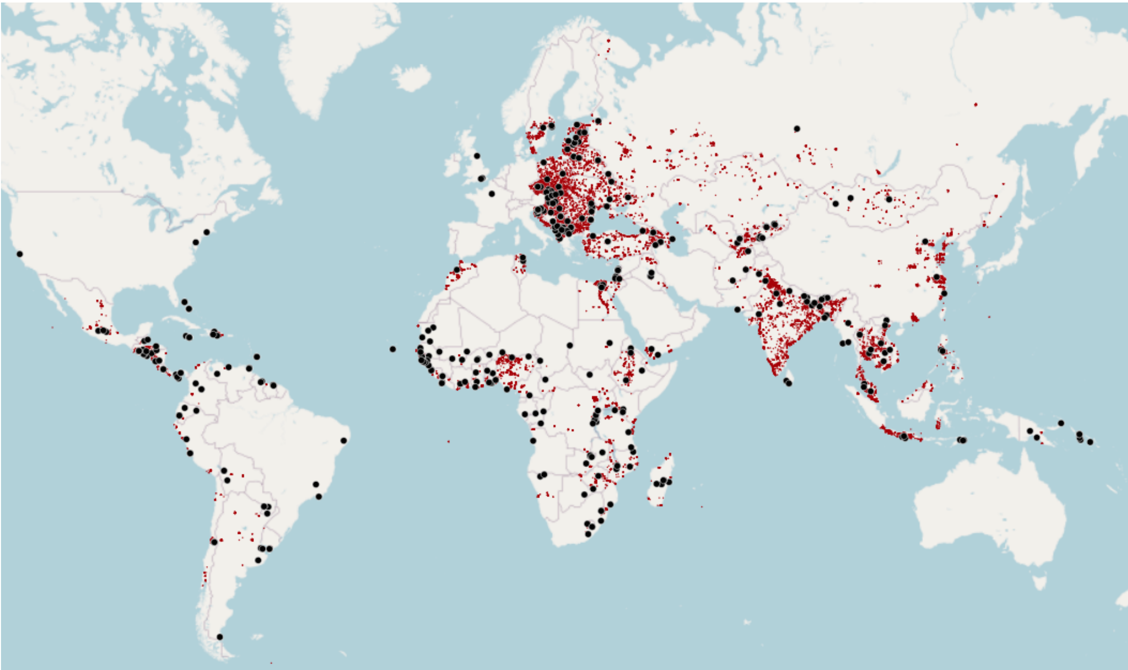

This map shows the geography of our sample. The small red dots represent firms, the large black dots leaders' birthplaces. Table B1 presents the list of countries and survey waves in our sample. There are around 25,000 African, 40,000 Asian, 20,000 European, 6,000 Middle American and 10,500 South American firms available in our main sample.

in many countries, especially less democratic ones, we do not observe leader transitions within our relatively short sample. Third, in cases where leaders were born in foreign countries, we do not identify any favored region. Taking into account these restrictions, our identifying variation comes from 25 countries in the baseline sample and from 33 countries in the regional sample.

Country characteristics In order to allow for comparisons between countries and for the interpretation of mean and aggregate values of monetary variables, we transform variables from local currency units to 2009 USD. For this transformation, we use period average exchange rates and GDP deflators from the World Bank's World Development Indicators. To study whether the effects of favoritism differ with respect to the political and institutional features of countries, we collect democracy index data from the V-Dem electoral democracy index, as well as data on perception of corruption from the World Banks Worldwide Governance Indicators.

Sample and summary statistics In total, there are around 100,000 and 150,000 enterprise surveys carried out in the geocoded and regional samples, respectively. However, the key variables we use have missing values to varying degrees. Additionally, to alleviate bias in our estimates from outliers, we exclude values that are outside three standard deviations of the calculated mean within an industry and country income level. For our baseline analysis, this leaves us with 82,000 to 94,000 firm-level observations, depending on the outcome we study. In the regional specification, we have between 126,000 and 142,000 observations.

Table B1 of the appendix lists the countries and survey years in our sample, with information on the number of firms per country and survey wave. We note the countries that contribute identifying variation in our two samples. Table B2 in the Appendix shows the summary statistics of the variables used in this paper.

2.2.2 Identification

Our empirical strategy exploits data on leader transitions and firm locations for identification in a difference-in-differences setup. We compare firms located in 'favored' areas in the sense that the current national leader was born in that area, to firms in the same area but in a time period when the current leader was not in office. Firms located in other non-favored areas but having similar observable characteristics, such as being in the same industry, serve as our control group.

2.2.2.1 The Spatial Extent of Favoritism

A central question of our empirical design relates to the spatial extent at which favoritism takes place. As we discussed in Section 2.2.1, our data measure the location of firms either by the geocoordinates of the firm, or by the region of their location as reported in the enterprise surveys. This lets us define treatment based on spatiality in a number of ways.

The first conceptual choice regards the use of regional boundaries versus the use of distances. We prefer the latter approach, utilizing the geocoded sample, as it allows us to study spatial effects around leaders' birthplaces at higher granularity and precision. This higher precision stems from two facts: First, depending on the shape and size as well as the location of the leader's birthplace within a region, treatment assignment based on regional boundaries will capture different firms. As an example, defining regional treatment for a leader born at the edge of an elongated region will assign firms in close proximity just across the border to the control group, whereas firms potentially far away on the other side of the region get assigned treatment status. Second, the region definitions within our data are not always consistent across time, and do not always coincide with administrative regions - which themselves might shift over time.[5] There are however two upsides to defining

[5] We can also manually map firms based on their geocoordinates into regional boundaries. We employ this technique in Section 2.3.3.1, where we overlay the world with a 0.5×0.5 degree grid layer and allocate the firms into cells. This enables the use of a stable granular region fixed effect and mirrors the geographic boundaries used in the prior favoritism literature with night lights. In the same vein, we could consistently map firms into administrative regions. However, both approaches do not alleviate the first drawback we mention, while removing the largest upside of the regional treatment definition, namely the larger sample.

treatment based on regions. Since we don't require information on the geocoordinates of the firms, we can utilize the full sample. Furthermore, if favoritism does take place through regional policies that are channeled at the level of administrative regions, this definition would capture the effect most precisely. Given this discussion, we start by studying firms whose exact geolocations are available, where we can identify treatment effects over granular distances. We then replicate our results on the larger sample with the regional specification to obtain complementary evidence.

Within the distance-based specification another conceptual question pertaining to the spatial extent of favoritism arises. Does favoritism travel equally far in countries of different sizes and shapes? For example, using larger distances for the treatment assignment might be adequate in a larger or less compactly shaped country, while the same distance might cover the majority of a smaller or more compact country. To address this issue, our baseline specification uses a country-specific distance measure that incorporates a country's shape and size. Inspired by Harari (2020), we construct the measure by overlaying each country's shape with a fine-grained point layer, and calculate the average Euclidean distance between a random selection of 10% of these points over 100 repetitions. The measure will vary across countries for two reasons: A larger country has a larger point layer, which on average leads to the randomly drawn points being further apart, thereby increasing the measure. The measure also increases the more a country's shape diverges from the most compact shape - a circle. We visualize these concepts in Figure B1.[6] An alternative approach to this is to define a fixed distance across countries, which we conduct as a robustness check.

2.2.2.2 Geocoded data

In our baseline we estimate a difference-in-differences model of the following form:

$$\begin{aligned} log(Outcome_{f,i,r,c,t}) =& \alpha + \beta^{km_c} \cdot LeaderArea_{l,c}^{km_c} \times Term_{l,c,t} + \\ & \gamma \cdot Controls_{f,t} + \tau_i + \mu_l^{km_c} + \lambda_r + \eta_{c,t} + \epsilon_{f,i,r,c,t} \end{aligned} \quad (2.1)$$

where $Outcome_{f,i,r,c,t}$ is the logarithm of either total sales, or the number of permanent employees. Our unit of observation is the firm f belonging to industry i located in region r of country c in year t.

β^{km_c} is our coefficient of main interest. It identifies the average treatment effect as the interaction of the dummy variables $LeaderArea_{l,c}^{km_c}$, which turn on if a firm is located within a country-specific kilometer radius km_c to the birthplace of leader l in country c, and the $Term_{l,c,t}$ dummy that indicates whether leader l is currently in office. We described the construction of our country-specific distance measure in the previous section. It is scaled such that in the median we match the area covered by the 0.5 $\times$ 0.5 degree pixels

[6]This measure can be interpreted as the average length of all hypothetical journeys through the country. We incorporate neither the degree of urbanization nor the ruggedness of the country's terrain in the calculation of the measure to maintain clarity on what is being captured.

commonly used in the favoritism literature.[7] This is the case for a radius of approximately 31km, which amounts to $\frac{1}{11}$ of our measure. Our results do not rely on choosing this particular share, as we demonstrate in the Appendix B.3.

Since favoritism might directly, or indirectly through spill-overs, affect firms farther away than the specified distance, we choose a cut-off distance between the treatment and control area. Firms falling into this area are excluded, which allows us to minimize diluting the control group with firms close enough to still be somewhat affected by the treatment. We again utilize the country-specific distance measure, and exclude firms between the treatment distance of $\frac{1}{11}$ and $\frac{1}{7}$ of the measure. Around 3% of observations fall into this area. We again verify robustness to alternative choices of the cut-off distance in the Appendix C.

$Controls_f$ is a vector of firm-specific control variables including the age of the firm, and its ownership shares belonging to foreigners, or to the public sector. τ_i, $\mu_l^{km_c}$, λ_r and $\eta_{c,t}$ are industry, leader area, region and country-by-time fixed effects, respectively. The error term is captured by $\epsilon_{f,i,r,c,t}$. We cluster the error term at the level of treatment following the arguments laid out by Abadie et al. (2017), which in the baseline estimation amounts to leader area by year. In the Appendix Table B3 we show the robustness of the estimated standard errors under alternative clustering strategies.

2.2.2.3 Regional data

We also estimate a version of Equation (2.1), where the treatment is defined based on the birth region of the leader. The equation is as follows:

$$\begin{aligned} log(Outcome_{f,i,r,c,t}) = & \alpha + \beta \cdot LeaderRegion_{r,c} \times Term_{l,c,t} + \\ & \gamma \cdot Controls_{f,t} + \tau_i + \lambda_r + \eta_{c,t} + \epsilon_{f,i,r,c,t} \end{aligned} \tag{2.2}$$

where the treatment status of a firm is defined by $LeaderRegion_{r,c}$ which is a dummy variable indicating whether any national leader was born in region r or not.

2.2.2.4 Identifying assumptions

Our model compares firms located within areas or regions around leaders' birthplaces before and after leaders assume power, while controlling for firms belonging to the same industries but located further away from leaders' birthplaces. The main identifying assumption in this difference-in-differences setting is that the treatment and control groups follow parallel trends prior to treatment. In our case, this will be violated if, for example, faster developing regions are more likely to nominate a national leader.

We test this assumption in Section 2.3.1 by conducting an analysis that tests for effects

[7] We match on the median not the mean to not overweight particularly large or small countries. The value of each individual country is reported in the Appendix Table B1.

in leads and lags of the treatment variable. We do not find evidence that the outcome variables between treated and control firms are statistically significantly different from zero in the year preceding the nomination of the leader. This absence of significant pre-trends suggests no systematic bias coming from selection as long as the selection effect is captured by the observables, and assuming that the selection effect is homogeneous across regions, such that the average effect of the pre-trends does not mask potentially offsetting trends. Our evidence from this test is consistent with previous work that has used regional-level data to study regional favoritism and does not find evidence for the existence of pretrends (see, for example, Hodler and Raschky, 2014).

We further validate our baseline results by augmenting the baseline difference-in-differences design with a propensity score approach in Section 2.3.3.2. This exercise suggests that our results are driven neither by differential firm characteristics between treatment and control groups that potentially affect firm outcomes, nor by changes in the composition of groups over time in our repeated cross-sectional data. We also implement a permutation test in Section 2.3.3.3, which suggests that assigning placebo treatments randomly to areas across time and space only very rarely leads to similarly large treatment effects as the ones we find in our baseline.

Finally, we follow the literature on difference-in-differences design with heterogeneous treatment effects (de Chaisemartin and D'Haultfœuille, 2023; Roth et al., 2022), to verify the validity of our setup which involves multiple periods and variation in treatment timing. Given the inclusion of country-by-time fixed effects, and the availability of only few survey waves per country, our results are almost always obtained from comparing treated, and never or not yet treated groups within countries, rather than by making 'forbidden' comparisons between already-treated units. More formally, we execute the diagnostics command *twowayfeweights* by de Chaisemartin and D'Haultfœuille (2020) to investigate the issue of potentially problematic comparisons of early and late treated groups. The test reveals that the vast majority of ATTs receive positive weights, which sum to 1.08, a large multiple of the sum of negative weights of -0.08. This reassures the use of the standard two-way fixed effects estimation.

2.3 Empirical results

2.3.1 Baseline results

We start by studying the treatment effect of favoritism using the geolocation of firms. We present our baseline results in Table 2.1. The first column regresses log sales on the treatment variable and the fixed effects as well as key firm characteristics as control variables. The estimated coefficient is statistically significant and implies that firms located close to leaders' birthplaces experience a 14% increase in sales relative to firms in the other parts of the country. In the second column our dependent variable is the log total number of employees. Again, we observe highly significant positive effects of 8% on average. These

Table 2.1: Baseline results: treatment effects around leaders' birthplaces

	(1) Log Sales	(2) Log Employees	(3) Log Sales	(4) Log Employees	(5) Log Sales	(6) Log Employees
Treated area	0.1434***	0.0817**	0.1657***	0.0874**	0.1349***	0.0757**
	(0.0506)	(0.0366)	(0.0513)	(0.0390)	(0.0504)	(0.0363)
Year before treatment start			0.0462	-0.1080		
			(0.1598)	(0.1081)		
Year after treatment end					-0.2762***	-0.2111***
					(0.0836)	(0.0424)
Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Firm level controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	82,527	94,093	82,377	93,921	82,377	93,921
R-squared	0.6621	0.2647	0.6622	0.2648	0.6622	0.2648
F	200.2	228.6	160.5	181.9	163.3	189.3

The regressions are estimated using Equation 2.1. Treatment area is defined as $\frac{1}{11}$ of the country-specific distance measure. Dependent variables are specified in logarithms. The mean values of the dependent variables in levels are 7.6 million USD for sales, and 80 employees. USD is measured in 2009 nominal values. Firm level controls include firm age, the share owned by foreigners and the share owned by the public sector. Columns (3) and (4) include a dummy that identifies the year before the start of treatment, and columns (5) and (6) a dummy for the year following the treatment end. Sample size changes are due to varying availability of the outcomes as well as the fact that we exclude rare cases in which observations are prior to and at the same time post treatment, which can occur if of three consecutive leaders the first and third are born in close proximity. The results hold when we fix the sample. All regressions include fixed effects for leader circles, regions, industries, and country-by-years. Standard errors are clustered at the level of treatment.

effects represent a sales increase of \$1.1 million, and an employment increase of 6 workers for an average firm.

The magnitudes of the effects are substantial. Taking into account the number of firms operating in these areas, and the sum of their sales, we can calculate the aggregate effects of favoritism in our sample. The favoritism effect leads to an estimated aggregate sales increase of \$19.5 billion (in 2009 nominal USD). Hodler and Raschky (2014) calculate that leaders' regions have on average 1% higher GDP in the worldwide sample, but the effects can reach up to 9% in certain subsamples, such as in countries with weak political institutions.[8] We take their approach of mapping the effects on nighttime light to GDP

[8]Following Hodler and Raschky (2014), we study whether the effects of favoritism on firm sales are different across countries with different political institutions. In Table B5 we interact our treatment variable with the electoral democracy score from V-Dem, and with the measure of corruption control from the World Bank. We do not find a linear relation between these institutional measures and our treatment effect. However, when allowing for a quadratic relation, we find suggestive evidence for a concave relation. In autocratic settings, leaders with a very strong grip on power have little incentive to seek support through regional favoritism. Such incentives increase with more democratization, but eventually, as the level of democratic institutions are sufficiently developed to impose the necessary constraints, possibilities of excessive regional redistribution are eliminated. This result should be interpreted with caution, given that the identification of this interaction effect comes from variation across countries.

growth using the correlation coefficient of 0.8 between firm revenues and GDP growth, as estimated by Cravino and Levchenko (2017). In our case, the corresponding effect on the favored regions is 11% when transformed into GDP growth values. This estimate is larger than that of Hodler and Raschky (2014), but not implausible considering that our sample consists of many countries with weaker political institutions.

We conduct placebo estimations to ensure that our results are driven by leader transitions rather than existing trends in regions. Since we are using a difference-in-differences specification, we want to make sure that there are no pre-trends that potentially drive our results. We construct a placebo pre-treatment variable by assuming that the leadership transition took place the year prior to when it actually happened. We also create a post-treatment variable in the same fashion. We then re-estimate Equation (2.1) including these leads and lags. The results are presented in columns 3 to 6 in Table 2.1. The pre-treatment dummy does not correlate with firm sales or with employment in a significant way, confirming the prior literature's notion of a lack of systematic anticipation or other sources of pre-trends.

Due to the limited frequency of the firm-level data, we are however unable to identify these dynamic effects on an annual basis for longer periods.[9] We therefore go one step further and reassess whether the prior literature's finding of a lack of pre-trends, which is established with nightlight data, also holds for our setting and sample. We present the treatment effect of leader changes on nighttime luminosity for exactly the 25 countries and the time period that lends identifying variation in our main specification. The yearly frequency of the nightlight data enables us to plot an event study in Figure B3. The figure confirms the existing picture. There is no evidence of systematic pre-trends, but a sizeable significant treatment effect three years after the new leader first comes into power.

Table 2.1 additionally shows large negative effects in the year after treatment ends, indicating that the effects of favoritism on firm growth are not persisting after leaders leave office. This evidence suggests that regional favoritism does not serve as a 'big push' industrial policy, according to which large positive shocks can help firms to permanently change their growth trajectories (Murphy et al., 1989).

2.3.2 Spatial gradient of the favoritism effect

While we are agnostic about the exact area around the leaders' birthplaces which is affected by favoritism, we postulate the following hypothesis: The closer a firm is located to the leader's birthplace, the more likely it is going to receive favorable treatment. This means that for smaller distances we should estimate larger absolute point estimates at higher precision. However, documenting strong effects on only a small share of the economy might hold little aggregate implications. On the other hand, capturing favoritism at a

[9]Our data similarly constrain us from studying the question of whether favoritism increases with the years a leader is in office. In our case, variation in tenure would come from across rather than within leaders' tenure.

Figure 2.2: Treatment effects by distance to leaders' birthplaces

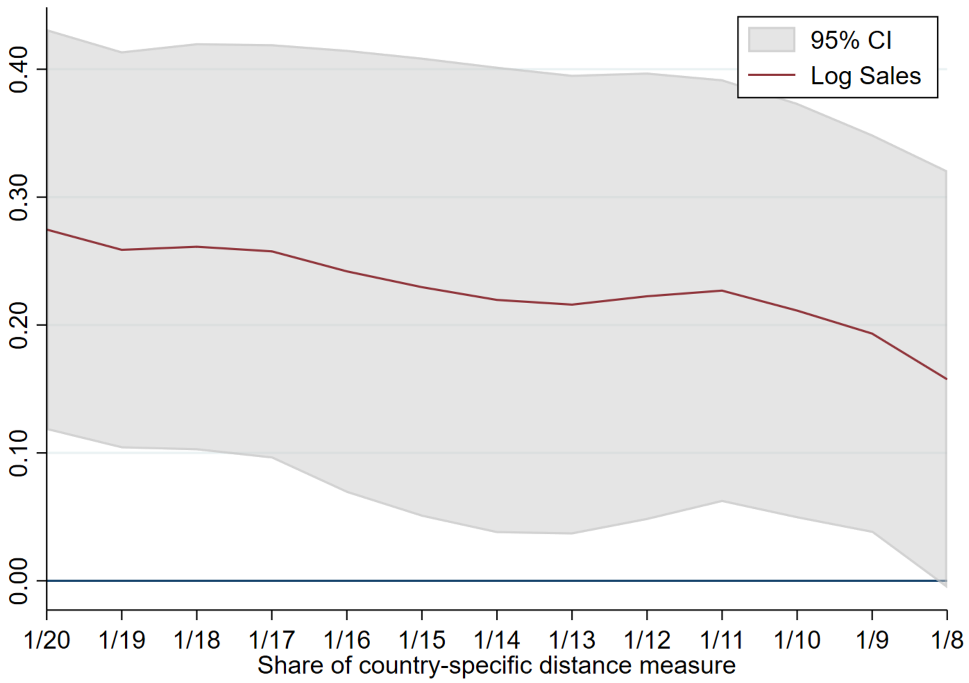

The regression is estimated using Equation 2.1. The red line plots the coefficient β^{km_c} estimated for each radius separately. The shaded area represents 95% confidence intervals. To keep the countries contributing to identification stable across the estimates we exclude five countries, which are marked in Table B1. The dependent variable is logarithm of total sales. All regressions include fixed effects for leader circles, regions, industries, and country-by-years. Standard errors are clustered at the level of treatment.

greater distance means a bigger share of the economy is implicated, but it will lead to more noisy results according to the hypothesis above.

We aim to test our hypothesis and reveal the spatial dimension of the treatment effect by running the estimation specified in Equation 2.1 over varying distances.[10] In Figure 2.2 we plot these treatment effects of favoritism over the varying distances to leaders' birthplaces for the logarithm of sales. As we hypothesized, we measure the strongest effects for areas very close to leaders' birthplaces, they decrease over distance, and become statistically not distinguishable from zero at a share of $\frac{1}{8}$ of our country-specific distance measure. In Figure B2 we document the same pattern of gradual decline for treatment defined by distances fixed across countries.

[10]The range of distances we display is informed by the necessity to keep the amount of countries contributing to the identification stable, such that changes in the effect size over distance can be interpreted as due to the spatial spread of favoritism, and not simply due to the sample composition changing.

Table 2.2: Alternative definitions of treated areas: regions, pixels, fixed radius

	(1) Log Sales	(2) Log Employees	(3) Log Sales	(4) Log Employees	(5) Log Sales	(6) Log Employees
Treated area	0.1308***	0.0609***	0.2378***	0.1667***	0.2139***	0.1404**
	(0.0389)	(0.0221)	(0.0447)	(0.0372)	(0.0749)	(0.0588)
Treatment Area defined by	Regions	Regions	Pixels	Pixels	50km Radius	50km Radius
Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Firm level controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	126,359	142,710	69,298	78,838	70,177	79,718
R-squared	0.6643	0.2626	0.6784	0.2833	0.6660	0.2582
F	654.4	826.9	34.97	347.1	129.0	148.4

The regressions are estimated with alternative definitions of the treatment area, each indicated at the bottom of the table. Dependent variables are specified in logarithms. Firm level controls include firm age, the share owned by foreigners and the share owned by the public sector. All regressions include fixed effects at the level of the respective treatment definition, regions, industries and country-by-years. Standard errors are clustered at the level of treatment.

2.3.3 Robustness tests

2.3.3.1 Definition of treated areas

As we discussed in Section 2.2.2.1, alternative choices in the definition of the treatment area are possible. In this section we assess the robustness of our baseline results by employing a number of these alternative specifications. We start by defining treatment by the region of leaders' birth rather than their exact birthplace. This allows us to utilize a substantially larger sample of firms, for which our data only indicate their location at the regional level. For the second alternative, we overlay countries' geographies with a fine grid layer of 0.5 x 0.5 degree pixels, and map firms into these grid cells.[11] Figure B4 of the Appendix visualizes this grid approach, which allows us to introduce granular pixel fixed effects to control for sub-regional time invariant confounding effects. The third approach we offer, fixes the radius of the treatment area to 50km for all countries, instead of relying on the country-specific measure. Table 2.2 collects the results. In all cases the evidence for positive and statistically significant effects is replicated.[12]

[11] At the equator 0.5 degree corresponds to roughly 55km. Results are also robust to a 1 degree specification.

[12] In an additional specification we interact the region treatment with the 50 km area treatment. Table B4 of the appendix shows the results. We find the strongest effects on firms that are located within a 50 km radius from the leader's birthplace, and at the same time belong to the leader's birth region.

Table 2.3: Comparison of baseline estimates with propensity score weighting estimates

	Treated Area		Observations	R-squared		F	
	Weighted	Unweighted		Weighted	Unweighted	Weighted	Unweighted
(1) Log Sales	0.2488***	0.2334***	69,352	0.6366	0.6639	8.846	193.6
	(0.0836)	(0.0707)					
(2) Log Employees	0.1096**	0.1405***	79,160	0.1145	0.2591	4.395	263.0
	(0.0523)	(0.0441)					

This table compares the treatment effects on the main outcomes estimated with unweighted (i.e. baseline) and weighted (propensity score) specifications of Equation 2.1. We restrict both specifications to the same sample. For the weighted specification, control variables are dropped, and instead the weights calculated according to Equation B1 are applied. The sample is trimmed to restrict the observations to the area of common support. Treatment area is defined as $\frac{1}{11}$ of the country-specific distance measure. Dependent variables are specified in logarithms. All regressions include fixed effects for leader circles, regions, and country-by-years. Standard errors are clustered at the level of treatment.

2.3.3.2 Propensity score weighting

Our difference-in-differences design leads to the identification of causal effects assuming that the group-specific pre-trends are parallel. Our analysis in the previous section did not find evidence for the existence of differential pre-trends. In this section, we provide a further robustness test by augmenting our difference-in-differences design with a propensity score approach. This exercise allows us to balance out observable differences between the treatment and control groups, thereby ruling out the possibility that the growth of firms in the treated area is driven by firm characteristics which differ systematically from the characteristics of control firms (Imbens, 2015).[13] This exercise also helps alleviate a second potential concern related to firm outcomes being driven by changes in the composition of the treatment and control groups over time. The sampling strategy of firm surveys is designed to make the data representative at the region level, such that, in principle, any compositional differences across the treatment and control groups over time would be the result of our treatment. However, given small sample sizes at the regional level, we nevertheless carry out this exercise.

One common shortcoming of this approach is that the choice of variables, as well as the functional form of the model used to calculate the propensity scores is under the discretion of the researcher. For this reason, we utilize the many firm characteristics available in our dataset in a data-driven machine learning approach. More specifically, we use random forests, an ensemble learning technique that averages the predictions of many individual decision trees, to calculate propensity scores (Lee et al., 2010; Zhao et al., 2016). We discuss the technical implementation of the random forest and calculation of

[13] An alternative approach is to include a long list of covariates. The advantage of our approach is that it is more data driven such that we do not need to take a stance on the importance of specific variables. Moreover, it allows for non-linear relationships between firm characteristics and outcome variables.

the propensity score weights in Appendix B.2. These weights help us make our treatment and control groups more similar in terms of the observable firm characteristics. Figure B8 shows the distribution of the standardized bias between the two groups before and after the application of the propensity score weights. The weighting shifts the distribution mass towards the center, indicating a substantial reduction in bias between the groups as captured by the observables.

In Table 2.3 we report the results of our difference-in-differences specification augmented by the propensity score weights. In order to draw comparisons to our baseline results, we re-estimate the baseline specification but restrict it to the same sample on which we run the weighted regressions. The two estimates are very similar in both size and precision for both outcome variables. These results reassure that our baseline results are neither driven by changes in the group composition across time, nor by differences in observable characteristics between the treatment and control groups.

2.3.3.3 Permutation test

We further address the direction of causality originating from leader transitions by conducting a placebo permutation analysis. Following Chetty et al. (2009), we perturb treatments randomly both across time and spatially. If leader transitions do drive the effects, we must see that they are a statistical rarity compared to the effects generated by the random permutations. To this end, we generate an empirical cumulative distribution function utilizing the grid-level estimation specification, and randomly assign each country with a treated pixel-year.[14] Originally treated observations and pixels with very few observations are dropped. We repeat this process to generate 5000 distinct estimates, and plot these in Figure B5 of the Appendix. The red line indicates the estimates of the correct treatment assignment on sales and employment for the grid-level specification. This exercise confirms that the result we find is indeed statistically rare. Furthermore, this test allows us to speak to the issue of serial correlation in difference-in-differences estimates raised by Bertrand et al. (2004). They state that, if uncorrected, serial correlation can lead to over-rejection of the null hypothesis in standard t-tests of difference-in-differences estimates. However, Figure B5 shows that, also in this non-parametric setting, the null hypothesis can be rejected at the 10% significance level.

2.3.3.4 Sensitivity of results to individual countries

We perform a jackknife-type exercise to test whether the average treatment effects we find are driven by strong favoritism effects emanating from individual countries. We re-estimate Equations 2.1 and 2.2, which are the regressions using geocoded and regional data, but successively dropping individual countries which provide identifying variation. Decreases (increases) in our coefficient of interest would indicate that the excluded country experienced a stronger (weaker) effect compared to the average country. Figure B6 of the

[14] Using the grid-level estimation has the upside of capturing equal sized areas for control and treatment groups over each permutation.

Appendix shows that changes to the average effects are small, and that they never lead to the average effect becoming statistically indistinguishable from zero. In specification 2.1 the largest change in the point estimate is not larger than four percentage points relative to the baseline effect, and in specification 2.2 this change is not larger than three percentage points relative to the baseline effect. Thus, we rule out that our findings are driven by individual countries.

2.4 Mechanisms

In order to shed light on the mechanisms behind our baseline results, we start by investigating whether the measured increases in sales and employment are accompanied by increases in productivity measures. We then assess whether regional favoritism affects the main sectors of the economy differentially. In the following sub-sections we study the role of government demand, of government regulatory policies, and of firm-level drivers of productivity in explaining our baseline favoritism effect.

2.4.1 Effects on productivity

From the information in the firm surveys, we construct three measures of firm productivity: Wage per worker, output per worker, and total factor productivity (TFP). We estimate TFP by regressing output in terms of sales on input factor costs and the net book value of land, buildings and machinery.[15] We then run Equation 2.1 with the residual from this regression and the other productivity measures as outcomes. Table 2.4 presents the results.

In Table 2.1 we found the size of the estimated coefficient for employment to be smaller than the coefficient for sales. Consistent with this, in columns (1) and (2) of Table 2.4 we find that firms in the treated area pay higher wages, and produce more output per capita. Column (3) shows that these firms not only grow in size, but also become more productive, as measured in terms of our revenue based total factor productivity measure. We are cautious of this last result, given that the estimate becomes statistically insignificant for many alternative choices of our distance measure as we show in Figure B12. Our further analysis of the mechanisms of favoritism in this section corroborates this caution, as we do not find patterns consistent with a productivity increase.

2.4.2 Sectoral results

We divide firms into the manufacturing and service sector. As we will discuss in Section 2.5, we expect redistributive policies implemented by the government to affect these two sectors differently. This is consistent with recent findings by Besley et al. (2021) who

[15]We sum up the costs for various input factors such as labor, raw materials, and intermediate goods, or electricity. As we use total sales as output in this regression, it constitutes as a revenue based TFP measure.

Table 2.4: Treatment effect on productivity outcomes

	(1) Log Wage	(2) Log Output per Worker	(3) TFP Residual
Treated area	0.0904***	0.0795***	0.0489*
	(0.0249)	(0.0236)	(0.0261)
Fixed effects	Yes	Yes	Yes
Firm level controls	Yes	Yes	Yes
Observations	77,946	81,735	68,318
R-squared	0.8293	0.7748	0.2988
F	40.38	61.36	77.50

The regressions are estimated using Equation 2.1. Treatment area is defined as $\frac{1}{11}$ of the country-specific distance measure. The mean values in levels are 7,420 USD in column (1), and 107,000 USD in column (2). USD is measured in 2009 nominal values. Firm level controls include firm age, the share owned by foreigners and the share owned by the public sector. All regressions include fixed effects for leader circles, regions, industries, and country-by-years. Standard errors are clustered at the level of treatment.

show that governments have less leverage to affect firms in the tradable versus the non-tradable sector. In particular, our model predicts that the non-tradable sector is likely to benefit more from redistributive policies. This prediction is similar and in line with the literature on the inflows of funds to developing countries from commodity booms, remittances, international aid, or borrowing. Such inflows increase household incomes, thus boosting consumption. The increased demand for tradable goods can be met by imports, while demand for non-tradable goods can only be satisfied with domestic production. Such episodes lead to relative increases in the prices of non-tradable goods (exchange rate appreciation), the reallocation of factors of production to the non-tradable sector, and deindustrialization. van der Ploeg (2011) provides a review of the resource curse literature and its implications. In a more recent study, De Haas and Poelhekke (2019) investigate the implications of natural resource booms and sectoral reallocation patterns while also using firm data from the Enterprise Surveys.

In Table 2.5 we include an interaction term between the treatment variable and a dummy variable for firms in the manufacturing sector. The results in column (1) show that manufacturing firms located around leaders' birthplaces benefit less from favoritism. Column (5) implies the same for measured TFP, in fact in favored areas the measured productivity growth is completely driven by service sector firms. Likewise we observe a large negative coefficient for output per worker in column (4), however it lacks the statistical precision to be deemed significantly different from zero. In column (3) we observe that wage growth is similar in both sectors indicated by the close to zero coefficient with a relatively small standard error. This result is consistent with the idea that there is high level of mobility of labor between the two sectors: Despite the fact that service sector

Table 2.5: Treatment effects by sector: manufacturing vs services

	(1) Log Sales	(2) Log Employees	(3) Log Wage	(4) Log Output per Worker	(5) TFP Residual
Treated area	0.2573***	0.1198**	0.0984***	0.1498***	0.1216***
	(0.0754)	(0.0546)	(0.0310)	(0.0478)	(0.0438)
Manufacturing	0.1602***	0.4058***	-0.1335***	-0.2403***	-0.2066***
	(0.0570)	(0.0363)	(0.0172)	(0.0703)	(0.0475)
Treated#Manufacturing	-0.1974*	-0.0772	-0.0111	-0.1089	-0.1364**
	(0.1031)	(0.0758)	(0.0334)	(0.0775)	(0.0673)
Fixed effects	Yes	Yes	Yes	Yes	Yes
Firm level controls	Yes	Yes	Yes	Yes	Yes
Observations	82,527	94,093	77,946	81,735	68,318
R-squared	0.6558	0.2500	0.8279	0.7696	0.2697
F	164.0	172.6	38.10	45.96	65.48

The regressions are estimated based on Equation 2.1, but include an interaction term between treatment and sectors. Treatment area is defined as $\frac{1}{11}$ of the country-specific distance measure. Dependent variables are specified in logarithms. Firm level controls include firm age, the share owned by foreigners and the share owned by the public sector. All regressions include fixed effects for leader circles, regions, industries, and country-by-years. Standard errors are clustered at the level of treatment.

firms experience higher growth, wage demands faced by firms in both sectors are similar, because both sectors compete for similar workers. In column (2) we document that there are no statistically significant sectoral differences in employment growth.

2.4.3 Government demand

In Table 2.6 we explore whether our baseline effect operates through the diversion of government demand towards firms in the favored regions. We consider the generation of additional government demand either through the public procurement system or through government owned firms more directly. Column (1) shows that firms located in proximity to leaders' birthplaces are 1.8% more likely than other firms to secure government contracts. The magnitude of this effect is substantial when compared to the mean probability of 17.8% of securing government contracts in our sample. In line with our sectoral results, column (2) presents evidence that this is driven by firms in the services sector. In columns (3) and (4), we then study whether sales and employment grow more in firms where the government has a partial ownership stake compared to privately owned firms. Our data provides weak evidence in support of this hypothesis. However, given that the Enterprise Surveys exclude firms which are fully government owned, we think about these estimates as lower bound effects. This interpretation will hold true as long as the government demand effect is more strongly present in firms fully rather than partially owned by the government.

Table 2.6: Government demand

	(1) Gov. contract secured?	(2) Gov. contract secured?	(3) Log Sales	(4) Log Employees
Treated Area	0.0179**	0.0337***	0.1387***	0.0745**
	(0.0077)	(0.0101)	(0.0507)	(0.0370)
Manufacturing		-0.0181**		
		(0.0079)		
Treated#Manufacturing		-0.0257**		
		(0.0121)		
Log employees	0.0285***	0.0302***		
	(0.0025)	(0.0026)		
Partial public ownership			0.8401***	0.7809***
			(0.2302)	(0.1232)
Treated#Partial public ownership			0.3879	0.4763***
			(0.2995)	(0.1745)
Fixed effects	Yes	Yes	Yes	Yes
Firm level controls	Yes	Yes	Yes	Yes
Observations	91,370	91,370	82,544	94,120
R-squared	0.1077	0.0975	0.6620	0.2650
F	49.68	44.84	173.5	206.5

The regressions are estimated using Equation 2.1, with logarithm of employees as an additional control variable to account for firm size. Treatment area is defined as $\frac{1}{11}$ of the country-specific distance measure. The mean values of the dependent variables in column (1) and (2) are 17.8%, in column (3) 7.6 million USD, and in column (4) 80 employees. Firm level controls include firm age, the share owned by foreigners and the share owned by the public sector. All regressions include fixed effects for leader circles, regions, and country-by-years, while (1), (3), and (4) also include industry fixed effects. Standard errors are clustered at the level of treatment.

2.4.4 Business environment

Next, we shift our attention to the supply rather than the demand side studied in the previous section, and investigate whether leaders use government regulatory policies to contribute to firm growth in their birth regions. The Enterprise Surveys ask questions regarding the constraints that firms face while doing business. Firms are asked to evaluate certain obstacles to their business on a five-point Likert scale. We center and normalize these variables and report the results in terms of standard deviations in Table 2.7.

In the first column, the dependent variable is the average of all business constraints. The estimated coefficient is positive and significant, indicating a worsening, not improving, business environment. In the following three columns, we study the more specific sources of business constraints. The results suggest that there is no change in the perceived institutional environment around leaders' birthplaces, but that the worsening business environment is driven by deficiencies in infrastructure and inputs.

On infrastructure, the result suggests that while leaders do divert resources to their

Table 2.7: Perceived business constraints

	Treated Area	Observations	R-squared	F
(1) Average	0.0850**	76,394	0.3959	11.41
	(0.0371)			
(2) Infrastructure	0.1190***	91,590	0.2907	17.50
	(0.0372)			
(3) Institutions	0.0140	79,775	0.3812	8.456
	(0.0436)			
(4) Inputs	0.0586***	88,522	0.2826	20.76
	(0.0226)			
(5) Land	0.0514**	90,616	0.2303	26.04
	(0.0200)			
(6) Finance	-0.0339	92,329	0.2022	49.78
	(0.0222)			
(7) Workforce	0.1251***	92,744	0.2330	26.82
	(0.0257)			

This table reports the treatment effect on firms' perceived business constraints. The regressions are estimated using Equation 2.1, with logarithm of employees as an additional control variable to account for firm size. Treatment area is defined as $\frac{1}{11}$ of the country-specific distance measure. Dependent variables are indices that have been centered at zero and normalized with a variance of one, with larger values indicating higher constraints. Average constraints in row 1 average the variable over business constraints related to infrastructure (2), institutions (3) and inputs (4). Input constraints are in turn an average over the constraints on land (5), finance (6), and workforce (7). All regressions include fixed effects for leader circles, regions, industries, and country-by-years. Standard errors are clustered at the level of treatment.

home region for example through generating higher government demand (Section 2.4.3), they do not promote sufficient infrastructure development to keep up with the increasing needs of the firms in these areas. This result is intuitive because infrastructure investments require planning and proper project implementation. Such activities require longer time horizons and more effort than, for example, simply awarding contracts to services firms in the favored areas. In this way, our results indicate that leaders are more likely to choose the latter option, or similar mechanisms to promote development in their home region. Infrastructure investments themselves can increase the incomes of local firms and workers, but do little to expand the infrastructure stock. Studies have shown that in the presence of limited absorptive capacity – in terms of skills, institutions, and management – countries are unable to translate every dollar of public investment into an additional dollar of capital stock (Presbitero, 2016).

On input constraints, the concept itself combines three components, the result for each of which are displayed in the last three columns of Table 2.7. From these regressions we observe that firms around leaders' birthplaces complain in particular about the lack of land and educated workforce, while the coefficient on the measure for access to finance is not significantly different from zero suggesting that leaders do not directly affect the capital market. The increasing complaints about lack of land make sense because this factor has

Table 2.8: Drivers of firm productivity

	Treated Area	N		Treated Area	N
Management			**Innovation**		
(1) Log Years of Manager's Experience	-0.0154	92,104	(6) R&D	0.0406***	74,303
	(0.0173)			(0.0151)	
			(7) New Processes	-0.0480*	63,214
Quality of Inputs				(0.0262)	
(2) % Workers with High School Degree	0.0101	68,945	(8) New Products	0.0613*	64,569
	(0.0133)			(0.0369)	
			(9) R&D controlling for new products	0..0225	69,136
(3) Formal Training	-0.0205*	93,450		(0.0171)	
	(0.0116)				
			(10) Technology Licensed from Abroad	0.0027	68,652
ICT Adoption				(0.0085)	
(4) Own Website	0.0089	93,698			
	(0.0089)		**Competition**		
(5) E-Mail Communication	0.0069	73,031	(11) Share Exports in Sales	0.0696	51,067
	(0.0122)			(0.5695)	

This table reports the treatment effects on firm's internal drivers of productivity. The regressions are estimated using Equation 2.1, with logarithm of employees as an additional control variable to account for firm size. Treatment area is defined as $\frac{1}{11}$ of the country-specific distance measure. All regressions include fixed effects for leader circles, regions, industries, and country-by-years. Standard errors are clustered at the level of treatment.

a fixed supply and does not increase proportionately with output, while the deteriorating perceptions about educated workforce suggest that the demand for labor exceeds the supply of skilled workers. This is consistent with increasing wage levels around leaders' birthplaces, as presented in Table 2.4. It is also worthwhile to note that, in the context of ethnic favoritism, Dickens (2018) shows that there is no increase in migration to the leader's ethnic region. It would therefore appear that adjustment is impaired by frictions to labor mobility. Specifically, tensions between ethnicities can be one factor hindering labor mobility within countries.

2.4.5 Drivers of firm productivity

Our baseline results show that firms located around leaders' birthplaces do not only grow in size, but that they also become more productive in terms of output per worker and measured TFP. However, given that both of these measures are based on nominal revenues, these measured productivity increases could be alternatively explained by increasing prices which we do not observe. Therefore, in order to better understand the question of whether, and if so how, favoritism leads to improvements in productivity, we adopt various drivers of firm productivity as comprehensively as possible, and test if firms located in favored areas improve on these measures.

We base our analysis on the review by Syverson (2011), and adopt ten measures from five broad categories of drivers of productivity. These are management practices, quality of inputs, adoption of ICT, research and development activities, and exports. Syverson (2011) also mentions that firm structure, and learning by doing effects can improve firm productivity, but we are unable to measure these components in our data.

Table 2.8 shows our estimates. Row (1) does not find evidence that firms in treated areas are managed by more experienced managers measured by the years of experience of working in the industry of the respective firms. Rows (2) and (3) study the role of firms' quality of inputs. There is no indication that firms in treated areas have a more educated workforce in terms of the share of workers with secondary school degrees, nor that these firms conduct formal training of their workforce. Rows (4) and (5) do not find evidence that firms in treated areas are more likely to adopt ICTs, as measured by firms having their own websites, or their use of emails when communicating with clients or suppliers. We then test the role of several variables measuring potential productivity improvements through innovation activity or adoption. In Row (6) we take note that firms in treated areas are significantly more likely to report any R&D expenditures than control firms. In rows (7) and (8), we study whether firms have introduced new products or processes. For new products we observe a positive and significant coefficient,[16] while for new processes a negative significant one. Our interpretation is that higher demand in the treated regions increases firms' incentives to introduce new products. However, this horizontal expansion does not necessarily imply improvements in efficiency, as process rather than product innovations are more likely to be associated with improved efficiency.[17] We then test in row (9) whether the increase in the likelihood to have reported any R&D expenses is driven by this vertical expansion of the firms' product portfolios. Indeed we find that controlling for the introduction of products that are new to the firm leads to an insignificant treatment effect on R&D.[18] In row (10), we do not find that firms in the treated area are more likely to adopt licensed technologies from abroad, which captures productivity improvements through technological diffusion from foreign countries. Finally, in row (11) of Table 2.8, we restrict our sample to manufacturing firms, and study whether they experience an increase in the share of sales coming from exports. Syverson (2011) warns that propensity of exporting is not necessarily a causal driver of productivity, but that it has been shown to be one of the most robust correlates of it. The direction of causality is not very important in our context, what is important is that this result, once again, does not show that firms in the treated area are more productive as far as productivity is correlated with export activity.

Given these null effects on this fairly comprehensive set of correlates of productivity, the explanation most consistent with our findings is that, despite the increases in measured

[16]This variable measures the introduction of products that are new to the firm, but not new to the market.

[17]For example, in the multi-product firm framework posited by Mayer et al. (2014) an exogenous increase in demand can lead the firm to expand its product scope without any improvement in productivity.

[18]This is not driven by the sample composition changing, as the treatment effect remains significant when restricting the sample to the subset with non-missing information on the introduction of new products but without including it as a control.

TFP, firms in fact do not become more productive. Instead, the treatment effects on our productivity measures rather reflect the change in local prices driven by the demand shock.

2.4.6 Size distribution of firms

In addition to the average effects of favoritism identified thus far, we are also interested in whether favoritism differently affects the size distribution of firms. Following Hsieh and Klenow (2009), in Figure 2.3 we present the distribution of firms in terms of total sales by plotting the approximated density of residuals from Equation 2.1 using Epanechnikov kernels. We separately plot the distribution of control and treated firms. If the favoritism effects were to change the distribution of firms, we would expect to observe substantial divergence in the density distribution of the two groups. This divergence is minimal, and therefore does not indicate a differential effect of favoritism across the size distribution of firms.[19] This result supports our assumptions in the following section, in which we model homogeneous firms.

Given that we have identified differential treatment effects for firms in the services and manufacturing sector in Table 2.5, we plot the distributions additionally for these sectors in Figure B7.[20]

2.5 Aggregate implications

Our empirical results are based on difference-in-differences estimations thus representing changes relative to the control group of firms in non-favored regions. As such we cannot draw conclusions regarding the aggregate effects of regional favoritism from these estimates directly. In this section, we propose a simple model that can fit the patterns we detected in the empirical section. In developing the model, we make a number of assumptions based on our empirical results, and as we describe the model, we motivate these assumptions by linking them to the related empirical findings. Then we calibrate the model to obtain a quantitative outcome for the aggregate economy resulting from the observed patterns of regional favoritism. Despite the fact that regional favoritism involves substantial reallocation of resources, we find that the aggregate impact is small. Our model includes minimal ingredients, but as we argue along the discussion of the quantitative results of the model, adding additional details will further mitigate the negative effect of favoritism.

[19] To test this hypothesis more formally, we use bootstrapping to construct a confidence interval of the ratio of the above mentioned residuals' standard deviations. The 95% confident interval of the ratio ranges between 0.980 and 1.001, thus suggesting that there are no statistically significant differences in the distributions between the control and treatment groups.

[20] Likewise the 95% confidence intervals of the bootstrapped ratios range from 0.960 to 1.006 for service sector firms and from 0.979 to 1.015 for manufacturing sector firms.

Figure 2.3: SIZE DISTRIBUTION OF TREATED AND UNTREATED FIRMS

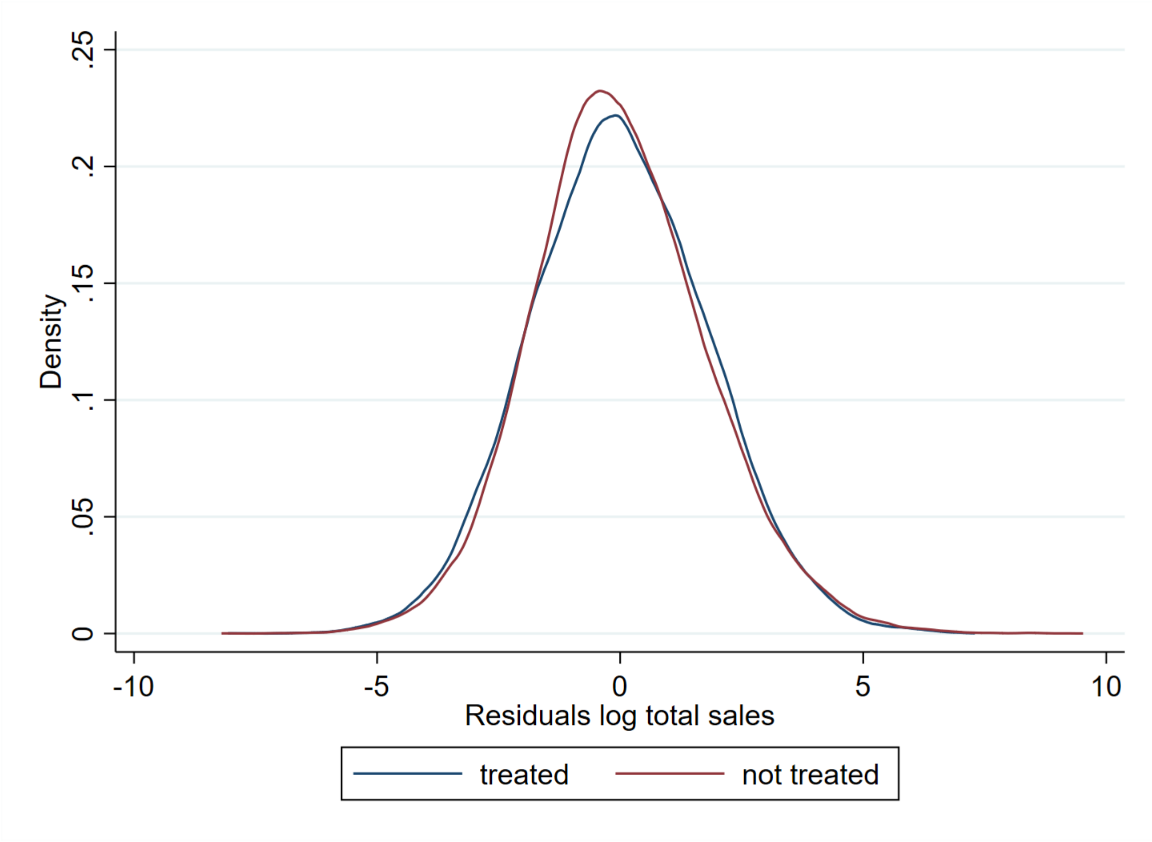

The figure plots the approximated density of residuals from Equation (2.1) with respect to logarithm of sales for the treatment and control group using Epanechnikov kernel estimator.

2.5.1 Framework

We consider a two-region and two-sector economy with perfectly competitive firms. Regions denoted $i \in \{h, a\}$ are the home region that receives subsidies τ_h and the rest of the country a that pays taxes τ_a to finance these subsidies. Positive values of τ_i denote taxes and negative values subsidies. We use the term taxes to refer to τ_i but this should not be taken literally because these taxes capture various wedges discussed by Restuccia and Rogerson (2008). Firms in both regions produce manufacturing goods (m) and services (s) $j \in \{m, s\}$. Manufacturing goods are traded across regions, whereas services are produced and consumed locally only. We assume that both regions use the same technologies with the same levels of productivity. Our data provide evidence in support of this assumption. We run regressions on outcomes that can proxy the average level of development (TFP, output per worker and wage), and include an indicator variable for areas which produced national leaders during the study period. Across all specification the estimated coefficients for the indicator variable turns out to be less than 10% and statistically not highly significant, which implies that the leader areas are not much wealthier compared to other places.[21, 22]

Production. We consider a representative firm that operates a simple production function

$$Y_{ij} = L_{ij}^{\alpha} \tag{2.3}$$

such that output Y_{ij} is produced using labor L_{ij}. Both regions are endowed with a fixed amount of homogeneous labor L_i, which is competitively allocated between sectors. Labor is perfectly mobile across sectors but immobile across regions. Our empirical results are consistent with a high level of labor mobility between sectors (Table 2.5), and low mobility between regions (Table 2.1). The mass of firms is proportional to the endowment of labor in each region. We do not introduce capital into the production function because our empirical results in Table 2.7 do not show any differential frictions in the capital market stemming from regional favoritism. Thus, to keep the model more tractable, we do not add capital. We will assume that the production function exhibits decreasing returns to scale ($\alpha < 1$), as in models with span of control. Another motivation for decreasing returns to scale assumption is that there are some fixed factors used in the production that do not adjust. This is consistent with the results in Table 2.7 where we showed that firms in favored regions perceive a worsening of the infrastructure and availability of land.

The firm's optimization problem can therefore be written as

$$\pi_{ij} = (1 - \tau_i) p_{ij} Y_{ij} - w_i L_{ij}, \tag{2.4}$$

where π_{ij} is the profit of the firm in region i and sector j, p_{ij} is the corresponding price and w_i the wage in region i. Perfect mobility between sectors implies that firms in both sectors face the same wage, for which we observe empirical evidence in Table 2.5 column (3). Since manufacturing goods are perfectly tradable between regions, their prices are the same in both regions and we normalize them to one ($p_{hm} = p_{am} = 1$)

Consumption. Both regions are populated by representative agents who derive utility by combining services (C_{is}) and manufacturing goods (C_{im}) given by $U_i = C_{im}^{\gamma} C_{is}^{1-\gamma}$. Agents earn wages by supplying labor inelastically, and receive the profits of the firms located in their region. The budget constraint is given by:

$$p_{is} C_{is} + C_{im} \leq w_i L_i + \pi_i, \tag{2.5}$$

[21] Our estimations include country-year fixed effects, and exclude observations for years and areas during which the respective leader was in office.

[22] At the end of this section we provide intuition on the outcomes if there were large differences between treated and non-treated areas in terms of productivity.

where π_i denotes total profits of firms in both sectors.

Market clearing. The equilibrium requires clearing in labor and goods markets

$$L_{hs} + L_{hm} = L_h, \quad L_{as} + L_{am} = L_a \tag{2.6}$$

$$C_{hs} = Y_{hs}, \quad C_{as} = Y_{as} \tag{2.7}$$

$$C_{hm} + C_{am} = Y_{hm} + Y_{am} \tag{2.8}$$

Finally, the government balances its books, which requires that the amount of tax collected in the non-home region equals the subsidies provided to the home region

$$\tau_h(p_{hs}Y_{hs} + Y_{hm}) + \tau_a(p_{as}Y_{as} + Y_{am}) = 0. \tag{2.9}$$

2.5.2 Discussion

The model yields several predictions that help us to understand the empirical results observed in Section 2.3. The key outcome of the model concerns the relationship between the tax rate and the relative allocation of labor between sectors. The model implies that the share of labor allocated to the services sector decreases with the tax rate.

$$\frac{\partial L_{is}}{\partial \tau_i} < 0. \tag{2.10}$$

Given that the home region receives a subsidy, and the non-home region pays taxes, this implies that a relatively larger share of labor in the home region will be allocated to the services sector. The intuition behind this result is rather simple. Since only the tradable good can be transferred across regions, the wedges introduced by the government require transfers from the non-home region. The relative supply of the tradable good in the home region increases because it receives transfers. As a result, it becomes optimal for firms in the home region to allocate relatively more resources to production in the services sector to meet consumer demand. Consequently, both regions will have relatively more resources allocated to one of the sectors compared to the economy without wedges. A concentration of resources in any of the sectors implies a lower level of marginal physical output in the presence of decreasing returns to scale. As a result, the implementation of taxes will generate aggregate losses in the economy.

Another prediction of the model concerns the effect of taxes on wages. Consistent with the empirical results documented in Table 2.4, wages decrease with taxes.

$$\frac{\partial w_i}{\partial \tau_i} < 0 \implies w_h > w_a. \tag{2.11}$$

In Section 2.3.1 we mentioned the possibility that regional favoritism can have long term effects in line with the 'big push' hypothesis. However, our model does not allow for such a possibility. The main reason for this is that our empirical results do not support this idea. It should also be mentioned that quantitative models of 'big push' are in their infancy (see Buera et al. 2021). Another important point is that we fit the data by varying wedges (τ) rather than the productivity terms of the production function which are fixed. In the latter case, it would be possible to obtain positive aggregate effects from regional favoritism. Again, the modeling choices are substantiated by the empirical results. Although in Table 2.4 we found some evidence in the increase in revenue-based TFP, the results in Table 2.5 show that this is completely driven by the non-tradable sector. Furthermore, in Section 2.4.2 we presented a series of results showing that the increases in sales in the favored regions are primarily driven by demand and not by productivity-enhancing activities. To be consistent with this evidence and fit the data, we model regional favoritism through changes in wedges rather than productivity terms.

2.5.3 Calibration

The qualitative discussion of the model predictions concluded that taxes generate net losses. In this section we use standard parameter values from the literature, and target some key moments from the empirical section to quantitatively asses the magnitude of taxation required to generate observed output differences, and to quantify associated output and welfare losses. We set the parameter governing the share of manufacturing goods consumption in developing economies to $\gamma = 0.30$ to generate an employment share of 30% in the manufacturing sector. As mentioned above, we assume that firms operate decreasing returns to scale technologies and set $\alpha = 0.85$ as in Restuccia and Rogerson (2008). We set the size of labor force in the home region (equivalently output in the undistorted economy) to 32% of total labor. This figure corresponds to the share of output produced by firms in the leader's region across our sample. Our key objective is to choose parameters τ_h and τ_a such that we can match the 14% total output increase in the home region, and make sure that the government's budget constraint (2.9) is satisfied. This value is taken from column (1) of Table 2.1.

Since both regions operate the same technologies, in the absence of wedges, both regions produce and consume exactly the same quantities per capita. In Table 2.9 we present changes in some key estimates relative to values for the economy without wedges. As already discussed, the share of labor allocated to the services sector in the home region increases. Quantitatively, this change is 10% (compared to 8% in our empirical results), while in the non-home region the corresponding figure decreases by -4.90%. Because in both regions labor is in fixed supply, the expansion of the services sector implies a decline in labor employed in the manufacturing sector, which is not consistent with our estimates in Table 2.5, where we did not find a decline in labor employed in the manufacturing sector. Introducing frictional labor mobility across regions, elastic labor supply, or rural-urban migration would allow us to address this issue. A model with these characteristics

Table 2.9: THE EFFECT OF DISTORTIONS ON FACTORS AND OUTPUT

	(1)	(2)	(3)	(4)	(5)	(6)
	L_{hs}	L_{as}	p_{hs}	w_h	Y	W
Changes in %	10.00	-4.90	5.56	14.47	-0.07	-0.04

The table displays the changes in percentages relative to the distortion-free economy. In column (5) Y refers to total output in the economy and in column (6) W refers to aggregate welfare in terms of consumption equivalents.

will mitigate the aggregate negative consequences of regional favoritism.

The following column displays the relative change in prices of non-tradable goods in the home region. There is a 5.56% increase in prices in the home region. In the data we do not observe prices and cannot compare them, but there was strong suggestive evidence that the price of non-tradable goods increases in treated circles. For example, in Table 2.5, we observed an increase in Y/L ratio only in the services sector. In our data, output is measured as price times quantity, and we do not have information on physical output. However, in Table 2.7 and 2.8 we did not find any supporting evidence for improvements in efficiency, so it is very likely that the Y/L ratio is driven by the increasing price of non-tradable goods. Column (4) displays the change in wages in the home region, which increase by about 14.5%. This figure exceeds our empirical estimate in column (1) of Table 2.4 but it is not far away. Additional features related to labor mentioned in the previous paragraph can improve the performance in this dimension. Overall, we find that this simple model performs relatively well in matching some key non-targeted moments. The fifth column displays the net loss in total real output, which amounts to 0.07% of annual output. In the last column we also report aggregate welfare changes, as measured in consumption equivalents.

Overall, despite substantial changes in output at the firm level, our model implies relatively small aggregate losses. Of course, our model is simple, but adding more features will not increase these losses because we have made a number of assumptions that work in the direction of generating output losses due to regional favoritism. For example, we assume a decreasing returns to scale technology, immobile labor across regions, and inelastic labor supply. Relaxing these assumptions would further shrink the negative effect of distortions on output and welfare. We also assumed that both regions have the same level of productivity. As explained above, our data provide weak evidence that the leader's region is slightly wealthier. If we incorporate this small difference and model migration with extreme value shock to location preferences, this can further decrease aggregate output losses. Output losses would have been greater if redistribution had been to less productive regions, but that scenario contradicts our data.

Another simplification is that we modeled an economy with a representative firm. Alternatively, we could add firm heterogeneity similar to Restuccia and Rogerson (2008), however, as we documented in Section 2.4.6 firm distribution is not affected by regional favoritism. If we were to model firm heterogeneity due to firm entry, aggregate losses would not increase. The general conclusion is that our simple model generates very small aggregate losses; however, given the empirical findings, adding the additional features discussed above would reduce rather than increase the aggregate losses.

2.6 Conclusions

Regional favoritism - that is, the geographic redistribution of resources within countries in favor of a political leader's home region - is a widespread phenomenon that is particularly prevalent in low and middle income countries. While evidence for regional favoritism has been extensively documented, its implications are not clearly understood. A commonly held normative view is that favoritism is necessarily a negative phenomenon that is fueled by corruption and other forms of rent seeking. However, preferential treatment of a region can also lead to higher welfare in the aggregate if, for example, leaders are well informed and are able to subsidize productive activities in the economy at the expense of more wasteful ones.

In this paper, we sought to solve this normative tradeoff by first identifying the micro effects of favoritism within a global sample of firms. We then quantified the macro effects of favoritism by feeding the estimated empirical parameters into a revised model of resource misallocation. Our empirical results suggest that firms located closer to leaders' birthplaces not only grow in size, but also become relatively more productive when measured by sales per worker, wages, and total factor productivity. While such improvements could potentially lead to higher growth for the entire country, this conclusion is not supported by our subsequent analysis. In particular, our evidence shows that this evolution of firms in favored regions is driven by a rapid expansion of the non-tradable sector, rather than substantial growth among manufacturing firms. Direct transfers to firms through public procurement contracts are one channel behind this effect. Importantly, these positive and economically substantial effects on firms are not sustainable and vanish after the leaders leave office.

We quantify that the net aggregate effects of the favoritism-based redistribution of resources between regions and sectors cost countries on average 0.07% of their output each year. We obtain a relatively small effect because on average leaders' home regions have similar levels of efficiency as the rest of the country. This means that resources are not redistributed towards less productive regions, which, if it were the case, would lead to larger aggregate losses.

Our results require several caveats. First, the regional favoritism we study may be an expression of various intentional and unintentional policies, including policies working on other forms of societal divides along ethnic, religious, or cultural lines. Future research could seek to disentangle the effects of these various policies. Second, owing to data constraints, we focus on leaders and ignore other systematically important national figures.

It would be potentially interesting to study regional favoritism in relation to other government figures. Third, future research could devote additional attention to the endogeneity of regions. Political leaders gain power often as a result of battles between complicated power structures, which may or may not reflect the underlying economic trends within specific regions. Although the evidence from our difference-in-differences framework assuages such concerns, our study remains a first pass. Fourth, we neglect the potential impact of favoritism on the entry and exit of firms, as well as its implications for firms in the informal and agricultural sectors. Since our survey data are not well equipped to explore these margins, future research may try to consolidate larger datasets, for example, from censuses or administrative sources, to better understand firm dynamics in general, and movements of firms and workers from informal and agricultural sectors more specifically.

References

Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey Wooldridge, "When Should You Adjust Standard Errors for Clustering?," 2017.

Amodio, Francesco, Giorgio Chiovelli, and Sebastian Hohmann, "The Employment Effects of Ethnic Politics," *CEPR Discussion Paper No. DP-14170*, 2019.

Anaxagorou, C, G Efthyvoulou, and V Sarantides, "Re-election Incentives, Ethnic Favoritism and the Subnational Allocation of Foreign Aid," *European Economic Review*, 2020, *127*, 1–32.

André, Pierre, Paul Maarek, Fatoumata Tapo et al., "Ethnic favoritism: Winner takes all or power sharing? Evidence from school constructions in Benin," *THEMA Working Paper No 2018-03*, 2018.

Asatryan, Z., T. Baskaran, C. Birkholz, and P. Hufschmidt, "Regional redistribution of mineral resource wealth in Africa," ZEW Discussion Paper 21-032 2021.

_ , _ , P. Hufschmidt, and A. Stöcker, "Regional favoritism and human capital accumulation in Africa," ZEW Discussion Paper 21-030 2021.

Bandyopadhyay, Sanghamitra and Elliott Green, "Roads and regional favoritism in Sub-Saharan Africa," Technical Report, LSE Working Paper No.19-195 2019.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-In-Differences Estimates?," *The Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Besley, Timothy and Hannes Mueller, "Predation, Protection, and Productivity: A Firm-Level Perspective," *American Economic Journal: Macroeconomics*, 2018, *10* (2), 184–221.

_ , Nicola Fontana, and Nicola Limodio, "Antitrust Policies and Profitability in Nontradable Sectors," *American Economic Review: Insights*, 2021, *3* (2), 251–65.

Brandt, Loren, Trevor Tombe, and Xiaodong Zhu, "Factor market distortions across time, space and sectors in China," *Review of Economic Dynamics*, 2013, *16* (1), 39 – 58.

Buera, Francisco J, Hugo Hopenhayn, Yongseok Shin, and Nicholas Trachter, "Big Push in Distorted Economies," Working Paper 28561, National Bureau of Economic Research March 2021.

Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padró i Miquel, "The value of democracy: Evidence from road building in Kenya," *American Economic Review*, 2015, *105* (6), 1817–51.

Chen, Yunsen, Jianqiao Huang, Hang Liu, and Weimin Wang, "Regional favoritism and tax avoidance: Evidence from China," *Accounting & Finance*, 2019, *58* (5), 1413–1443.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99* (4), 1145–77.

Chu, Jian, Raymond Fisman, Songtao Tan, and Yongxiang Wang, "Hometown favoritism and the quality of government monitoring: Evidence from rotation of Chinese auditors," *American Economic Journal: Economic Policy, forthcoming*, 2021.

Coppedge, Michael, John Gerring, Carl Henrik Knutsen, Staffan I. Lindberg, Jan Teorell, David Altman, Michael Bernhard, Agnes Cornell, M. Steven Fish, Lisa Gastaldi, Haakon Gjerløw, Adam Glynn, Allen Hicken, Anna Lührmann, Seraphine F. Maerz, Kyle L. Marquardt, Kelly McMann, Valeriya Mechkova,

Pamela Paxton, Daniel Pemstein, Johannes von Römer, Brigitte Seim, Rachel Sigman, Svend-Erik Skaaning, Jeffrey Staton, Aksel Sundtröm, Eitan Tzelgov, Luca Uberti, Yi ting Wang, Tore Wig, and Daniel Ziblatt, "V-Dem Codebook v11.1," 2021.

Cravino, Javier and Andrei A Levchenko, "Multinational firms and international business cycle transmission," *The Quarterly Journal of Economics*, 2017, *132* (2), 921–962.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–96.

_ **and Xavier D'Haultfoeuille**, "Difference-in-differences estimators of intertemporal treatment effects," *Review of Economics and Statistics*, 2024, pp. 1–45.

_ **and Xavier D'Haultfœuille**, "Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey," *The Econometrics Journal*, 2023, *26* (3), C1–C30.

Desmet, Klaus and Esteban Rossi-Hansberg, "Urban Accounting and Welfare," *American Economic Review*, 2013, *103* (6), 2296–2327.

Dickens, Andrew, "Ethnolinguistic favoritism in African politics," *American Economic Journal: Applied Economics*, 2018, *10* (3), 370–402.

Do, Quoc-Anh, Kieu-Trang Nguyen, and Anh N Tran, "One mandarin benefits the whole clan: Hometown favoritism in an authoritarian regime," *American Economic Journal: Applied Economics*, 2017, *9* (4), 1–29.

Dreher, Axel, Andreas Fuchs, Roland Hodler, Bradley C Parks, Paul A Raschky, and Michael J Tierney, "African leaders and the geography of China's foreign assistance," *Journal of Development Economics*, 2019, *140*, 44–71.

Fajgelbaum, Pablo D, Eduardo Morales, Juan Carlos Suarez Serrato, and Owen Zidar, "State Taxes and Spatial Misallocation," *The Review of Economic Studies*, 2018, *86* (1), 333–376.

Franck, Raphael and Ilia Rainer, "Does the leader's ethnicity matter? Ethnic favoritism, education, and health in Sub-Saharan Africa," *American Political Science Review*, 2012, *106* (2), 294–325.

Golden, Miriam and Brian Min, "Distributive Politics Around the World," *Annual Review of Political Science*, 2013, *16* (1), 73–99.

Haas, Ralph De and Steven Poelhekke, "Mining matters: Natural resource extraction and firm-level constraints," *Journal of International Economics*, 2019, *117*, 109–124.

Harari, Mariaflavia, "Cities in bad shape: Urban geometry in India," *American Economic Review*, 2020, *110* (8), 2377–2421.

Heckman, James J, Hidehiko Ichimura, and Petra E Todd, "Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme," *The review of economic studies*, 1997, *64* (4), 605–654.

Hodler, Roland and Paul A. Raschky, "Regional Favoritism," *The Quarterly Journal of Economics*, 2014, *129* (2), 995–1033.

Hopenhayn, Hugo A, "Firms, misallocation, and aggregate productivity: A review," *Annual Review of Economics*, 2014, *6* (1), 735–770.

Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, 2009, *124* (4), 1403–1448.

_ **and** _ , "Development Accounting," *American Economic Journal: Macroeconomics*, 2010, *2* (1), 207–223.

i Gerard Padró Miquel et al., "The control of politicians in divided societies: The politics of fear," *Review of Economic Studies*, 2007, *74* (4), 1259–1274.

Imbens, Guido W, "Matching methods in practice: Three examples," *Journal of Human Resources*, 2015, *50* (2), 373–419.

Kramon, Eric and Daniel N Posner, "Ethnic favoritism in education in Kenya," *Quarterly Journal of Political Science*, 2016, *11* (1), 1–58.

Lee, Brian K, Justin Lessler, and Elizabeth A Stuart, "Improving propensity score weighting using machine learning," *Statistics in medicine*, 2010, *29* (3), 337–346.

Luca, Giacomo De, Roland Hodler, Paul A. Raschky, and Michele Valsecchi, "Ethnic favoritism: An axiom of politics?," *Journal of Development Economics*, 2018, *132*, 115–129.

Martinez-Bravo, Monica and Leonard Wantchekon, "Political Economy and Structural Transformation: Democracy, Regulation and Public Investment," STEG Pathfinding Paper 2021.

Mayer, Thierry, Marc J. Melitz, and Gianmarco I. P. Ottaviano, "Market Size, Competition, and the Product Mix of Exporters," *American Economic Review*, 2014, *104* (2), 495–536.

Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny, "Industrialization and the Big Push," *Journal of Political Economy*, 1989, *97* (5), 1003–1026.

Presbitero, Andrea F., "Too much and too fast? Public investment scaling-up and absorptive capacity," *Journal of Development Economics*, 2016, *120*, 17–31.

Ranasinghe, Ashantha, "Property rights, extortion and the misallocation of talent," *European Economic Review*, 2017, *98*, 86–110.

Restuccia, Diego and Richard Rogerson, "Policy distortions and aggregate productivity with heterogeneous establishments," *Review of Economic Dynamics*, 2008, *11* (4), 707–720.

_ and _ , "The causes and costs of misallocation," *Journal of Economic Perspectives*, 2017, *31* (3), 151–74.

Roth, Jonathan, Pedro HC Sant'Anna, Alyssa Bilinski, and John Poe, "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature," *arXiv preprint*, 2022.

Stuart, Elizabeth A, Haiden A Huskamp, Kenneth Duckworth, Jeffrey Simmons, Zirui Song, Michael E Chernew, and Colleen L Barry, "Using propensity scores in difference-in-differences models to estimate the effects of a policy change," *Health Services and Outcomes Research Methodology*, 2014, *14* (4), 166–182.

Syverson, Chad, "What determines productivity?," *Journal of Economic literature*, 2011, *49* (2), 326–65.

van der Ploeg, Frederick, "Natural Resources: Curse or Blessing?," *Journal of Economic Literature*, 2011, *49* (2), 366–420.

Zhao, Peng, Xiaogang Su, Tingting Ge, and Juanjuan Fan, "Propensity score and proximity matching using random forest," *Contemporary clinical trials*, 2016, *47*, 85–92.

B Appendix

B.1 Additional tables and figures

Table B1: Sample description

Country	Year	Firms	Distance measure	Country	Year	Firms	Distance measure
Afghanistan	2008	535	617	Costa Rica	2010	538	153
	2014	410		Croatia†	2007	633	225
Albania†	2007	304	125		2013	360	
	2013	360			2019	404	
	2019	377		Czech Republic*	2009	250	205
Angola	2006	425	730		2013	254	
	2010	360			2019	502	
Argentina*	2006	1063	1286	Côte d'Ivoire	2009	526	347
	2010	1054			2016	361	
	2017	991		DRC	2006	340	1087
Armenia	2009	374	143		2010	359	
	2013	360			2013	529	
Azerbaijan	2009	380	201	Djibouti	2013	266	92
	2013	390		Dominican Republic**	2010	360	133
Bahamas	2010	150	242		2016	359	
Bangladesh	2013	1442	266	Ecuador	2006	658	383
Barbados	2010	150	12		2010	366	
Belarus	2008	273	307		2017	361	
	2013	360		Egypt	2013	2897	539
	2018	600			2016	1814	
Belize	2010	150	104	El Salvador*	2006	693	101
Benin	2016	150	273		2010	360	
Bhutan	2015	253	137		2016	719	
Bolivia**	2006	613	728	Estonia*	2009	273	136
	2010	362			2013	273	
	2017	364			2019	360	
Botswana	2006	342	517	Eswatini	2006	307	74
	2010	268			2016	150	
Brazil	2009	1802	1948	Ethiopia§	2011	644	752
Bulgaria	2007	1015	213		2015	848	
	2009	288		Gambia	2006	174	109
	2013	293			2018	151	
	2019	772		Georgia*	2008	373	197
Burkina Faso	2009	394	391		2013	360	
Burundi	2006	270	115		2019	581	
	2014	157		Ghana	2007	494	300
Cambodia	2016	373	267		2013	720	
Cameroon	2009	363	552	Guatemala	2006	522	226
	2016	361			2010	590	
Chad	2018	153	773		2017	345	
Chile	2006	1017	1330	Guinea	2006	223	349
	2010	1033			2016	150	
China	2012	2700	1327	Guinea Bissau	2006	159	123
Colombia*	2006	1000	810	Guyana	2010	165	336
	2010	942		Honduras**§	2006	436	272
	2017	993			2010	360	

Continued on next page

† Identifying variation in geocoded sample only.
* Identifying variation in both samples.
** Identifying variation in region sample only.
§ Dropped in Figure 2.2.

Table B1 – continued from previous page							
Country	**Year**	**Firms**	**Distance measure**	**Country**	**Year**	**Firms**	**Distance measure**
Honduras**§	2016	332	272		2019	150	
Hungary**	2009	291	220	Morocco	2013	407	720
	2013	310			2019	1096	
India	2014	9281	1350	Mozambique	2007	479	710
Indonesia	2009	1444	1713		2018	601	
	2015	1320		Myanmar*	2014	632	676
Iraq	2011	756	479		2016	607	
Israel	2013	483	480	Namibia	2006	329	663
Jamaica	2010	376	53		2014	580	
Jordan	2013	573	226	Nepal	2009	368	337
	2019	601			2013	482	
Kazakhstan	2009	544	225	Nicaragua**	2006	478	227
	2013	600			2010	336	
	2019	1446			2016	333	
Kenya*	2007	657	500	Niger	2017	151	784
	2013	781		Nigeria**	2007	1891	608
	2018	1001			2014	2676	
Kosovo†§	2009	269	499	North Macedonia*	2009	366	101
	2013	202			2013	360	
	2019	271			2019	360	
Kyrgyz Republic*	2009	235	373	Pakistan	2013	1247	781
	2013	270		Panama	2006	604	248
	2019	360			2010	365	
Lao PDR*	2009	360	457	Papua New Guinea	2015	65	526
	2012	270		Paraguay**§	2006	613	469
	2016	368			2010	361	
	2018	332			2017	364	
Latvia†	2009	271	182	Peru**	2006	632	804
	2013	336			2010	1000	
	2019	359			2017	1003	
Lebanon†	2013	561	70	Philippines**	2009	1326	642
	2019	532			2015	1335	
Lesotho	2016	150	120	Poland*	2009	455	339
Liberia	2017	151	219		2013	542	
Lithuania**	2009	276	173		2019	1369	
	2013	270		Romania	2009	541	323
	2019	358			2013	540	
Madagascar	2009	445	487	Russia	2009	1004	2918
	2013	532			2012	4220	
Malawi	2014	523	330		2019	1323	
Malaysia	2015	1000	889	Rwanda	2006	212	105
Mali	2007	490	890		2019	360	
	2010	360		Senegal	2007	506	281
	2016	185			2014	601	
Mauritania	2006	237	663	Serbia*	2009	388	203
	2014	150			2013	360	
Mexico**	2006	1480	1152		2019	361	
	2010	1480		Sierra Leone	2017	152	168
Moldova**§	2009	363	159	Slovak Republic*	2009	275	167
	2013	360			2013	268	
	2019	360			2019	429	
Mongolia*	2009	362	975	Slovenia†	2009	276	101
	2013	360			2013	270	
	2019	360			2019	409	
Montenegro*	2009	116	83	Solomon Islands	2015	151	281
	2013	150		South Africa	2007	937	756

Continued on next page

Table B1 – continued from previous page							
Country	**Year**	**Firms**	**Distance measure**	**Country**	**Year**	**Firms**	**Distance measure**
South Sudan	2014	738	595		2013	762	
Sri Lanka	2011	610	139	Ukraine**	2008	851	558
Sudan	2014	662	849		2013	1002	
Suriname	2010	152	236		2019	1337	
	2018	233		Uruguay	2006	621	261
Sweden	2014	600	564		2010	607	
Tajikistan	2008	360	310		2017	347	
	2013	359		Uzbekistan*	2008	366	624
	2019	352			2013	390	
Tanzania	2006	419	623		2019	1239	
	2013	813		Venezuela	2010	320	692
Thailand	2016	1000	579	Vietnam**	2009	1053	606
Timor-Leste	2015	126	106		2015	996	
Togo	2016	150	206	Yemen	2010	477	476
Trinidad and Tobago	2010	370	49		2013	353	
Tunisia	2013	592	299	Zambia*	2007	484	616
Turkey	2008	1152	608		2013	720	
	2013	1344			2019	601	
	2019	1663		Zimbabwe	2016	600	415
Uganda	2006	563	323				

Table B2: Summary statistics

Geocoded sample	N	Mean	Std. Dev.	p5	p95
Treated area	101350	0.16	0.37	0	1
Year before treatment start	101350	0.019	0.14	0	0
Year after treatment start	101350	0.011	0.10	0	0
Total sales in 2009 USD	87218	7597616	63214844	12045	24197024
Num. full-time employees	99707	79.6	223	5	320
Output per employee in 2009 USD	86300	106982	1622484	1154	258941
Wage in 2009 USD	82360	7420	53922	195	23362
TFP residual	72333	0.0095	1.38	-1.8	2.4
Firm age	100047	18.7	15.5	3	49
Firm share owned private foreign	100025	7.00	23.6	0	90
Firm share owned public	100070	0.68	6.61	0	0
Government contract secured?	98287	0.18	0.38	0	1
Avgerage of constraints	81644	31.6	20.5	1.7	68.3
Infrastructure constraints	98627	33.8	28.2	0	87.5
Institutional constraints	85401	30.3	22.6	0	70
Input constraints	95075	30.2	23.0	0	75
Obstacle land	97548	24.5	31.4	0	100
Obstacle finance	99345	34.1	32.0	0	100
Obstacle inadequately educated workforce	99788	31.9	31.2	0	100
Years of experience top manager	98826	18.0	11.2	3	40
Share employees completed high school	73101	0.65	0.35	0.02	1
Formal Training for employees	100383	0.38	0.48	0	1
Firm has own website	100995	0.53	0.50	0	1
Firm communicates via email	78932	0.75	0.43	0	1
Firm spent on R&D excl. market research	80057	0.22	0.41	0	1
New product / service last 3 years?	95133	0.36	0.48	0	1
New / improved process last 3 years?	93444	0.36	0.48	0	1
Firm licensed technology from foreign firm	74001	0.15	0.36	0	1
Share of sales: direct exports	99605	7.64	21.9	0	70
V-Dem electoral democracy index	101350	0.49	0.22	0.09	0.92
Scaled WB Control of Corruption percentile	100447	0.36	0.21	0	1
Region sample	N	Mean	Std. Dev.	p5	p95
Treated region	148593	0.16	0.37	0	1
Total sales in 2009 USD	129050	8121428	172838953	11797	23715758
Num. full-time employees	146365	77.6	214.5	5	306
Output per employee in 2009 USD	127761	129963	4156877	1187	246908
Wage in 2009 USD	123875	7475	63795	207	22143
TFP residual	109796	0.0084	1.31	-1.6	2.3

Table B3: Overview of results using alternative clustering approaches

	(1) LA-Y	(2) C-Y & LA-Y	(3) R-Y & LA-Y	(4) S & LA-Y	(5) C-S-Y & LA-Y
Log(Sales)	.0506	.0523	.0510	.0130	.0548
Log(Employees)	.0366	.0371	.0370	.0162	.0437
# of Cluster 1	556	198	890	46	877
# of Cluster 2		556	556	556	556

The table showcases changes to the main estimates' standard errors from Equation 2.1 using other clustering approaches. The nomenclature is as follows: 'C' stands for 'Country', 'S' for 'Sector', 'Y' for 'Year', 'R' for 'Region' and 'LA' for 'Leader Area'. Column (1) thus lists standard errors for clustering of leader area by year - our main specification for comparability.

Figure B1: Example how the country-specific distance measure varies with size and shape

(a) Circle, area ≈ 36 (b) Circle, area ≈ 72 (c) Rectangle, area = 36

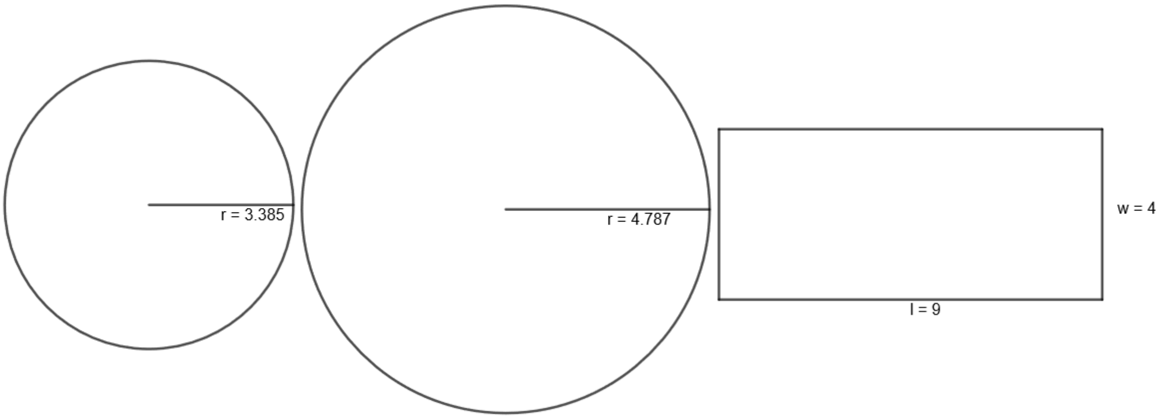

Distance measure = 278 Distance measure = 393 Distance measure = 310

The figure showcases conceptually how our country-specific distance measure varies across countries of different sizes and geographic outlines. We created stylized geographic forms to which we apply the same algorithm as described in Section 2.2.2.1 to calculate the measure. Moving from figure (a) to figure (b) we keep the same circular shape, but double the area, and consequently the distance measure increases substantially. On the other hand going from figure (a) to figure (c) we keep the area constant, but change the shape to a rectangle. The larger distance measure of figure (c) reflects the decrease in compactness.

Figure B2: Treatment effects by fixed distance to leaders' birthplaces

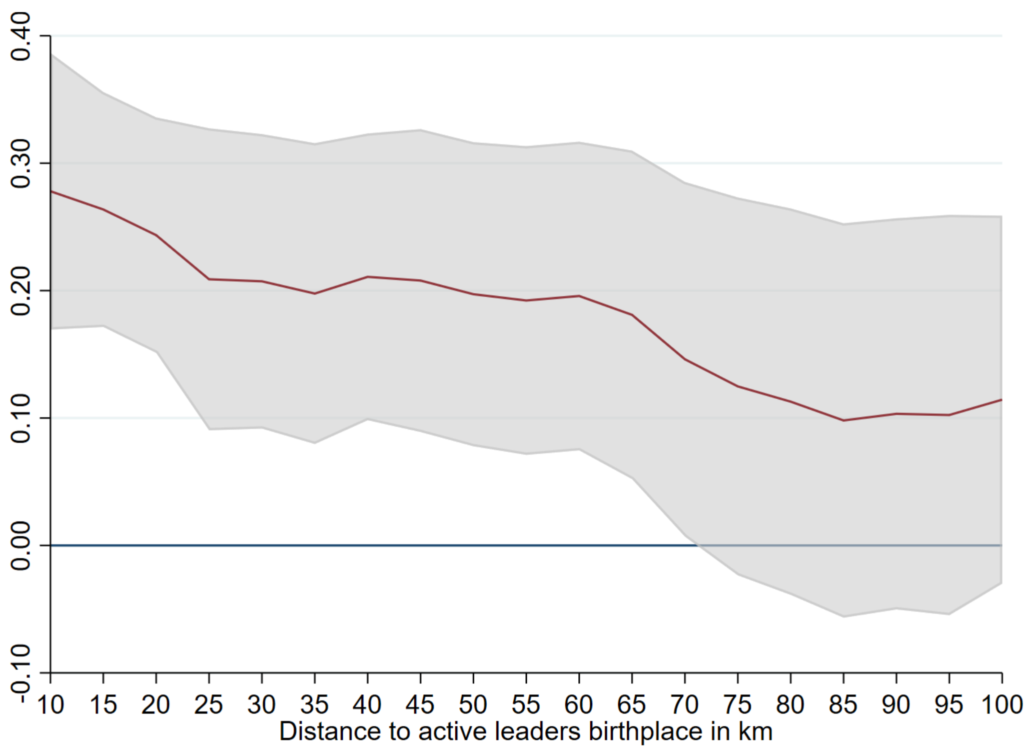

Logarithm total sales

In the figure, the red line plots the coefficient β^{km} estimated for each radius around the leaders' birthplaces' stated on the x-axis separately. Firms located in a circle of 10 km have on average nearly 30% higher sales than similar firms located further away. These effects decrease by distance, and become indistinguishable from zero beyond 70 km from leaders' birthplaces. The regression is estimated using equation 2.1. We drop eleven countries such that the estimates are identified by a stable set of countries over all distances. The shaded area represents 95% confidence intervals. The dependent variable is total sales and is specified in logarithm. All regressions include fixed effects for leader circles, regions, industries, and country-by-years. Standard errors are clustered at the level of treatment.

Table B4: SPATIAL VERSUS REGIONAL TREATMENT EFFECTS

	(1) Log sales	(2) Log employees
Treated area in leader region	0.1658*** (0.0533)	0.0967** (0.0389)
Treated area <u>not</u> in leader region	0.0390 (0.0817)	0.0126 (0.0640)
Fixed effects	Yes	Yes
Firm level controls	Yes	Yes
Observations	82,527	94,093
R-squared	0.6621	0.2647
F	161.3	186.1

The regressions are estimated using Equation 2.1. In this specification we interact the spatial and regional definition of treatment. Dependent variables are specified in logarithms. Firm level controls include firm age, the share owned by foreigners and the share owned by the public sector. All regressions include fixed effects for leader circles, regions, industries, and country-by-years. Standard errors are clustered at the level of treatment.

Table B5: Treatment effects by institutional setting

	(1) Log sales	(2) Log sales	(3) Log sales	(4) Log sales
Treated Area	-0.0565	-1.0897**	0.0401	-0.3871***
	(0.1701)	(0.4842)	(0.1146)	(0.1365)
Treated#V-Dem electoral democracy index	0.3293	4.0786**		
	(0.2890)	(1.7230)		
Treated#(V-Dem electoral democracy index)^2		-3.1328**		
		(1.4590)		
Treated#Control of Corruption			0.2252	2.3898***
			(0.2359)	(0.6027)
Treated#(Control of Corruption)^2				-2.2891***
				(0.6376)
Fixed effects	Yes	Yes	Yes	Yes
Firm level controls	Yes	Yes	Yes	Yes
Observations	82,527	82,527	81,697	81,697
R-squared	0.6621	0.6621	0.6632	0.6632
F	160.7	135.9	155.6	133.2

The regressions are estimated using Equation 2.1 augmented by interacting the treatment variable with the V-Dem electoral democracy index and the control of corruption index from the World Banks Worldwide Governance Indicators. The former index seeks to answer the question 'to what extent is the ideal of electoral democracy in its fullest sense achieved' by aggregating a number of relevant sub-indices. It ranges from 0 (low) to 1 (high). The aggregation encompasses both the idea of a weakest link argument and partial compensation between the sub-indices (Coppedge et al., 2021). The latter index is also an aggregate of a number of sources' perception of corruption. It is expressed as a percentile rank and scaled to the 0 (worst rank) to 1 (best rank) interval. Firm level controls include firm age, the share owned by foreigners and the share owned by the public sector. All regressions include fixed effects for leader circles, regions, industries, and country-by-years. Standard errors are clustered at the level of treatment.

Figure B3: Treatment effect on nightlight luminosity - difference-in-differences as introduced in de Chaisemartin and D'Haultfoeuille (2024)

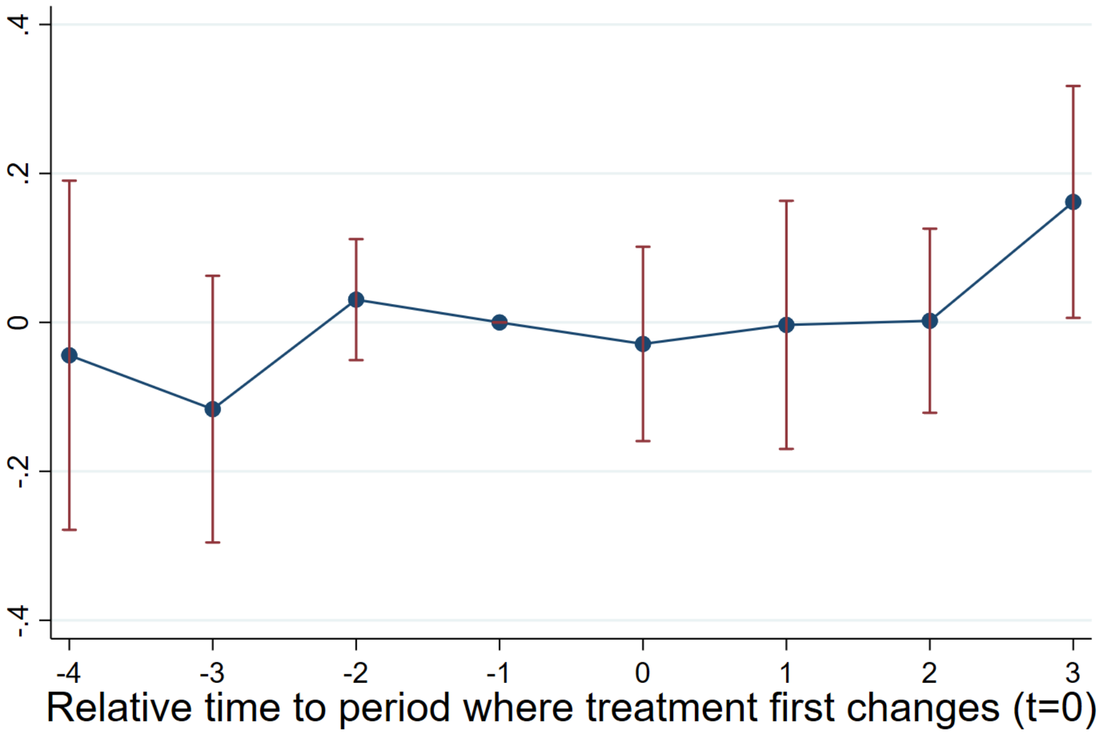

We construct shape files of the treatment and control areas based on the country specific distance measures for the countries with identifying variation in our baseline specification. Control areas are constructed by utilizing admin-1 region shape files and subtracting the treated and exclusion areas from them. We then extract average night lights for these areas from the Earth Observation Group's extended annual DMSP nighttime lights time series using R's exact_extract function. The figure presents the estimated treatment effects relative to the year before treatment status switches for the first time, comparing treatment status switchers to non-switchers (de Chaisemartin and D'Haultfoeuille, 2024). The estimates are unbiased under heterogeneous and dynamic effects, which is a potentially more prominent issue given the yearly frequency of the nightlight data. We include region fixed effects as control variables and cluster at the group level.

Figure B4: Example of a 0.5 x 0.5 degree grid layer over India

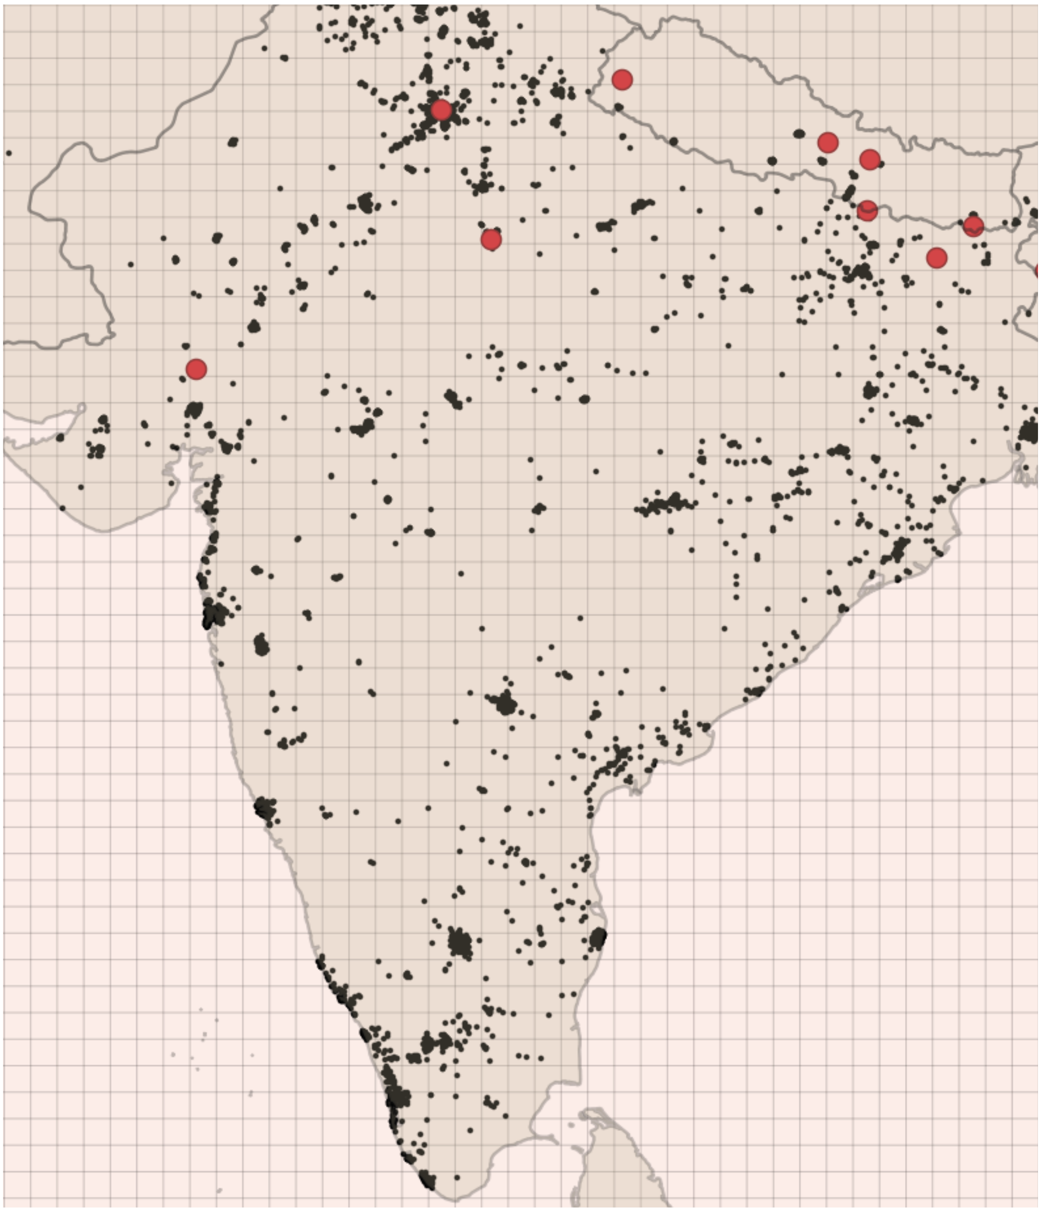

This map serves as a visual example of the grid-layer over India. The grid is spanned by 0.5 x 0.5 degree pixels across the world. The small black dots represent firms. The large red dots represent leader birthplaces.

Figure B5: Permutation test: effect of placebo treatment

(a) Log sales

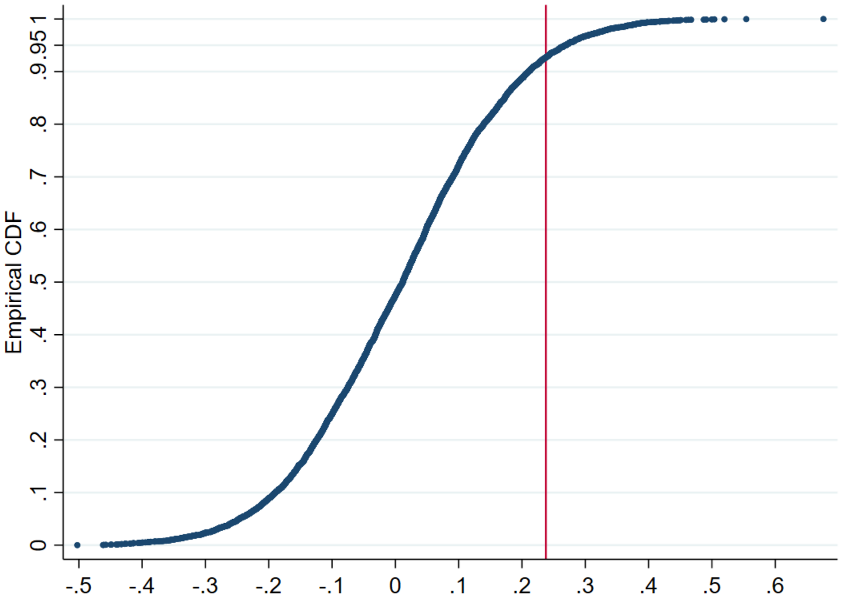

(b) Log employees

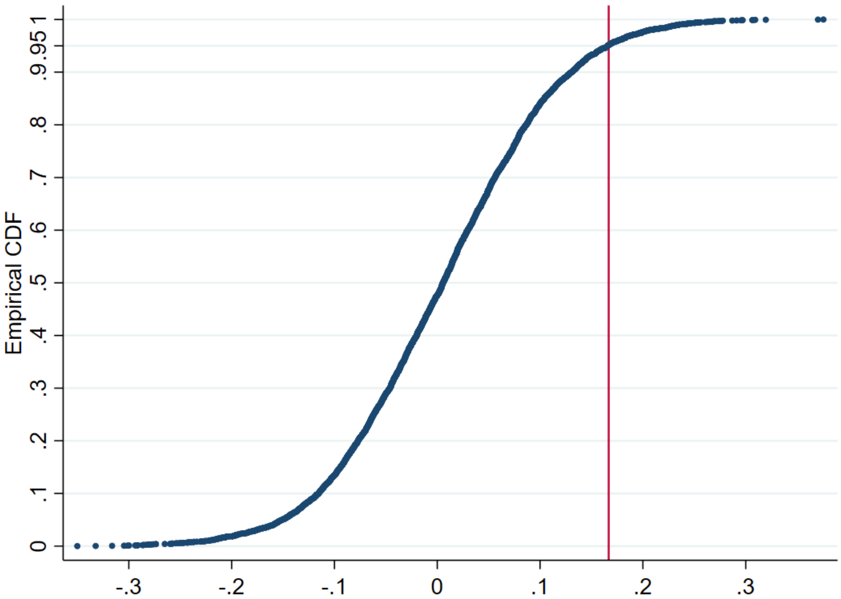

The figure depicts the cumulative distribution of 5000 placebo estimates of the permuted treatment effect. The estimates are derived from the grid-level specification with size 0.5 x 0.5 degrees, where in each country on permutation a random grid cell receives treatment status. The vertical red lines show the magnitude of the actual treatment effects from columns (3) and (4) of Table 2.2.

Figure B6: Changes to the average treatment effect when dropping countries with identifying variation one-by-one

(a) Specification 2.1, geocoded data

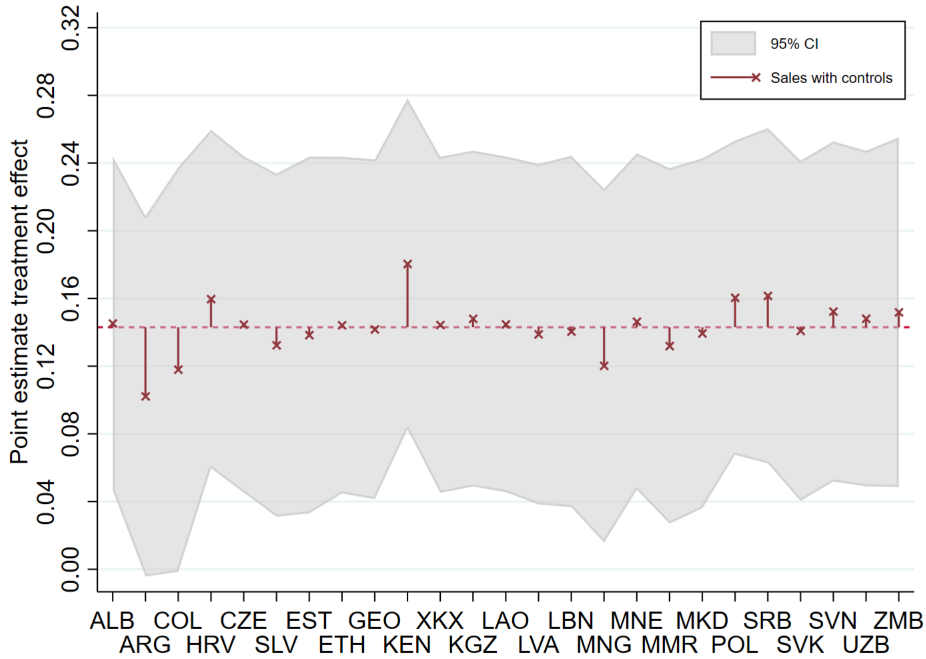

(b) Specification 2.2, regional data

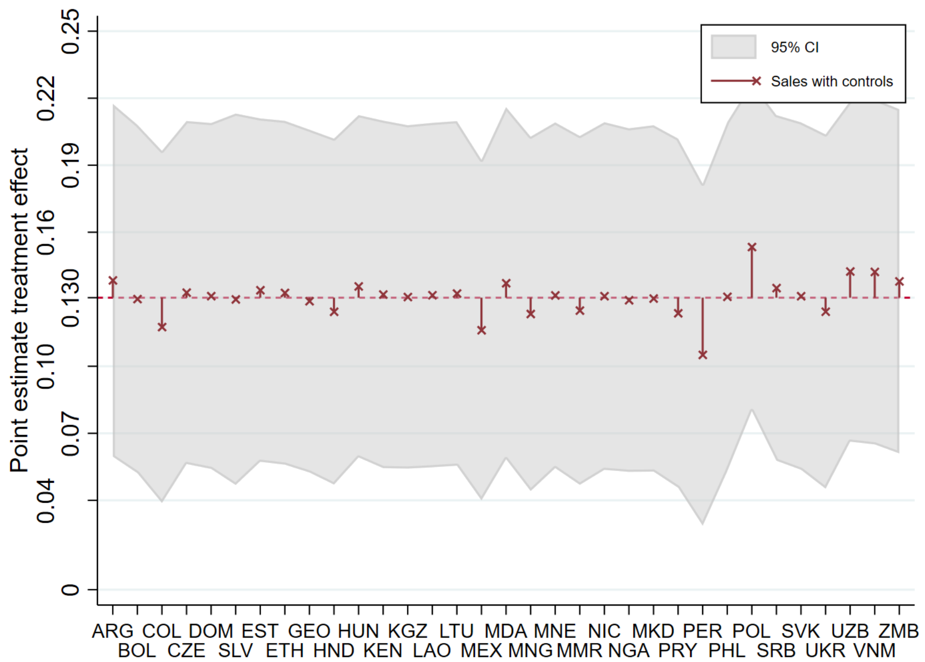

The x-axis lists the 3-letter ISO 3166 country code of the country that is dropped from the estimation for the respective estimate. The red line depicts the average effect of the corresponding unrestricted samples from Tables 2.1 and 2.2.

Figure B7: Size distribution of treated and untreated firms by sector

(a) Only services firms

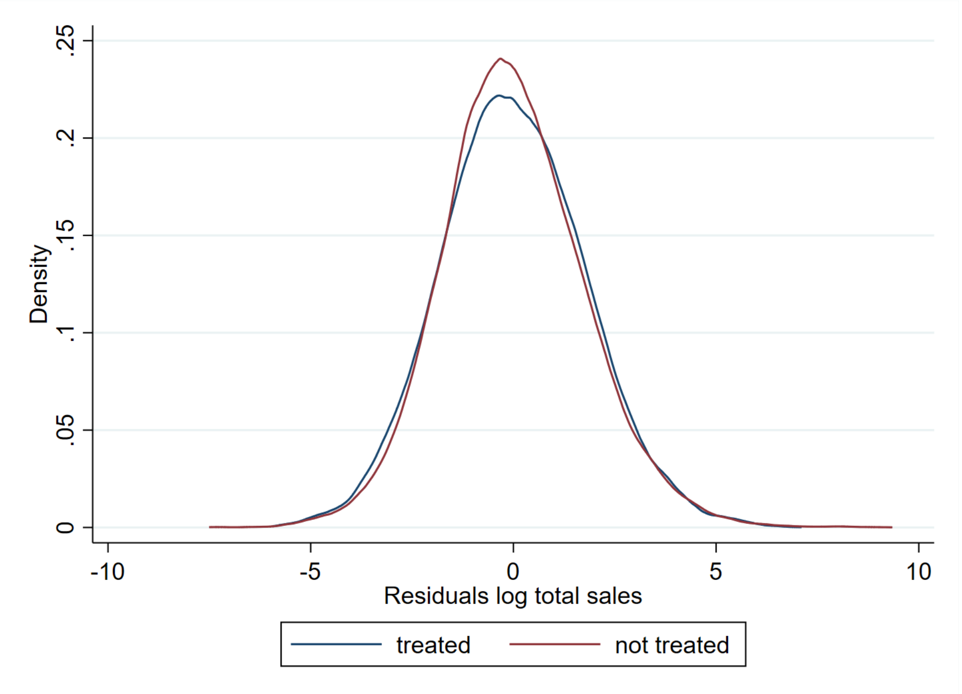

(b) Only manufacturing firms

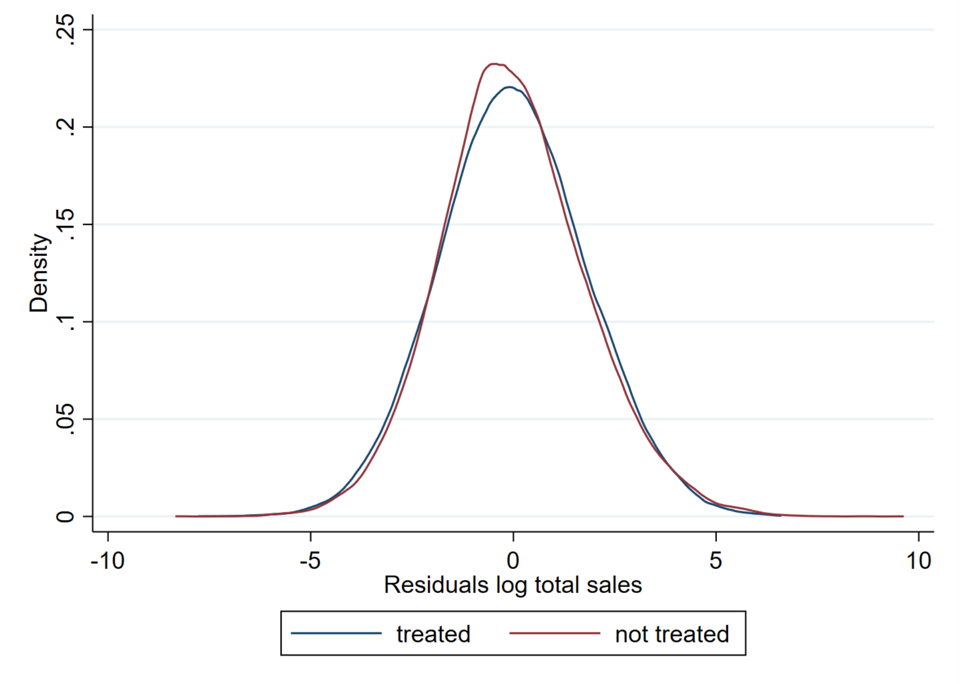

The figure plots the approximated density of residuals from Equation (2.1) with respect to logarithm of sales for the treatment and control group restricting the sample to service sector firms in panel (a) and to manufacturing sector firms in panel (b) using Epanechnikov kernel estimator.

B.2 Implementation of random forest

Random forests operate by averaging over a number of unique uncorrelated decision trees. Each individual decision tree splits the data based on a number of randomly selected variables at each node to purify the data. That is, at each node the data are partitioned into groups based on the observations' similarity in terms of the randomly selected variables. Decision trees reach their terminal nodes once no further purification of a given data partition can be reached. These terminal nodes then determine our estimated propensity scores as the share of observations belonging to the treatment group at that node for the subjects present.

There are two main parameters that establish the generation of the random forest. The first is the number of trees to be grown. Figure B9 shows that the prediction error rate of our forest is stable after 100 trees; however, to be extra diligent, we grow 500 trees. The second parameter is the number of randomly sampled variables available to split the data at each tree node. In Figure B10 we investigate its optimal value by starting from a value of 2, and gradually showing the response of the prediction error rate. At 20, the error rate has virtually converged to a stable value, which we therefore set as the parameter.

All firm level variables with less than 20% missing values that are not our regression outcomes are fed into the random forest algorithm. Zhao et al. (2016) demonstrate that random forests can perform well with variables missing even up to 40% of values. We let the algorithm classify firms into the following groups: the not yet treated, the treated and the never treated. We do this to adopt a weighting scheme similar to the one suggested by Stuart et al. (2014) that specifically accounts for a difference-in-differences design with cross-sectional data. The weights are calculated as follows:

$$w_i = \frac{p_1(X_i)}{p_g(X_i)} \tag{B1}$$

where firms' weight w_i is equal to the predicted probability of being in group 1 given the observed covariates X_i over the predicted probability to be in the group they are actually in. Group 1 consists of the not yet treated. Firms in the other groups receive a weight that is proportional to the predicted probability of them being in group 1, relative to the predicted probability of them being in the group they actually belong to.

In figure B11 we visualize the distribution of the predicted probabilities to belong to any of the groups for the not yet treated, the treated, and the never treated by kernel density estimation. This serves to evaluate the overlap and common support hypothesis. First we exclude observations with probabilities close to 0 or 1 of belonging to any group to avoid perfect predictability given a set of covariates. Then we trim the observations to the area of common support following the approach of Heckman et al. (1997). We drop areas where the estimated densities of the kernel estimator are below a threshold of 0.01.

Figure B8: Distribution of standardized % bias across covariates between treated and untreated observations

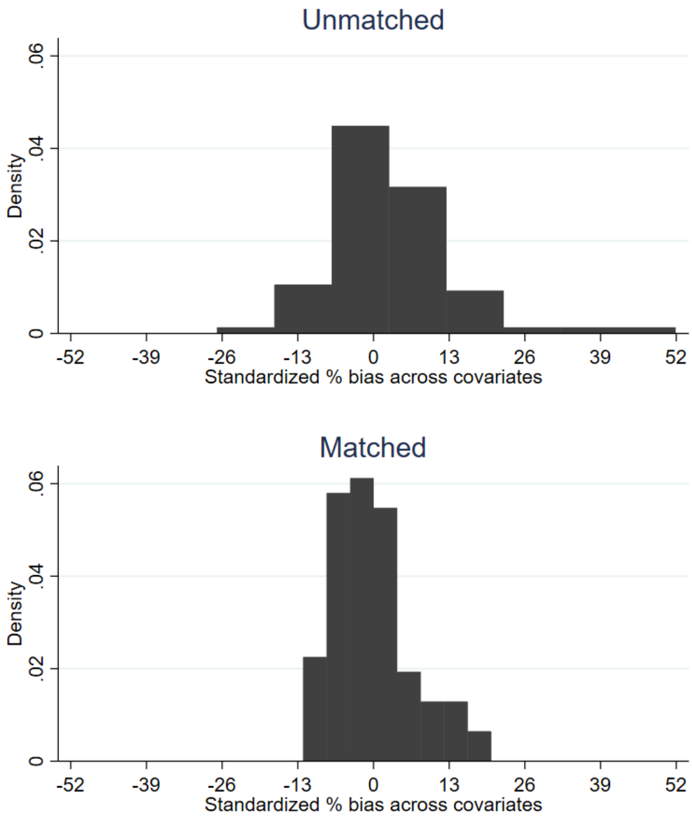

The figure shows the reduction in the standardized bias of firm level covariates between treatment and control firms by showing the standardized % bias before applying the propensity score weights derived from the random forest approach in the upper panel and after the weighting in the lower panel.

Figure B9: Random forest accuracy over the number of trees grown

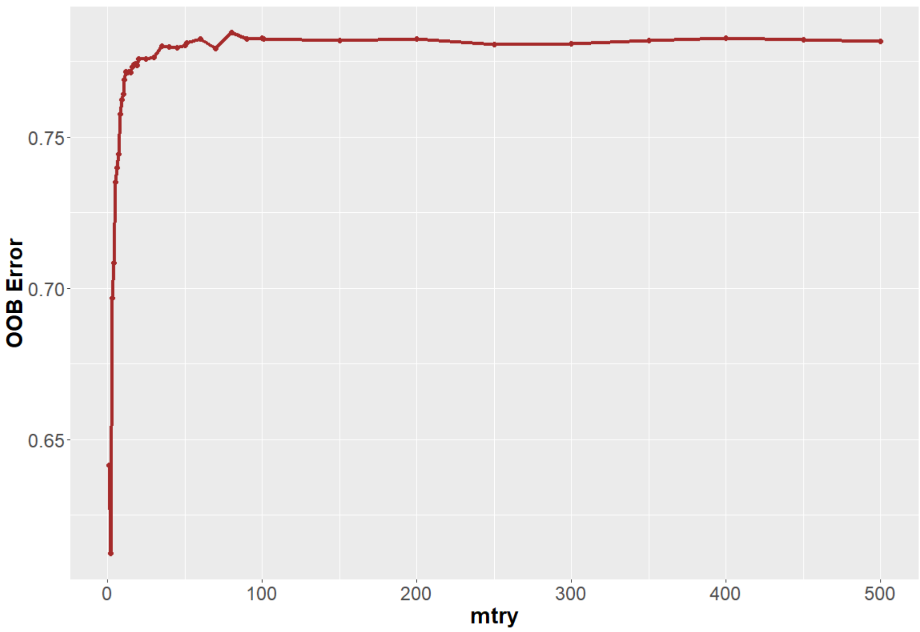

The figure shows the out-of-bag (OOB) prediction performance of the random forest on the y-axis over increasing numbers of independent random trees on the x-axis.

Figure B10: Random forest OOB error rate over number of variables used to split at each tree node

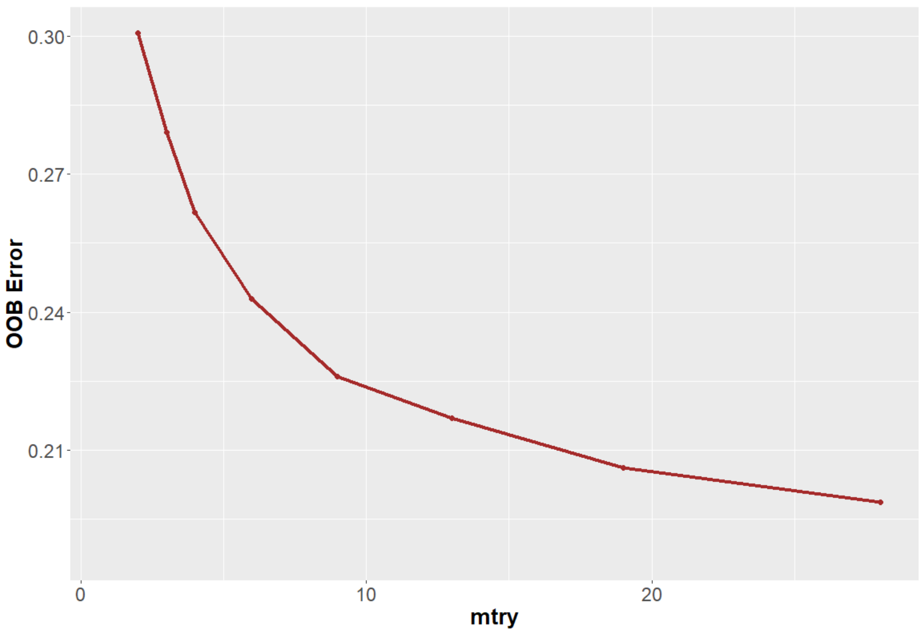

The figure shows the out-of-bag (OOB) prediction performance of the random forest on the y-axis over increasing numbers of variables used to split the data at each tree node on the x-axis.

Figure B11: Distribution of the predicted probability to belong to groups 1 to 3 given the observed covariates by group

(a) Predicted probability to belong to "not yet treated"

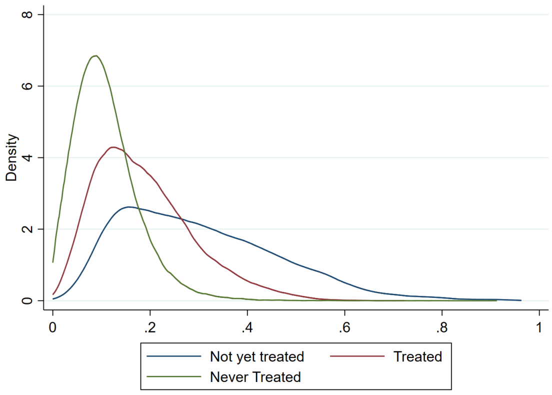

(b) Predicted probability to belong to "treated"

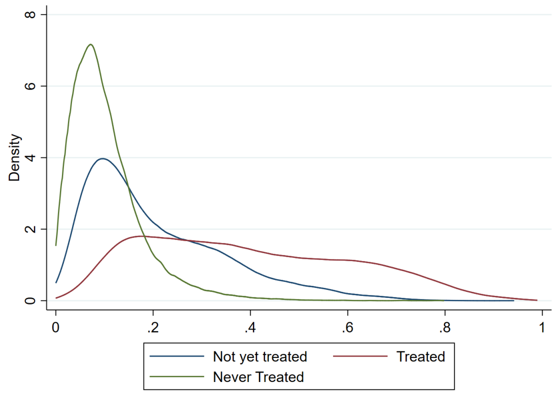

(c) Predicted probability to belong to "never treated"

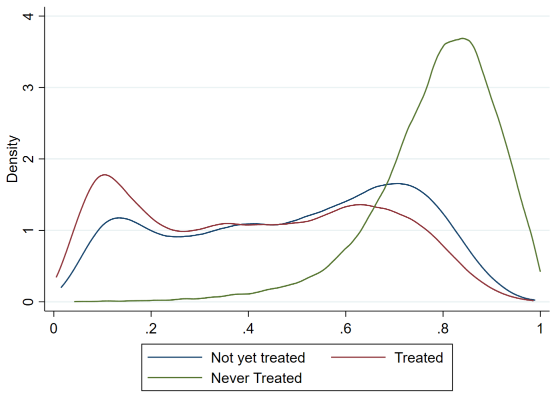

The figure shows distributions of predicted probabilities to belong to the group of "not yet treated" observations in panel (a), the group of "treated" observations in panel (b) and the group of "never treated" observations in panel (c) given the observed covariates separated by actual group status. The red line belongs to the "treated" group, the blue line the "not yet treated" group and the green line to the "never treated" group.

B.3 Robustness of results to alternative specifications

In this section we showcase the robustness of our results to a host of alternative choices for the share of the distance measure as well as the distance cut-off. We plot the distribution of all 90% confidence intervals of the treatment effects for distances from $\frac{1}{8}$ to $\frac{1}{20}$ and for cutoffs of $\frac{1}{5}$, $\frac{1}{6}$, $\frac{1}{8}$ and $\frac{1}{9}$ for our main outcomes sales and employment, as well as the productivity outcomes. We also re-estimate our baseline table fixing the sample size such that the number of observations is stable.

Figure B12: Alternative choices of the distance measure and cutoff threshold

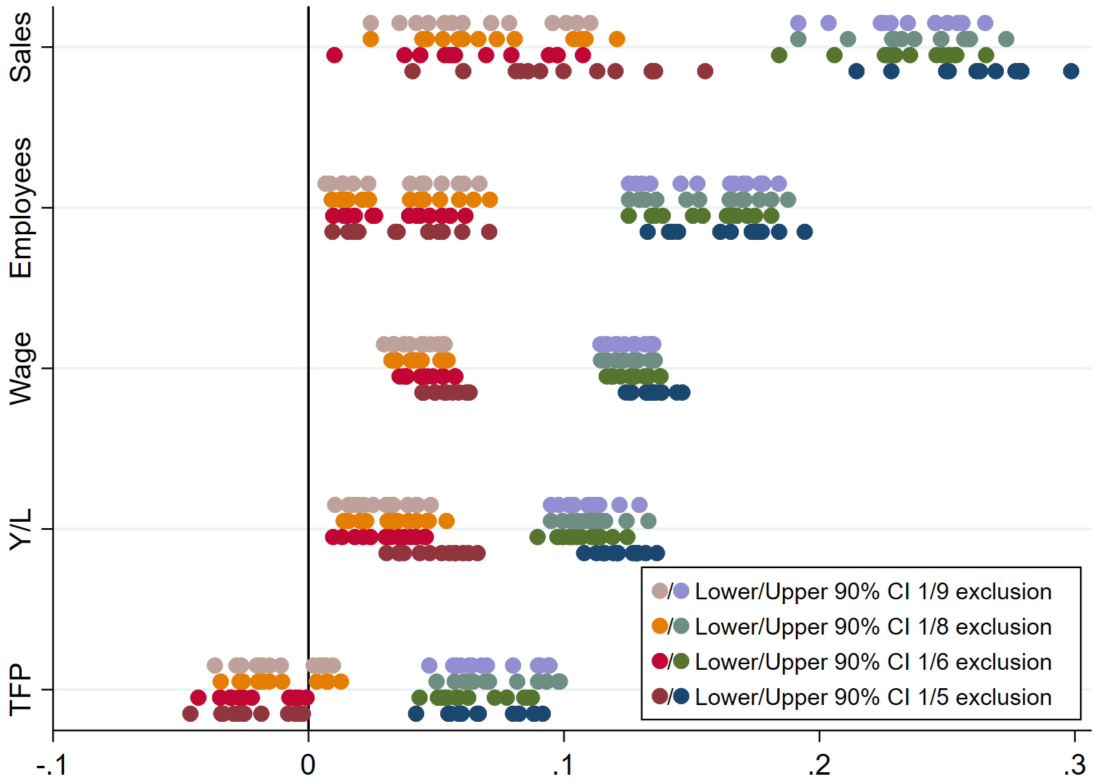

The figure shows the upper and lower bounds of 90% confidence intervals for alternative specifications of the baseline specification. Each pair of points represents a pairing of an alternative treatment distance, going from $\frac{1}{8}$ to $\frac{1}{20}$ of the country-specific distance measure, with an alternative cutoff distance of $\frac{1}{5}$, $\frac{1}{6}$, $\frac{1}{8}$, or $\frac{1}{9}$ of the country-specific distance measure. Outcomes are noted on the y-axis, the vertical black line marks zero.

Table B6: Baseline results: fixed sample size

	(1) Log Sales	(2) Log Employees	(3) Log Sales	(4) Log Employees	(5) Log Sales	(6) Log Employees
Treated area	0.1811***	0.0969**	0.1839***	0.0873**	0.1562***	0.0747*
	(0.0524)	(0.0387)	(0.0536)	(0.0415)	(0.0533)	(0.0392)
Year before treatment start			0.0317	-0.1094		
			(0.1600)	(0.1170)		
Year after treatment end					-0.2579***	-0.2298***
					(0.0881)	(0.0485)
Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Firm level controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	81,587	81,587	81,587	81,587	81,587	81,587
R-squared	0.6661	0.2719	0.6661	0.2719	0.6661	0.2720
F	189.5	198.3	151.7	158.4	154.2	166.4

The regressions are estimated using Equation 2.1 replicating Table 2.1, but restricting the sample such that the number of observations is stable. Treatment area is defined as $\frac{1}{11}$ of the country-specific distance measure. Dependent variables are specified in logarithms. Firm level controls include firm age, the share owned by foreigners and the share owned by the public sector. All regressions include fixed effects for leader circles, regions, industries, and country-by-years. Standard errors are clustered at the level of treatment.

Favoritism by the Governing Elite *

Zareh Asatryan † Thushyanthan Baskaran ‡ Carlo Birkholz §
Patrick Hufschmidt ¶

Abstract

This article extends the study of regional favoritism beyond primary political leaders to ministers of the national cabinets, offering the first comprehensive analysis of the broader governing elites' influence on this type of resource allocation. We hand collect a novel dataset of more than 15,900 georeferenced birthplaces of national cabinet members from 1992 to 2016 and find in a staggered difference-in-differences design that birthplaces of ministers exhibit approximately 9% higher nighttime luminosity, with the strongest effects observed in Africa, for powerful ministerial portfolios, and in contexts with weak institutions and high corruption. While long-serving ministers attract some migration to their home regions, population levels seem to slightly decline once ministers leave office. These findings highlight the systemic nature of favoritism within the governing elite and its dependence on political power and institutional constraints.

Keywords: Favoritism, elite capture, spatiality, luminosity, population, democracy.
JEL: D72, H72, H77, R11.

*We thank participants of the 79th Annual Congress of the International Institute of Public Finance, the 2023 German Development Economics Conference, the 2023 meeting of the European Public Choice Society, the 16th RGS Doctoral Conference, the 1st ZEW-ifo Young Scholar Political Economy Workshop, the University of Siegen Brown Bag Seminar and the University of Göttingen Development Economics Chair Seminar for valuable feedback. We gratefully acknowledge financial support from the German Research Foundation (DFG) within the Project "Regional Favoritism and Development" (Grant no. 423358188 / BA 496716-1).

†ZEW Mannheim
‡Ruhr University Bochum
§University of Mannheim, ZEW Mannheim
¶TU Dortmund

3.1 Introduction

In economics, public choice theory cautions the government as a self-interested actor. One notable manifestation of this self-interest is the tendency of political leaders to favor certain regions in the allocation of public resources over others for private motives. Empirical studies have documented this phenomenon, known as regional favoritism, focusing on primary political leaders. Evidence of such favoritism has, for example, been observed in increased night light intensity (Hodler and Raschky, 2014), in firm performance (Asatryan et al., 2021a) and foreign aid allocation (Dreher et al., 2021a). However, these analyses are centered on primary rulers, thereby providing an incomplete perspective. The ability to engage in, and the extent of regional favoritism is unlikely to be solely a reflection of individual choices; instead, it is shaped by the broader political dynamics within the governing elite.

For this reason, this paper extends the analysis of regional favoritism beyond primary rulers to the broader governing elite, which we define in the context of this paper as including all national cabinet members. We then seek to answer three research questions: First, do ministers engage in regional favoritism and, if so, how does the extent compare to primary rulers? While ministers typically wield less power than a country's primary leader and may therefore lack the ability to direct resources at a comparable magnitude, they may encounter less public scrutiny, potentially enabling them to redirect resources to their home regions more freely. Second, how do characteristics of specific ministerial portfolios, such as their prestige and power, influence regional favoritism? Third, what role do institutions play in moderating the effects of regional favoritism? Although stronger democratic institutions might constrain politicians' ability to channel resources toward personal interests, they may simultaneously incentivize such behavior as a means of securing electoral support.

Several prominent examples of regional favoritism at the level of ministers underscore the importance of understanding the dynamics within the governing elite. The political trajectory of the Rajapaksa family in Sri Lanka offers a particularly illustrative case. Mahinda Rajapaksa, who ascended to power as Prime Minister in 2004 and later as President in 2005, consolidated influence as the primary ruler by also controlling key ministries, including Defence, Finance, and the Ministry of Highways, Ports and Shipping. During his tenure, the Rajapaksa family entrenched itself deeply within government, with three of Mahinda's brothers assuming powerful positions: one as Minister of Economic Development, another as Secretary to the Ministry of Defense, and a third as Speaker of Parliament. This concentrated authority and decision-making within a close-knit group from Sri Lanka's Southern Province resulted in a surge of major infrastructure projects in the region, including the construction of the Mahinda Rajapaksa International Cricket Stadium, the Mattala Rajapaksa International Airport, and the Magampura Mahinda Rajapaksa Port, the latter poised to become the country's largest. The family's grip on power remained firm through various reshuffles of ministerial roles among the brothers until mass protests in March 2022 directed against the Rajapaksa family saw them ousted.

Building on this anecdote, we systematically investigate regional favoritism by the governing elite. To this end, we compiled a dataset of hand-collected and geo-referenced birthplaces of national cabinet members globally. Our sample spans from 1992 to 2016 and includes geo-coordinates for approximately 15,900 unique cabinet member birthplaces. We describe this dataset in detail in Section 3.2, and plan to make it publicly available as part of the Political Leaders' Affiliation Database (PLAD) to serve as input for research on regional favoritism specifically, but more generally to support a wider range of geospatial studies in political economy.

Our empirical strategy leverages the timing of ministers' appointments and the geographical spread of their birthplaces. Using satellite imagery, we compare nightlight intensity and population numbers for small geographical units (0.5 x 0.5 degree pixels, where 0.5 degrees correspond to about 55km at the equator) before and after a minister assumes office. Areas, or pixels, as we call them interchangeably, that have never or not yet been home to a minister serve as the control group. To cope with the biases identified in the recent differences-in-differences literature (for a synthesis, see (Roth et al., 2023)), we implement estimators capable of addressing the shortcomings of traditional two-way fixed effects models. Specifically, we implement the estimator proposed by Callaway and Sant'Anna (2021) (**CS**) to estimate the persistent effects of an area having been the birthplace of a minister, and the estimator developed by de Chaisemartin and D'Haultfœuille (2020) (**CH**) to estimate the dynamic effects of treatment status switches.[1]

Our main finding is an aggregate increase in nighttime light intensity of roughly 9% for minister pixels, indicating regional favoritism effects of ministers surpassing those previously documented for primary leaders. A sub-sample analysis by continent reveals that these effects are driven primarily by countries in Africa, with no statistically significant average effect observed for Asia, Europe, or the Americas. The dynamic treatment effects show that the favoritism effect intensifies with ministers' tenure. We find no evidence of persistent positive population growth in ministerial regions. The **CS** estimates show a population decline of 1% and 2% globally, while the **CH** dynamic effects indicate no significant changes for shorter tenures, but positive effects for longer ones. Taken together, these results suggest that consistent favoritism over an extended period can trigger migration responses; however, when ministers leave office and favoritism ends, these regions experience relative population declines compared to untreated areas.

To explore the mechanisms underlying these findings, we incorporate individual-level data on ministers from WhoGov and country-level data on corruption and institutions from Transparency International and Freedom House. We find that greater political power, measured by the prestige of a minister's portfolio, is associated with stronger favoritism effects. Notably, finance and foreign ministers drive the results of the highest prestige

[1]We refer to the estimates of **CS** as persistent effects, as treated units remain so once they receive treatment, while we call the estimates of **CH** dynamic effects, as treatment status can switch on and off.

category, indicating that access to domestic and foreign financial capital plays a critical role in allocating resources toward favored regions. Additionally, we examine heterogeneity by institutional context. Our baseline results are primarily driven by autocratic, more corrupt, and less industrialized countries. In contrast, in democratic settings, ministers appear more constrained in their ability to redistribute resources to favored regions.

Our paper contributes to the evolving literature on regional favoritism. The seminal work by Hodler and Raschky (2014) demonstrates that regions connected to the primary ruler exhibit greater economic activity, as measured by nighttime luminosity. Hodler and Raschky (2014) also find that such favoritism does not have a persistent effect once the leader steps down. Asatryan et al. (2021a) show that firms particularly in the non-tradable sector located in favored regions experience higher sales and employment, primarily due to short-term government demand. However, neither firms' perceived business environment nor metrics related to productivity of firms improve. The induced allocation towards non-tradable firms thus leads to small aggregate output losses in the economy, due to diminishing marginal returns.

A series of papers investigates favoritism specifically on the African continent. Dreher et al. (2021b,a) show that for home regions of primary rulers, the allocation of Chinese aid is subject to favoritism, and that favored regions appear to benefit in terms of local economic development measured by nighttime luminosity. World Bank aid does not exhibit the same pattern; however, new evidence suggests that the allocation of Western aid is not absent of regional favoritism, as birthplaces of leaders' spouses attract substantially more aid from European donors, the United States, and China (Bomprezzi et al., 2024). Asatryan et al. (2021c) find that the economic benefits of mine openings are concentrated in leaders' birth regions, but only in autocratic regimes. Further, Asatryan et al. (2021b) show that exposure to favoritism during adolescence increases human capital for men and women co-ethnic to the primary leader later in life. Closest to our work, Widmer and Zurlinden (2022) examine 36 African countries and find reduced infant and neonatal mortality in regions linked to health ministers, especially for rural or uneducated mothers. Our study significantly expands the geographic and temporal scope of the sample of ministers and focuses on broader economic outcomes. In the context of India, Khalil et al. (2021) study ministers as well, however at the sub-national level. They document that constituencies represented by a chief minister see a 13% increase in luminosity, although this effect is stronger in constituencies outside the chief minister's home region.

Related research explores the mechanisms of favoritism, often within single-country contexts. Burgess et al. (2015) find that during autocratic periods in Kenya, regions inhabited by co-ethnics of the president receive more road spending, while democracy shifts favoritism toward less visible mechanisms, such as educational transfers. More evidence on mechanisms of regional favoritism emerges from a diverse set of countries: public sector employment in Germany (Baskaran and da Fonseca, 2021), infrastructure in Vietnam (Do et al., 2017), public transfers in Italy (Carozzi and Repetto, 2016), and lending of the

European Investment Bank across regions of Europe (Asatryan and Havlik, 2020).

Our paper also intersects with the broader literature on politician selection and accountability (Barro, 1973; Besley and Coate, 2003; Besley, 2005; Maskin and Tirole, 2004; Alesina and Tabellini, 2007, 2008; Francois et al., 2015). In our paper, we are also interested in understanding which factors restrict politicians from engaging in rent-seeking activities and whether a particular selection of ministers, for instance women, behave differentially.[2]

Finally, our paper is connected to the literature on the spatial implications of distributive politics. Neoclassical models of distributive politics propose that office-motivated politicians have strong incentives to allocate disproportionate public resources to electorally important geographies (Weingast et al., 1981), such as core, swing, or politically aligned districts (Cox and McCubbins, 1986; Cox, 2010; Albouy, 2013; Baskaran and Hessami, 2017). The spatial distortion we investigate focuses on birth regions, which could reflect electoral motivations, rent-seeking behavior, or intrinsic preferences for one's place of origin.

The remainder of this paper is structured as follows: Section 3.2 introduces the newly collected data on ministers' birthplaces and other data sources used in the analysis. Section 3.3 describes our baseline empirical strategy. Section 3.4 presents the main results, while Section 3.5 explores the mechanisms. Section 3.6 concludes.

3.2 Data

3.2.1 Birthplaces of the governing elite

A key contribution of our paper is the manual collection and geo-coding of the birthplaces of national cabinet members. We will publicly share the full dataset as part of the Political Leaders' Affiliation Database (https://www.plad.me/). Our dataset identifies governing elites using the WhoGoV database, which itself covers all countries with populations exceeding 400,000, totaling 177 countries and spanning the years 1966 to 2016. To date and to our knowledge, this is the largest global dataset on ministers and cabinets, containing information on 50,197 cabinet members. In addition to their names, this dataset contains variables documenting the years they were in power, official position, years of birth and death, party affiliation, ministerial portfolio, and several other information (Nyrup and Bramwell, 2020).

Building on this, we conducted desk research to extend the data with two geographic dimensions: birthplace and birth region. We identified 15,931 birthplaces and 17,066 birth regions of cabinet members in 141 countries. Table C1 provides an overview of the countries and time periods covered. Figure 3.1 displays the spatial distribution of all identified birthplaces on a global map, while Figures C1, C2, C3, C4, and C5 in the

[2]We turn to this particular aspect in Section 3.5.5.

Appendix zoom in to each of the continents. The maps reveal considerable variation in the spatial distribution of cabinet members' birthplaces, with natural gaps in areas such as the Australian Outback, the Amazon Rainforest, the Sahel Desert, and the Himalayan Mountain Range.

This dataset offers several advantages for studying spatial political economy questions compared to conventional datasets that focus solely on primary rulers. First, the inclusion of a more comprehensive set of the governing elite enables to study and control for dynamics within the elite. Second, the larger sample size of cabinet members provides greater identifying variation across both time and space. Third, the dataset allows researchers to leverage cross-sectional variation in political portfolios, prestige levels, and ideology. Fourth, it facilitates the examination of within-unit variation, capturing not only changes over a minister's tenure but also changes between different ministerial portfolios.

Given the availability of the nighttime light data, described in detail below, we restrict our main estimation sample to the period from 1992 to 2016. During this time, we identify 13,951 birthplaces for 27,238 cabinet members. Naturally, finding reliable information on birthplaces becomes more challenging for earlier time periods. From 1992 to 2000, the detection rate drops to 48.1%, compared to the overall average of 51.2%, but increases to 53.9% after 2000. Similarly, it might be easier to identify the birthplaces of ministers who hold more powerful and prestigious positions. Using Nyrup and Bramwell (2020)'s classification of ministerial portfolios into high, medium, and low prestige categories (see Table C2), we find that birthplaces of high-prestige ministers are detected at a rate of 54.6%. However, detection rates for medium- and low-prestige categories are similar, at 47.2% and 47.8%, respectively.

The high average detection rates and fairly small differences across time periods and prestige levels suggest that selection bias in the detection of birthplaces is unlikely to introduce significant systematic bias into our estimates. Biased estimates would require a correlation between birthplace detection and either nighttime light intensity or the treatment effect. For example, if birthplaces in faster-growing areas were more likely to be identified, our estimates could be upward biased. However, our event study estimations do not show differential pretrends (see Figure 3.2), which speaks against this being the case.

3.2.2 Luminosity data

We use nighttime luminosity as a proxy for local economic development at the local level (Alesina et al., 2016; Hodler and Raschky, 2014; Michalopoulos and Papaioannou, 2016; Bruederle and Hodler, 2018; Martínez, 2022). These data are derived from satellite images of Earth at night, captured by the US Air Force (USAF) Defense Meteorological Satellite Program Operational Linesman System (DMSP-OLS). The original imagery is processed by the National Oceanic and Atmospheric Agency (NOAA) and released as raster datasets.

Figure 3.1: Birthplaces of cabinet members

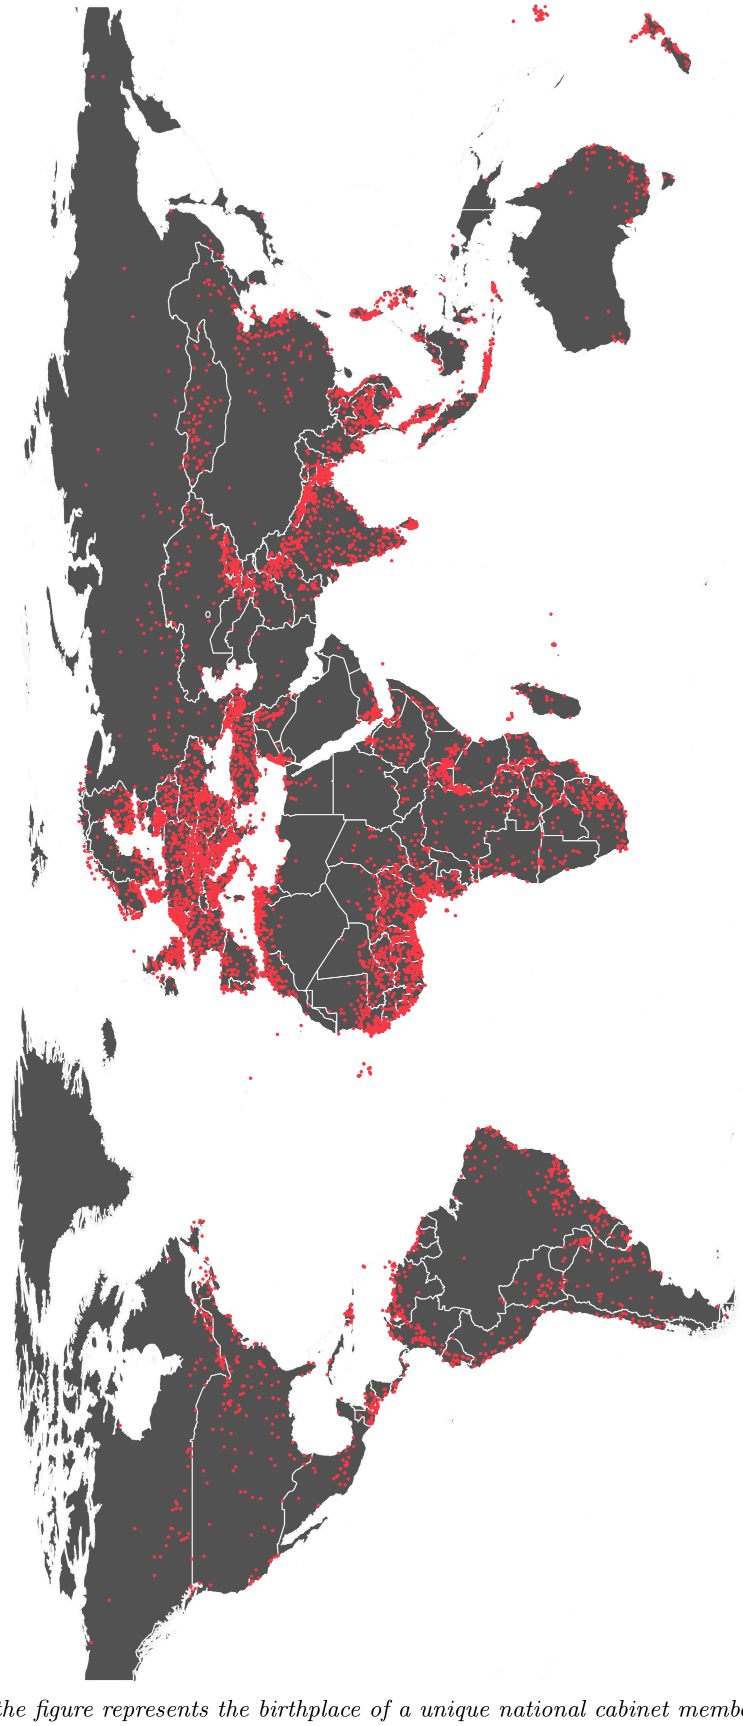

Each red dot in the figure represents the birthplace of a unique national cabinet member in our sample.

We use annual composites collected from satellites F10, F12, F14, F15, F16, and F18, where ephemeral lights (e.g., fires and flaring) and data from nights affected by clouds, moonlight, or other glare are excluded. The images are available at a resolution of 30 arc-seconds (about 0.86 square kilometers at the equator) for years after 1992. Each pixel of the dataset stores a 6-bit digital value ranging from 0 to 63 representing the average light intensity, higher values implying that a pixel emanates more light (Henderson et al., 2012). While the initial release of stable light data ended in 2013, the series was extended to 2021 using improved algorithms and data from satellites F15 and F16 (Ghosh et al., 2021).

3.2.3 Population data

We obtain population data from the WorldPop Project (WorldPop and CIESIN, 2018), which provides annual gridded population estimates as raster files for 2000-2020. Population values per pixel of the WorldPop data are derived from official census data and various other input data sources, such as location and extent of settlements, roads, land cover, building maps, satellite nightlights, vegetation, topography, health facility locations, and refugee camps. Stevens et al. (2015) shows methodological details regarding the random forest regression tree-based mapping approach that is used to produce the gridded pixel data at spatial resolutions of 1 km and 100 m.

3.2.4 Further data sources

In Section 3.5, which explores the mechanisms of the observed baseline effects, we incorporate data on democracy and civil rights (Freedom House, 2019) and corruption (Transparency International, 2022). These country-level variables allow us to examine how institutional settings influence the effects of regional favoritism.

3.2.5 Combining all data on a grid

We bring all data sources together on a global grid of 0.5×0.5 degree squares.[3] These grid cells, or "pixels", are intersected with country borders to identify within which country a particular cell is located. Border cells spanning multiple countries are excluded, resulting in a final sample of 1,189,560 cells for 1992-2016.

For each grid cell, we calculate yearly measures of economic development and population by overlaying the grid over the raster datasets for nightlights and population described above. Luminosity data are aggregated by computing the mean of the values of the night light raster image pixels that fall within the boundaries of each of the 0.5×0.5 degree grid cells. For population, we proceed identical except that we calculate the sum of the values of each cell of the population raster that falls within the boundaries of each of the

[3] At the equator 0.5 degrees correspond to about 55km.

0.5 x 0.5 degree grid cells. Both outcomes are plotted for a sample year in Figures C6 and C7.

Finally, to define the treatment status, we perform a spatial join of this grid with our geocoded cabinet member dataset. In our sample, approximately 6% of cells are at some point home to a minister, while only 0.6% are the birthplace of a primary ruler. About 8.5% of the pixels that at some point host a minister are also home to a primary ruler during the sample period. To account for potential confounding effects of this dual treatment, our regressions control for the presence of a primary ruler. Conversely, 84.8% of leader pixels are also home to a minister, underscoring the importance of analyzing regional favoritism within the broader governing elite rather than focusing solely on primary leaders. On average, ministers remain in power for 4.2 years. Pixels can be the home of more than one active minister simultaneously, in fact the average number of ministers during treatment is 1.55.[4]

3.3 Empirical strategy

3.3.1 Staggered difference-in-differences

A recent series of papers analyzes the inference question when treatment is staggered across units over time and has discovered that the two-way fixed effects estimator (TWFE) may not be an unbiased estimator of the average treatment effect on the treated (ATT) when treatment effects occur at different point in time and are heterogeneous. Many authors suggest alternative estimators and provide diagnostic tools to reveal potential bias (Baker et al., 2022; Borusyak et al., 2022; Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020; Sun and Abraham, 2021).

The canonical difference-in-differences models involve two periods and two groups. The untreated group never participates in the treatment, and the treated group becomes treated in the second period. However, using the canonical models in cases where there are more than two time periods and where different units can become treated at different times, already treated units may serve as control group for later treated units because their treatment status is constant over time. An important finding is that every group acts as a control group at some point in time. If treatment effects vary over time, the estimated coefficients may be biased. Goodman-Bacon (2021) proves that the usual fixed effects estimator yields a weighted average of all possible pairs of the underlying TWFE estimator. In particular, the Goodman-Bacon Decomposition shows that when treatment effects are not homogeneous, some of these weights may be negative.

In other words, the TWFE is not robust to treatment effect heterogeneity, as relatively comparing newly treated units to already treated units adjusts the path of outcomes for newly treated units by the path of outcomes for already treated units. However, this

[4]This treatment intensity can be captured by the **CH** estimator but not by the **CS** estimator - providing another rationale for employing both.

is not the path of untreated potential outcomes, it includes treatment effect dynamics. As a result, these dynamics appear in the coefficient of the treatment dummy, making it difficult to give a convincing causal interpretation. Callaway and Sant'Anna (2021) show in simulations that examples exist where the effect of participating in the treatment is positive for all units in all time periods, but the TWFE estimation results indicate a negative effect of participating in the treatment.

With multiple treatment timings (appointments to ministerial positions) across units (cells in countries) and potentially heterogeneous treatment effects, as countries are heterogeneous in size and cabinets are heterogeneous regarding political power, our setting calls for an empirical design that addresses the previously described estimation pitfalls. We use the **CS** estimator as our main specification, which we will introduce below. In Section 3.4.3 we supplement it with the **CH** estimator for robustness, as well as its ability to allow for treatment status switches.

3.3.2 Specification

In our main specification, treatment occurs at the time a minister is appointed. For the pixel containing the birth place of the respective minister, the treatment status switches on, and, because of the properties of the estimator, remains on. In other words, we measure the treatment effect of ever having been the birthplace of a minister during the sample period. The control group are all remaining pixels of our sample. Callaway and Sant'Anna (2021) propose numerous ways to aggregate group-time average treatment effects. We use the aggregation methods simple and dynamic as defined in the *did* R package. Both procedures are outlined in the following.

The ATT in setups with multiple treatment groups and multiple time periods can be formalized by:

$$ATT(g,t) = E[Y_t(g) - Y_t(0)|G_g = 1]. \tag{3.1}$$

The $ATT(g,t)$ represents the average treatment effect for pixels that are members of a particular group g[5] at a particular time period t.

Consider the average effect of receiving treatment, separately for each group. This can be denoted as:

$$\theta_S(g) = \frac{1}{T-g+1}\sum_{t=g}^{T} \mathbf{1}\{g \leq t\} ATT(g,t). \tag{3.2}$$

$\theta_S(g)$ is the average effect of receiving the treatment among units in group g, across their post-treatment periods. There are T total time periods, where t in our setting is

[5]Groups are defined by treatment timing. For example, a pixel that is a birth place of cabinet member that came into power in the year 1996 belongs to $g = 1996$.

yearly $t = 1, ..., T$. The parameter $\theta_S(g)$ allows to emphasize treatment effect heterogeneity with respect to treatment adoption time. Furthermore, it is fairly straightforward to further aggregate $\theta_S(g)$ to receive an overall effect parameter that is easy to interpret:

$$\theta_S^O = \sum_{g \in G}^{T} \theta_S(g) P(G = g | G \leq T). \tag{3.3}$$

θ_S^O is the average effect of receiving the treatment for units (pixels) in group g as defined in equation 3.2. θ_S^O first calculates the average effect for each group (across all time periods). Then it averages these effects together across groups to summarize the total average effect of receiving the treatment. Hence, θ_S^O is the average effect of participating in the treatment for all units that ever received treatment. In this regard, its interpretation is the same as the ATT in the traditional DiD setup with two periods and two groups.

As shown, the simple aggregation method is an intuitive approach. It yields a weighted average of all group-time average treatment effects with weights proportional to group size. This type of aggregation circumvents the negative weights problem that might occur in two-way fixed effects regressions. Therefore, it is a straightforward summary statistic of the overall effect of receiving the treatment in the context of multiple time periods and variation in treatment timing. However, this simple aggregation has the tendency to overestimate the effect of early-treated groups simply because more of them exist during post-treatment periods. Therefore, we also implement a dynamic approach, as outlined next.

In our application, there is a large number of groups and time periods and we are interested in understanding treatment effect dynamics. A common approach to analyze these dynamics is to aggregate group-time effects into an event study plot. We do this by computing average effects across different lengths of exposure to the treatment and plot the results.

Let e be event-time, i.e., $e \cdot t - g$ captures the years passed since treatment was adopted. A way to aggregate the group-time average treatment effect $ATT(g, t)$ to highlight treatment effect dynamics with respect to e is given by:

$$\theta_D(e) = \sum_{g \in G}^{T} \mathbf{1}\{g + e \leq T\} P(G = g | G + e \leq T) ATT(g, g + e). \tag{3.4}$$

$\theta_D(e)$ is the aggregated parameter of interest for our event study. It captures the average effect of a pixel having a birthplaces of a ministers e years after the treatment was adopted across all pixels that are ever observed to have birthplace of a minister for specifically e years. In this specification, the "on impact" average effect of receiving the treatment appears at $e = 0$. This aggregation avoids the drawbacks associated with the dynamic TWFE specification discussed in the previous section. The overall effect is then calculated by averaging the effect of the treatment across all positive lengths of exposure.

A methodological challenge is that regions or pixels that are connected to the governing elite may be systematically different from other polygons. For example, ministers might be more likely to originate from more urbanized parts of their respective countries who's economic activity grows faster irrespective of hosting a minister. As such, comparing pixels that were connected to a cabinet member with all other (not yet treated) pixels may lead to biased estimates. We do not find any indication for such differential growth trends in the event studies we present with our results.

To further address this concern, we incorporate covariates in our event study estimations. In particular, we utilize a matrix of covariates that includes country dummies and controls for leader birthplaces. We use the default doubly robust approach of the *did* R command to compute group-time average treatment effects. This procedure allows us to verify if the results hold after conditioning on these pre-treatment covariates.[6]

3.4 Empirical results

3.4.1 Luminosity in minister pixels

In Table 3.1 we present the aggregate effect of being a minister's birthplace on the intensity of nightlight a pixel emits from our baseline specification. The aggregation of the group-time specific effects follows the two procedures outlined in Section 3.3.2. In column (1) we show the aggregate effect for our full sample which spans countries around the world. Both aggregation methods result in sizeable significant effects, suggesting aggregate increases between 5% and 9% of nighttime light intensity after ministers come into power. In their seminal paper Hodler and Raschky (2014) estimate a baseline effect of 3.8% increased nighttime light intensity in leaders' birth regions.

There are a number of potential reasons for the larger effect sizes that we measure: First, the sample compositions have a large overlap, but are not identical. This is true for the countries included, but particularly for the time periods. As Hodler and Raschky show a strong interaction effect with leader tenure, i.e. effects start becoming statistically different from zero only in year 14, our longer study period might capture more long tenures. Second, the unit of study in our estimations is the pixel level, and thereby more granular than the region level employed by Hodler and Raschky. Third, the use of **CS** differences-in-differences estimators addresses the issues of staggered treatment adoption discussed in Section 3.3.1. In countries with multiple switches of primary rulers during the sample period, potentially harmful comparisons of treated and already treated pixels might arise in a standard difference-in-differences design. However, this methodological issue is clearly more pronounced with the many more treatments we observe for ministers.

[6]The *did* package requires that covariates are time-invariant. For time varying variables, the did package sets the value of the covariate to be equal to the value of the covariates in the base period. In the post-treatment periods the base period is the period immediately before observations in a particular group receive the treatment, and in pre-treatment periods the base period is the period immediately before the current period.

Fourth, our results imply large treatment effects for ministers. If minister cabinet changes typically coincide with changes of the primary ruler, then not controlling for minister birth pixels dilutes the control group and downward biases the estimate. Additionally, leaders' birthplaces are very often also the birthplace of a minister such that not controlling for their presence may confound the estimated effect. Fifth, ministers may be more strongly incentivized and better able to exert favoritism toward their birthplaces. For example, they might rely more on regional political support, while at the same time being under less public scrutiny.

In columns (2) to (5) of the table, we present the results for sub-samples of individual continents. We observe strong heterogeneity of effects between the continents. We find that the effect is driven by African countries, as the other subsamples have small estimates that are not significantly different from zero. Part of these differences might be driven by the fact that nightlights as a measure will behave differently across the continents in our sample. For instance, already very strongly electrified countries in Europe may have a different potential to become brighter. Furthermore, it is likely that the institutional setting mediates the size of the effect. We turn to this aspect in Section 3.5.2.

Table 3.1: Treatment effects in minister birth pixels: nightlights

	Dependent variable: luminosity				
Aggregation method	(1) World	(2) Africa	(3) Europe	(4) Asia	(5) Americas
simple	0.054***	0.144***	-0.006	-0.010	-0.017
	(0.015)	(0.032)	(0.044)	(0.036)	(0.029)
dynamic	0.094***	0.187***	0.009	0.027	-0.07
	(0.022)	(0.043)	(0.051)	(0.052)	(0.047)
Observations	957,350	209,900	324,825	250,550	172,075

The dependent variable average nighttime light intensity is specified in logarithmic form. The method of aggregation *simple* is defined by Equation 3.3, and the *dynamic* aggregation is defined by Equation 3.4. To limit the duration of a plausible treatment effect on the outcome, we constrain the *dynamic* aggregation to 20 post-treatment periods. All estimations include covariates identifying the birth pixels of a country's primary rulers, capital cities, and a country-specific factor variable. Stars denote significance levels at 10%(*), 5%(**), and 1%(***). Heteroscedasticity-robust standard errors, which are clustered at the pixel level, are presented in parentheses.

We are also interested in the time dynamics of the effects we measure, as tenure showed to be an important factor in Hodler and Raschky (2014). To this end, we plot the group aggregates by distance to treatment start in an event study type plot in Figure 3.2. We observe a slowly increasing effect over the first ten years after a minster comes into power for the global sample, which again is driven by the African sub-sample (see Figure C8). For the other continents, the line plotting the aggregated coefficients remains fairly flat and statistically insignificant. The steady increase over the years is in line with the notion that ministers are diverting resources and differentially benefit their home regions more, the longer they stay in power. The figures let us also investigate the existence of pre-trends. If minister pixels were substantially different from non-minister pixels, or if ministers coming into power could be anticipated and elicit a change of nightlights, this should lead

Figure 3.2: Dynamic treatment effects in minister birth pixels

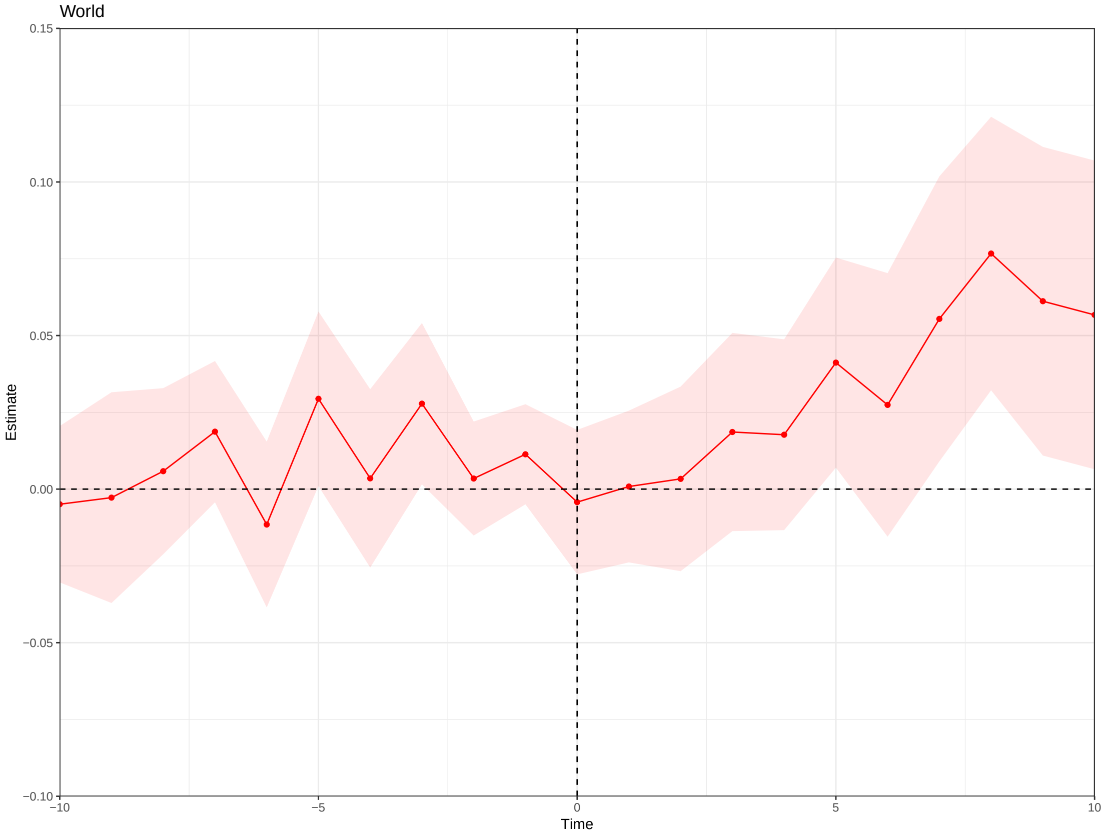

The figure shows an event study based on the Callaway and Sant'Anna (2021) difference-in-differences estimator, relating birth places of ministers in power to luminosity at the grid level. The red shaded areas on the plot represent 95 percent confidence intervals, with standard errors clustered at the cell level.

to significant effects in the time periods prior to them getting into office. None of the samples in Figure C8 displays a pattern that is consistent with this narrative.

3.4.2 Population in minister pixels

Nightlight intensity is by design a very broad measure and naturally raises the question: What is actually happening on the ground? In this section, we turn to another measure that lets us keep the large scale nature of our study, but sheds some light on this question. As we lay out in Section 3.2.3, we construct pixel-year population sums from the WorldPop Project data. We run our baseline specification employing this measure as the outcome variable.

Table 3.2 presents the results. For the world sample, we observe small negative effects that are statistically significant. Our results suggest an aggregate population decline between 1% and 2% in the minister birth pixels compared to the control group. For Africa and Europe, we find no effects. There is a smaller negative effect in the Asia sub-sample and a surprisingly large negative effect for the Americas sub-sample, that drives the world result.

Table 3.2: Treatment effects in minister birth pixels: population

	Dependent variable: population				
Aggregation method	(1) World	(2) Africa	(3) Europe	(4) Asia	(5) Americas
simple	-0.014***	-0.008	0.009	-0.013*	-0.053***
	(0.004)	(0.007)	(0.008)	(0.007)	(0.007)
dynamic	-0.028***	-0.016	0.008	-0.039***	-0.074***
	(0.007)	(0.011)	(0.010)	(0.013)	(0.013)
Observations	664,343	148,291	229,857	169,609	116,586

The dependent variable total population in the pixel is specified in logarithmic form. The method of aggregation *simple* is defined by Equation 3.3, and the *dynamic* aggregation is defined by Equation 3.4. To limit the duration of a plausible treatment effect on the outcome, we constrain the *dynamic* aggregation to 20 post-treatment periods. All estimations include covariates identifying the birth pixels of a country's primary rulers, capital cities, and a country-specific factor variable. Stars denote significance levels at 10%(*), 5%(**), and 1%(***). Heteroscedasticity-robust standard errors, which are clustered at the pixel level, are presented in parentheses.

We offer two interpretations for this finding: First, nomination of a minister and subsequent favoring of one ethnic or political group could increases out-group tensions leading to migration responses of the disfavored group. Second, the negative estimates could be a result of the not-switching treatment status of the estimator, meaning population decreases are driven by places that are no longer home of an active minister. As they no longer receive the benefits from being home to a high-ranking public official, firms and people relocate leading to population decline compared to the control group. Since treatment units over time eventually lose their active ministers, the persistent effect measured by the estimator captures this decline.

Indeed, in Section 3.4.3 we find evidence in support of this last point. When we estimate the average treatment effect based only on units with active treatment status, we find positive effects on population for ministers with long tenure. This implies that the negative effects of the estimator that measures persistent effects are driven by places where ministers loose their office.

Overall, our interpretation of the population results is that the regional favoritism effect we estimate in the nightlights appears to not induce persistent growth of the local population. However, taken together with the results from the **CH** estimator below, regional favoritism may lead to short-term migration responses.

3.4.3 Dynamic versus persistent effects

In this section, we test the robustness of our results to the use of alternative estimators. In particular, we benchmark our baseline results from the **CS** estimator against the canonical TWFE estimator and the **CH** estimator. Because the latter allows for treatment status switches, we can speak more to the persistence of the treatment effect, as well as the role that minister tenure plays by comparing its results to our baseline, that captures the persistent effect of ever having been the birthplace of a minister. We plot the results of all three estimators for nightlights in Figure 3.3 and for population in Figure 3.4.

Figure 3.3: Robustness of dynamic treatment effects in minister birth pixels: nightlights

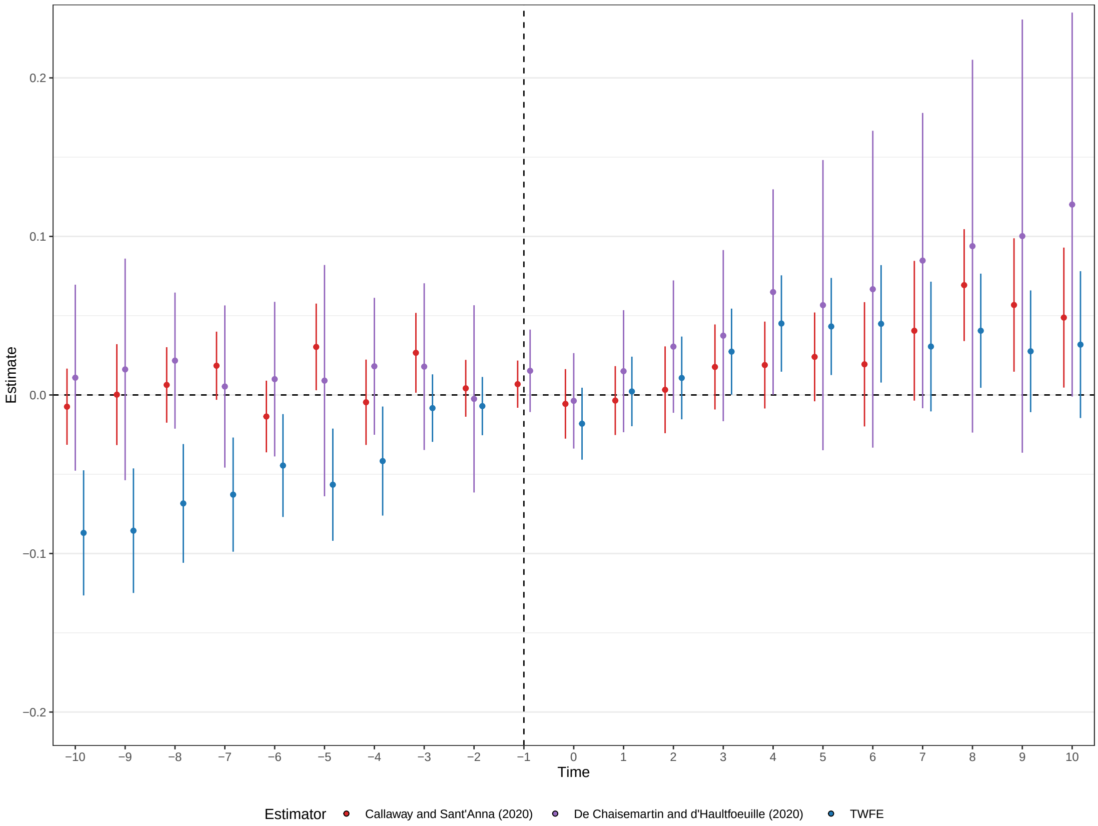

The figure displays event-studies that examine the relationship between the birthplaces of ministers in power and the (logarithm of) nightlight output at the grid-level. The estimators utilized include the dynamic version of the TWFE model (blue), Callaway and Sant'Anna (2021) (red), and De Chaisemartin and d'Haultfoeuille (2020) (purple). These estimates were computed using the did2s R package. Cell-level covariates include the birth pixels of a country's primary rulers, capital cities, and a country-specific factor variable. Comparison groups were defined by the default settings: not-yet treated and never-treated entities (cells). The x-axis represents time, measured in years, with the vertical reference line indicating the reference period. The bars on the plot represent 95 percent confidence intervals, with standard errors clustered at the cell level.

There are two core findings that we want to highlight. First, the classical TWFE displays significant pretrends for both our baseline analyses. Estimates from the TWFE are very likely to be biased in our setting with a strongly staggered and potentially heterogeneous treatment effect. Second, the **CH** estimates are consistently more positive and come with larger standard errors attached the further away from treatment. Both findings are in line with expectations, as the treatment effect in these specifications will be estimated only against observations with active and ongoing treatment. That means that for the late dynamic treatment effects the number of still treated observations goes down as ministers drop out of office, and naturally the precision of the estimates decreases. For the same reason it is sensible that the measured effects are more positive vis-à-vis our

Figure 3.4: Robustness of dynamic treatment effects in minister birth pixels: population

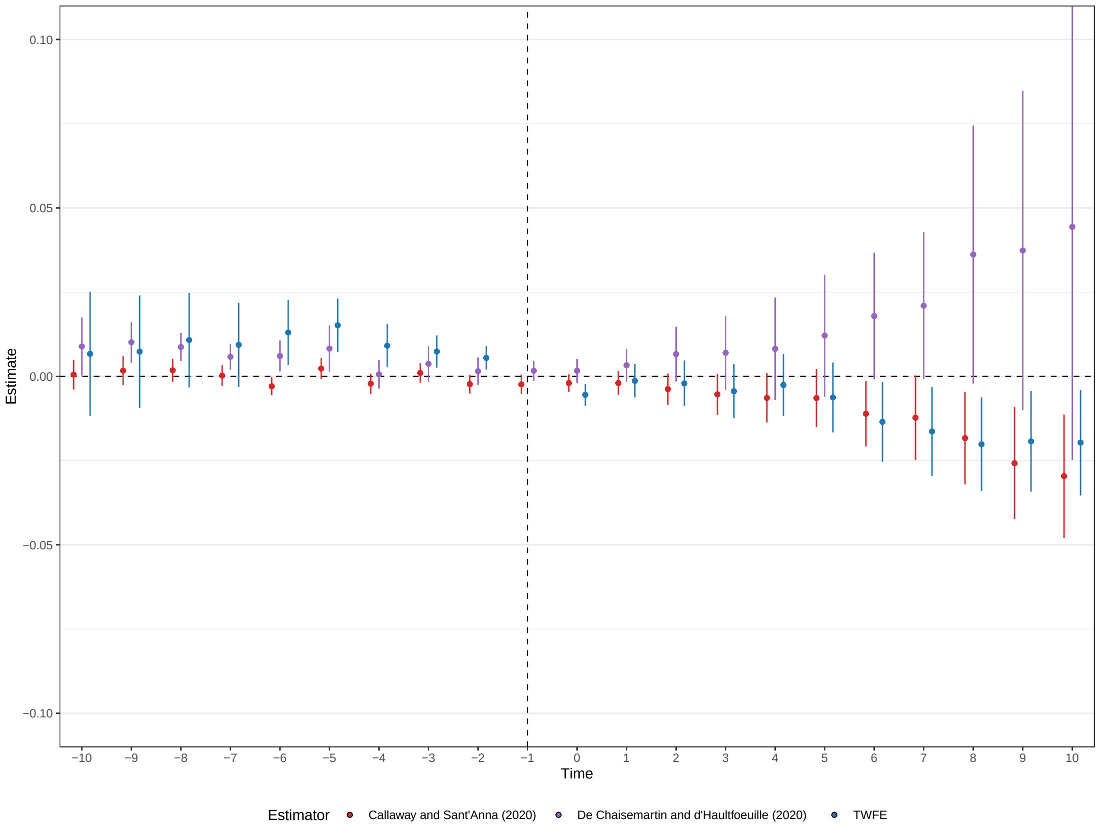

The figure displays event-studies that examine the relationship between the birthplaces of ministers in power and the (logarithm of) population at the grid-level. The estimators utilized include the dynamic version of the TWFE model (blue), Callaway and Sant'Anna (2021) (red), and De Chaisemartin and d'Haultfoeuille (2020) (purple). These estimates were computed using the did2s R package. Cell-level covariates include the birth pixels of a country's primary rulers, capital cities, and a country-specific factor variable. Comparison groups were defined by the default settings: not-yet treated and never-treated entities (cells). The x-axis represents time, measured in years, with the vertical reference line indicating the reference period. The bars on the plot represent 95 percent confidence intervals, with standard errors clustered at the cell level.

baseline estimator that estimates the persistent effect of ever having been treated and as such combines the effects of still treated and not anymore treated observations.

3.5 Mechanisms

In this section, we leverage cross-sectional variation in ministers' characteristics and cross-country variation in institutional settings to explore mechanisms that could explain the baseline effects.

3.5.1 Prestige levels and portfolios

The extent to which public officials can redirect resources to their birthplaces may depend on the perceived power of their office or the specific characteristics of the portfolio they oversee. We test both hypotheses below. First, we redefine the treatment variables of our main specification according to three prestige levels high, medium and low as classified by Nyrup and Bramwell (2020). The prestige category to which each portfolio belongs is detailed in Table C2.

In these specifications, we estimate the effect of a pixel ever having been the birthplace of a high-, medium- or low-prestige minister compared to all other pixels, including birthplaces of ministers from the other two categories. To account for the effects from ministers of the respective other prestige categories, we include dummy variables in the covariates matrix identifying their birthplace pixels.

Table 3.3: TREATMENT EFFECTS IN MINISTER BIRTH PIXELS BY PRESTIGE LEVEL

		Dependent variable: luminosity				
Aggregation method	Prestige	World	Africa	Europe	Asia	Americas
simple	High	0.057	0.119***	-0.109	0.048	-0.018
		(0.206)	(0.041)	(0.162)	(0.054)	(0.025)
	Medium	0.039**	0.117***	-0.079	0.013	-0.012
		(0.017)	(0.034)	(0.062)	(0.037)	(0.029)
	Low	-0.041	0.022	0.024	-0.224**	-0.105***
		(0.053)	(0.085)	(0.071)	(0.112)	(0.037)
dynamic	High	0.067	0.146***	-0.076	0.067	-0.031
		(0.310)	(0.051)	(0.140)	(0.105)	(0.031)
	Medium	0.074***	0.149***	-0.053	0.042	-0.027
		(0.024)	(0.039)	(0.060)	(0.074)	(0.033)
	Low	-0.029	0.018	0.093	-0.212**	-0.143***
		(0.053)	(0.083)	(0.092)	(0.099)	(0.051)
Observations	High	967,000	216,300	325,750	251,550	173,400
	Medium	960,675	212,225	324,925	250,875	172,650
	Low	973,625	221,150	326,100	252,125	174,250

The dependent variable average nighttime light intensity is specified in logarithmic form. The method of aggregation *simple* is defined by Equation 3.3, and the *dynamic* aggregation is defined by Equation 3.4. To limit the duration of a plausible treatment effect on the outcome, we constrain the *dynamic* aggregation to 20 post-treatment periods. All estimations include covariates identifying the birth pixels of a country's primary rulers, capital cities, and a country-specific factor variable. Stars denote significance levels at 10%(*), 5%(**), and 1%(***). Heteroscedasticity-robust standard errors, which are clustered at the pixel level, are presented in parentheses. Dynamic treatment effects are displayed in Figures C10, C11, and C12.

The results in Table 3.3 indicate that portfolios categorized as high and medium prestige are the primary drivers of the observed effects. In the global sample, we find positive treatment effects for both categories, but the effects for the high-prestige category are not statistically significant. This is likely due to the smaller sample size of high-prestige ministers and pronounced effect heterogeneities across continents. For pixels linked to medium-prestige ministers, we observe significant ATTs of 3.9% and 7.4%. By contrast, treatment effects for low-prestige ministers are negative and statistically insignificant.

When we disaggregate the analysis by continent, we find, consistent with the baseline results, that Africa drives the positive average effects in the global sample. We observe particular large estimates for the high- and medium-prestige portfolios for African countries. In African countries, the treatment effects are particularly large for high- and medium-prestige portfolios. In other continents, the results are less pronounced, insignificant, and occasionally negative for low-prestige ministers.

These findings suggest that political power, as proxied by the prestige associated with a ministerial portfolio, plays a significant role in enabling regional favoritism. However, beyond the general importance of a ministerial position, specific portfolio characteristics may also influence the ability to direct preferential transfers. To investigate this, we drill-down into the high-prestige category, which comprises four key portfolios: "defense, military & national security", "foreign relations", "finance, budget & treasury", and "government, interior & home affairs". Analogously to the specification for prestige levels, we redefine the treatment variables to include only ministers from these portfolios while controlling for the presence of any other minister, including those from the remaining high-prestige portfolios.

The results, presented in Table 3.4, align with the aggregated findings for the high-prestige category in Table 3.3. Generally, we observe positive treatment effects, with significant results for defense and foreign ministers in the global sample. The continent-level analysis reveals that these effects are primarily driven by Asia and Africa. In African countries, finance ministers also exhibit significant treatment effects, with estimates of 13.3% and 15.3% for foreign and finance ministers, respectively (dynamic, column 2). In Asia, foreign ministers show an effect of 19.3% (dynamic, column 4). These findings highlight that ministries with access to (flexible) financial resources are particularly influential in enabling regional favoritism.

The significant effects for foreign ministers resonate with the literature on the capture of foreign aid by primary leaders (Dreher et al., 2021b; Bomprezzi et al., 2024). This literature shows that foreign aid often benefits the birth regions of primary rulers or, depending on lender scrutiny, their spouses' birth regions. It is plausible that foreign ministers, by overseeing aid flows, can similarly channel resources toward their own birthplaces.

Finally, we observe very strong treatment effects, ranging from 17.1% to 29%, for defense ministers in Asia. This sub-sample includes countries such as Myanmar and Thailand, where the military has direct or de facto control for extended periods of our sample, as well as nations like Cambodia, Indonesia, Pakistan, and the Philippines, where the armed forces play significant roles in politics.

Table 3.4: TREATMENT EFFECTS IN MINISTER BIRTH PIXELS OF HIGH PRESTIGE PORTFOLIOS

		Dependent variable: luminosity				
Aggregation method	Portfolio	World	Africa	Europe	Asia	Americas
simple	Defense	0.093*	0.086	0.051	0.171***	-0.073
		(0.051)	(0.102)	(0.064)	(0.064)	(0.068)
	Foreign	0.086**	0.080	-0.005	0.075	0.071
		(0.038)	(0.062)	(0.094)	(0.067)	(0.246)
	Finance	0.056	0.112*	0.119	0.007	-0.060
		(0.073)	(0.059)	(0.135)	(0.084)	(0.057)
	Interior	0.030	0.040	0.116	0.158	-0.006
		(0.043)	(0.046)	(0.099)	(0.115)	(0.058)
dynamic	Defense	0.116*	0.074	0.119	0.290***	-0.137
		(0.058)	(0.163)	(0.090)	(0.105)	(0.094)
	Foreign	0.102**	0.133*	0.098	0.193**	0.243
		(0.044)	(0.076)	(0.070)	(0.085)	(0.190)
	Finance	0.083	0.153***	0.143	-0.058	-0.070
		(0.093)	(0.069)	(0.155)	(0.100)	(0.063)
	Interior	0.042	0.065	0.117	0.236	-0.014
		(0.048)	(0.050)	(0.115)	(0.159)	(0.065)
Observations	Defense	973,550	221,400	326,000	252,050	174,100
	Foreign	972,775	220,800	326,000	252,100	173,875
	Finance	973,350	221,075	326,050	252,150	174,075
	Interior	973,300	220,900	326,175	252,200	174,025

The dependent variable average nighttime light intensity is specified in logarithmic form. The method of aggregation *simple* is defined by Equation 3.3, and the *dynamic* aggregation is defined by Equation 3.4. To limit the duration of a plausible treatment effect on the outcome, we constrain the *dynamic* aggregation to 20 post-treatment periods. All estimations include covariates identifying the birth pixels of a country's primary rulers, capital cities, and a country-specific factor variable. Stars denote significance levels at 10%(*), 5%(**), and 1%(***). Heteroscedasticity-robust standard errors, which are clustered at the pixel level, are presented in parentheses. Dynamic treatment effects are displayed in Figures C13, C14 and C15.

3.5.2 Democracy versus autocracy

Next we investigate whether the institutional context mediates the effects we measured in the baseline specification. We interact the treatment variables of our main specification with a dummy indicating democratic and autocratic country-years according to the Freedom House classification. The treatment then occurs when the first autocratic (democratic) minister in our sample comes into office, while adding a dummy that indicates the existence of a democratic (autocratic) minister at any other time. We thus estimate the effect of having ever been the birth place of a minister in an autocratic or democratic regime on the nightlight intensity emitted by a pixel, compared to the other pixels. An alternative approach would be to split the sample into autocratic and democratic country-years. When comparing the two options, we choose the one that preserves the largest sample, as sample size reductions, and specifically the imbalance they introduce to the panel structure, impose additional restrictions on the estimator.

Table 3.5: TREATMENT EFFECTS IN MINISTER BIRTH PIXELS BY INSTITUTIONAL SETTING

		Dependent variable: luminosity				
Aggregation method		World	Africa	Europe	Asia	Americas
simple	Autocracy	0.085***	0.106***	0.056	0.019	-0.022
		(0.027)	(0.035)	(0.044)	(0.038)	(0.037)
	Democracy	-0.034	0.079	-0.036	0.011	-0.024
		(0.024)	(0.061)	(0.041)	(0.054)	(0.024)
dynamic	Autocracy	0.116***	0.141***	0.085*	0.045	-0.035
		(0.031)	(0.051)	(0.048)	(0.052)	(0.038)
	Democracy	-0.011	0.229*	-0.022	0.190	-0.026
		(0.047)	(0.131)	(0.068)	(0.116)	(0.036)
Observations	Autocracy	961,350	211,325	325,600	250,850	173,575
	Democracy	971,875	221,750	325,500	251,850	172,775

The dependent variable average nighttime light intensity is specified in logarithmic form. The method of aggregation *simple* is defined by Equation 3.3, and the *dynamic* aggregation is defined by Equation 3.4. To limit the duration of a plausible treatment effect on the outcome, we constrain the *dynamic* aggregation to 20 post-treatment periods. All estimations include covariates identifying the birth pixels of a country's primary rulers, capital cities, and a country-specific factor variable. Stars denote significance levels at 10%(*), 5%(**), and 1%(***). Heteroscedasticity-robust standard errors, which are clustered at the pixel level, are presented in parentheses. Dynamic treatment effects are displayed in Figure C16 and Figure C17.

Table 3.5 shows the results. We take note of two findings: For autocratic settings, we measure large positive effects. The effects are statistically significant for the world and the African and European sub-sample. For democratic settings, we observe a close to zero result for the full sample. The sub-sample analysis reveals some tentative evidence for sizeable positive effects in democratic countries of the African and Asian continent, however both come with large standard errors attached to them.

Conceptually it is not unambiguous which institutional setting should come up with the larger effects. We think of the institutional context as a mediator that affects both the possibility to engage in regional favoritism, as well as the incentives to do so. While autocratic ministers might be less constrained to engage in favoritism than their democratic counterparts, they might face a lower incentive to share rents broadly, as they face less electoral competition. Our results in this section then suggest that the restrictive features of some democracies in our samples dominate these electoral incentives, giving rise to the stronger observable effects in autocratic settings.

3.5.3 OECD versus non-OECD

Membership in the Organisation for Economic Co-operation and Development (OECD) serves as an indication of a country's economic and political development. Typically, OECD members are high-income economies with a high Human Development Index. These nations are typically democratic, with market-based economies, regulated by international standards and norms set by the OECD. Thus, an OECD membership not only indicates economic prosperity, but also reflects a country's commitment to democratic principles, free market practices, and global policy cooperation.

Table 3.6: Treatment effects in minister birth pixels by OECD membership

	Dependent variable: luminosity	
Aggregation method	OECD	Non-OECD
simple	-0.092	0.062***
	(0.061)	(0.018)
dynamic	-0.054	0.093***
	(0.059)	(0.024)
Observations	60,350	897,000

The dependent variable average nighttime light intensity is specified in logarithmic form. The method of aggregation *simple* is defined by Equation 3.3, and the *dynamic* aggregation is defined by Equation 3.4. To limit the duration of a plausible treatment effect on the outcome, we constrain the *dynamic* aggregation to 20 post-treatment periods. All estimations include covariates identifying the birth pixels of a country's primary rulers, capital cities, and a country-specific factor variable. Stars denote significance levels at 10%(*), 5%(**), and 1%(***). Heteroscedasticity-robust standard errors, which are clustered at the pixel level, are presented in parentheses.

We divide our sample of countries into OECD and non-OECD members and reestimate the baseline equations. Table 3.6 reports the results. We detect evidence for regional favoritism as indicated by nightlight intensity only in non-OECD countries. This finding serves as a further puzzle piece that points to the role robust institutions play in constraining the ability of politicians to redistribute resources to their birth places. OECD countries generally have stronger institutions and governance structures, as well as higher levels of transparency, all of which can help deter regional favoritism.

It is important to note that nighttime light luminosity may not serve as an effective measure of economic development in industrialized countries, such as OECD-countries.

This observation does not dismiss the validity of nighttime light luminosity as a global indicator; rather, it suggests that its interpretative power may be limited in the context of highly developed economies.

Using luminosity as an indicator of economic activity is particularly useful in developing countries. However, in industrialized nations, it might not serve as an accurate measure for several reasons. First, these countries typically have widespread and uniformly high illumination, making it challenging to spot differences in economic activity based only on nightlight data. This is exacerbated by the fact that nightlight data in dense urban areas is top coded. Second, energy efficiency measures and regulations against light pollution can further reduce the perceived nightlight output. Third, significant service and digital sectors in these countries may not correlate with high nightlight output. Therefore, while nightlight output may be useful in certain contexts, it may not accurately represent economic development in industrialized countries (Gibson et al., 2021).

3.5.4 Corruption

To gain further insights into the mediating role of the institutional setting, we use the Corruption Perception Index (CPI) developed by Transparency International. measures perceived levels of public sector corruption globally, combining various indices based on surveys of businesspeople and assessments by country experts. These assessments rate countries on their perceived corruption levels, providing a composite score. The CPI scores range from 0 to 100, where 0 indicates a country is perceived as highly corrupt, and 100 reflects very low perceived corruption levels.[7]

We divide our sample into two groups based on the CPI threshold of 50, categorizing countries as either more or less corrupt. The results, summarized in Table 3.7, reveal that minister pixels in more corrupt countries are significantly more likely to exhibit regional favoritism. This evidence suggests that corruption substantially influences regional resource allocation among the governing elite.

The results are summarized in Table 3.7. Our analysis reveals that regions with higher perceived corruption are more likely to exhibit regional favoritism (dynamic, column 2). This evidence suggests that corruption may considerably influence resource allocation among the ruling elite, particularly in environments with less robust institutions. Interestingly, the coefficients in column 1 indicate potential reverse favoritism in less corrupt countries. However, these countries are primarily industrialized nations where nighttime light, as a proxy for economic development, may have limitations, as discussed earlier. Therefore, this finding should be interpreted cautiously.

[7]It is important to note that the CPI measures perceptions of corruption rather than actual levels. Since corruption typically occurs behind closed doors and is difficult to observe directly, perception-based measures provide a reasonable second-best for the measurement of actual corruption, as surveys and expert assessments are likely highly correlated with true corruption levels.

Table 3.7: Treatment effects in minister birth pixels by corruption

	Dependent variable: luminosity	
Aggregation method	Less corrupt	More corrupt
simple	-0.198	0.059***
	(0.149)	(0.015)
dynamic	-0.296***	0.099***
	(0.115)	(0.023)
Observations	24,900	932,450

The dependent variable average nighttime light intensity is specified in logarithmic form. The method of aggregation *simple* is defined by Equation 3.3, and the *dynamic* aggregation is defined by Equation 3.4. To limit the duration of a plausible treatment effect on the outcome, we constrain the *dynamic* aggregation to 20 post-treatment periods. All estimations include covariates identifying the birth pixels of a country's primary rulers, capital cities, and a country-specific factor variable. Stars denote significance levels at 10%(*), 5%(**), and 1%(***). Heteroscedasticity-robust standard errors, which are clustered at the pixel level, are presented in parentheses.

3.5.5 Women ministers

Prior literature showed that policy makers' gender can interact in various ways with the outcome of their governance (see Hessami and da Fonseca (2020) for a comprehensive review). Hence, we ask: Do men and women ministers engage in regional favoritism? We redefine the treatment variables in our main specification based on the gender of ministers and estimate the potential impact on a pixel of ever having been the birthplace of a woman minister compared to all other pixels, including birthplaces of men ministers. These specifications add dummy variables that identify the birthplace pixels of men ministers to our covariates matrix.

Table 3.8: Treatment effects in female minister birth pixels

	Dependent variable: luminosity				
Aggregation method	(1) World	(2) Africa	(3) Europe	(4) Asia	(5) Americas
simple	-0.050	-0.003	-0.141	-0.148**	0.004
	(0.033)	(0.038)	(0.123)	(0.065)	(0.033)
dynamic	-0.074	-0.015	-0.155	-0.174*	-0.047
	(0.043)	(0.052)	(0.160)	(0.077)	(0.039)
Observations	974,050	221,875	325,925	252,150	174,100

The dependent variable average nighttime light intensity is specified in logarithmic form. The method of aggregation *simple* is defined by Equation 3.3, and the *dynamic* aggregation is defined by Equation 3.4. To limit the duration of a plausible treatment effect on the outcome, we constrain the *dynamic* aggregation to 20 post-treatment periods. All estimations include covariates identifying the birth pixels of a country's primary rulers, capital cities, and a country-specific factor variable. Stars denote significance levels at 10%(*), 5%(**), and 1%(***). Heteroscedasticity-robust standard errors, which are clustered at the pixel level, are presented in parentheses.

The results in Table 3.8 suggest that women ministers do not engage in regional favoritism. This finding aligns with literature positing that greater representation of women enhances institutional quality by reducing corruption (Hessami and da Fonseca, 2020). However, since women's participation at the level of ministers is low, only around 10% in our sample are women, inference is based on much less identifying variation.

3.6 Conclusion

This paper demonstrates that ministers possess the ability to and do actively engage in regional favoritism. To quantify: The largest increases in nighttime light intensity that we measured between 9.4% and up to 18.7% in the African sub-sample, translate into average local GDP growth of 2.8% to 5.6% using the correlation of 0.3 suggested by Henderson et al. (2012).

We also find that tenure is an important factor, as the effects of regional favoritism intensify with longer exposure to treatment. Furthermore, cross-sectional heterogeneity analyses reveal that the most powerful ministers, particularly those with direct control over budgets, drive the baseline effects. Regional favoritism is more pronounced in autocratic and corrupt countries, where weaker institutional frameworks provide greater opportunities for resource misallocation. In contrast, stronger democratic institutions appear to restrain ministers from channeling resources disproportionately toward their birthplaces.

These findings suggest that policy interventions aimed at strengthening institutional checks and balances, enhancing oversight of powerful ministers, particularly those controlling easily misdirected funds such as foreign aid, and enforcing term limits could help mitigate regional favoritism. There is also tentative evidence that increasing the representation of women in ministerial positions may further reduce the extent of favoritism.

Beyond its immediate findings, this paper highlights the value of the geocoded cabinet member dataset collected for this study. The dataset offers unique advantages for studying political economy questions where the spatial distribution of the governing elite is relevant. Compared to conventional data focusing solely on primary rulers, this dataset allows for a more comprehensive analysis by incorporating the dynamics within the governing elite; providing greater identifying variation across time and space due to the larger sample size of cabinet members; enabling cross-sectional variation in political portfolios, prestige levels, and ideology; and facilitating within-unit variation, capturing not only changes over a minister's tenure but also transitions between different ministerial portfolios.

There remain many open questions and avenues for future research in this field. While nighttime light data provide a globally consistent measure of local economic activity, future studies could benefit from incorporating alternative measures, such as land classifications derived from high-resolution daytime satellite imagery, which may better capture economic activity in urban and highly developed areas. Additionally, a systematic investigation into the factors determining why certain regions produce ministers while others do not could address potential concerns regarding reverse causality. Although the evidence from our event study suggests reverse causality in our sample is not a concern on average, a deeper understanding of regional political participation and its determinants would be valuable.

References

Albouy, David, "Partisan Representation in Congress and the Geographic Distribution of Federal Funds.," *Review of Economics and Statistics*, 2013, *95* (1), 127–141.

Alesina, A., S. Michalopoulos, and E. Papaioannou, "Ethnic inequality," *Journal of Political Economy*, 2016, *124*, 428–488.

Alesina, Alberto and Guido Tabellini, "Bureaucrats or Politicians? Part I: A Single Policy Task," *American Economic Review*, March 2007, *97* (1), 169–179.

_ and _ , "Bureaucrats or politicians? Part II: Multiple policy tasks," *Journal of Public Economics*, April 2008, *92* (3), 426–447.

Asatryan, Z., T. Baskaran, C. Birkholz, and D. Gomtsyan, "Favoritism and Firms: Micro evidence and macro implications," 2021. Published: ZEW Discussion Paper 21-031.

_ , _ , P. Hufschmidt, and A. Stöcker, "Regional favoritism and human capital accumulation in Africa," 2021. Published: ZEW Discussion Paper 21-030.

Asatryan, Zareh and Annika Havlik, "The political economy of multilateral lending to European regions," *The Review of International Organizations*, 2020, *15* (3), 707–740. Publisher: Springer.

_ , Thushyanthan Baskaran, Carlo Birkholz, and Patrick Hufschmidt, "Regional Redistribution of Mineral Resource Wealth in Africa," SSRN Scholarly Paper ID 3820533, Social Science Research Network, Rochester, NY 2021.

Baker, Andrew C., David F. Larcker, and Charles C. Y. Wang, "How much should we trust staggered difference-in-differences estimates?," *Journal of Financial Economics*, May 2022, *144* (2), 370–395.

Barro, R. J., "The control of politicians: an economic model," *Public Choice*, 1973, *14*, 19–42.

Baskaran, T. and M. L. da Fonseca, "Appointed public officals and local favoritism: evidence from the German States," *Journal of Urban Economics*, 2021, *124*, 103354.

_ and Z. Hessami, "Political alignment and intergovernmental transfers in parliamentary systems: evidence from Germany," *Public Choice*, 2017, *171*, 75–98.

Besley, T., "Political Selection," *Journal of Economic Perspectives*, 2005, *19*, 43–60.

Besley, Timothy and Stephen Coate, "Elected versus Appointed Regulators: Theory and Evidence," *Journal of the European Economic Association*, September 2003, *1* (5), 1176–1206.

Bomprezzi, Pietro, Axel Dreher, Andreas Fuchs, Teresa Hailer, Andreas Kammerlander, Lennart C Kaplan, Silvia Marchesi, Tania Masi, Charlotte Robert, and Kerstin Unfried, "Wedded to Prosperity? Informal Influence and Regional Favoritism," 2024.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event Study Designs: Robust and Efficient Estimation," April 2022.

Bruederle, Anna and Roland Hodler, "Nighttime lights as a proxy for human development at the local level," *PloS one*, 2018, *13* (9), e0202231. Publisher: Public Library of Science.

Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padró i Miquel, "The Value of Democracy: Evidence from Road Building in Kenya," *American Economic Review*, jun 2015, *105* (6), 1817–1851.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, December 2021, *225* (2), 200–230.

Carozzi, Felipe and Luca Repetto, "Sending the pork home: Birth town bias in transfers to Italian municipalities," *Journal of Public Economics*, feb 2016, *134*, 42–52.

Cox, Gary W., "Swing voters, core voters, and distributive politics," in Alexander S. Kirshner, Elisabeth Jean Wood, Ian Shapiro, and Susan C. Stokes, eds., *Political Representation*, Cambridge: Cambridge University Press, 2010, pp. 342–357.

_ **and Mathew D. McCubbins**, "Electoral Politics as a Redistributive Game," *The Journal of Politics*, May 1986, *48* (2), 370–389. Publisher: The University of Chicago Press.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, September 2020, *110* (9), 2964–2996.

Do, Quoc-Anh, Kieu-Trang Nguyen, and Anh Tran, "One mandarin benefits the whole clan: hometown favoritism in an authoritarian regime," *American Economic Journal: Applied Economics*, 2017, *9*, 1–29.

Dreher, A., A. Fuchs, R. Hodler, B. C. Parks, P. A. Raschky, and M. J. Tierney, "Is favoritism a threat to chinese aid effectiveness? A subnational analysis of chinese development projects," *World Development*, 2021, *139*, 105291.

Dreher, Axel, Andreas Fuchs, Bradley Parks, Austin Strange, and Michael J Tierney, "Aid, China, and growth: Evidence from a new global development finance dataset," *American Economic Journal: Economic Policy*, 2021, *13* (2), 135–174.

Francois, Patrick, Ilia Rainer, and Francesco Trebbi, "How Is Power Shared in Africa?," *Econometrica*, 2015, *83* (2), 465–503. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA11237.

Freedom House, "Freedom in the world 2019," 2019. Published: www.freedomhouse.org.

Ghosh, Tilottama, Kimberly E. Baugh, Christopher D. Elvidge, Mikhail Zhizhin, Alexey Poyda, and Feng-Chi Hsu, "Extending the DMSP Nighttime Lights Time Series beyond 2013," *Remote Sensing*, January 2021, *13* (24), 5004. Number: 24 Publisher: Multidisciplinary Digital Publishing Institute.

Gibson, John, Susan Olivia, Geua Boe-Gibson, and Chao Li, "Which night lights data should we use in economics, and where?," *Journal of Development Economics*, March 2021, *149*, 102602.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, December 2021, *225* (2), 254–277.

Henderson, J. V., A. Storeygard, and D. N. Weil, "Measuring economic growth from outer space," *American Economic Review*, 2012, *102*, 994–1028.

Hessami, Zohal and Mariana Lopes da Fonseca, "Female political representation and substantive effects on policies: A literature review," *European Journal of Political Economy*, 2020, *63*, 101896.

Hodler, Roland and Paul A. Raschky, "Regional favoritism," *The Quarterly Journal of Economics*, May 2014, *129* (2), 995–1033.

Khalil, Umair, Mandar Oak, and Sundar Ponnusamy, "Political favoritism by powerful politicians: Evidence from chief ministers in India," *European Journal of Political Economy*, January 2021, *66*, 101949.

Martínez, L. R., "How much should we trust the dictator's GDP growth estimates?," *Journal of Political Economy*, 2022, *forthcoming.*

Maskin, Eric and Jean Tirole, "The Politician and the Judge: Accountability in Government," *American Economic Review*, September 2004, *94* (4), 1034–1054.

Michalopoulos, S. and E. Papaioannou, "The long-run effects of the scramble for Africa," *American Economic Review*, 2016, *106*, 1802–1848.

Nyrup, Jacob and Stuart Bramwell, "Who Governs? A New Global Dataset on Members of Cabinets," *American Political Science Review*, November 2020, *114* (4), 1366–1374. Publisher: Cambridge University Press.

Roth, Jonathan, Pedro HC Sant'Anna, Alyssa Bilinski, and John Poe, "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

Stevens, Forrest R., Andrea E. Gaughan, Catherine Linard, and Andrew J. Tatem, "Disaggregating Census Data for Population Mapping Using Random Forests with Remotely-Sensed and Ancillary Data," *PLOS ONE*, February 2015, *10* (2), e0107042. Publisher: Public Library of Science.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, December 2021, *225* (2), 175–199.

Transparency International, "Corruption Perceptions Index," `https://www.transparency.org/en/cpi` 2022. Data retrieved for multiple years.

Weingast, B., K. Shepsle, and C. Johnsen, "The political economy of costs and benefits: a neoclassical approach to distributive politics," *Journal of Political Economy*, 1981, *89* (4), 642–664.

Widmer, Philine and Noémie Zurlinden, "Ministers Engage in Favoritism Too," *Journal of Public Economics*, September 2022, *213*, 104707.

WorldPop and CIESIN, "Global High Resolution Population Denominators Project," `http://www.worldpop.org` 2018. Funded by The Bill and Melinda Gates Foundation (OPP1134076). Collaboration between School of Geography and Environmental Science, University of Southampton; Department of Geography and Geosciences, University of Louisville; and Département de Géographie, Université de Namur.

C Appendix

C.1 Additional data description

Table C1: Countries and years of collected birthplaces

Country	Continent	Years	Country	Continent	Years
Algeria	Africa	1966-2016	South Africa	Africa	1966-2016
Angola	Africa	1975-2015	South Sudan	Africa	2012-2016
Benin	Africa	1966-2016	Sudan	Africa	1989-2012
Botswana	Africa	1967-2015	Eswatini	Africa	1986-2012
Burkina Faso	Africa	1966-2016	Tanzania	Africa	1978-2016
Burundi	Africa	1966-2016	Togo	Africa	1966-2015
Cameroon	Africa	1966-2016	Tunisia	Africa	1966-2016
Cape Verde	Africa	1976-2017	Uganda	Africa	1966-2016
Central African Republic	Africa	1966-2015	Zambia	Africa	1966-2016
Chad	Africa	1966-2016	Zimbabwe	Africa	1981-2016
Comoros	Africa	1966-2016	Afghanistan	Asia	1990-2018
Djibouti	Africa	1977-2016	Armenia	Asia	1992-2016
Egypt	Africa	1966-2016	Azerbaijan	Asia	1992-2016
Equatorial Guinea	Africa	1978-2016	Bangladesh	Asia	1972-2016
Eritrea	Africa	1993-2006	Bhutan	Asia	1973-2014
Ethiopia	Africa	1966-2016	Cambodia	Asia	1967-2016
Gabon	Africa	1966-2016	China	Asia	1982-2014
Gambia	Africa	1966-2015	Georgia	Asia	1992-2016
Ghana	Africa	1966-2014	India	Asia	1980-2018
Guinea	Africa	1966-2016	Indonesia	Asia	1991-2016
Côte d'Ivoire	Africa	1987-2013	Iraq	Asia	1990-2015
Kenya	Africa	1966-2016	Israel	Asia	1970-2017
Lesotho	Africa	1967-2016	Jordan	Asia	1975-2016
Liberia	Africa	1991-2014	Kazakhstan	Asia	1992-2016
Libya	Africa	1991-2015	Kyrgyz Republic	Asia	1992-2016
Madagascar	Africa	1984-2016	Lao PDR	Asia	1966-2016
Malawi	Africa	1966-2014	Lebanon	Asia	1967-2014
Mali	Africa	1969-2016	Malaysia	Asia	1968-2016
Mauritania	Africa	1972-2016	Mongolia	Asia	1972-2016
Mauritius	Africa	1966-2016	Myanmar	Asia	1986-2016
Morocco	Africa	1966-2014	Nepal	Asia	1990-2016
Mozambique	Africa	1975-2016	Pakistan	Asia	1966-2016
Namibia	Africa	1990-2015	Philippines	Asia	1966-2018
Niger	Africa	1974-2016	Sri Lanka	Asia	1990-2016
Nigeria	Africa	1966-2016	Tajikistan	Asia	1992-2015
Congo	Africa	1966-2017	Thailand	Asia	1966-2016
Rwanda	Africa	1966-2015	Timor-Leste	Asia	2002-2018
São Tomé Príncipe	Africa	1975-2017	Turkey	Asia	1966-2016
Senegal	Africa	1966-2015	Uzbekistan	Asia	1992-2015
Sierra Leone	Africa	1966-2014	Vietnam	Asia	1976-2016
Somalia	Africa	1969-2016	Yemen	Asia	1966-2016

Continued on next page

Table C1 – continued from previous page					
Country	**Continent**	**Years**	**Country**	**Continent**	**Years**
Albania	Europe	1990-2016	Ukraine	Europe	1992-2016
Austria	Europe	1990-2018	United Kingdom	Europe	1990-2018
Belarus	Europe	2006-2014	Canada	North America	1990-2018
Belgium	Europe	1990-2018	Costa Rica	North America	1973-2018
Bosnia and Herzegovina	Europe	1992-1998	Dominican Republic	North America	1982-2018
Bulgaria	Europe	2006-2016	El Salvador	North America	1966-2015
Croatia	Europe	2006-2016	Guatemala	North America	1990-2016
Czech Republic	Europe	2006-2016	Honduras	North America	1990-2015
Denmark	Europe	1990-2019	Mexico	North America	1990-2018
Estonia	Europe	2006-2018	Nicaragua	North America	1979-2015
Finland	Europe	1990-2018	Panama	North America	1995-2015
France	Europe	1991-2017	United States	North America	1991-2018
Germany	Europe	1990-2018	Argentina	South America	1989-2018
Greece	Europe	1981-1993	Bolivia	South America	1986-2016
Hungary	Europe	2006-2016	Brazil	South America	1990-2016
Italy	Europe	1990-2018	Chile	South America	1990-2018
Lithuania	Europe	1992-2016	Colombia	South America	1973-2018
Moldova	Europe	1992-2016	Ecuador	South America	1990-2016
Montenegro	Europe	1997-2016	Guyana	South America	1984-2016
North Macedonia	Europe	1995-2016	Paraguay	South America	1990-2016
Norway	Europe	1990-2018	Peru	South America	1984-2016
Netherlands	Europe	1990-2018	Suriname	South America	1979-2018
Poland	Europe	1990-2016	Trinidad and Tobago	South America	1978-2018
Portugal	Europe	1990-1992	Uruguay	South America	1990-2015
Romania	Europe	1990-2016	Venezuela	South America	1990-2016
Russia	Europe	1992-2015	Australia	Oceania	1990-2018
Slovak Republic	Europe	1993-2016	Fiji	Oceania	1992-2016
Slovenia	Europe	1992-2016	New Zealand	Oceania	1990-2018
Spain	Europe	1990-2015			
Sweden	Europe	1990-2016			

Figure C1: Birthplaces of cabinet members in Africa

Each red dot in the figure represents the birthplace of a unique national cabinet member in our sample for Africa.

Figure C2: Birthplaces of cabinet members in Asia

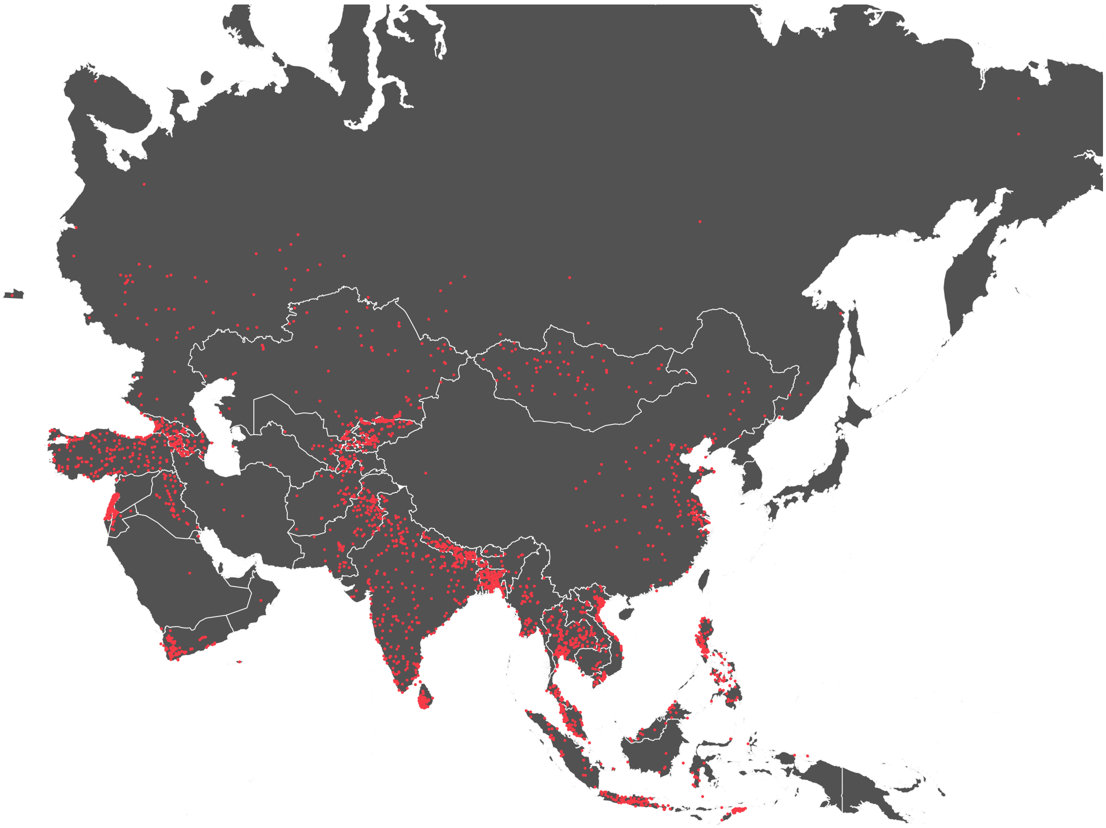

Each red dot in the figure represents the birthplace of a unique national cabinet member in our sample for Asia.

Figure C3: Birthplaces of cabinet members in Europe

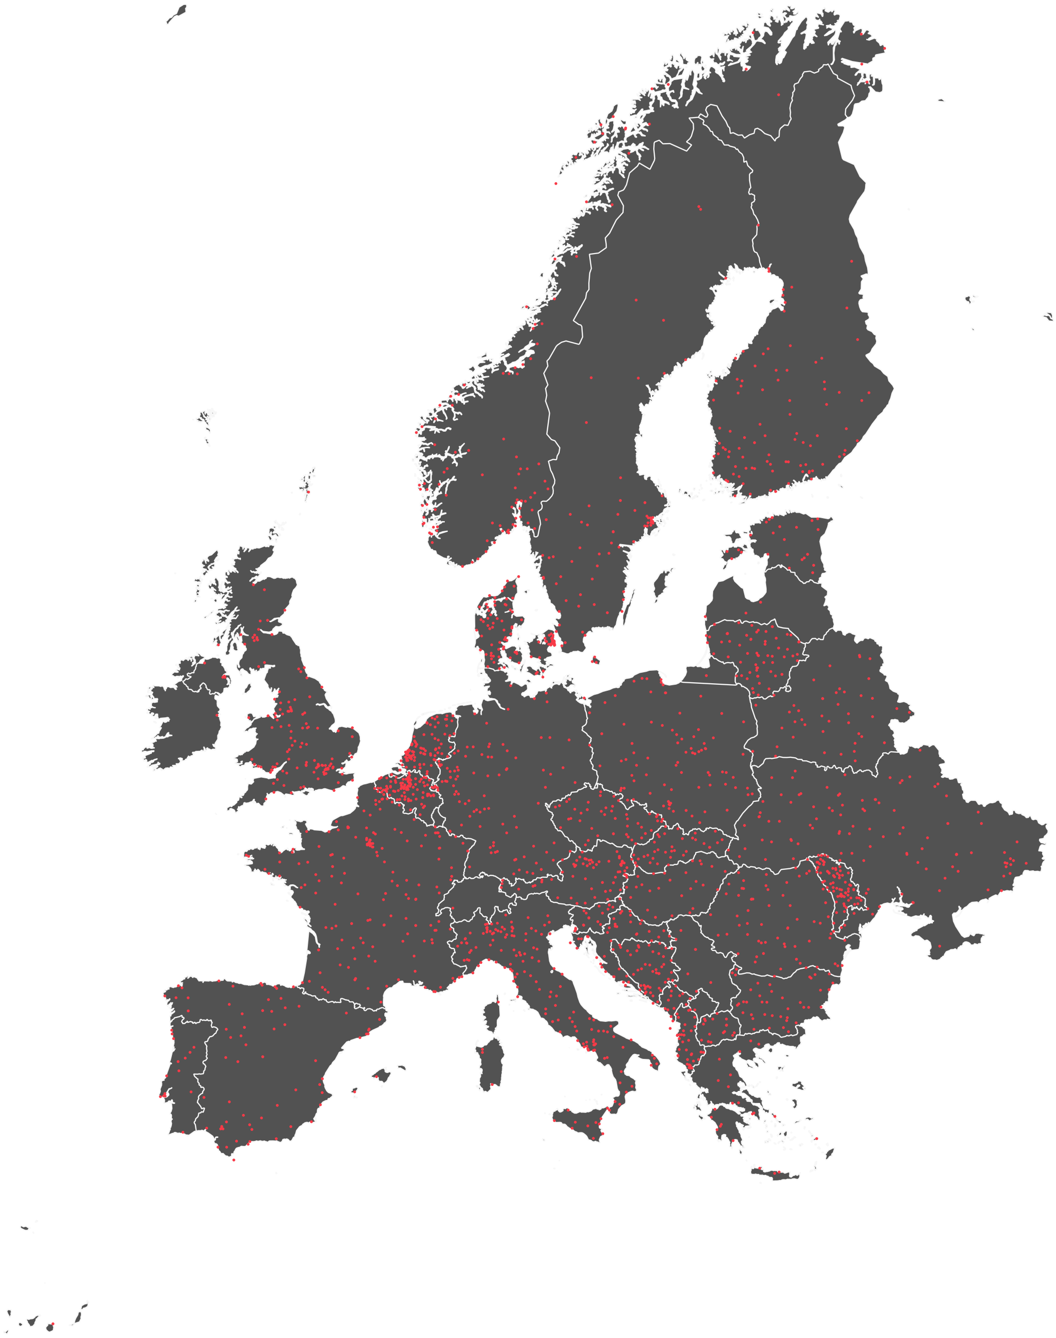

Each red dot in the figure represents the birthplace of a unique national cabinet member in our sample for Europe.

Figure C4: Birthplaces of cabinet members in the Americas

Each red dot in the figure represents the birthplace of a unique national cabinet member in our sample for the Americas.

Figure C5: Birthplaces of cabinet members in Oceania

Each red dot in the figure represents the birthplace of a unique national cabinet member in our sample for Oceania.

Table C2: Portfolios and prestige level categories

Portfolio	Prestige
Defense, Military & National Security	High
Foreign Relations	High
Government, Interior & Home Affairs	High
Finance, Budget & Treasury	High
Agriculture, Food, Fisheries & Livestock	Medium
Audit, Oversight & Internal Affairs	Medium
Civil Service	Medium
Communications & Information	Medium
Construction & Public Works	Medium
Correctional Services & Police	Medium
Culture & Heritage	Medium
Education, Training & Skills	Medium
Energy	Medium
Enterprises, Companies & Business	Medium
Environment	Medium
Executive & Legislative Relations	Medium
Foreign Economic Relations	Medium
General Economic Affairs	Medium
Health & Social Welfare	Medium
Housing	Medium
Industry & Commerce	Medium
Justice & Legal Affairs	Medium
Labor, Employment & Social Security	Medium
Medium Local Government	Medium
Planning & Development	Medium
Political Reform	Medium
Properties & Buildings	Medium
Religion	Medium
Regional	Medium
Tax, Revenue & Fiscal Policy	Medium
Transport	Medium
Ageing & Elderly	Low
Children & Family	Low
Immigration & Emigration	Low
Minorities	Low
Science, Technology & Research	Low
Sports	Low
Tourism	Low
Veterans	Low
Without Portfolio	Low
Women	Low
Youth	Low

Source: Nyrup and Bramwell (2020)

Figure C6: Grid of the world displaying mean nightlight intensity

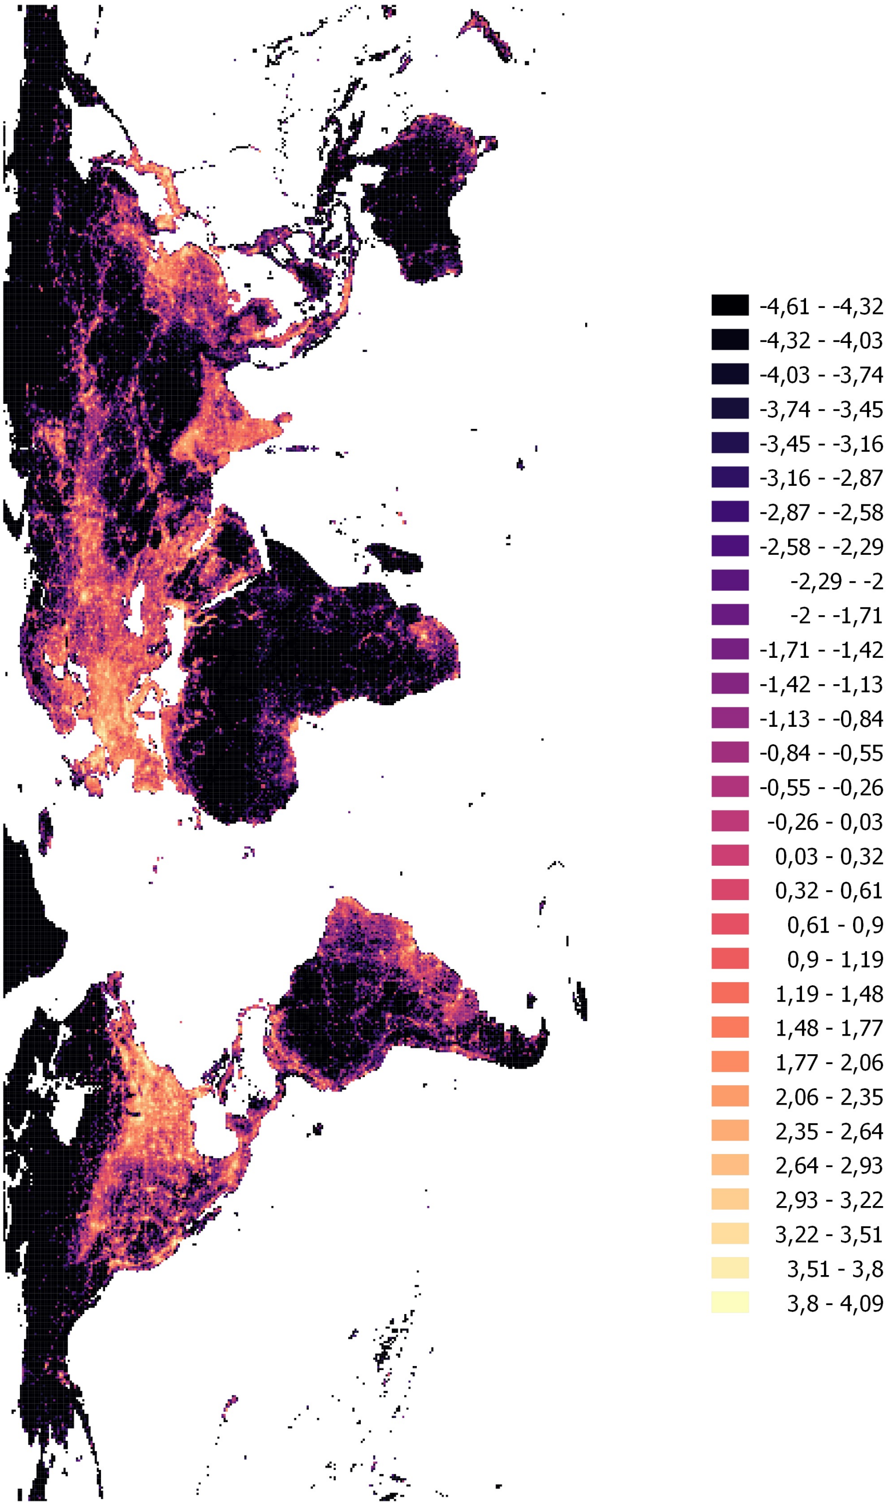

The figure shows (the logarithm) of mean night light output for the period of our sample (1992-2016). The values for the pixels were computed by extracting information from the night light raster files based on the grid of the world utilized in our empirical analysis. For this process we used the exactextractr R package. Brighter cells indicate higher nighttime light intensity. The corresponding values are tabulated in the legend.

Figure C7: Grid of the world displaying mean population

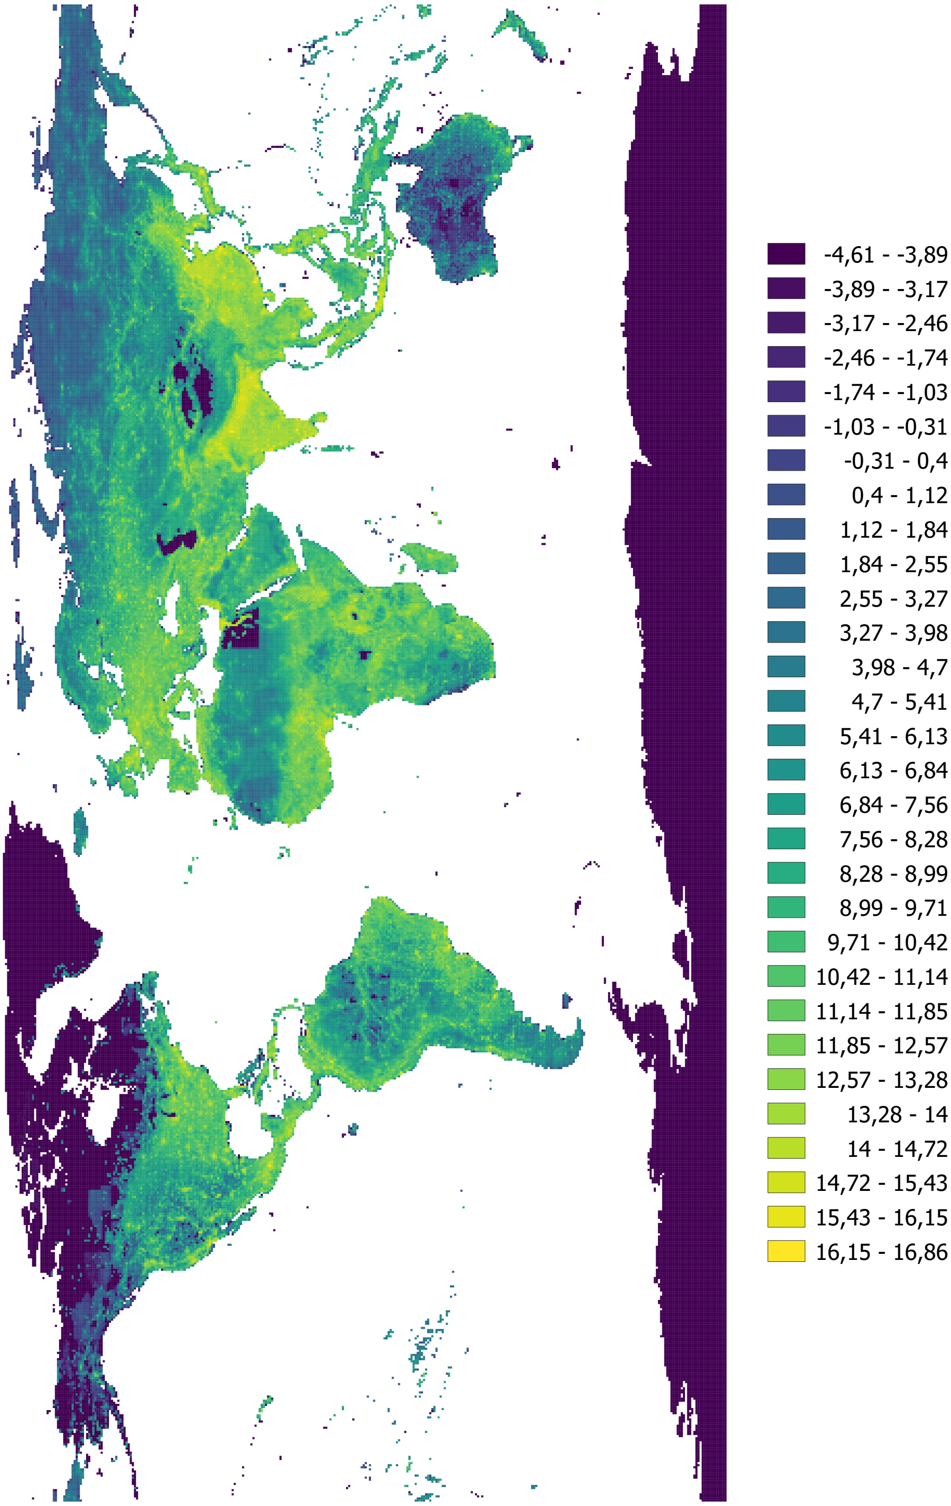

The figure shows (the logarithm of) sum population for the period (2000-2016). The values for the pixels were computed by extracting information from the population raster files based on the grid of the world utilized in our empirical analysis. For this process we used the exactextractr R package. Brighter cells indicate higher population numbers. The corresponding values are tabulated in the legend.

C.2 Additional results: baseline

Figure C8: Dynamic treatment effects in minister birth pixels: nightlights

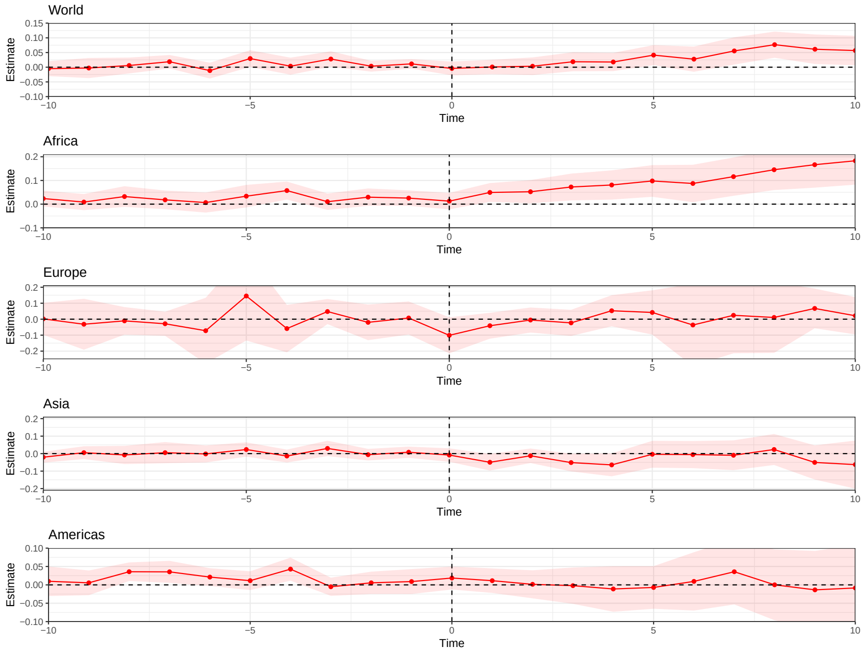

The figure shows an event study of the baseline specification for the outcome log nightlight intensity by continent. The red shaded areas on the plot represent 95 percent confidence intervals.

Figure C9: Dynamic treatment effects in minister birth pixels: population

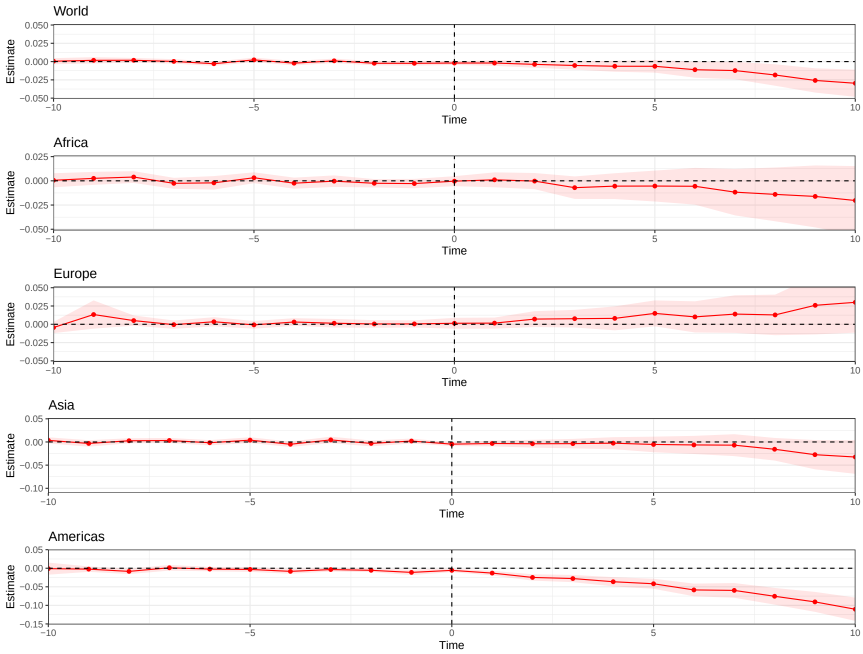

The figure shows an event study of the baseline specification for the outcome log total population by continent. The red shaded areas on the plot represent 95 percent confidence intervals.

C.3 Additional results: mechanisms

C.3.1 Prestige levels

Figure C10: Dynamic treatment effects in high prestige minister birth pixels

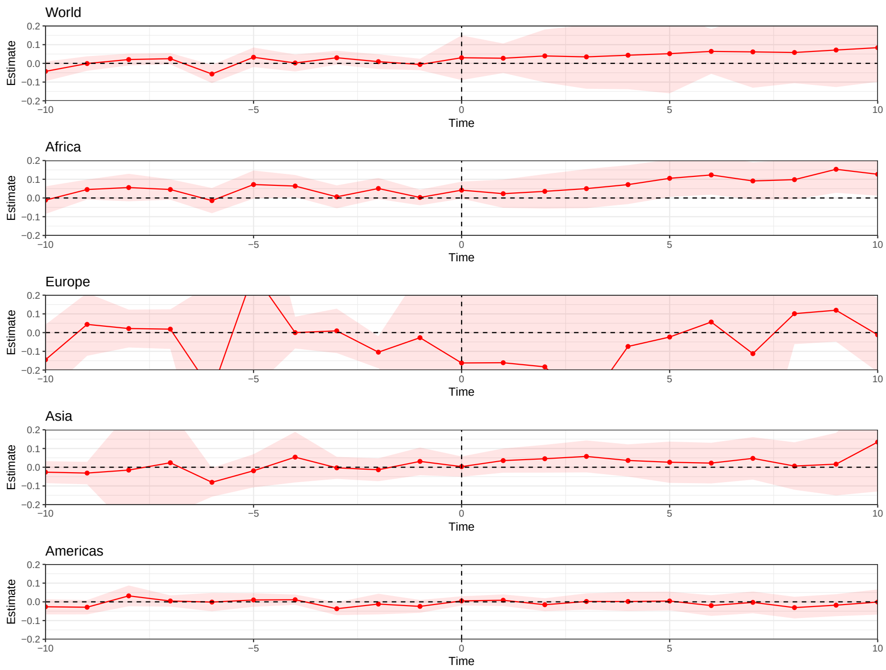

The figure shows an event study for the outcome log nightlight intensity where treatment is restricted to cover high prestige ministers only. Dummy variables identifying all non-high prestige minister birth pixels are include as controls. The red shaded areas on the plot represent 95 percent confidence intervals.

Figure C11: Dynamic treatment effects in medium prestige minister birth pixels

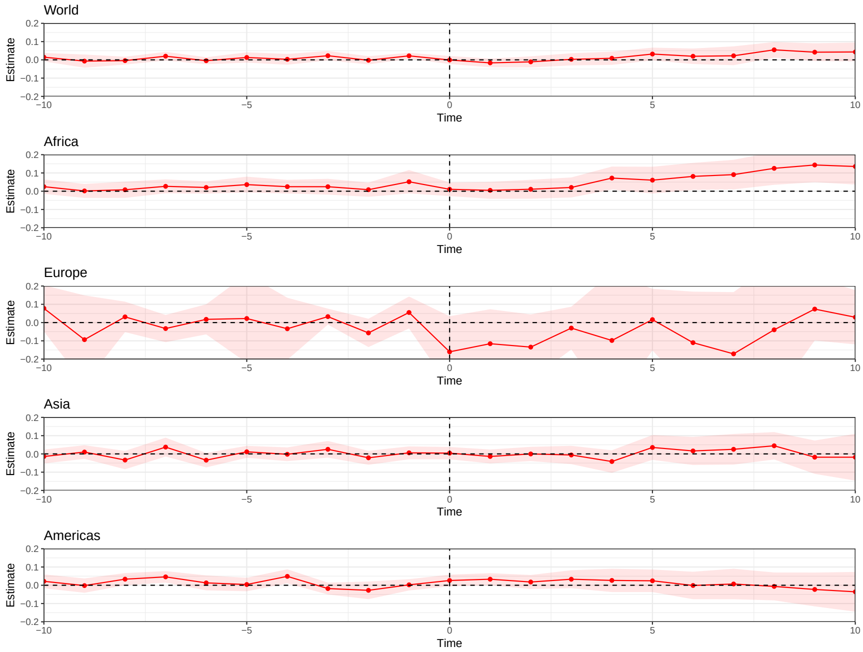

The figure shows an event study for the outcome log nightlight intensity where treatment is restricted to cover medium prestige ministers only. Dummy variables identifying all non-medium prestige minister birth pixels are include as controls. The red shaded areas on the plot represent 95 percent confidence intervals.

Figure C12: Dynamic treatment effects in low prestige minister birth pixels

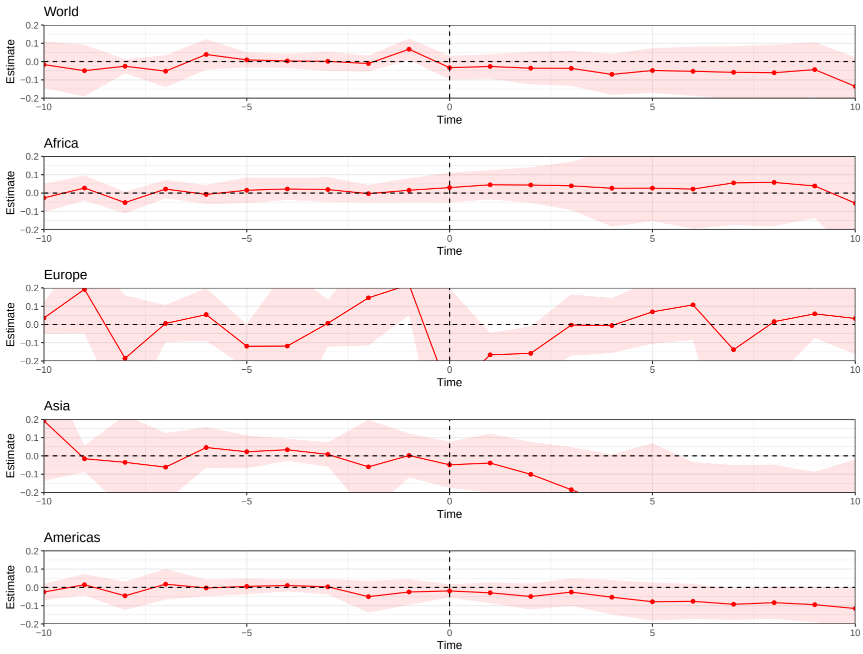

The figure shows an event study for the outcome log nightlight intensity where treatment is restricted to cover low prestige ministers only. Dummy variables identifying all non-low prestige minister birth pixels are include as controls. The red shaded areas on the plot represent 95 percent confidence intervals.

C.3.2 Ministerial portfolios

Figure C13: Dynamic treatment effects in defense minister birth pixels

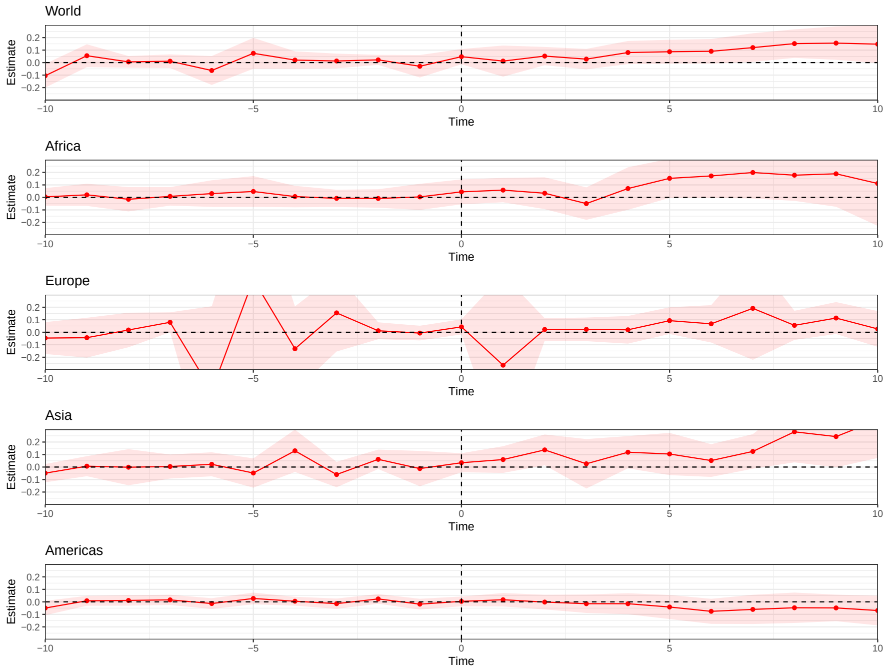

The figure shows an event study for the outcome log nightlight intensity where treatment is restricted to cover defense ministers only. Dummy variables identifying all non-defense minister birth pixels are include as controls. The red shaded areas on the plot represent 95 percent confidence intervals.

Figure C14: Dynamic treatment effects in foreign minister birth pixels

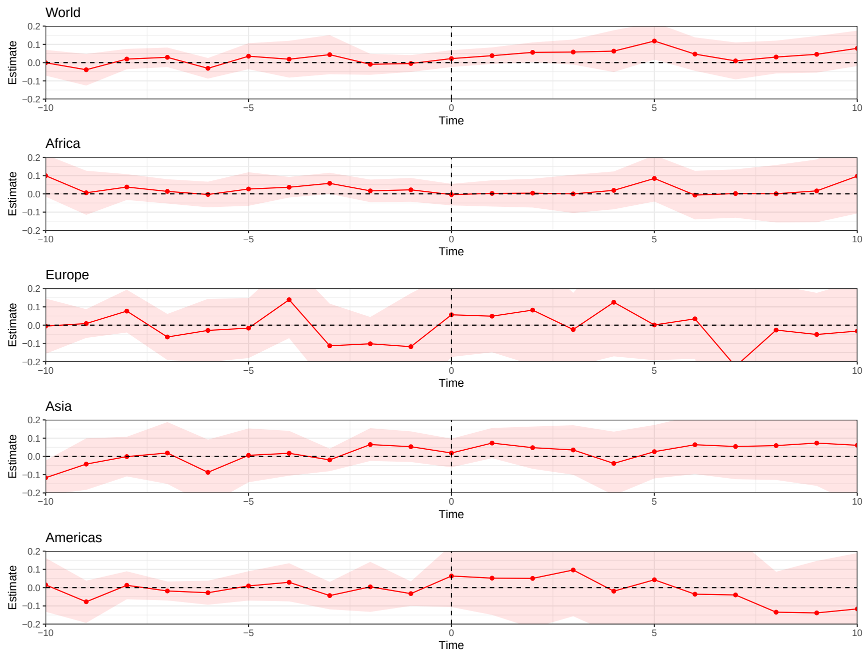

The figure shows an event study for the outcome log nightlight intensity where treatment is restricted to cover foreign ministers only. Dummy variables identifying all non-foreign minister birth pixels are include as controls. The red shaded areas on the plot represent 95 percent confidence intervals.

Figure C15: Dynamic treatment effects in finance minister birth pixels

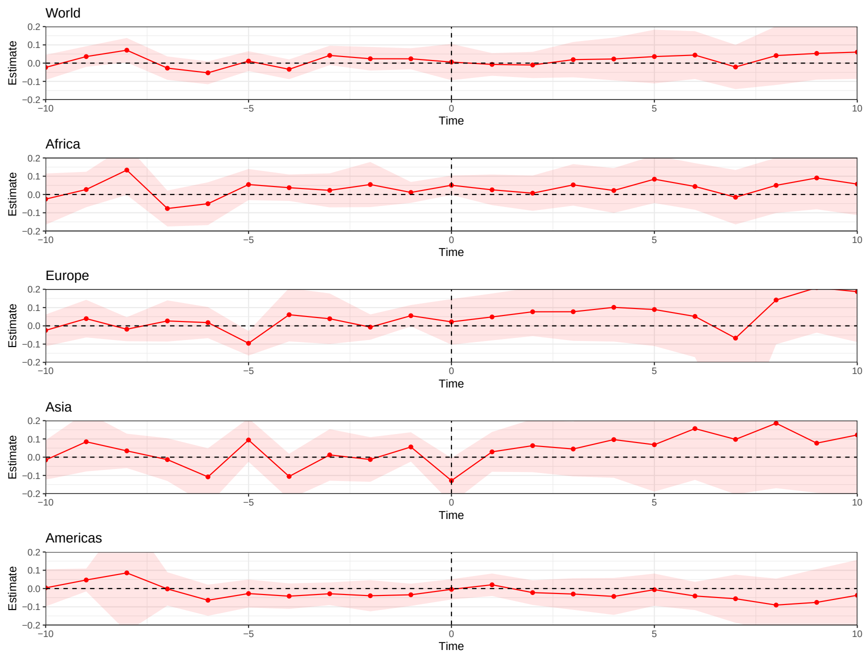

The figure shows an event study for the outcome log nightlight intensity where treatment is restricted to cover finance ministers only. Dummy variables identifying all non-finance minister birth pixels are include as controls. The red shaded areas on the plot represent 95 percent confidence intervals.

C.3.3 Institutions

Figure C16: Dynamic treatment effects: ministers in autocracies

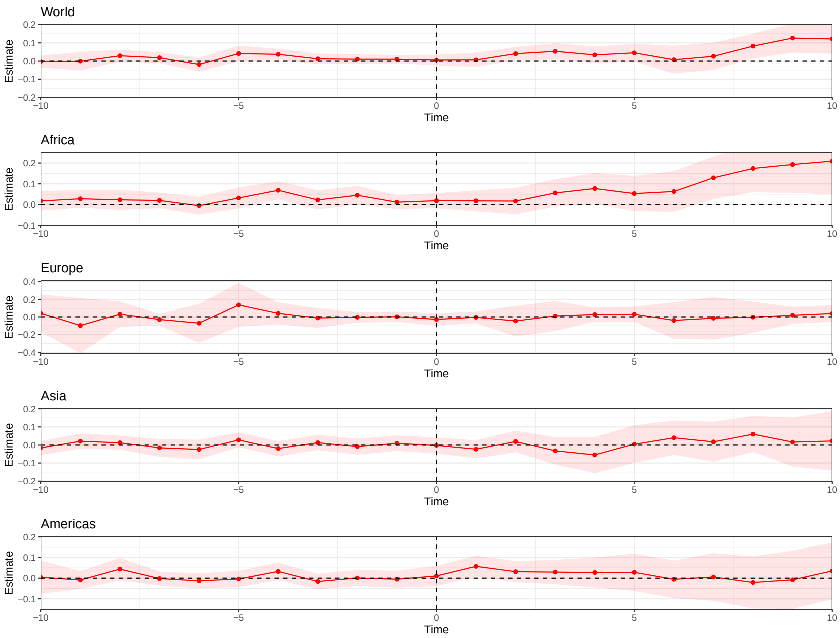

The figure shows an event study for the outcome log nightlight intensity where treatment is restricted to cover ministers in autocracies only. Dummy variables identifying all non-autocratic minister birth pixels are include as controls. The red shaded areas on the plot represent 95 percent confidence intervals.

Figure C17: Dynamic treatment effects: ministers in democracies

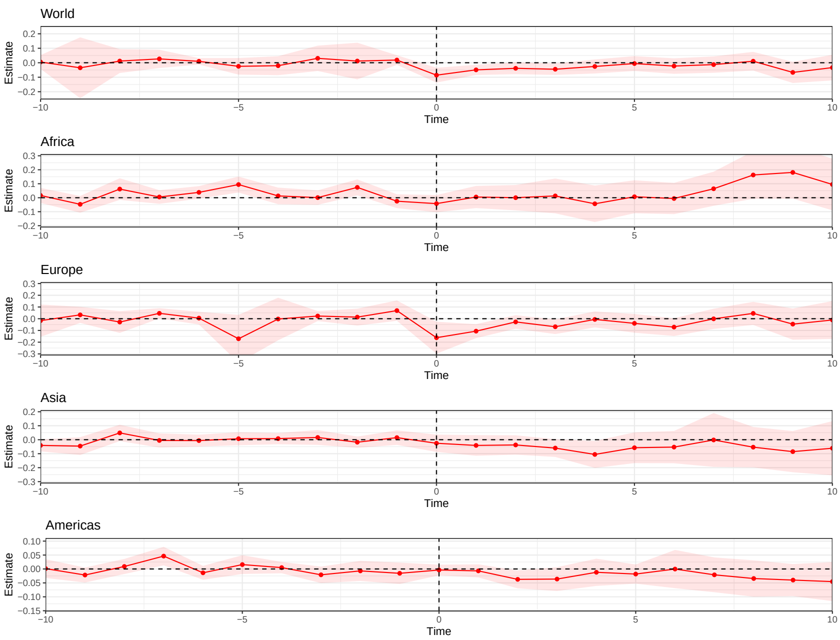

The figure shows an event study for the outcome log nightlight intensity where treatment is restricted to cover ministers in democracies only. Dummy variables identifying all non-democratic minister birth pixels are include as controls. The red shaded areas on the plot represent 95 percent confidence intervals.

The Regional Economics of Mineral Resource Wealth in Africa

Zareh Asatryan [*] Thushyanthan Baskaran [†] Carlo Birkholz [‡]
Patrick Hufschmidt [§]

Abstract

We study the regional economics of mineral resource activity in Africa. Using geocoded data on mine openings and closures in Africa, we document that mining regions experience local economic booms while a mine is in operation. We then explore how mineral resources affect non-mining regions. Non-mining regions might be affected by mining activity due to deliberate government policies (e.g., regional redistribution) or due to various inadvertent country-level macroeconomic adjustments (e.g., Dutch Disease type effects or declining institutional quality). Our results suggest that mineral resources have heterogeneous effects on non-mining regions. Politically important regions benefit economically, while generic non-mining regions are, in general, worse off. Exploring mechanisms, we find that these spatial patterns arguably emerge both due to deliberate government policies as well as Dutch-Disease-style macroeconomic adjustments that harm regions specializing in sectors other than mining.

Key words: Mineral resources, spillovers, luminosity, favoritism, Africa.

JEL: H77, O13, R12.

[*]ZEW Mannheim
[†]Ruhr University Bochum
[‡]University of Mannheim, ZEW Mannheim
[§]TU Dortmund

4.1 Introduction

For many countries, mineral resources are an indispensable source of income. In 2017, the share of mineral resource rents in GDP was as high as 28% in Mongolia or 14% in the Democratic Republic of Congo. On average, mineral resource rents constituted 0.5% of World GDP.[1] In fact, these numbers may even understate the importance of the mineral resource sector. In addition to pure rents, the sector contributes to the national economy by providing employment to millions of formal and informal workers (Ericsson and Löf, 2019).

Even though the macroeconomic importance of the mineral resource sector is undeniable, its local economic implications across different geographies – both within mining and non-mining regions – are not well understood. One strand of the relevant literature explores the economic implications of mineral resources for mining regions (Cust and Poelhekke, 2015).[2] However, the evidence remains ambiguous. Some studies find that mineral resources have positive short-run effects on local economic development and household income (Michaels, 2011; Loayza et al., 2013; Allcott and Keniston, 2017; Feyrer et al., 2017; Mamo et al., 2019; Benshaul-Tolonen, 2019; de la Sierra, 2020), while others find evidence for environmental and societal damages (James and Aadland, 2011; Aragón and Rud, 2013; Kotsadam and Tolonen, 2016), with potentially adverse long-run economic effects.[3]

Besides their economic effects within mining regions, another important question regarding the spatial economic implications of mineral resources is how they affect economic outcomes in non-mining regions. It is important to study how mineral resources affect such regions in order to better understand why – in many cases – mineral resources are associated with aggregate economic decline and political instability rather than prosperity, and why countries with significant mining revenues perform worse across a range of standard measures for welfare (nutrition, literacy, life expectancy) than their non-resource neighbors (Chuhan-Pole et al., 2017). Yet, there is almost no literature on this question.[4]

The extraction of mineral resources can affect non-resource regions for two main reasons. First, national governments, which are typically the primary claimant of mineral resource revenues (Brosio and Singh, 2014), could redistribute mining revenues spatially

[1]Mineral resource rents are defined as the difference between the value of production for a stock of minerals at world prices and their total costs of production. Data are taken from the World Bank's World Development Indicators.

[2]This literature follows a long strand of research using cross-country data (Sachs and Warner, 1995; van der Ploeg, 2011)

[3]For example, the resource sector may crowd out other sectors that could be more viable in the longer term (Cust and Viale, 2016). Mineral resource wealth could also depress regional incomes over a longer horizon if children and young adults drop out of education (human capital accumulation) to work in the mineral resource sector (Ahlerup et al., 2019).

[4]Marginally related papers are Huang et al. (2022), who study how the regional redistribution of resource rents influences urbanization and structural transformation across regions and Hodler et al. (2023), who study how the interplay between the location of mines and the spatial distribution of ethnic groups influences the likelihood of conflicts. Although related, these two papers do not explore the entire spectrum of spatial economic implications of mineral resources.

either explicitly through intergovernmental transfer schemes or implicitly through general government spending and nationally provided public goods. The way in which natural resource revenues are shared between resource-producing and non-producing regions is thus often politically contentious and a major source of conflicts (Fearon and Laitin, 2003).

Second, mineral resources can cause inadvertent macroeconomic adjustments that affect different subnational regions in a heterogeneous fashion. For example, the Dutch Disease literature suggests that when resource-rich regions experience booms, non-resource regions can be negatively affected by adverse terms of trade effects (Corden, 1984). Another possible macroeconomic consequence of mineral resource wealth is general political instability or declining quality of government (Humphreys et al., 2007; Maystadt et al., 2014; Berman et al., 2017).[5] For non-resource regions, such adverse effects might outweigh the (potentially) positive effects of inter-regional transfers funded by higher resource revenues. In fact, as there are typically more non-resource than resource regions in a given country, such ex- and implicit negative spillovers of mining activity into non-resource regions might result in negative aggregate effects at the country-level, even if minerals induce booms locally.

This paper is the first to shed light on the broader spatial implications of mineral resources. It explores how and why mineral resources affect various types of non-mining regions by combining cross-country data with spatially disaggregated micro data on Africa. Specifically, we use geo-referenced data on the operation of mines for different minerals combined with luminosity data over the period 1992-2013, and study how the operation of mines affects luminosity in mining and non-mining regions in African countries.

With respect to non-mining regions, we define three types: capital cities, birth regions of national leaders, and generic non-mining regions. Capital cities and leaders' birth regions are natural candidates for regional favoritism by national governments. Policy makers may favor these two types of regions because they themselves or close acquaintances could directly benefit from any disproportionate resource allocations. In contrast, generic (non-mining) regions are the remainder of the country and, depending on circumstances, might or might not benefit from more mineral activity. For example, they might receive additional resources through intergovernmental equalization. However, as discussed above, they might be worse off due to adverse macroeconomic effects.

Implementing several variants of difference-in-difference designs at the grid-level, we find that the opening of a mine leads to visible and persistent economic booms in the mining regions. This effect is centered on the mine location and is observable up to a distance of 30km. The effect is robust across a range of sensitivity tests.

To explore how mines affect non-mining regions, we exploit cross-country variation in the number of operating mines. We find that an additional mine anywhere in the country increases luminosity in the capital city. This suggests that a share of the proceeds from mineral resources are shifted from mining regions to capital cities. We also observe that

[5] See Aragón et al. (2015) for a more extensive discussion of possible macroeconomic channels.

additional mines increase luminosity in the birth region of the current national leader when the country in question is under autocratic rule. However, we observe no such effect in democracies. This evidence suggests that the distribution of mineral resource wealth is subject to regional favoritism towards leaders' birth regions in non-democratic settings, but not when countries are democratic. This finding is in line with the cross-country evidence suggesting that the natural resource curse only emerges in countries with weak institutions (Collier and Hoeffler, 2009; Boschini et al., 2007).

We also observe that the opening of additional mines in a given country decreases luminosity in generic non-mining regions compared to similar regions in adjacent countries. Exploring this finding further, we find that the number of conflict events increases in generic non-mining regions when mineral activity expands elsewhere. This result suggests that governments might use mineral revenues to fund conflicts in other (non-mining) parts of the country. We also find that mines have weaker economic effects on luminosity in regions inhabited by ethnicities that are "politically weak", i. e., considered to be discriminated against in their country, and more positive effects in regions with politically powerful ethnicities. Overall, these results suggest that deliberate government actions can provide a partial explanation for the adverse effects of mineral resource activity on generic non-mining regions.

However, macroeconomic adjustments appear to be important as well. Although we do not observe a decline in the institutional quality at the country level when mineral resource activity expands, subnational regions specializing in manufacturing, agriculture, or the hospitality industry are ostensibly worse off (in contrast to regions that have a relatively large mining sector). These results indicate further that exchange rate adjustments due to increased mineral resource exports disadvantage non-mining regions and contribute to a decline in economic activity – a finding in line with the Dutch Disease literature.

In addition to the general literature on the economic implications of mineral resources discussed above, this paper primarily contributes to a recent literature that studies the link between mineral resources and intergovernmental transfers.[6] Existing studies tend to focus on the implications of intergovernmental transfers funded by resource revenues allocated specifically to mining regions. For example, Cust and Ridwan (2014) discuss evidence from Indonesia that fiscal transfers related to oil production boost local GDP in oil-producing regions. However, the direct effect of project investments appears to be small. Caselli and Michaels (2013) study the effect of oil revenue windfalls in Brazil on municipalities that benefit from fiscal sharing rules. Though in their study, only municipalities close to offshore production facilities are considered.

This paper is also related to the emerging literature on regional favoritism. The seminal contribution by Hodler and Raschky (2014) uses night-time luminosity as a proxy for local economic development to show that national leaders favor their birth towns. Subsequent

[6]Thereby, we also implicitly contribute to the broader literature on fiscal federalism (Wildasin, 1997; Baskaran, 2012).

contributions explore regional favoritism in more detail. For example, Do et al. (2017) find that bureaucrats in Vietnam favor their hometowns in infrastructure investments. In institutionally more mature settings, Baskaran and da Fonseca (2021) show that German state ministers allocate more state employment to their place of residence, while Asatryan and Havlik (2020) show similar home bias in the allocation of loans in Europe. Another closely related paper is Dreher et al. (2021). This paper adopts an empirical strategy similar to ours to study the subnational economic effects of Chinese aid flows. Specifically, it identifies subnational impacts of country-level variation in Chinese aid flows and shows that Chinese aid tends to have positive local economic effects.

Another related strand of literature explores the importance of artisanal or small-scale mining. Bazillier and Girard (2017) show that in Burkina Faso, artisanal mining can have positive local effects. For example, an increase in the gold price increases the consumption level of households who live near artisanal mines. Industrial mines, on the other hand, have no effect on local consumption. Pokorny et al. (2019) report similar results on artisanal and industrial mining. One reason for these findings might be that the proceeds from artisanal mines are harder to tax than industrial mines, and thus less likely to be redistributed to non-mining regions.

4.2 Background

4.2.1 Mineral resources and their exploitation in Africa

Minerals are materials with economic value in or on the Earth's crust. They can be extracted and applied as inputs for various productive uses, including industrial applications. Mineral resources of significant value typically belong to the state and their proceeds often constitute a large fraction of public revenues (besides contributing to overall GDP).

Minerals are exploited either by state-owned corporations or private firms that have acquired a license from the government and thus pay royalties or are taxed according to production (Land, 2009). In Africa, most countries tend to rely on private investors due to limited domestic mining capacity (Laporte and Quatrebarbes, 2015). The government taxation of mineral resources and the cost of licenses are therefore the main means by which African governments tap into resource rents.

The rules by which revenues are shared between governments and private corporations vary between countries. Due to the idiosyncrasies of the mineral resource sector, corporations and investors often receive unique tax treatments. The share of the resource rent that accrues to the public sector thus varies depending on such factors as global market conditions or the bargaining power of governments. It is estimated that the rent captured by governments can range from 25% to 65% (Land, 2009).[7]

4.2.2 Regional distribution of mineral resource revenues

Governments are complex organizations in which power is shared vertically and horizontally. Power is shared vertically between the national and subnational governments and horizontally among various subnational governments. This raises the question of how the resource rents to which the "government" is entitled are shared between government units.

The traditional literature on fiscal federalism suggests that resource rents should accrue to the national government (Oates, 1999). National governments are better equipped to deal with the inherent volatility of resource revenues, which are, for example, subject to global demand shocks. Subnational governments may also lack sufficiently competent staff and the absorptive capacity to make adequate use of resource rents. Consequently, subnational governments might make inefficient investments or waste resources on vanity projects (Brosio and Singh, 2014).

Countries in Africa broadly follow the recommendations given in the traditional fiscal federalism literature.[8] While national constitutions tend to make only vague statements about ownership, proclaiming that resources belong to the "people" or the "state", both legislation that adds detail to the constitution and administrative reality suggest that it is the national governments that are the first claimant of any resource revenues.[9] Therefore, it is the national government that typically negotiates with private corporations and decides on their tax treatment.

There are various taxes that national governments levy on mineral resources, ranging from income and profit taxes, to royalties and licensing fees, sales and excise taxes, VAT on goods and services, and stamp duties (Otto, 2001).[10] The share of the rent that accrues to the national government is then distributed across governmental units, and thus either explicitly or implicitly across different regions of the country. Whether this regional distribution takes place according to pre-determined rules or in a discretionary fashion depends on the institutional arrangements in a given country.

Many African countries have formal revenue sharing schemes by which regional gov-

[7] According to Laporte and Quatrebarbes (2015), whether governments (and other domestic stakeholders) receive their "fair" share of resource rents is up for debate. In any case, how resource rents should be shared between governments and private investors is a question that is separate from how rents should be shared between different tiers of government and across the various regions within the country. Basic economic theory suggests that pure rents can be fully captured by the government (e.g., by a lump-sum tax), but in practice this is not feasible due to, e. g., international production capacity constraints.

[8] Of course, the true reason why revenues are assigned in this way may not be the normative prescriptions of this literature but political expediency.

[9] Control over natural resources is typically only allocated to the local government tier if their economic importance is small (Brosio and Singh, 2014). With the recent wave of decentralization, the ownership of resources has been partially transferred to subnational governments, but national governments generally continue to be the main claimants (Brosio and Singh, 2014). See Table D19 for details on the constitutional arrangements with respect to the ownership of mineral resources in several African countries.

[10] Depending on the country in question, regional and local governments may be allowed to tax resource rents to a degree using such taxes.

ernments participate in national government revenues (both from resource rents and other revenue sources).[11] A general feature of revenue sharing in Africa is that resource producing regions receive a relatively large share of rents (Brosio and Singh, 2014). This is often justified as compensation for environmental or societal damages caused by mining activities. There are, however, also cases where the distribution of resource revenues is not tied to where the natural resources are produced. Chuhan-Pole et al. (2017), for example, note that in Tanzania, a fiscally highly centralized country, public revenues accrue exclusively to the central government and are then allocated according to priorities not related to the location of the mines. Arora et al. (2017) report that in Ghana, about 80% of resource revenues are retained by the (national) government and used for general budget support.

Accordingly, non-producing regions also often receive a share of the resource rents through formal revenue sharing mechanisms. One normative reason for why non-mining regions should benefit from resource rents is a standard insurance argument. Ex-ante, it is unclear which regions will have valuable mineral resources. Under the veil of ignorance, it is thus welfare-increasing to share resource revenues. However, national governments also allocate resource revenues discretionarily to non-resource regions. Such discretionary transfers may be justifiable from a normative point of view if they are granted to achieve desirable economic, societal, or environmental goals. On the other hand, discretionary transfers might be granted to pursue narrow political or electoral agendas or due to clientelism and favoritism. In general, there is a lack of evidence about the motives that would lead national governments to redistribute resource rents to non-mining regions.

Besides explicit transfers that affect subnational revenues, national governments may also use a share of the resource rents to fund national expenses. National government spending typically has a distributional consequence across space, with some regions benefiting more than others (Berry et al., 2010). In general, it is opaque and dependent on idiosyncratic country-specific political and institutional circumstances which regions will benefit from any increase in national spending, either due to newly discovered mineral resource wealth or other reasons (Reingewertz and Baskaran, 2020).

Although intergovernmental flows of public funds might appear as an appealing outcome to study the wider spatial economic effects of mineral revenues, it is difficult to identify these effects using only fiscal data. Without knowledge of the entirety of transfers that flow from the national to subnational governments (vertical transfers) and between subnational governments (horizontal transfers) and a spatial disaggregation of national government spending, it is not possible to arrive at an accurate assessment. In addition, intergovernmental transfer schemes in Africa are generally not sufficiently developed to transparently transfer resource rents between regions (Brosio and Singh, 2014). Transfers as such must also not have broad positive welfare effects if funds are wasted locally

[11] Indeed, Fjeldstad et al. (2014) point out that local governments in Africa, with the possible exception of South Africa, rely heavily on central government transfers to fund their expenses. One notable example is Botswana, where rural councils receive 92% and urban areas receive 62% of their revenues from the central government.

due to mismanagement (Brollo et al., 2013; Standing and Hilson, 2013). Finally, mineral revenues could affect the local economy of non-mining regions for reasons other than the intergovernmental redistribution of public funds, notably macroeconomic adjustments that have heterogeneous impacts across geographies. For these reasons, in this paper we rely on night-time luminosity as a catch-all measure of the economic implications of the redistribution of mineral resource rents beyond the mining regions.

4.3 Data

4.3.1 Grid

Most analyses below are conducted at the grid level. To this end, we overlay the African continent with a grid of 0.5 × 0.5 degree cells (0.5 degrees correspond to about 55km at the equator) (Berman et al., 2017). We then intersect this grid with a map of country borders to identify within which country a particular cell is located. We then drop from this grid all cells that are located in more than one country. The final sample consists of 9,068 cells over the period 1992-2013 (see Figure 4.1).[12]

4.3.2 Mineral resource data

Previous research has used mine openings and closures to approximate the amount of mineral resource revenues that accrues to the government (Kotsadam and Tolonen, 2016; Knutsen et al., 2017). Another strand of the literature has also used variation in mineral resource prices (Berman et al., 2017).

In this paper, we rely on mine openings and closures as our main approach, but also explore the effects of price variations. We use openings and closures as our baseline approach since variation in prices is less suitable in our context. First, it is difficult to estimate effects for non-mining regions when a country has several minerals and prices for each mineral evolve differently (for some minerals price data is also unavailable). That is, it is difficult to assess whether overall mineral revenues increase or decline without knowledge of export volumes for each mineral. Second, temporary variation in prices might induce only short-term resource reallocations rather than long-term adjustments, which in turn would persistently affect luminosity in non-mining regions.

We obtain historical data on mineral resource activity in Africa from MinEx Consulting, which is a private mining consulting company. The database contains a comprehensive list of *significant and unique* deposits in Africa. In terms of minerals, the database covers all commodities except bulk minerals (i. e., coal, iron ore, bauxite, potash and phosphate). In terms of size, MinEx Consulting estimates that its data covers 99% of all giant-sized de-

[12] In the regressions, sample sizes are typically slightly smaller, primarily due to missing data on luminosity.

posits, 95% of all major deposits, 70% of moderate deposits and 50% of minor deposits.[13] The information on mineral resource deposits and activity is derived from various (propitiatory and free) sources and is more comprehensive and up to date than most alternative datasets.[14]

[13] The thresholds for precious metals are: Minor $\geq$ 0.03 Moz Au (millions of ounces gold) equivalents, Moderate $\geq$ 0.32 Moz Au equivalents, Major $\geq$ 2.24 Moz Au equivalents, Giant $\geq$ 11.18 Moz Au equivalents, Supergiant $\geq$ 80.00 Moz Au equivalents. For other minerals, the thresholds are: Minor $\geq$ 0.03 Mt Cu (millions of megatonnes of copper) equivalents, Moderate $\geq$ 0.32 Mt Cu equivalents, Major $\geq$ 2.45 Mt Cu equivalents, Giant $\geq$ 18.97 Mt Cu equivalents, Supergiant $\geq$ 35.00 Mt Cu equivalents.

[14] In particular, in contrast to the freely available data from the U.S. Geological Survey (`https://data.doi.gov/dataset/mineral-operations-of-africa-and-the-middle-east`), the version of the MinEx database available to us provides information on mineral resource activity up to 2015 and includes more detail on mines, in particular their startup and shutdown dates. See Section D.1 in the appendix for more details on the MinEx data.

Figure 4.1: Mining regions across Africa

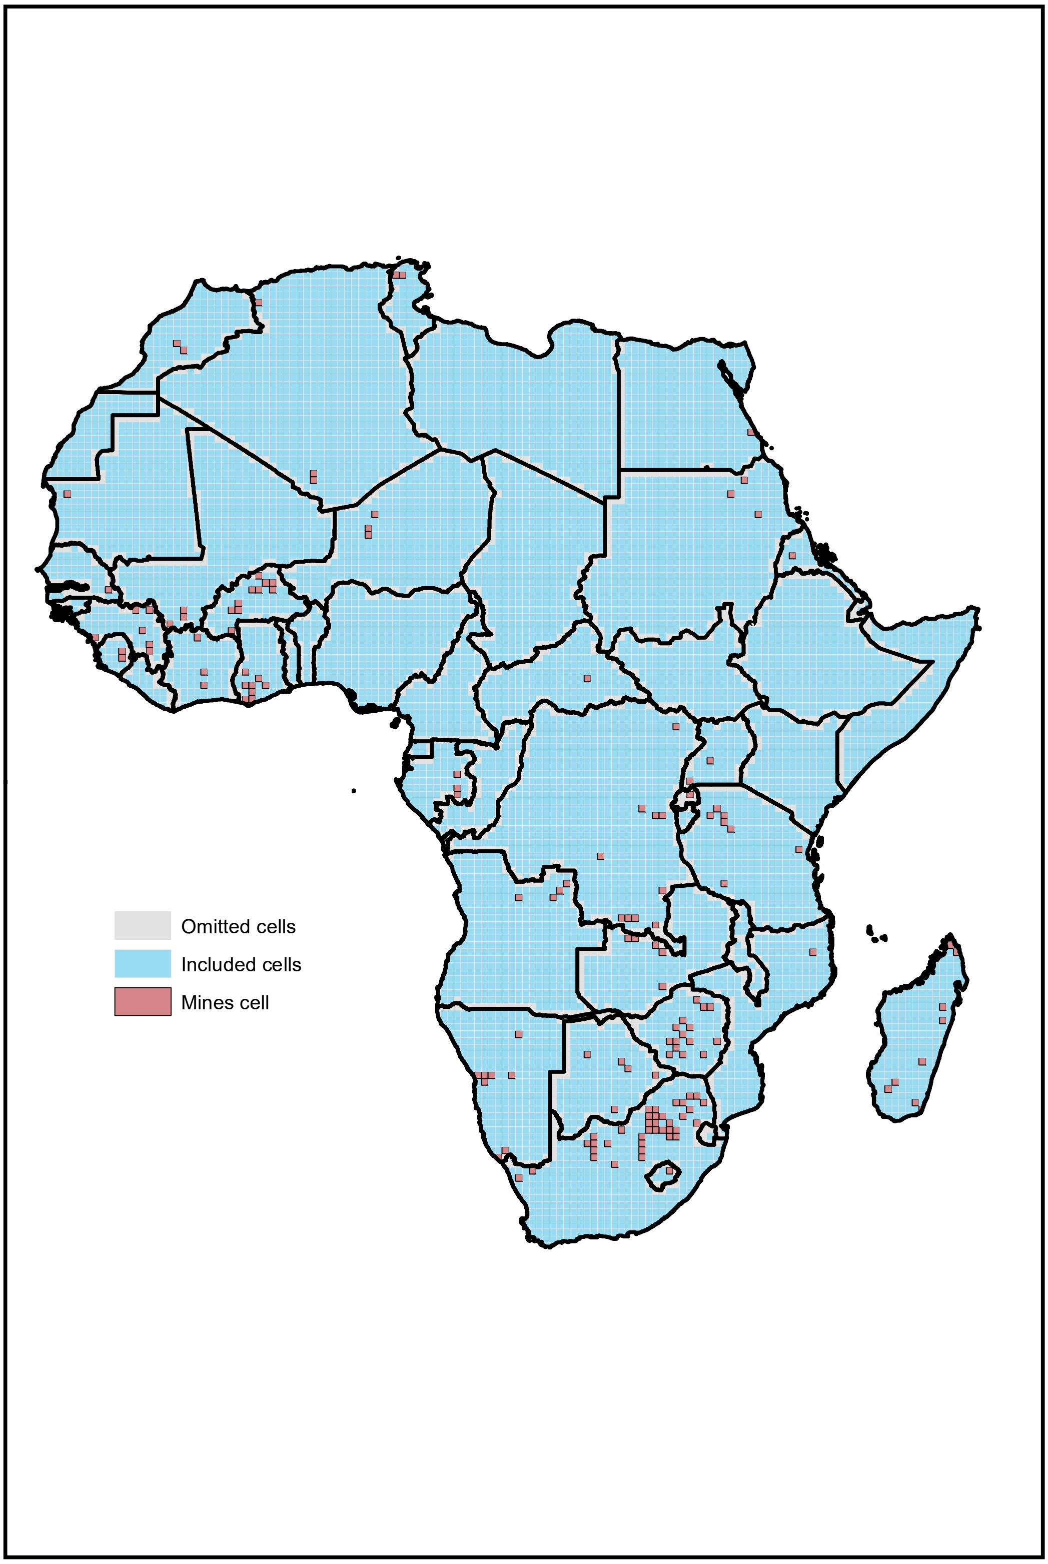

The figure shows the location of mineral deposits included in our estimation sample plotted to the grid of Africa.

The database lists 519 mines of (non-bulk) minerals that were potentially in operation at least for one year during the 1992-2013 period. Of these 519 mines, we lack information on the startup or shutdown date for 228 mines. We omit these mines from the analysis, which leaves us with a final sample of 291 mines of any size, a coverage 56% of all mines potentially in operation.[15] However, the actual sample coverage relevant for the sample period is likely higher than these 56%. For example, of the 115 mines with information on the year of discovery but no information on the startup year, 100 had been discovered before 1992 (the beginning of the sample period). 76 mines had been discovered even before 1950. As such, it is unlikely that many of the mines with missing startup dates were started up during the sample period. Similarly, of the 90 mines that were recorded as closed during the sample period but for which we lack information on the exact shutdown date (but have information on the startup date), 56 were started up before 1950. It is unlikely that many of these mines were closed during the sample period. Second, mines with missing information on startup or shutdown dates tend to be smaller and thus in all likelihood less economically consequential.

Table 4.1 lists the type of minerals and the number of respective mines included in our sample. The most common mineral in the sample is gold, which makes up 46% of all mines. Diamond and copper mines also constitute large share of all mines.

Table 4.1: Types of mines and their frequency in the sample

Mine	**Freq.**	**Percent**
Andalusite	1	0.352
Asbestos	3	1.056
Chromium	1	0.352
Cobalt	1	0.352
Copper	26	9.155
Diamonds	48	16.90
Flourine	1	0.352
Gold	131	46.13
Lead	2	0.704
Manganese	7	2.465
Mineral Sands	5	1.761
Nickel	9	3.169
Platinum Group Elements (PGE)	13	4.577
Platinum	2	0.704
Ruby	3	1.056
Sapphire	7	2.465
Tin	3	1.056
Tungsten	2	0.704
Uranium	10	3.521
Zinc	9	3.169

The table shows the type of minerals used in the estimations below and their frequency in the sample.

We project the latitude and longitude coordinates of the 291 mines in our sample onto

[15] The number of mines that contribute to the regressions is slightly lower, i. e., 284 mines, because some mines were closed exactly in 1992 or opened up and closed in the same year.

the grid included in our sample; see Figure 4.1. Subfigure (a) of Figure 4.2 shows the number of operational mines across Africa included in each year of our sample period. Subfigure (b) and (c) show the number of mine openings and closings per year in our sample, respectively. There is significant variation in mining activity. It is apparent that openings generally outnumber closures, particularly in the second half of the sample period.

Figure 4.2: Number of mines over time

(a) Total

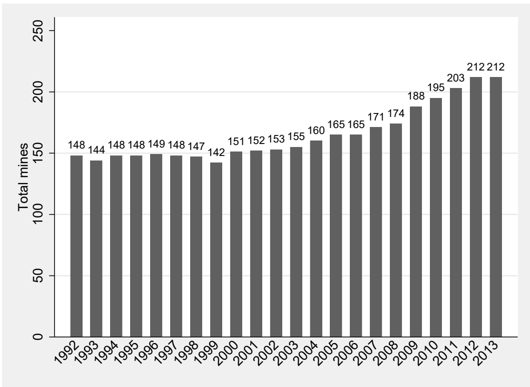

(b) Openings

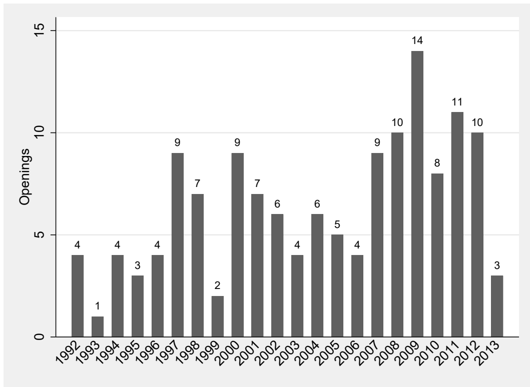

(c) Closures

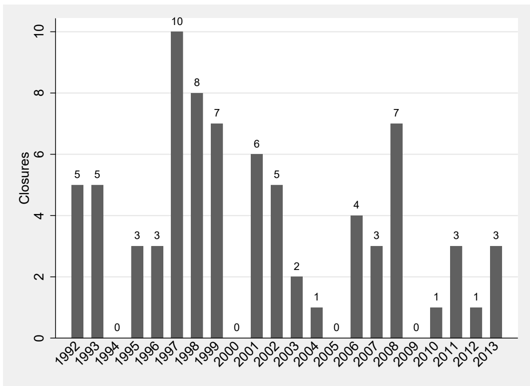

The figure shows for mines with non-missing data on openings and closures the total number of mines (subfigure a), the number of mine openings (subfigure b), and the number of closures (subfigure c) in each year during the sample period.

We also obtain prices for most of the minerals listed in Table 4.1 from the World Bank, the IMF, and the US Geological Survey (USGS). See Section D.2 in the appendix for details on the price data for each of the minerals.

4.3.3 Luminosity data

Following previous literature, we use nighttime luminosity as a proxy for economic activity at the local level (Alesina et al., 2016; Hodler and Raschky, 2014; Michalopoulos and Papaioannou, 2016; Martínez, 2022). This data is based on images of the earth at night obtained by satellites of the US Air Force (USAF) Defense Meteorological Satellite Program Operational Linesman System (DMSP-OLS). The original imagery is processed by the National Oceanic and Atmospheric Agency (NOAA) and released to the public as raster datasets.

The raster datasets consist of annual average stable night lights between 8.30pm to 10pm and are available at a resolution of 30 arc-seconds (about 0.86 square kilometer at the equator) for all years after 1992. Each pixel of the dataset stores a digital value ranging from 0 to 63 indicating the amount of average light of an area covering 30 arc-seconds. Higher values imply that a pixel emanates more light (Henderson et al., 2012).

To obtain cell-level measure of economic development, we overlay the grid of cells over the raster datasets. We then calculate the area mean of the digital values of each cell with size 30 arc-seconds that falls within the boundaries of each of the 0.5×0.5 degree cells.[16]

While the DMSP-OLS luminosity data has been used as a proxy for local economic activity in previous research, it has important limitations (Gibson et al., 2021). First, it may suffer from measurement error. For example, light in dimly lit areas might not be detected by the satellites due limited dynamic range of their sensors. Second, luminosity might not reflect economic activity if the satellites mainly pick up ephemeral phenomena that emit light at night (e. g., forest fires). However, Bruederle and Hodler (2018) find luminosity accurately reflects human well-being and thus local economic development by comparing night lights with individual-level outcomes as reported by survey data from 29 African countries. Similar results are reported by Määttä et al. (2022). In any case, we address concerns regarding sensor sensitivity in a robustness test and also verify the suitability of luminosity as proxy for economic activity by relating luminosity to other

[16]Table D13 in the Appendix provides summary statistics of the luminosity data and all further data used below.

proxies for local economic development (see Section D.3 in the Appendix).[17]

4.3.4 Capital regions

We retrieve information on the location of national capitals from the CEPII's GeoDist database (Mayer and Zignago, 2011). To identify cells that belong to capital regions, we draw a buffer of 10km size across the longitude and latitude coordinates for capital cities noted in the CEPII's GeoDist database. All cells that fall within each buffer are indicated as capital regions; see Figure D5 in the Appendix.[18]

4.3.5 Leader regions

We use information on the birth cities of national leaders from the Archigos database (Goemans, 2016); in cases where information on birth regions was missing, we collect this information ourselves. We then geocode the birth cities using ArcGIS.

To identify cells covering a leader's birth region, we draw, as for capital cities, a buffer of 10km around each leader's birth city's longitude and latitude coordinates. We then classify all cells that fall within this buffer as a leader region (see Figure D6).

4.3.6 Other data

We obtain further data for robustness tests and extensions, notably proxies for the level of democracy in a country from Freedom House, proxies for institutional quality as well as gross mineral revenues from the World Bank's World Development Indicators, survey-based proxies for local economic development from the Development and Health Surveys (DHS), data on population counts from WorldPop[19], data on conflicts from the Armed Conflict Location and Event Data Project (ACLED), and data on employment per economic sector from the Integrated Public Use Microdata Series (IPUMS). We discuss these data further below in more detail as the need arises.

[17]Another more specific concern with using luminosity as a proxy for local economic development in mining regions is that mineral activity itself can emanate light at night. For example, producers could install floodlights to enable production at night or to reduce road hazards in mining regions. Mines could also be illuminated at night to prevent illegal mining activity by artisanal miners. Technically, the effect of mines we estimate below represents a composite effect of such direct effects of mining activity on luminosity and the broader welfare effects. However, such alternative reasons for light at night in mining regions are unlikely to influence our estimates substantially. The largest mine by area in the world (Hull Rust open pit mine in the US) has an area of 8.1 km^2, which is significantly smaller than the grid cells. In addition, we also find further below that mineral resource activity increases lights in cells that neighbor mining cells but have no mines themselves.

[18]We draw a buffer around the geographic coordinates to capture capitals that are close to the sea. Without a buffer-based approach, the longitude and latitude coordinates might be projected slightly outside of the range of the African land cells by our GIS software due to projection inaccuracies. Note that capitals falling entirely into cells at the border to another country are dropped as well when we remove border pixels in our baseline specification.

[19]Source is WorldPop and CIESIN (2018).

4.4 Mines and economic activity in mining regions

4.4.1 Empirical model

We start out by exploring the local effect of mines within mining regions. The baseline model to estimate the local effect of mines is as follows:

$$y_{it} = \alpha_i + \sum_c \gamma_t \times c + \text{Mine}_{i,t} + \epsilon_{i,t}, \tag{C1}$$

with y_{it} the log of the mean of luminosity[20] in cell i in year t, α_i cell fixed effects, γ_t year fixed effects and c country dummies (i. e., we include country-specific year fixed effects), and $Mine_{i,t}$ a dummy that is 1 when there is at least one operating mine in cell i in year t and 0 else.

The key identifying assumption is that the opening and closure of mines is exogenous to local economic trends. Naturally, the setup and operation of a mine in a given locality is not a random event. We report in the Appendix event-studies to further validate the parallel trends assumption and to assess whether potential violations are economically meaningful (see Figures D2 and D3). In any case, cell fixed effects can account for time-constant geographical or environmental characteristics of different geographies. Similarly, country-specific year fixed effects account for country-level developments that might be correlated with startups or closures of mines.

In addition, we estimate an extension of Equation C1 where we interact the mines dummy with the price of the respective mineral. These specifications make use of variation in (global or US) mineral resource prices on top on the variation in opening and closures of mines.[21]

4.4.2 Main result

In Table 4.2, we report the results from estimating Equation C1. To evaluate statistical significance, we rely on heteroscedasticity-robust standard errors. Standard errors are also clustered at the cell-level.[22]

In model (1), we collect the baseline estimate and observe large positive effects. An operating generic mine increases luminosity by about 86% in the mining cell. In model (2)-(5), we estimate variations of Equation C1 where we focus on mines of different sizes. That is, the mine dummy is only 1 if a cell has a mine of the indicated size. We find that larger mines have stronger effects on luminosity. Overall, these results suggest that

[20]We add +0.01 to each grid's value of mean luminosity to avoid missing values once we take the log.

[21]We drop cells where more than one mineral is extracted in these specifications. For minerals for which we have no prices, the interaction variable is set to 0.

[22]For the baseline estimates, we replicate the estimations with Conley-standard errors as well as different units of clustering in Table D17 in the appendix.

Table 4.2: Mineral resources and economic activity in mining regions

	(1: All mines)	(2: ≥ Major mines)	(3: ≥ Giant)	(4: < Major)	(5: < Moderate)	(6: Prices)
Mine	0.858***	1.078***	1.371***	0.410***	0.178*	
	(0.116)	(0.156)	(0.280)	(0.121)	(0.095)	
Mineral price						0.105***
						(0.018)
Cell FE	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Cells	9053	9053	9053	9053	9053	9053
N	199138	199138	199138	199138	199138	199138

This table collects difference-in-differences regressions that relate mineral resource activity (operating mines) to luminosity at the grid-level for all of Africa (0.5 × 0.5 decimal degree cells). In this specification, we study whether mines increase luminosity in mining regions. The dependent variable is the log of mean light output in each cell. The independent variable in models (1)-(5) is a dummy variable that is 1 if a cell had an operating mine of a certain size as indicated in the column header. The independent variable in model (6) is the contemporanous log price of the mineral that is extracted in a given cell. The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

mineral resource activity induces significant local economic booms.[23]

In model (6), we report the results from the specification where we include an additional interaction between the dummy for operating mines with the price of the respective mineral. The estimate for the interaction effect is positive and significant, and suggests a 0.1% increase in night time luminosity when global mineral prices for the respective mineral being produced in the cell increase by 1%. Thus, both the existence of mines as well as higher mineral prices appear to increase luminosity.

Building on the specification in Equation C1, we also explore how far the local economic effect of mines spread. More specifically, we adapt Equation C1 by including dummies for cells within certain distance bands around mines.[24]

The results are collected in Table 4.3. We find that after the opening of a mine, luminosity increases the most in the mining cells and declines successively in cells that are further away. Mines cease to have a noticeable effect on luminosity in cells that are about 30km away from a mine.[25]

[23] Among the 519 deposits, MinEx classifies 246 as major, giant or supergiant. Among these 246 deposits that were of at least major size, we lack information on startup or shutdown dates for only 47 mines, leaving us with a sample of 199 at least major mines (coverage 81%). As such, another advantage of studying the economic implications of major mines is the better sample coverage.

[24] However, one concern with these specifications is that it becomes increasingly unclear whether regions within the distance bands should be considered as treatment or control units. Non-mining cells that are close to mines likely constitute appropriate counterfactuals for mining cells and thus their inclusion in the control groups strengthens identification. On the other hand, such cells are also those that are subject the most to spillovers.

[25] We report a replication of the baseline specifications after dropping the mining cells in Table D16 in the Appendix. In addition, we report a number or further robustness tests and extensions in the Appendix: in Section D.4.1, we study the effect of mine discoveries rather than openings; in Section D.4.2, we report results from event-studies with mine openings and discoveries; in Section D.4.3, we focus on mines located in border regions to refine our identification strategy; in Section D.4.4, we explore the effect of different types of minerals on local economic development; in Section D.4.5, we study the implications of the Extractive Industries Transparency Initiative (EITI). The results reported in these sections are in line with the baseline results.

Table 4.3: Mineral resources and economic activity in mining regions – geographical spillovers

	(1)	(2)	(3)	(4)	(5)
Mine	0.887***	0.893***	0.897***	0.901***	0.895***
	(0.116)	(0.116)	(0.116)	(0.116)	(0.116)
10 km	0.303***	0.310***	0.314***	0.317***	0.312***
	(0.084)	(0.084)	(0.084)	(0.085)	(0.085)
20-30 km		0.093**	0.098**	0.102**	0.096**
		(0.041)	(0.041)	(0.042)	(0.043)
30-50 km			0.031	0.035	0.029
			(0.037)	(0.038)	(0.039)
50-100 km				0.009	0.003
				(0.021)	(0.023)
100-200 km					-0.011
					(0.015)
Cell FE	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	Yes	Yes
Cells	9056	9056	9056	9056	9056
N	199230	199230	199230	199230	199230

This table collects difference-in-differences regressions that relate mineral resource activity (operating mines) to luminosity at the grid-level for all of Africa (0.5 × 0.5 decimal degree cells). In this specification, we study whether mines increase luminosity in neighboring cells. The dependent variable is the log of mean light output in each cell. The independent variables are dummies that are one if a cell had an operating mine within the indicated distance. The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

4.5 Mines and economic activity in non-mining regions

How does mineral resource activity (anywhere in a given country) affect economic outcomes in non-mining regions? We study this question for three types of non-mining regions.

First, we focus on generic non-mining regions (which have no special connection to the government and its leader). While such regions might benefit from mineral resource activity elsewhere, e. g., through institutionalized intergovernmental transfer schemes, they could also be disadvantaged due to, e. g., adverse macroeconomic developments or other, more deliberate government actions.

Second, we study capital regions. It is plausible that policy makers use a share of the natural resource revenues to spend on amenities in the capital region as they or their "friends and family" could directly benefit from such investments, or that resource discoveries cause urbanization (Huang et al., 2022).

Finally, we focus on leaders' birth regions. Leaders' birth regions might benefit for several reasons if the leader has access to additional resources, ranging from an innate, altruistic desire to benefit his or her birth place to more parochial and political considerations, such as fostering local political support (Baskaran and da Fonseca, 2021).

4.5.1 Mineral resources and luminosity in generic regions

4.5.1.1 Estimation model

To study the effect of mines on generic non-mining regions, we implement a cross-country design with border regions . We focus on border regions as standard cross-country designs (i. e., regressions with entire countries as units of observation) would likely produce biased estimates in our context due to unobserved heterogeneity. For example, countries facing economic difficulties might be more likely to search for and grant permissions to new mining operations (Maddala, 1999).

The idea underlying the border design is that economic trajectories in border regions in neighboring countries are relatively similar in the absence of country-specific economic shocks such as the opening of new mines. That is, the identifying assumption is that year-specific effects in border regions of different countries are similar. If this assumption holds, a disproportionate increase in luminosity in the border regions of a country where a new mine starts up, compared to the border regions of neighboring countries, can be ascribed to increased mineral resource activity.

We calculate average luminosity in cells that neighbor border cells but are located in foreign countries as described in Section D.4.3 in the Appendix. See also Figure D7 in the Appendix for the cells included in the border design sample.

Using the border-cell sample, we estimate the following model as our preferred specification:

$$y_{i,t} - \overline{y}_{i,t} = \alpha_i + \gamma_t + \beta \text{Mines}_{c,t} + \epsilon_{i,t}, \tag{C2}$$

with $y_{i,t}$ the log of the mean of luminosity in border cell i in year t and $\overline{y}_{i,t}$ the log of the mean luminosity averaged across border cells located in neighboring countries within a 250km distance. α_i are cell fixed effects, γ_t are year fixed effects (which are not country-specific), and $\text{Mines}_{c,t}$ is the number of mines in the country where cell i is located in year t.

4.5.1.2 Results

The results are collected in models (1a)-(1b) of Table 4.4. We continue to rely on heteroscedasticity- and cluster-robust standard errors (with cells as the unit of clustering) to evaluate significance.

We find that an additional mine anywhere in the country decreases luminosity in generic border regions relative to luminosity in border cells located in neighboring countries. Specifically, the difference between both sets of cells declines by about 1%. Overall, these results suggest that non-mining regions are disadvantaged by additional mining operations elsewhere in the country. We explore in Section 4.7 below whether they are disadvantaged due to deliberate government actions or inadvertent country-level developments.

4.5.2 Mineral resources and luminosity in capital regions

4.5.2.1 Empirical model

Our preferred specification to study the effect of additional mineral resources on luminosity in capital regions is as follows:

$$y_{i,t} = \alpha_i + \sum_c \gamma_t \times c + \beta \text{Mines}_{c,t} \times \text{Capital}_i + \epsilon_{i,t}, \tag{C3}$$

where i indicates grid cells and t time periods. $y_{i,t}$ is the log of the mean of luminosity in cell i in year t.[26] α_i are cell fixed effects, $\gamma_t \times c$ are country-specific fixed effects, $Capital_i$ is a dummy variable that is 1 for cells that cover capital cities. $Mines_{c,t}$ is a count variable indicating the number of mines in country c and in year t.

Note that the capital region dummy is perfectly collinear with the cell fixed effects (as capitals did not move during the sample period).[27] The number of mines in a country in year t is also perfectly collinear with the country-specific year fixed effects. However, the interaction between the number of mines variable and the capital region dummy variable is neither collinear with the cell fixed effects (since the number of mines varies over time) nor the country-specific year fixed effects (since capital regions constitute only a small fraction of the country). With this specification, we thus compare how luminosity evolves in capital regions when the number of mines increases relative to other regions in the country.

The identifying assumption for this specification is that the variation in the number of mines in a given year throughout the country is orthogonal to unobserved variables in capital regions. This appears to be a reasonable assumption as the startup or closures of mines across the country are plausibly unrelated to specific developments in capital regions (recall that we account for country-wide developments with the country-year fixed effects).

4.5.2.2 Results

The results from estimating Equation C3 are collected in models (2a)-(2c) of Table

[26] As in Equation C1, we add +0.01 to each grid cells' value of mean luminosity to avoid missing values once we take the log.

[27] We leave out the capital of South Sudan, Juba, as this country has only officially existed since 2011.

4.4. To evaluate significance, we again rely on on heteroscedasticity- and cluster-robust standard errors (with cells as the unit of clustering).

Model (2a) reports results with only country and year fixed effects. Capital regions are on average more brightly lit than other regions, which is plausible given that they are typically (substantially) richer than the remainder of the country (African Development Bank, 2015). The number of mines in the country is negatively correlated with luminosity. This may indicate that (additional) mines have adverse aggregate effects.

The variable of interest is the interaction between capital regions and the number of mines. In model (2a), this interaction is insignificant, which suggests that when the number of mines changes, luminosity in capitals does not evolve differently than in other regions. However, model (2a) is a basic specification which, in particular, does not account for cell-specific characteristics other than whether a cell is part of the capital city or not. For example, luminosity might increase more in less developed regions over-time due to catch-up effects, rendering any effects of mine openings on capital cities' luminosity empirically undetectable without further adjustments to the specification.

Indeed, the results in model (2b) suggest that once we add cell fixed effects to account for time-constant differences between cells, the interaction between the capital cities and mines variable turns positive and significant. This suggest that capital regions benefit more from mine openings than other regions in the country. In model (2c), we replace the generic year fixed effects with country-specific year fixed effects. The results are virtually identical to those reported in model (2b).

Overall, the results imply that a share of mineral resource rents are shifted to the national capital. Capitals experience an increase in luminosity of about 5.1% when an additional mine opens, compared to other regions in the country. This effect may come about, on the one hand, because the national government uses additional resource revenues to provide various public goods in the capital. Alternatively, since the national elite resides in the capital, it may also be that the increase in luminosity is due to higher private incomes (and subsequent trickle down effects). In this case, this increase in luminosity in the nation's capital may indicate that at least part of the resource rents is siphoned off and mostly benefits the ruling elites and individuals with relevant connections.

4.5.3 Mineral resources and luminosity in leaders' birthregions

4.5.3.1 Estimation model

Another region within a country that might experience a disproportionate increase in luminosity when aggregate mineral activity picks up is the birth region of the national leader. Previous evidence suggests that (national) leaders engage in regional favoritism and allocate disproportionate resources to their homelands when in power (Hodler and Raschky, 2014; Asatryan et al., 2022, 2021). The motives for such behavior could range

Table 4.4: Non-mining regions and mineral resources

	Panel A– Generic regions		Panel B– Capital cities			Panel C– Leader regions		
	(1a)	(1b)	(2a)	(2b)	(2c)	(3a)	(3b)	(3c)
Mines	-0.011***	-0.011***	-0.014***	-0.015***		-0.013***	-0.014***	
	(0.003)	(0.003)	(0.002)	(0.002)		(0.002)	(0.002)	
Capital city			2.959***					
			(0.192)					
Capital × Mines			0.008	0.062***	0.051***			
			(0.011)	(0.012)	(0.012)			
Leader						1.655***	0.068**	0.052*
						(0.140)	(0.033)	(0.029)
Leader × Mines						-0.005	0.001	0.001
						(0.007)	(0.002)	(0.002)
Country FE	Yes	-	Yes	-	-	Yes	-	-
Cell FE	No	Yes	No	Yes	Yes	No	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	-	Yes	Yes	-
Country-Year FE	No	No	No	No	Yes	No	No	Yes
Countries	48	48	50	50	48	50	50	48
Cells	5041	5041	9058	9058	9056	9058	9058	9056
N	110902	110902	199274	199274	199230	199274	199274	199230

This table collects results for specifications following Equation C2 (Panel A), Equation C3 (Panel B), and Equation C4 (Panel C) that relate mineral resource activity (operating mines) to luminosity at the grid-level across Africa (0.5 × 0.5 degree pixels). We study whether mines lead to an increase in luminosity in generic (border) regions, capitial cities, and leader birth regions. The dependent variable in models (1a)-(1b) is the difference in log mean luminosity in cell i and log mean luminosity in neighboring cells located in other countries. The dependent variable in models (2a)-(3c) is the log of mean light output in each cell. The variable of interest in models (1a)-(1b) is the number of mines in a country in year t. The variable of interest in models (2a)-(2c) is the interaction between a dummy indicating cells that cover capital regions and a count variable indicating the number of mines in a country in year t. The variable of interest in models (3a)-(3c) the interaction between a dummy indicating cells that cover leaders' birth regions and a count variable indicating the number of mines in a country. Only mines with available data on startup and shutdown dates are included. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

from an innate, emotional attachment to their birth regions, to political calculus (e. g., birth regions might be the power base of a leader), or demands for public resources from regional elites of the birth region that the leader might be unable to resist (Baskaran and da Fonseca, 2021).

Our preferred specification to study this question is as follows:

$$y_{i,t} = \alpha_i + \sum_c \gamma_t \times c + \delta \text{Leader}_{i,t} + \beta \text{Mines}_{c,t} \times \text{Leader}_i + \epsilon_{i,t}, \tag{C4}$$

where all indices and variables except the leader dummy are defined as in Equation C3. $\text{Leader}_{i,t}$ is a dummy that is one when a cell covers the birth region of the current national leader and zero otherwise. Note that unlike capital cities, leaders' birth regions can vary over time due to regime changes. As such, the dummy indicating leaders' birth region is included in our preferred specification despite the cell fixed effects.

4.5.3.2 Results

The results from estimating Equation C4 are collected in models (3a)-(3c) of Table 4.4. As before, we rely on heteroscedasticity- and cluster-robust standard errors (with cells as the unit of clustering) to evaluate significance.

In model (3a), the dummy capturing leaders' birth regions is positive and highly significant. This indicates that leaders, on average, originate from richer parts of the country. As in the regressions above on capital cities, the aggregate number of mines is negatively correlated with luminosity, suggesting that mineral resources have detrimental aggregate effects. The interaction between leaders' birth regions and number of mines is insignificant.

The effect remains insignificant when we replace the country fixed effects with cell fixed effects (model 3b) and the generic year fixed effects with country-specific year fixed effects (model 3c). Overall, leaders' birth regions do not benefit more from additional mines than other regions. One reason for this might be that leaders' birth regions might be prioritized even in the absence of additional mineral resource revenues. When additional resources become available due to the opening of mines, they can then be used for purposes other than birthtown favoritism, e. g., for public goods in capital cities or siphoned off for private gains.[28]

[28] We report a number of robustness tests and extensions on the results reported in this section in the Appendix. First, we report results with Conley-standard errors as well as different units of clustering for the baseline specifications in Table D18 in the Appendix. Second, we study in Section D.5.1 whether the effects we identify are stronger (or weaker) for mines that MinEx classifies as at least "major". Third, we explore the effect of different types of mines in Section D.5.2. Fourth, we study whether the EITI has potentially affected non-mining regions in Section D.5.3. Overall, the results reported in the Appendix are in line with the baseline findings reported in this section.

4.6 Heterogeneity by country-level characteristics

There are two important potential sources of heterogeneity regarding the spatial effects of mineral resources. First, the type of political regime, notably whether a country is a democracy or autocracy. It is conceivable that autocratic countries, in particular, fully extract the proceeds of mines out of the mining regions, leaving the mining regions with none of the benefits of higher mineral resource activity. Alternatively, autocratic countries might give a smaller weight to an equitable spatial distribution of mining proceeds and instead allow mining regions to retain a larger fraction of the rents to minimize discontent in these regions.

With respect to non-mining regions, previous research suggests that (regional or ethnic) favoritism is more prevalent in non-democracies (Burgess et al., 2015). Capital regions may also benefit more in non-democratic settings if a lack of accountability and oversight enables national elites to capture more resource rents (Libman, 2013).[29]

The second important source of heterogeneity is the level of corruption in a given country. Similar to the potential heterogeneity with respect to the type of political regime, it is possible that mining regions benefit less in more corrupt countries from the extraction of their minerals. Favored non-mining regions, in turn, could benefit more if expropriated mineral rents are shifted to these.

4.6.1 Type of political regime

To study the effect heterogeneous effect of mines on mining regions, we interact the mining dummy in Equation C1 with dummies for whether or not a country is a democracy. We define dummy variables for democracy and autocracy, respectively, depending on whether the Freedom House index classifies a country as fully democratic or not in a given year (House, 2019). Using these two dummies, we estimate separate coefficients for the intra-regional effect of mines in autocratic and democratic countries, respectively.

The results are collected in model (1) of Table 4.5. We find that the local effects of mining do not differ substantially between more and less democratic countries. In both types of regimes, mineral resource activity induce a local boom of similar size.

To explore the heterogeneous effect of mines on generic non-mining regions, we interact the number of mines variable in Equation C2 with the dummy variables for whether a country is a democracy or autocracy. The results are collected in model (2) of Table 4.5. We observe that the effect of mineral activity does not depend on the extent of democracy. Generic non-mining regions are equally worse off from mining activity elsewhere in both democracies and autocracies.

To study heterogeneous effects on capital cities, we interact the interaction between the

[29]Libman (2013), for example, shows that the effect of natural resources on regional economic performance depends on the level of *subnational* democracy using variation in the level of subnational democracy across the Russian Federation.

Table 4.5: Mineral resources and luminosity in non-mining regions – heterogenous effects by regime type

	(1– Mining)	(2– Border)	(3– Capitals)	(4– Leaders)
Mine × More democratic	0.715***			
	(0.130)			
Mine × Less democratic	0.911***			
	(0.125)			
Mines × More democratic		-0.010***		
		(0.003)		
Mines × Less democratic		-0.012***		
		(0.003)		
Capital × More democratic			0.052***	
			(0.013)	
Capital × Less democratic			0.051***	
			(0.013)	
Leader × More democratic				0.002
				(0.002)
Leader × Less democratic				0.007***
				(0.002)
Country FE	-	-	-	-
Cell FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	-	-
Country-Year FE	Yes	No	Yes	Yes
Countries	46	46	46	46
Cells	8956	4968	8959	8959
N	182368	101684	182460	182460

This table reports an extension of the results in column (1) of Table 4.2 and columns (1b), (2c), and (3c) of Table 4.4. In these extensions, we explore heterogenous effects according to the level of democracy in a country. We interact the respective variables of interest with dummies for whether a country in a given year is a full democracy or not according to the Freedom House index. The dependent variable in column (1),(3), and (4) is the log of mean light output in each cell. In column (2), the dependent variable is the difference in log mean luminosity in cell i and log mean luminosity in neighboring cells located in other countries. The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

capital cities' dummy and the number of mines variable in Equation C3 further with the democracy and autocracy dummies. As per model (3) of Table 4.5, we find no difference in the effect of additional mining activity on capital regions between democracies and autocracies. In both political regimes, capital regions appear to benefit by similar amounts.

We adopt a similar approach for leaders' birth regions and collect the results in model (4) of Table 4.5. For these regions, we observe heterogeneous effects. Specifically, we find that in autocracies, one additional mine increases luminosity in leaders' birth regions by about 0.7%. No similar effects exists for democracies. As discussed above, this result is consistent with previous findings on the importance of democracy for the prevalence of regional or ethnic favoritism.

4.6.2 Level of corruption

To explore heterogeneous effects by the level of corruption on luminosity within mining-regions, we interact the mining dummy in Equation C1 with dummies for whether or not a country is (relatively) corrupt in a given year. We define a country as "more corrupt" in a given year if its score for the "control of corruption" index from the World Governance Indicators is below the median index value in a given year, and vice versa. Using these two dummies, we estimate separate coefficients for more and less corrupt countries.

The results are in model (1) of Table 4.6. We again find that the local effects of mines do not vary significantly between more and less corrupt countries.

To study whether the effects of mineral resources across the three types of non-mining regions differ by the level of corruption in a country, we again estimate separate interactions between the variables of interest and dummies for whether or not a country is more and less corrupt.

The results are collected in models (2)-(4) of Table 4.6. We find that capital cities benefit in both corrupt and non-corrupt settings. However, leaders' birth regions only experience an increase in luminosity when the country is relatively corrupt. Similarly, generic regions only witness a decline in luminosity in corrupt countries. These results suggest that the level of corruption rather than the political regime determines how mineral revenues affect non-mining regions, in particular generic non-mining regions.

Table 4.6: Mineral resources and luminosity in non-mining regions – heterogenous effects by the level of corruption

	(1– Mining)	(2– Border)	(3– Capitals)	(4– Leaders)
Mine × Less corruption	0.966***			
	(0.135)			
Mine × More corruption	0.739***			
	(0.147)			
Mines × Less corruption		0.002		
		(0.003)		
Mines × More corruption		-0.025***		
		(0.003)		
Capital × Less corruption			0.035***	
			(0.011)	
Capital × More corruption			0.056***	
			(0.016)	
Leader × Less corruption				0.001
				(0.001)
Leader × More corruption				0.025***
				(0.009)
Country FE	-	-	-	-
Cell FE	Yes	Yes	Yes	Yes
Year FE	-	Yes	-	-
Country-Year FE	Yes	No	Yes	Yes
Countries	46	46	46	46
Cells	8956	4968	8959	8959
N	125091	69759	125156	125156

This table reports an extension of the results in column (1) of Table 4.2 and columns (1b), (2c), and (3c) of Table 4.4. In these extensions, we explore heterogenous effects according to the level of democracy in a country. We interact the respective variables of interest with dummies for whether a country in a given year has a high or low level of corruption according to the "control of corruption" index from the World Governance Indicators. The dependent variable in column (1),(3), and (4) is the log of mean light output in each cell. In column (2), the dependent variable is the difference in log mean luminosity in cell i and log mean luminosity in neighboring cells located in other countries. The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

4.7 Mechanisms

In the estimates reported in Section 4.4 and 4.5, we observe positive effects of mines on mining regions. These positive effects are visible up to 30km away from the location of the mine. We also observe positive effects for capital regions. There are also positive effects on leaders' birth regions as long as a country is an autocracy. In contrast, we observe negative effects on generic regions.

The extensions we have explored suggest the following mechanisms for these findings. First, that mines have spillovers for up to 30km suggests a genuine expansion in economic activity. That capital regions and – if a country is not a democracy – leader regions benefit from more mineral resource activity indicates that the government deliberately redistributes mineral resource revenues from the mining regions to regions that are politically important.

What is less clear is why luminosity declines in generic non-mining regions. We thus explore in the following the mechanisms behind our baseline findings further, focusing in particular on generic non-mining regions.

4.7.1 Conflicts

One potentially important consequence of mines is an uptick in conflicts. In mining regions, conflicts could emerge for control of the mineral resources. However, non-mining regions could also witness more conflicts, for example if mineral resources are used to fund conflicts elsewhere. Conflicts, in turn, might have negative consequences for economic development.

In Table 4.7, we explore the effect of mines on the incidence of conflicts. For this, we replace in Equation C1 luminosity with the number of conflicts in cell i and year t. It appears that – in our sample – mines do not lead to more conflicts.[30] In fact, for giant mines, we even observe a small negative and statistically significant effect. Overall, the insignificant effect of mines on conflicts within mining regions is consistent with their positive effect on luminosity.

In Table 4.8, we explore the effect of mines on conflicts in non-mining regions. We find no effects in capital regions. However, it appears conflicts decline in leaders' birth regions if mineral resource activity expands elsewhere in the county (while the estimate is insignificant, it has a relatively large p-value). At the same time, the number of conflicts increases in generic regions across the country. Both findings can be interpreted in conjunction with the results found for luminosity. That is, the increase in luminosity in leaders' birth regions in non-democratic countries (model 4 in Table 4.5) could be explained by a decline in conflicts, for example due to leaders using mineral resource rents to pacify competing local groups.

Similarly, the decline in luminosity in generic non-mining regions (models (1a)-(1b) of Table 4.4) could be in part explained by an increase in conflicts, possibly because leaders make use of the resource revenue to fund conflicts elsewhere. Of course, another possible interpretation of this finding is that it is the deterioration of economic conditions in non-mining regions (e. g., due to macroeconomic adjustments) that is responsible for the increase in conflicts. This question is an exiting area for future research. In any

[30] Note that Arezki et al. (2015) report similar results.

Table 4.7: Mineral resources and conflicts in mining regions

	(1: All mines)	(2: $\geq$ Major mines)	(3: $\geq$ Giant)	(4: < Major)	(5: < Moderate)	(6: Prices)
Mine	-0.037	-0.074	-0.036***	0.036	0.107	
	(0.038)	(0.054)	(0.012)	(0.034)	(0.095)	
Mineral price						-0.003
						(0.004)
Cell FE	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Cells	9063	9063	9063	9063	9063	9063
N	199360	199360	199360	199360	199360	199360

This table collects difference-in-differences regressions that relate mineral resource activity (operating mines) to the number of conflicts at the grid-level for all of Africa (0.5 $\times$ 0.5 decimal degree cells). In this specification, we study whether mines affect conflict incidents in mining regions. The dependent variable is the sum of conflict events in each cell. The independent variable in models (1)-(5) is a dummy variable that is 1 if a cell had an operating mine of a certain size as indicated in the column header. The independent variable in model 6 is the contemporanous price of the mineral that is extracted in a given cell. The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

case, there appears to be a clear link between between country-wide aggregate mineral resources, conflicts, and economic developments in non-mining regions.

4.7.2 Ethnic redistribution

To further understand to what extent the spatial patterns we observe are due to intentional decisions of policy makers – notably inter-regional redistribution of mining revenues – we explore how luminosity responds to mineral resource activity based on the spatial distribution of ethnic groups. Specifically, we study whether regions inhabited by ethnic groups that are politically powerful, i. e., that are in positions of power at the national level, respond differently to mineral resources than regions inhabited by ethnic groups that are politically weak (i. e., discriminated against).

We classify the grid cells in each year as “strong” if they are inhabited by at least one group whose “power status” is coded by Vogt et al. (2015) as “Monopoly” or “Dominant”. Similarly, we classify all cells inhabited by at least one group coded as “Discriminated” as weak. We classify all other grid cells as neutral. Cells can be both weak and strong if they are inhabited by two or more groups with opposite “power status”. Cells can also change their classification over time (see Figure 4.3 for a visual representation of cells' power status).

We then explore whether mines lead to a larger increase in luminosity if the region surrounding a mine is inhabited by an ethnic group that is politically strong and vice versa. Such a pattern would indicate that the national government deliberately redistributes more resources away from mineral resource regions if they are inhabited by ethnicities that are discriminated against. The corresponding empirical model we use to explore this question is an extension of Equation C1. Specifically, we add an interaction of the mines dummy with a dummy for whether or not a cell is inhabited by a politically strong or weak ethnic group, respectively.

Figure 4.3: Grid of Africa matched to political power of ethnic groups

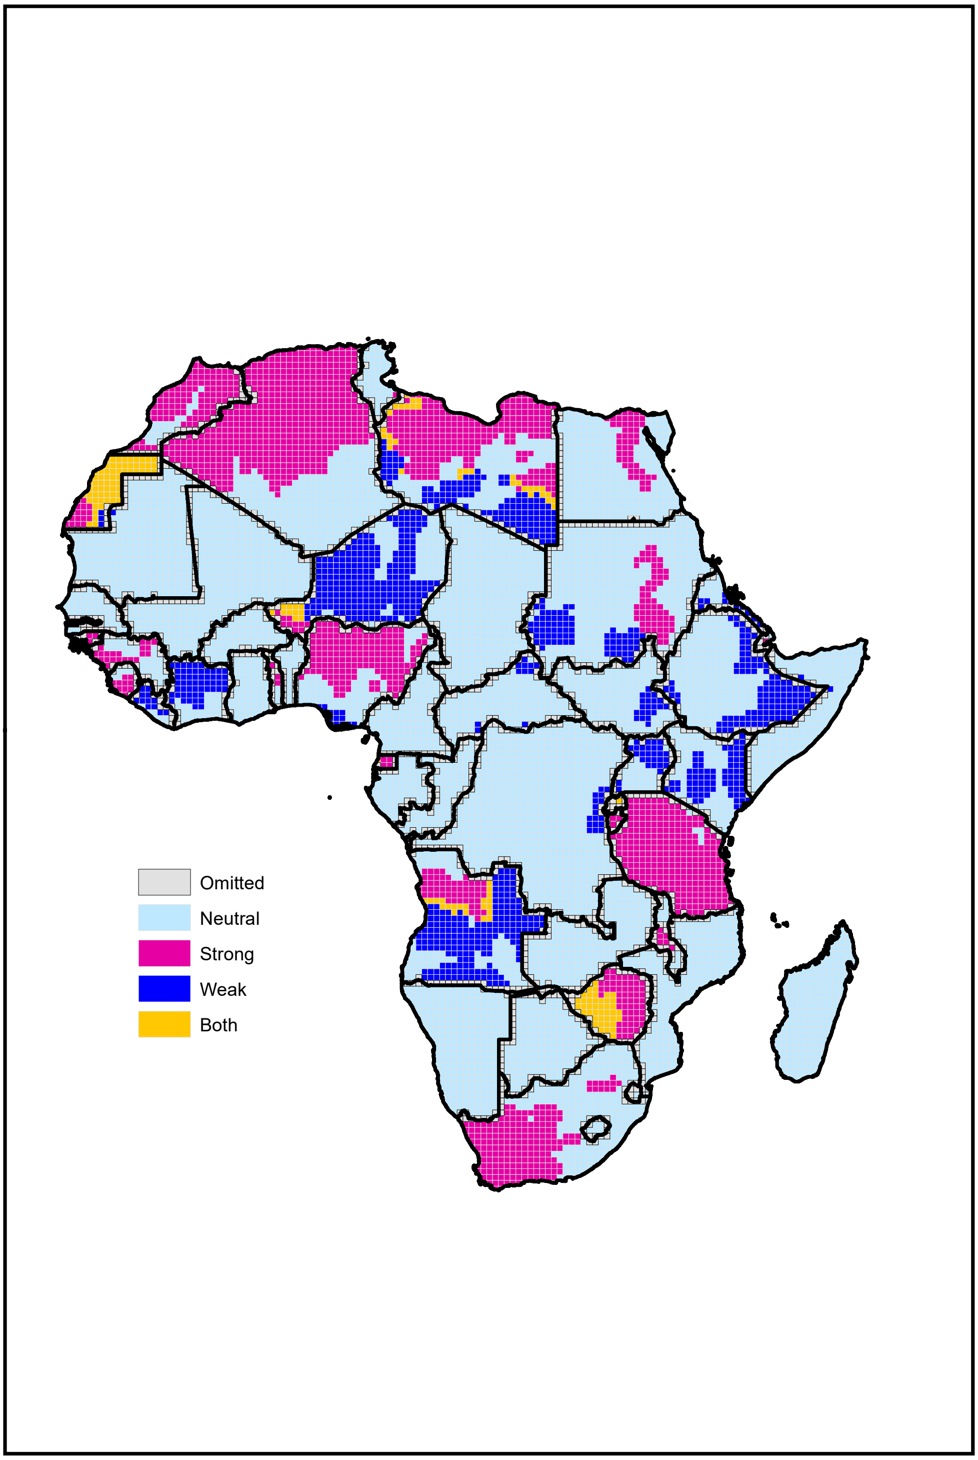

The figure shows a grid over Africa. It indicates all cells that were inhabited by at least one ethnic groups that was politically powerful (i. e. that was classified as having a "monopoly" on power or as politically "dominant" in the ethnic power dataset by Vogt et al. (2015)) for at least for one year during the sample period. It also indicates all cells that were politically weak (i. e., that were inhabited by at least one ethnic group classified as "discriminated" for at least one year during the sample period). Some cells have both politically strong and politically weak ethnic groups. All other cells are indicated as neutral.

Table 4.8: MINERAL RESOURCES AND CONFLICTS IN NON-MINING REGIONS

	(1– Border)	(2– Capitals)	(3– Leaders)
Mines	0.020***		
	(0.005)		
Capital × Mines		0.011	
		(0.019)	
Leader × Mines			-0.004
			(0.002)
Country FE	-	-	-
Cell FE	Yes	Yes	Yes
Year FE	Yes	-	-
Country-Year FE	No	Yes	Yes
Countries	48	48	48
Cells	5050	9066	9066
N	111100	199452	199452

This table reports results that explore the effect of mineral resources on conflicts in non-mining regions. In column (1), the dependent variable is the difference in sum of conflicts events in cell i and the mean of the sum of conflict events in neighboring cells located in other countries. The dependent variable in column (2)-(3) is the sum of conflict events in each cell in a given year. We study whether mines lead to an increase in conflict events in generic regions (column (1)), in capital regions (column (2)), and in leader's birth regions (column (3)). The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

We collect the results in Table 4.9. They suggest that the effect of mines is substantially smaller, in fact almost zero, when the mining region is inhabited by a politically weak ethnic group. This finding indicates that national governments intentionally redistribute resources spatially, in particular that they redistribute more resources away from mining regions when they are inhabited by politically weak ethnicities.

In Table 4.10, we explore how generic non-mining regions respond to mineral resources in view of their ethnic markup. We focus on generic non-mining regions as capital regions and leaders' birth regions are arguably politically well represented in all circumstances. Specifically, we interact in Equation C2 the number of mines variable with dummies for whether or not a cell was inhabited by a politically weak or strong ethnicity, respectively.

Table 4.10 suggests that generic regions inhabited by politically weak groups experience a large decline in luminosity when mineral resource activity expands. In contrast, no such decline is observable in regions that are inhabited by politically strong groups. These results reaffirm that, at least in part, the spatial implications of mineral resource activity emerge due to deliberate redistribution by the national government.

Table 4.9: Mines and luminosity in mining regions inhabited by politically strong and weak ethnicities

	(1)	(2)	(3)
Mine	0.821***	0.846***	0.842***
	(0.125)	(0.125)	(0.125)
Strong	-0.045		-0.047
	(0.036)		(0.036)
Strong × Mine	-0.078		0.030
	(0.120)		(0.128)
Weak		0.126***	0.126***
		(0.028)	(0.028)
Weak × Mine		-0.611***	-0.625***
		(0.218)	(0.229)
Cell FE	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes
Cells	7281	7281	7281
N	158547	158547	158547

This table collects difference-in-differences regressions that relate mineral resource activity (operating mines) to luminosity at the grid-level for all of Africa (0.5 × 0.5 decimal degree cells). In this specification, we study whether mines increase luminosity in mining regions less if they are inhabited by politically weak regions and vice versa. The dependent variable is the log of mean light output in each cell. The independent variables are a dummy variable that is 1 if a cell had an operating mine and an interaction variable that is one if a cell is inhabited by a contemporaneously politically weak or strong ethnicity, respectively. The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

Table 4.10: Generic non-mining regions and mineral resources inhabited by politically strong and weak ethnicities

	(1)	(2)
Strong	-0.416***	0.023
	(0.101)	(0.031)
Weak	0.067	0.056**
	(0.098)	(0.028)
Mines × Strong	0.004	-0.000
	(0.004)	(0.001)
Mines × Weak	-0.138***	-0.024***
	(0.030)	(0.006)
Country FE	Yes	-
Cell FE	No	Yes
Year FE	Yes	Yes
Country-Year FE	No	No
Countries	44	44
Cells	4077	4077
N	88427	88427

This table collects results for specifications following equation C2 that relate mineral resource activity (operating mines) to luminosity at the grid-level across Africa (0.5 × 0.5 degree pixels). We study whether the aggregate number of mines affects luminosity in generic non-mining regions. We additionally distinguish between border regions inhabited by "politcally strong" and politically weak" ethnicities. The dependent variable is the difference between the log of mean light output in each border cell and the average of the log mean light output in neighboring cells in foreign countries. The independent variable is a count variable indicating the number of mines in a given country in year t. Only mines with available data on startup and shutdown dates are included. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

4.7.3 Country-level institutional quality

An alternative channel through which mineral resources might negatively affect economic outcomes in generic non-mining regions is a general decline in institutional quality. If mineral resources result in corruption and political instability, it is plausible that this will negatively affect economic development across the country, and in particular in generic non-mining regions. This channel would thus imply that luminosity declines in generic non-mining regions also due to inadvertent country-level adjustments, in addition to those direct government actions as discussed above.

To explore this channel, we relate the number of mines in a country to country-level institutional quality using standard cross-country regressions. More specifically, we use four proxies for institutional quality: whether or not a country is democratic, whether it has high levels of corruption, whether it has high levels of government effectiveness, and whether it has high levels of political stability.

The results are collected in Table 4.11. Overall, we find no strong associations between mines and institutional quality. As such, the reason why mines diminish economic activity in generic non-mining regions does not appear to be a decline in aggregate institutional quality.

Table 4.11: Mines and country-level institutional quality

	(1– Democracy)	(2– Corruption)	(3– Government effectiveness)	(4– Political stability)
Mines	-0.006	0.018	0.016*	-0.034
	(0.010)	(0.018)	(0.009)	(0.028)
Country FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Countries	47	47	47	47
N	993	680	680	680

This table collects results for specifications that relate measures for institutional quality to country-level mineral resource activity. The dependent variable in model 1 is a dummy for whether or not a country is a full democracy according to the Freedom House index. The dependent variable in model 2 is the control of corruption score from the Worldwide Governance Indicators. The dependent variable in model 3 is the government effectiveness score from the Worldwide Governance Indicators. The dependent variable in model 4 is the political stability and absence of violence/terrorism score from the Worldwide Governance Indicators. Higher scores correspond to better outcomes in the Worldwide Governance Indicators. Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the country.

4.7.4 Terms of trade adjustments

Another channel through which mineral resources could cause economic harm to non-mining regions is terms of trade adjustments. Regions that specialize in sectors other than mining might be harmed if the exchange rate appreciates when mineral resource exports increase. For example, exports of manufacturing or agricultural products might decline, in turn hurting regions that specialize in these sectors. Similarly, regions with a large hospitality industry that caters to foreign tourists might witness lower demand. These effects would be in line with the Dutch Disease literature.[31]

To study this channel, we make use of census data from IPUMS. The IPUMS census data indicate the location of respondents up to the second tier (GEOLEV2) of government (e.g., districts). Using all available census waves and countries within Africa, we calculate the share of respondents in each tier 2 region employed in (i) mining, (ii) manufacturing, (iii) agriculture, and (iv) in the hospitality industry (hotel and restaurants) over the sample period (i. e., we take the over-time average for regions that have data in multiple census waves). We then estimate interaction models of the following form:

$$y_{i,t} = \alpha_i + \sum_c \gamma_t \times c + \beta \text{Mines}_{c,t} \times \text{Sector share}_i + \epsilon_{i,t}, \tag{C5}$$

with i indicating different tier-two regions, α_i region fixed effects, γ_t year fixed effects and c country dummies (i. e., we include country specific year fixed effects), and Mines the number of mines in a country. Sector share is the share of respondents in each region that works in each of the four sectors discussed above.

The interaction between the number of operating mines in the country and the sector share is the variable of interest and captures whether regions that specialize in either of the four sectors are better or worse off if mineral resource activity expands in the country. Figure D4 in the Appendix shows the industry shares across regions with available data. The data is available for 20 countries. Note that the variable capturing the share of each sector is perfectly collinear with the region fixed effects.

We collect the results form estimating Equation C5 in Table 4.12. As expected, we observe that an uptick in mining activity increases light output in regions where a larger share of the respondents is employed in the mining sector (model 1). The interaction effect between the number of mines and the share of respondents working in the mining sector is positive and significant. In contrast, the interaction between the number of mines and the share of respondents working in manufacturing is negative and significant at the 10% level

[31]The recent empirical evidence on the Dutch Disease in general, and specifically in Africa is mixed. Harding and Venables (2016) find that natural resource revenues generally decrease exports by the non-resource sectors. Asiamah et al. (2022) find similar evidence specifically for sub-Saharan Africa. Cust et al. (2022) show that Dutch Disease effects are relevant in Africa and that one important channel for its emergence is the public sector. On the other hand, Pegg (2010) argues that while Botswana exhibits many symptoms consistent with the Dutch Disease, the prime reasons for these symptoms are not those that are discussed in the Dutch Disease literature.

(model 2), indicating that additional mineral resource activity harms regions specializing in manufacturing. We also observe negative (albeit insignificant) coefficients for regions specializing in agriculture (model 3) and the hospitality industries (model 4).

Overall, these estimates indicate that an expansion in mineral resource activity has negative effects on regions that specialize in sectors other than mining. As discussed, one obvious reason for this is terms of trade adjustments that harm non-mining regions.

Table 4.12: SPATIAL EFFECT OF MINES BY SECTORAL SPECIALIZATION

	(1– Mining)	(2– Manufacturing)	(3– Agriculture)	(4– Hospitality)
Mines × Mining share	2.350***			
	(0.690)			
Mines × Manufacturing share		-0.665*		
		(0.341)		
Mines × Agriculture share			-0.062	
			(0.040)	
Mines × Hospitality share				-1.318
				(1.320)
Country FE	-	-	-	-
Cell FE	Yes	Yes	Yes	Yes
Year FE	-	-	-	-
Country-Year FE	Yes	Yes	Yes	Yes
Countries	20	20	20	20
Cells	1734	1734	1734	1734
N	38148	38148	38148	38148

This table collects difference-in-differences regressions that relate mineral resource activity (operating mines) to luminosity at the tier 2 regional level (GEOLEV2) for all of Africa. In these specifications, we study whether mines lead to an increase in luminosity in regions with a larger (i) mining sector, (ii) manufacturing sector, (iii) agricultural sector, and (iv) hospitality sector. The dependent variable is the log of mean light output in each cell. The variable of interest is the interaction between the over-time average of the share of census respondents employed in each sector and a count variable indicating the number of mines in a country in year t. Only mines with available data on startup and shutdown dates are included. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the region.

4.8 Conclusion

We study how mineral resources affect economic conditions both in mining and non-mining regions across the African continent by combining cross-country data and subnational variation. Using nighttime luminosity as proxy for the local economy, we find that mines expand economic activity in mining regions. The economic implications of mineral resources on non-mining regions, on the other hand, are heterogeneous: (i) generic non-mining regions are in general disadvantaged, (ii) capital regions benefit from mineral resource activity anywhere in the country, and (iii) leaders' birth regions benefit in autocratic regimes.

This second set of results suggests that mineral resources have important economic implications beyond the mining regions. While the improvement in economic conditions in capitals and leaders' birth regions suggest that the government engages in regional favoritism and redistributes mining revenues to politically important regions, the explanation for the decline in economic activity in generic non-mining regions is less obvious. Exploring mechanisms for this result, we found that conflicts increase in generic non-mining regions when mineral resource activity expands, a finding consistent with the interpretation that national governments partially use mining revenues to fund conflicts elsewhere in the country. We also find evidence for terms of trade adjustments that disadvantage regions specializing in sectors other than mining. As such, the decline in economic activity within generic non-mining regions can be explained by both deliberate government policies as well as inadvertent macroeconomic adjustments due to increased mineral exports.

Overall, these findings advance our understanding of how mineral resources affect the spatial distribution of economic activity. They underscore, in particular, that mineral resources not only affect mining regions, but, through various channels, non-mining regions as well. As the effects of mineral resources on non-mining regions have been neglected in the literature, one important conclusion from this paper is that such broader spatial effects should be taken into account in future research on the economic implications of mineral resources.

References

African Development Bank, “Africa Economic Outlook 2015: regional development and spatial inclusion,” 2015.

Ahlerup, P., T. Baskaran, and A. Bigsten, “Gold mining and education: a long-run resource curse in Africa?,” *Journal of Development Studies*, 2019, *56*, 1745–1762.

Alesina, A., S. Michalopoulos, and E. Papaioannou, “Ethnic inequality,” *Journal of Political Economy*, 2016, *124*, 428–488.

Allcott, Hunt and Daniel Keniston, “Dutch Disease or agglomeration? The local economic effects of natural resource booms in modern America,” *Review of Economic Studies*, 07 2017, *85* (2), 695–731.

Aragón, F. M. and J. P. Rud, “Natural resources and local communities: evidence from a Peruvian gold mine,” *American Economic Journal: Economic Policy*, 2013, *5*, 1–25.

_ , P. Chuhan-Pole, and B. C. Land, “The local economic impacts of resource abundance: what have we learned?,” Policy Research Working Paper 7263 2015.

Arezki, R., S. Bhattacharyya, and N. Mamo, “Resource discovery and conflict in Africa: what do the data show?,” CSAE Working Paper WPS/2015-14 2015.

Arora, N., J. Ghose, and S. K. Bakshi, “Managing revenue from natural resources: a multi-country analysis of sharing resource revenue with sub-national levels,” Working Paper 2017.

Asatryan, Z., T. Baskaran, C. Birkholz, and D. Gomtsyan, “Favoritism and Firms: micro evidence and macro implications,” CESifo Working Paper No. 9797 2022.

_ , _ , P. Hufschmidt, and A. Stöcker, “Regional favoritism and human capital accumulation in Africa,” ZEW Discussion Paper 21-030 2021.

Asatryan, Zareh and Annika Havlik, “The political economy of multilateral lending to European regions,” *The Review of International Organizations*, 2020, *15* (3), 707–740.

Asiamah, Oliver, Samuel Kwaku Agyei, Bossman Ahmed, and Ellen Animah Agyei, “Natural resource dependence and the Dutch Disease: Evidence from sub-Saharan Africa,” *Resources Policy*, 2022, *79*, 103042.

Baker, Andrew C., David F. Larcker, and Charles C. Y. Wang, “How much should we trust staggered difference-in-differences estimates?,” *Journal of Financial Economics*, May 2022, *144* (2), 370–395.

Baskaran, T. and M. L. da Fonseca, “Appointed public officals and local favoritism: evidence from the German States,” *Journal of Urban Economics*, 2021, *124*, 103354.

Baskaran, Thushyanthan, “Soft budget constraints and strategic interactions in subnational borrowing: evidence from the German States, 1975-2005,” *Journal of Urban Economics*, 2012, *71*, 114–127.

Bazillier, R. and V. Girard, “The gold digger and the machine. Evidence from the artisanal and industrial gold rush in Burkina Faso,” Mimeo 2017.

Benshaul-Tolonen, A., “Local industrial shocks and infant mortality,” *Economic Journal*, 2019, *129*, 1561–1592.

Berman, N., M. Couttenier, D. Rohner, and M. Thoening, “This mine is mine! How minerals fuel conflicts in Africa,” *American Economic Review*, 2017, *107*, 1564–1610.

Berry, C. R., B. C. Burden, and W. G. Howell, “The president and the distribution of federal spending,” *American Political Science Review*, 2010, *104*, 783–799.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting event study designs: robust and efficient estimation," April 2022.

Boschini, A. D., J. Pettersson, and J. Roine, "Resource curse or not: a question of appropriability," *Scandinavian Journal of Economics*, 2007, *109*, 593–617.

Brollo, F., T. Nannicini, R. Perotti, and G. Tabellini, "The political resource curse," *American Economic Review*, 2013, *103*, 1759–1796.

Brosio, G. and R. J. Singh, "Revenue sharing of natural resources in Africa: reflections from a review of international literature," World Bank Background Paper 90252 2014.

Bruederle, Anna and Roland Hodler, "Nighttime lights as a proxy for human development at the local level," *PloS one*, 2018, *13* (9), e0202231.

Bulte, E. H., R. Damania, and R. T. Deacon, "Resource intensity, institutions, and development," *World Development*, 2005, *33*, 1029–1044.

Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padro i Miquel, "The Value of Democracy: Evidence from Road Building in Kenya," *American Economic Review*, 2015, *105* (6), 1817–1851.

Callaway, B. and P. H. C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 2021, *225*, 200–230.

Campbell, G., "Blood diamonds: tracing the deadly path of the world's most precious stones.," 2002. Westview Press.

Caselli, F. and G. Michaels, "Do oil windfalls improve living standards? Evidence from Brazil," *American Economic Journal: Applied Economics*, 2013, *5*, 208–238.

Chuhan-Pole, P., A. L. Dabalen, and B. C. Land, "Mining in Africa," World Bank Report 2017.

Collier, P. and A. Hoeffler, "Testing the Neocon Agenda: democracy in resource-rich societies," *European Economic Review*, 2009, *53*, 293–308.

Conley, T.G., "GMM estimation with cross sectional dependence," *Journal of Econometrics*, 1999, *92*, 1–45.

Corden, W. M., "Booming sector and Dutch Disease economics: survey and consolidation," *Oxford Economic Paper*, 1984, *36*, 359–380.

Cust, J. and C. Viale, "Is there evidence for a subnational resource curse," Natural Resource Governance Institute Policy Paper 2016 2016.

_ **and D. R. Ridwan**, "The economic spillovers from resource extraction: a partial resource blessing at the subnational level?," CREA Discussion Paper 2014-08 2014.

_ **and S. Poelhekke**, "The local economic impacts of natural resource extraction," *Annual Review of Resource Economics*, 2015, *7*, 251–268.

Cust, James, Shantayanan Devarajan, and Pierre Mandon, "Dutch Disease and the public sector: how natural resources can undermine competitiveness in Africa," *Journal of African Economies*, 09 2022, *31*, i10–i32.

Dahlberg, S., S. Holmberg, B. Rothstein, N. Alvarado Pachon, and S. Axelsson, "Quality of Government Basic Dataset," University of Gothenburg: The Quality of Government Institute 2020.

de la Sierra, R. S., "On the origins of the state: stationary bandits and taxation in Eastern Congo," *Journal of Political Economy*, 2020, *128*, 32–74.

Do, Quoc-Anh, Kieu-Trang Nguyen, and Anh Tran, "One mandarin benefits the whole clan: hometown favoritism in an authoritarian regime," *American Economic Journal: Applied Economics*, 2017, *9*, 1–29.

Dreher, A., A. Fuchs, R. Hodler, B. C. Parks, P. A. Raschky, and M. J. Tierney, "Is favoritism a threat to chinese aid effectiveness? A subnational analysis of chinese development projects," *World Development*, 2021, *139*, 105291.

Ericsson, M. and O. Löf, "Mining's contribution to national economies between 1996 and 2016," *Mineral Economics*, 2019, *32*, 223–250.

Fearon, J. D. and D. D. Laitin, "Ethnicity, insurgency, and civil war," *American Political Science Review*, 2003, *97*, 75–90.

Fetzer, T., "Can workfare programs moderate violence? Evidence from India," STICERD Working Paper 2014.

Feyrer, J., E. T. Mansur, and B. Sacerdote, "Geographic dispersion of economic shocks: evidence from the fracking revolution," *American Economic Review*, 2017, *107*, 1313–1334.

Fjeldstad, O.-H., G. Chambas, and J.-F. Brun, "Local government taxation in Sub-Saharan Africa," CHR. Michelsen Institute Working Paper 2014.

Gibson, J., S. Olivia, G. Boe-Gibson, and C. Li, "Which night lights data should we use in economics, and where?," *Journal of Development Economics*, 2021, *149*, 102602.

Global Witness, "Digging in corruption: fraud, abuse and exploitation in Katanga's copper and cobalt mines," Global Witness Report 2006.

Goemans, H. E., "ARCHIGOS: a data set on leaders 1875-2015, Version 4.1," https://www.rochester.edu/college/faculty/hgoemans/data.htm 2016.

Gollin, D., R. Jedwab, and D. Vollrath, "Urbanization with and without industrialization," *Journal of Economic Growth*, 2016, *21*, 35–70.

Guidolin, M. and E. La Ferrara, "Diamonds are forever, wars are not: is conflict bad for private firms?," *American Economic Review*, 2007, *97*, 1978–1993.

Harding, T. and A. J. Venables, "The implications of natural resource exports for nonresource trade," *IMF Economic Review*, 2016, *64*, 268–302.

Henderson, J. V., A. Storeygard, and D. N. Weil, "Measuring economic growth from outer space," *American Economic Review*, 2012, *102*, 994–1028.

Hodler, R., P. Schaudt, and A. Vesperoni, "Mining for peace," Cesifo Working Paper 10207 2023.

Hodler, Roland and Paul Raschky, "Regional Favoritism," *Quarterly Journal of Economics*, 2014, *129* (2), 995–1033.

House, Freedom, "Freedom in the world 2019," www.freedomhouse.org 2019.

Hsiang, S. M., "Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America.," *PNAS*, 2010, *107*, 15367–72.

Huang, Q., V. W. Xie, and W. You, "Resource rents, urbanization, and structural transformation," Mimeo 2022.

Humphreys, M., J. D. Sachs, and J. E. Stiglitz, "Introduction: what is the problem with natural resource wealth?," in "Escaping the resource curse," New York: Columbia University Press, 2007.

Isham, J., M. Woolcock, L. Pritchett, and G. Busby, "The varieties of resource experience: natural resource export structures and the political economy of economic growth," *World Bank Economic Review*, 2005, *19*, 141–174.

James, A. and D. Aadland, "The curse of natural resources: an empirical investigation of U.S. counties," *Resource and Energy Economics*, 2011, *33*, 440–453.

Khan, T. S., T. T. T. Nguyen, R. Schodde, and F. Ohnsorge, "From Commodity Discovery to Production," World Bank Policy Research Working Paper No. 7823 2016.

Knutsen, C. H., A. Kotsadam, E. H. Olsen, and T. Wig, "Mining and local corruption in Africa," *American Journal of Political Science*, 2017, *61*, 320–334.

Kotsadam, A. and A. Tolonen, "African mining, gender, and local employment," *World Development*, 2016, *83*, 325–339.

Land, B., "Capturing a fair share of fiscal benefits in the extractive industry," *Transnational Corporations*, 2009, *18*, 157–174.

Laporte, B. and Quatrebarbes, "What do we know about the sharing of mineral resource rent in Africa?," FERDI Working Paper 126 2015.

Libman, A., "Natural resources and sub-national economic performance: does sub-national democracy matter?," *Energy Economics*, 2013, *37*, 82–99.

Loayza, N., A. Mier y Teran, and J. Rigolini, "Poverty, inequality, and the local natural resource curse," IZA discussion paper 7226 2013.

Määttä, I., T. Ferreira, and C. Leßmann, "Nighttime lights and wealth in very small areas: Namibian complete census versus DHS data," *Review of Regional Science*, 2022, *42*, 161–190.

Maddala, G. S., "On the use of panel data methods with cross-country data," *Annales d' Économie et de Statistique*, 1999, *55/56*, 429–448.

Mamo, N., S. Bhattacharyya, and A. Moradi, "Intensive and Extensive Margins of Mining and Development: Evidence from sub-Saharan Africa," *Journal of Development Economics*, 2019, *139*, 28–49.

Martínez, L. R., "How much should we trust the dictator's GDP growth estimates?," *Journal of Political Economy*, 2022, *130.*

Mayer, T. and S. Zignago, "Notes on CEPII's distances measures: The GeoDist database," CEPII WP No. 2011 − 25 2011.

Maystadt, J.-F., G. De Luca, P. G. Sekeris, and J. Ulimwengu, "Mineral resources and conflicts in DRC: a case of ecological fallacy?," *Oxford Economic Papers*, 2014, *66*, 721–749.

Michaels, G., "The long term consequences of resource-based specialization," *Economic Journal*, 2011, *121*, 31–57.

Michalopoulos, S. and E. Papaioannou, "The long-run effects of the scramble for Africa," *American Economic Review*, 2016, *106*, 1802–1848.

Oates, W., "An essay on fiscal federalism," *Journal of Economic Literature*, 1999, *37* (3), 1120–1149.

Otto, J. M., "Fiscal decentralization and mining taxation," Mimeo (World Bank Group Mining Deartment) 2001.

Pegg, Scott, "Is there a Dutch Disease in Botswana?," *Resources Policy*, 2010, *35* (1), 14–19.

Pokorny, Benno, Christian von Lübke, Sidzabda Djibril Dayamba, and Helga Dickow, "All the gold for nothing? Impacts of mining on rural livelihoods in Northern Burkina Faso," *World Development*, 2019, *119*, 23–39.

Reingewertz, Y. and T. Baskaran, "Distributive spending and presidential partisan politics," *Public Choice*, 2020, *186*, 65–85.

Roth, J., P. H. C. Sant'Anna, A. Bilinski, and A. Poe, "What's trending in difference-in-differences," arXiv:2201.01194 2022.

Sachs, J. D. and A. M. Warner, "Natural resource abundance and economic growth," Mimeo (Harvard University) 1995.

Standing, A. and G. Hilson, "Distributing mining wealth to communities in Ghana," Report (U4 Anti-Corruption Resource Centre) 2013.

van der Ploeg, F., "Natural Resources: curse or blessing?," *Journal of Economic Literature*, 2011, *49*, 366–420.

Vogt, M., N.-C. Bormann, S. Rüegger, L.-E. Cederman, P. Hunziker, and L. Girardin, "Integrating data on ethnicity, geography, and conflict: the ethnic power relations data set family," *Journal of Conflict Resolution*, 2015, *59*, 1327–42.

Wildasin, D., "Externalities and bailouts: hard and soft budget constraints in intergovernmental fiscal relations," World Bank Policy Research Working Paper No. 1843 1997.

WorldPop and CIESIN, "Global High Resolution Population Denominators Project," `http://www.worldpop.org` 2018. Funded by The Bill and Melinda Gates Foundation (OPP1134076). Collaboration between School of Geography and Environmental Science, University of Southampton; Department of Geography and Geosciences, University of Louisville; and Département de Géographie, Université de Namur.

D Appendix

D.1 Further information on MinEx mining data

MinEx states regarding the compilation of the database the following:

"MinEx's Deposit Database has been built up over several years and is based on information sourced from company public reports (Annual Reports, press releases and NR 43-101 studies etc), technical and trade journals (such as Economic Geology, Northern Miner and Mining Journal), Government Files (from the various Geological Surveys) and personal communications with key people in the industry. The data is current as of end-June 2015".

Furthermore, the data used in this paper was compiled:

"from MinEx's main database which contains information on over 55,000 mineral deposits across a wide range of metals. A large number of these deposits are smaller than "Minor" - and as such are of limited commercial interest".

The (primary) minerals produced in the 519 significant deposits in operation during the sample period are:

Andalusite, Asbestos, Barium, Chromium, Cobalt, Copper, Diamonds, Flourine, Fluorite, Gold, Lead, Manganese, Mineral Sands, Nickel, PGE, Platinum, Rare Earths, Ruby, Sapphire, Silver, Sulphur, Tin, Tungsten, Uranium, Vermiculite, Zinc.

D.2 Data on mineral prices

We obtain data on mineral prices from different sources. Whenever available, we rely on price data provided by the World Bank (i. e., World Bank Commodity Price Data – The Pink Sheet). For cobalt, we use data from the IMF Primary Commodity Prices Database.[32] For the remaining minerals – for which no price data is available from the World Bank nor the IMF – we use data on US unit prices (rather than world prices) from the US Geological Survey. For several minerals, no price data is available from any of these public sources. See Table D14 for further details regarding the price data and Table D15 for summary statistics.

D.3 Suitability of luminosity as proxy for local economic development

Is luminosity an accurate proxy for local economic development? To study this question, we compare cell-level luminosity with (i) cell-level population, (ii) access to electricity within a cell, and (iii) respondents wealth within a cell.

The population data is obtained from WorldPop. Specifically, data on population counts are available annually for the period 2000-2020 as raster files. We use raster files with a resolution of 1 km and aggregate the population counts to the level of the 0.5×0.5 grid cells.

The WorldPop data is based on recent official census population data and various other input data sources, such as location and extent of settlements, roads, land cover, building maps, vegetation, topography, health facility locations, and refugee camps. However, it also takes satellite nightlights into account, implying some degree of built in correlation between night lights and population counts.

Access to electricity and wealth are based on survey data from the Demographic and Health Surveys (individual recode). The DHS data is available only for selected countries. The surveys are also conducted every few years (depending on the country). We construct the access to electricity variable from a dummy variable indicating an individual's response to whether or not they have access to electricity. The household wealth variable is based on the wealth index provided by the DHS, which aggregates different responses regarding wealth-related survey questions.

Figure D1 plots luminosity as well as the three further proxies for economic development on the grid of Africa. Specifically, we plot for each outcome the cell-level average over the sample period. Visually, there appears to be a correlation between luminosity and the other three outcomes. In Table D1, we provide more formal evidence by regressing mean luminosity on the three other proxies (all variables averaged over the sample period). We consistently find positive and statistically significant associations.

Even though luminosity and other proxies for local development appear strongly re-

[32] https://www.imf.org/en/Research/commodity-prices.

lated, there remains the further concern that this relationship might not emerge in dimly lit areas. In particular, the satellites might not detect small changes in light due to the low sensitivity of the sensors. If dimly lit areas are not randomly dispersed throughout the country, this form of measurement error might lead to bias. To explore the sensitivity of the baseline results to this issue, we report replications of the baseline estimates after dropping all cells where the value of mean luminosity is zero at least twice during the sample period *and* positive at least twice during the sample period. These cells oscillate between zero and positive lights, possibly because they are dimly lit and the satellite sensors are not sufficiently sensitive to detect light in all years during the sample period.

The results for the effects within mining regions are collected in Table D2. While the sample of cells is, as discussed, smaller due to the omission of arguably dimly lit cells, we find that the results are similar to the baseline. In Table D3, we collect the corresponding results for non-mining regions. As before, we find that they, too, are similar to the baseline results.

Table D1: Relationship between luminosity and other cell-level proxies for economic activity

	(1: Population)	(2: Electricity)	(3: Wealth)
Mean luminosity	0.240***	0.049***	0.116***
	(0.009)	(0.002)	(0.005)
N	5196	2257	2121

This table collects results from regressing three different proxies for economic activity on mean luminosity. The dependent variable in model (1) is log population in cell i, in model (2) the share of respondents with access to electricit, and in model (3) the average value of the DHS's wealth index. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity-robust standard errors in parentheses.

Table D2: Mineral resources and economic activity in mining regions – drop dimly lit cells

	(1: All mines)	(2: ≥ Major mines)	(3: ≥ Giant)	(4: < Major)	(5: < Moderate)	(6: Prices)
Mine	0.423***	0.528***	0.524*	0.202**	0.063	
	(0.103)	(0.144)	(0.295)	(0.081)	(0.112)	
Mineral price						0.051***
						(0.017)
Cell FE	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Cells	7272	7272	7272	7272	7272	7272
N	159956	159956	159956	159956	159956	159956

This table collects difference-in-differences regressions that relate mineral resource activity (operating mines) to luminosity at the grid-level for all of Africa (0.5 × 0.5 decimal degree cells). In this specification, we study whether mines increase luminosity in mining regions. The dependent variable is the log of mean light output in each cell. The independent variable in models (1)-(5) is a dummy variable that is 1 if a cell had an operating mine of a certain size as indicated in the column header. The independent variable in model 6 is the contemporanous price of the mineral that is extracted in a given cell. The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. The sample with which these results are produced omits cells that have a value for mean light of zero in at least two years *and* positive values for mean light in at least two years during the sample period. It is possible that cells oscillate between no lights and positive lights because they are dimly light and the satellite sensors are not sufficiently sensitive to detect the light they emanate in certain years. With this table, we explore whether the baseline results are robust to this type of measurement error. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

Table D3: Mineral resources and luminosity in non-mining regions – drop dimly lit cells

	(1– Border)	(2– Capitals)	(3– Leaders)
Mines	-0.016***		
	(0.003)		
Capital × Mines		0.067***	
		(0.014)	
Leader × Mines			0.001
			(0.002)
Country FE	-	-	-
Cell FE	Yes	Yes	Yes
Year FE	Yes	-	-
Country-Year FE	No	Yes	Yes
Countries	48	47	47
Cells	3991	7275	7275
N	87802	160048	160048

This table reports a replication of the baseline results after dropping cells that have a value for mean light of zero in at least two years and positive values for mean light in at least two years during the sample period. It is possible that cells oscillate between no lights and positive lights because they are dimly light and the satellite sensors are not sufficiently sensitive to detect the light they emanate in certain years. With this table, we explore whether the baseline results are robust to this type of measurement error. In model (1), the dependent variable is the difference in log mean luminosity in cell i and log mean luminosity in neighboring cells located in other countries. The dependent variable in models (2)-(3) is the log of mean light output in each cell. The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

Figure D1: OUTCOMES AT THE CELL-LEVEL

(a) Luminosity

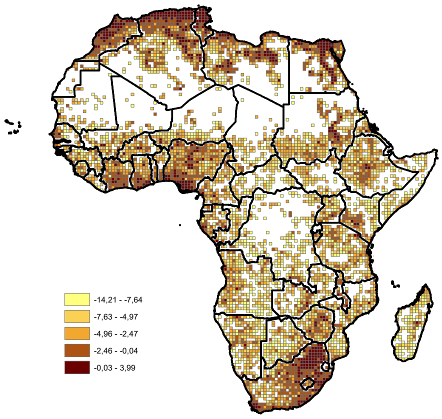

(b) Population

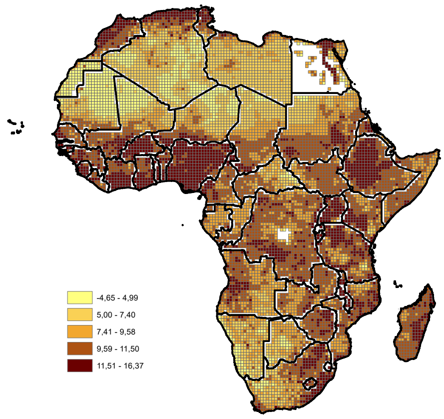

(c) Electricity

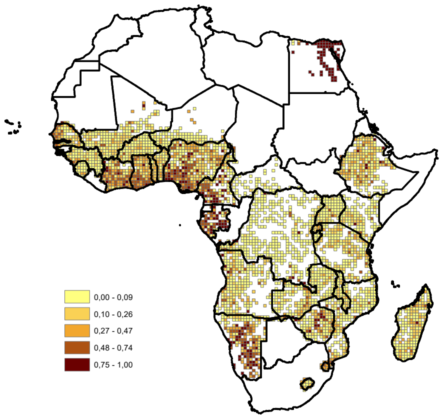

(d) Wealth

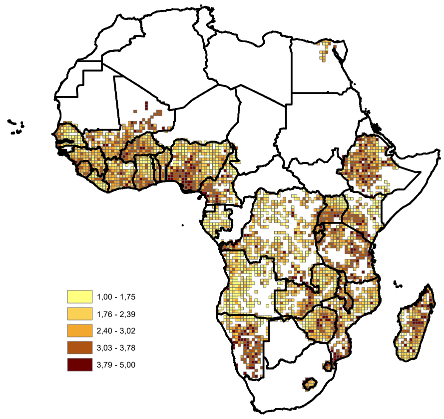

The figure shows four outcomes plotted to the cell-level. Subfigure (a) reports the log of each cell's mean luminosity averaged over the sample period. Subfigure (b) reports the log of each cell's population count averaged over the sample period. Subfigure (c) reports the share of respondents located in each cell with access to electricity averaged over all available waves of the DHS. Subfigure (d) reports the average value for the wealth index reported by the DHS averaged over all available waves of the DHS surveys.

D.4 Additional results for mining regions

D.4.1 Mine discoveries

One concern with using operating mines as the main independent variable as in Equation C1 is that mines might be more likely to be started up in regions that are economically prospering for other reasons. Such regions might receive infrastructure investments (roads, trains, etc.), which might make it easier to extract and transport mineral resources.

To explore whether the focus on mine openings biases the estimated coefficients significantly, we re-estimate the baseline results in model (1) of Table 4.2 after replacing the dummy for operating mines with a dummy for discovered mines. One disadvantage of this specification is, however, that not all discovered mines are ultimately exploited and the time from discovery to exploitation is endogenous to world prices (Khan et al., 2016).

The results are collected in Table D4. We estimate two specifications. First, we use a dummy that is one for a cell for all years after a mine discovery (model 1). Second, we use a dummy that is one after a mine is discovered but changes to zero when the mine is closed (model 2). In both specifications, we estimate a significantly positive coefficient for the mine discovery dummy.

Table D4: MINERAL RESOURCES AND LUMINOSITY AFTER MINE DISCOVERY

	(1)	(2)
Discovered mine	0.243***	
	(0.051)	
Discovered mine (not closed)		0.245***
		(0.051)
Cell FE	Yes	Yes
Country-Year FE	Yes	Yes
Cells	9056	9056
N	199230	199230

This table collects difference-in-differences regressions that relate mineral resource discovery to luminosity at the grid-level for all of Africa (0.5 × 0.5 decimal degree cells). In this specification, we study whether the discovery of mines increases luminosity in mining regions in subsequent years. The dependent variable is the log of mean light output in each cell. The independent variable in model (1) is a dummy that is one after a mine has been discovered until the end of the sample period. The independent variable in model (2) is a dummy that is one after a mine has been discovered until the closure of the first mine located within a particular cell. The sample includes mines with available information on the date of discovery in the MinEx database. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

D.4.2 Event-studies

We estimate event-studies to gauge how the treatment effect varies over time. In this context, we also account for recent research suggesting that canonical difference-in-differences models estimated with two-way fixed effects (TWFE) – as our baseline specification in Equation C1 – produce estimates that cannot necessarily be interpreted as average treatment effects (ATT) if treatments are staggered (occur at different points in time) and treatment effects exhibit group- or event-time specific heterogeneity (Baker et al., 2022; Borusyak et al., 2022; Roth et al., 2022). The TWFE-estimator can be problematic under these conditions as it makes "invalid" comparisons between units in the treatment group that are treated at different points in time. In our context, this implies that cells already treated for several periods can serve as a counterfactual for cells that are treated only recently.

Callaway and Sant'Anna (2021) suggest an approach that circumvents the invalid comparisons made by TWFE. In their approach, separate group- and event-time specific coefficients are estimated by choosing for each group- and event-time specific coefficient unique control groups that avoid "invalid" comparisons, i. e., avoid units that had already been treated in $t-n$ in the control group for units that are treated in some period t. These group- and event-time specific coefficients can then be aggregated to produce either group-specific ATTs (one coefficient for each group over the entire sample period) or event-time specific ATTs (one coefficient for each time period across all groups – i. e., an event-study). The group- and event-time specific ATTs can also be jointly aggregated to produce the standard ATT.

We focus on heterogeneity by event-time e and aggregate group-specific average treatment effects $ATT(g,t)$ at each calendar year t as follows to produce an event-study:

$$\theta_D(e) = \sum_{g\in G}^{T} \mathbf{1}\{g+e \leq T\} P(G=g|G+e \leq T) ATT(g, g+e). \tag{C6}$$

Adjusting the notation of Callaway and Sant'Anna (2021) slightly, $t = 1, ..., T$ denotes in this model a calendar year with T the total number of years in the sample. G is the year when a treated cell (i. e., a cell with at least one mine opening) becomes treated first. Individual treated cells are grouped into different treatment groups g depending on when they are treated first, i. e., their value of $g \in \{G_{min}, ..., G_{max}\}$. Then, $e = t - g$ captures the years passed since treatment occurred first for group g. $P(G = g|G + e \leq T)$ denotes the size of group g among all groups treated in event-year e and serves as weights in the aggregation.[33]

$\theta_D(e)$ is the aggregated parameter of interest for the event-study. Following the notation introduced above, this parameter captures the (weighted) average effect of a cell having a mine in operation e years after the treatment was adopted across all cells that

[33] See footnote 9 in Callaway and Sant'Anna (2021).

are ever observed to have a mine in operation for exactly e years.

Figure D2 plots the event-study coefficients from the aggregation described in Equation C6. The plots suggest a noticeable increase in luminosity after the opening of a mine. Luminosity peaks after around 5 years, but continues to be higher than in the pre-treatment period thereafter. These results are consistent with our baseline estimates. This is presumably because most cells are never treated in our setting (have no mine openings), which implies the potentially "invalid" comparisons between cells treated at different points in time likely influence the overall ATT only marginally.

We also observe a small increase in luminosity already in t-2, suggesting some degree of treatment anticipation. This is presumably due to using the official opening date of a mine to define the onset of treatment. Even before the official opening of a mine, there is likely activity ongoing due to construction and in-migration of workers. An slight uptick in luminosity a few years before the official opening is thus expected.

To explore this further, we report in Figure D3 results from an event study where we use the discovery date of a deposit rather than the opening of a mine to define the onset of a treatment. We observe similar treatment effects as in Figure D3 but no pre-trends.

Figure D2: Event-study on mine openings and luminosity

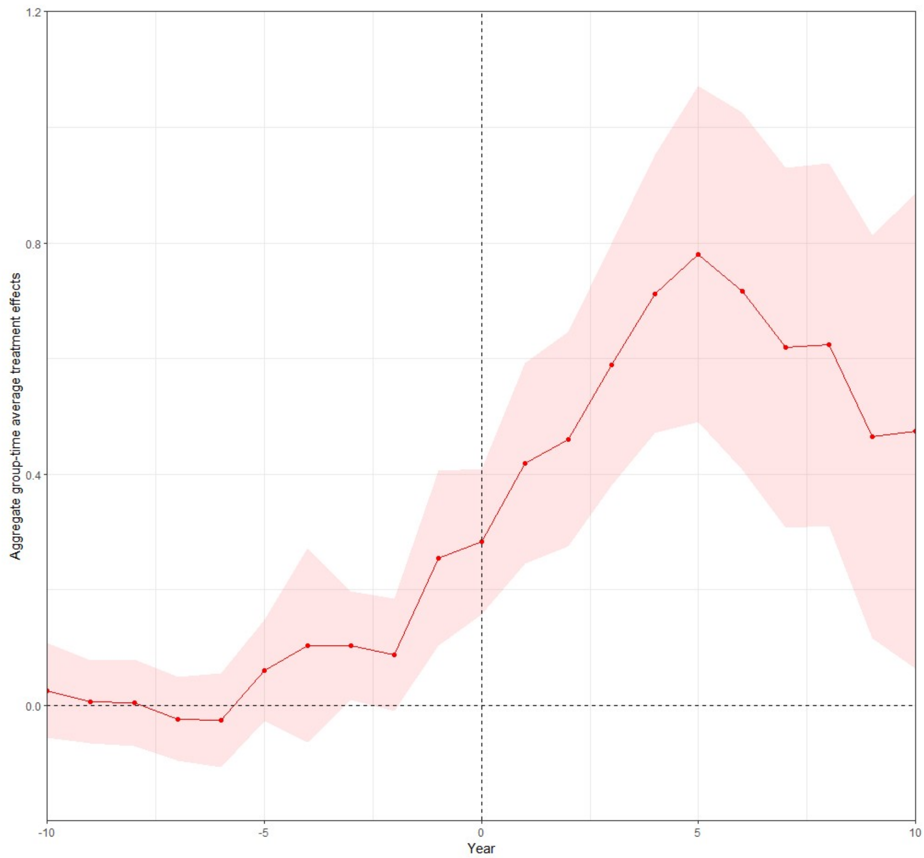

The figure shows an event-study based on the Callaway and Sant'Anna difference-in-differences with multiple time periods estimator, relating mine openings to luminosity at the grid-level.

Figure D3: Event-study on mine discoveries and luminosity

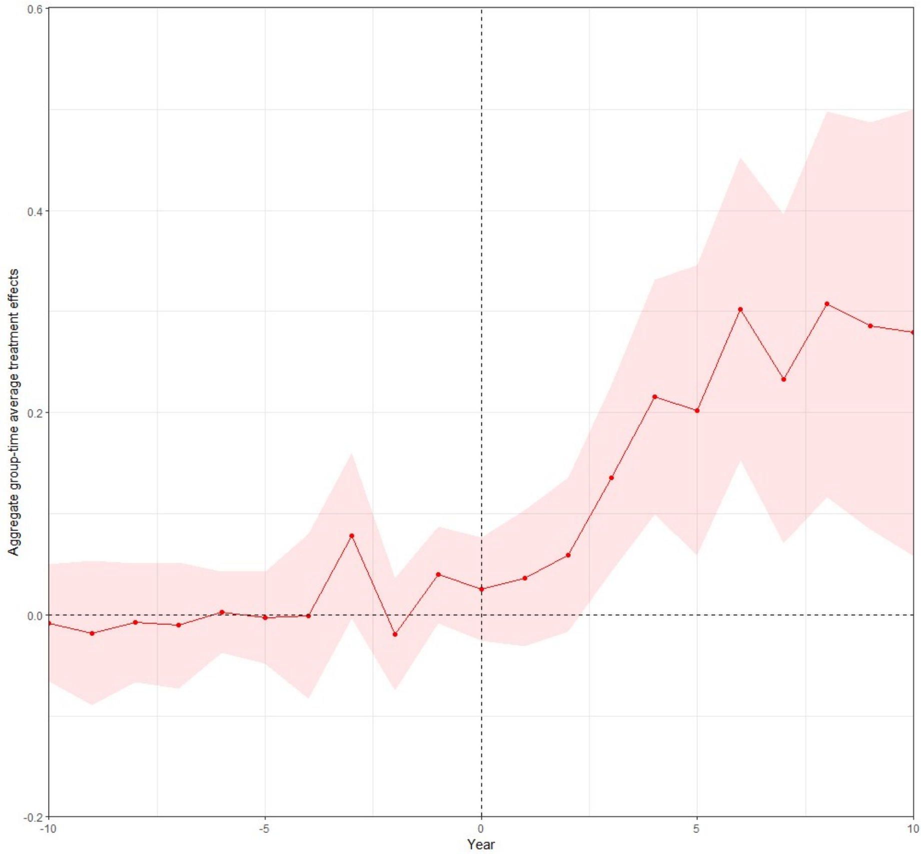

The figure shows an event-study based on the Callaway and Sant'Anna difference-in-differences with multiple time periods estimator, relating mine discoveries to luminosity at the grid-level.

D.4.3 Mines in border regions

To refine the identification strategy of Equation C1, we limit the sample to cells located close to the border. As in the regressions for generic non-mining regions (Section 4.5.1), focusing on border regions is one way to account for country-level trends that might have heterogeneous effects across regions within the same country. More specifically, cells on opposite sides of the border are likely to experience similar trends once country-specific year effects have been partialled out.

We first classify each cell according to whether it is adjacent to its country's border. We use a distance of 250km as the threshold; see Figure D7 for a map that indicates these cells. Next, we calculate for each border cell in a given country the average luminosity of all cells within a 250km distance that are located in a different country. This average luminosity across neighboring foreign cells constitutes the counterfactual for each domestic index cell. Figure D8 shows how the log of mean luminosity in index cells and their neighbors evolve during the sample period. The trends are remarkably similar, suggesting that cells in neighboring countries can provide a reasonable counterfactual for the border cells in a given country.

Table D5: Mineral resources and economic activity in mining regions – sample restricted to border regions

	(1: All mines)	(2: ≥ Major mines)	(3: ≥ Giant)	(4: < Major)	(5: < Moderate)	(6: Prices)
Mine	0.837***	0.975***	1.133***	0.490***	0.289**	
	(0.128)	(0.166)	(0.278)	(0.159)	(0.131)	
Mineral price						0.098***
						(0.018)
Cell FE	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Cells	6108	6108	6108	6108	6108	6108
N	134350	134350	134350	134350	134350	134350

This table collects difference-in-differences regressions that relate mineral resource activity (operating mines) to luminosity at the grid-level for all of Africa (0.5 × 0.5 decimal degree cells). In this specification, we study whether mines increase luminosity in mining regions while limiting the sample to border cells. The dependent variable is the log of mean light output in each cell. The independent variable in models (1)-(5) is a dummy variable that is 1 if a cell had an operating mine of a certain size as indicated in the column header. The independent variable in model 6 is the contemporanous price of the mineral that is exploited in a given cell. The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

We collect the results in Table D5. Note that the number of cells is substantially smaller in this sample. Nevertheless, the results are similar to the baseline estimates. In particular, the estimate for an increase in the price of a mineral is highly significant (model 6).

D.4.4 Types of minerals

In this section, we study whether treatment effects vary according to the type of mineral that is exploited in a mine. It is possible that mineral resources differ in how much local economic activity they induce, either because they generate different amounts of revenue or because they require a different mix of capital and labor for exploitation. Also, previous research indicates that especially “point” resources such as gemstones and gold, i. e., resources that are concentrated at one point, often induce violent conflicts due to a rapacity effect (Isham et al., 2005; Bulte et al., 2005). “Point” resources have also been found to generate only relatively small local employment effects (Gollin et al., 2016).

We focus in the following on gold, diamonds, copper and all remaining minerals. Gold, diamonds, and copper are the three most common minerals in the MinEx Consulting database (see Table 4.1).[34]

We estimate Equation C1 by replacing the generic mine dummy with dummies for gold, diamonds, copper, and all other types of mines, respectively. The results are collected in Table D6. We find that while the magnitudes of the coefficients vary slightly, all types of mines have positive local economic effects. Luminosity is significantly higher when a cell has at least one operating mine, irrespective of the type of mineral exploited in that mine.

[34] The remaining minerals are: Andalusite, Asbestos, Barium, Chromium, Cobalt, Flourine, Fluorite, Lead, Manganese, Mercury, Mineral Sands, Nickel, PGE, Platinum, Rare Earths, Rubies, Sapphires, Silver, Sulphur, Tantalum, Tin, Tungsten, Uranium, Vermiculite, and Zinc.

Table D6: Mineral resources and economic activity in mining regions – heterogenous effects by type of mineral

	(1: All mines)	(2: ≥ Major mines)	(3: ≥ Giant)	(4: < Major)	(5: < Moderate)	(6: Prices)
Panel A: Gold						
Mine	0.821***	1.005***	1.164***	0.282	0.387**	
	(0.151)	(0.182)	(0.323)	(0.187)	(0.186)	
Mineral price						0.123***
						(0.022)
Panel B: Diamonds						
Mine	0.483	0.526	1.191*	0.223**	0.117***	
	(0.317)	(0.423)	(0.618)	(0.090)	(0.036)	
Mineral price						0.035
						(0.022)
Panel C: Copper						
Mine	1.192***	1.273***	1.211	0.863	no mines of this size	
	(0.401)	(0.474)	(0.794)	(0.621)		
Mineral price						0.151***
						(0.045)
Panel D: Other						
Mine	0.885***	1.373***	1.876***	0.475***	0.028	no common price
	(0.219)	(0.387)	(0.701)	(0.175)	(0.114)	

This table reports a replication of the baseline results for selected minerals. Specifically, we consider country-level variation in gold mines (Panel A), diamond mines (Panel B), copper mines (Panel C), and all other mines (Panel D). Other mines include the following minerals: Andalusite, Asbestos, Barium, Chromium, Cobalt, Flourine, Fluorite, Lead, Manganese, Mercury, Mineral Sands, Nickel, PGE, Platinum, Rare Earths, Rubies, Sapphires, Silver, Sulphur, Tantalum, Tin, Tungsten, Uranium, Vermiculite, and Zinc. The dependent variable is the log of mean light output in each cell. The independent variable in models (1)-(5) is a dummy variable that is 1 if a cell had an operating mine of a given type. The independent variable in model (6) is the contemporanous log price of the mineral in question. Note that there are no copper mines of size smaller than moderate (Panel C, model 5). There is also no price regression for "other mines" as there is no common price (Panel D, model 6). The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

D.4.5 Extractive Industries Transparency Initiative and mining regions

One important recent initiative for the mineral resource sector is the Extractive Industries Transparency Initiative (EITI). The purpose of the EITI as per its mission statement is to promote "the understanding of natural resource management", to "strengthen public and corporate governance and accountability", and to "provide the data to inform policy-making and multi-stakeholder dialogue".[35] As such, the EITI might have enabled mining regions to benefit more from the extraction of mineral resources.

The EITI was established in 2003, but countries have joined the initiative at different points in time. Using this variation, we explore whether the local economic effects of mineral resources varies before and after a country has joined this initiative. For this, we collect information on the dates at which different countries have joined the EITI (see Table D7) and estimate an extension of model C1 where we interact the mineral resource dummy with a dummy for whether a country has joined the EITI in a given year.

We collect the results in Table D8. We find a positive effect of the EITI on local

[35] See `https://eiti.org/our-mission`.

economic development in most specifications. It appears that mining regions did indeed benefit from the EITI. However, note that these results do not establish a causal effect as the adoption date of the EITI was not random.

Table D7: Year of EITI membership by country

Country	**Year of membership**
Burkina Faso	2009
Cameroon	2007
Central African Republic	2008
Chad	2010
Cote D'Ivoire	2008
Democratic Republic of the Congo	2007
Ethiopia	2014
Gabon	2021
Ghana	2007
Guinea	2007
Liberia	2008
Madagascar	2008
Malawi	2015
Mali	2007
Mauritania	2007
Mozambique	2009
Niger	2020
Nigeria	2007
Republic of the Congo	2007
Senegal	2013
Seychelles	2014
Sierra Leone	2008
Sao Tome and Principe	2008
Tanzania	2009
Togo	2010
Uganda	2020
Zambia	2009

This table shows the year countries joined the Extractive Industries Transparency Initiative. Note that countries which joined after 2013 are treated as non-EITI countries in the regressions as our sample stops in 2013.

D.5 Additional results for non-mining regions

D.5.1 Major mines and luminosity in non-mining regions

In the baseline regressions, we neglected the specific size of a mine when assessing its implications on non-mining regions. However, it is plausible that larger mines generate more revenues and are thus more consequential for luminosity in non-mining regions (e. g., due to more resources available for inter-regional redistribution). We thus re-estimate the baseline models with the number of mines that are classified as at least of "major" size by MinEx Consulting.

The results are collected in Table D9 and are virtually identical to the baseline es-

Table D8: Mineral resources and economic activity in mining regions – the effect of the EITI

	(1: All mines)	(2: $\geq$ Major mines)	(3: $\geq$ Giant)	(4: < Major)	(5: < Moderate)	(6: Prices)
Mine	0.732***	0.871***	1.096***	0.422***	0.176*	
	(0.108)	(0.148)	(0.267)	(0.127)	(0.106)	
Mineral price						0.089***
						(0.017)
After EITI	0.372***	0.610***	0.660***	-0.035	0.008	0.048***
	(0.128)	(0.172)	(0.224)	(0.160)	(0.211)	(0.018)
Cell FE	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Cells	9053	9053	9053	9053	9053	9053
N	199138	199138	199138	199138	199138	199138

This table collects difference-in-differences regressions that relate mineral resource activity (operating mines) to luminosity at the grid-level for all of Africa (0.5 × 0.5 decimal degree cells). In this specification, we study whether mines increase luminosity in mining regions more or less after the introduction of Extractive Industries Transparency Initiative (EITI). The dependent variable is the log of mean light output in each cell. The independent variables in models (1)-(5) is a dummy variable that is 1 if a cell had an operating mine of a certain size as indicated in the column header and an interaction variable that is 1 once a country adopts the EITI. The independent variable in model 6 is the contemporanous price of the mineral that is extracted in a given cell and an interaction variable that is one once a country adopts the EITI. The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

timates. Additional major mines increase luminosity in capital regions and decrease luminosity in border regions. In fact, the estimates are slightly larger than in the baseline specifications with all mines.

D.5.2 Types of minerals and luminosity

Governments might be more likely to regionally redistribute the proceeds from some minerals than others. For example, gold or diamonds might come with greater (international) scrutiny compared to less salient minerals. It might be more difficult for governments to siphon off the proceeds from such minerals, reducing the opportunity costs of using them for redistribution towards non-resource regions. On the other hand, if "point" resources indeed induce violent conflicts, they might in turn lead to a deterioration of the macroeconomic environment and thus negatively affect non-resource regions. Differences in extraction technology might also necessitate joint ventures with foreign companies. For such minerals, the government might have fewer pure rents available to redistribute across regions. Minerals might also vary in how easily their extraction can be taxed (de la Sierra, 2020). Finally, different countries might specialize in the production of different minerals, and these countries might vary in how they redistribute mineral resources spatially for idiosyncratic reasons.

Due to such considerations, we hypothesize that different types of mines will vary in their economic implications on non-mining regions. We thus study heterogeneity in the effect of different minerals by focusing, as in Section D.4.4, on gold, diamonds, copper and all remaining minerals.

We estimate the preferred specifications as outlined in Section 4.5 for border regions,

Table D9: MINERAL RESOURCES AND LUMINOSITY IN NON-MINING REGIONS – MINES CLASSIFED AT LEAST AS MAJOR

	(1– Border)	(2– Capitals)	(3– Leaders)
Major mines	-0.020***		
	(0.003)		
Capital × Major Mines		0.055***	
		(0.019)	
Leader × Major Mines			0.002
			(0.002)
Country FE	-	-	-
Cell FE	Yes	Yes	Yes
Year FE	Yes	-	-
Country-Year FE	No	Yes	Yes
Countries	48	48	48
Cells	5041	9056	9056
N	110902	199230	199230

This table reports a replication of the results in columns (1b), (2c), and (3c) of Table 4.4 with only mines classifed as at least major by MinEx Consulting. In column (1), the dependent variable is the difference in log mean luminosity in cell i and log mean luminosity in neighboring cells located in other countries. The dependent variable in column (2)-(3) is the log of mean light output in each cell. The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

capital regions, and leaders' birth regions. The results are collected in Table D10. They suggest some, but rather limited heterogeneity.

Gold (Panel A) and diamond (Panel B) mines induce in an increase in luminosity in the capital region. Copper mines display a positive coefficient (Panel C), even though it is relatively small and insignificant. The remaining mines, too, result in an increase in luminosity in capital regions (Panel D). For leader regions, we observe again as in the baseline specification, no significant effects for any of the four sets of minerals. That is, there is no mineral-specific heterogeneity. For border regions, we find that it is mainly diamond and copper mines that reduce luminosity in generic regions. Gold mines have insignificant effects. Other types of mines even appear to increase luminosity in generic regions.

Different explanations can be put forward for the result that only copper and diamond mines reduce luminosity in generic non-mining regions. Countries that opened up new copper mines during the sample period could be larger or less prosperous, leaving less room for the government to shift resource rents broadly to generic non-mining regions. Indeed, the country that contributes a substantial share of the variation for copper mines regressions is the Democratic Republic of the Congo (DRC). There is widespread concern about the implications of copper mining in the DRC on corruption and good governance (Global Witness, 2006). With respect to diamond mines, it is known that their extraction has an intricate relationship with violent conflicts (Campbell, 2002; Guidolin and La Ferrara, 2007). Through this channel, diamond mines could reduce luminosity in non-mining regions.

Table D10: Mineral resources and luminosity in non-mining regions – types of minerals

	Panel A: Gold			Panel B: Diamonds		
	(1– Border)	(2– Capital)	(3– Leader)	(4– Border)	(5– Capital)	(6– Leader)
Mines	-0.001			-0.037***		
	(0.003)			(0.008)		
Capital × Mines		0.043***			0.106	
		(0.015)			(0.067)	
Leader × Mines			0.005			0.007
			(0.004)			(0.009)
Countries	48	48	48	48	48	48
Cells	5041	9056	9056	5041	9056	9056
N	110902	199230	199230	110902	199230	199230
	Panel C: Copper			**Panel D: Other**		
	(7– Border)	(8– Capital)	(9– Leader)	(10– Border)	(11– Capital)	(12– Leader)
Mines	-0.053***			0.069***		
	(0.006)			(0.009)		
Capital × Mines		0.022			0.061***	
		(0.023)			(0.011)	
Leader × Mines			0.008			0.001
			(0.011)			(0.004)
Countries	48	48	48	48	48	48
Cells	5041	9056	9056	5041	9056	9056
N	110902	199230	199230	110902	199230	199230

This table reports a replication of the results in columns (1b), (2c), and (3c) of Table 4.4 for selected minerals. Specifically, we consider country-level variation in gold mines (Panel A), diamond mines (Panel B), copper mines (Panel C), and all other mines (Panel D). All other mines include the following minerals: Andalusite, Asbestos, Barium, Chromium, Cobalt, Flourine, Fluorite, Lead, Manganese, Mercury, Mineral Sands, Nickel, PGE, Platinum, Rare Earths, Rubies, Sapphires, Silver, Sulphur, Tantalum, Tin, Tungsten, Uranium, Vermiculite, and Zinc. The dependent variable in columns (1), (4), (7), and (10) is the difference in log mean luminosity in cell i and log mean luminosity in neighboring cells located in other countries. In columns (2)-(3), (5)-(6), (8)-(9), and (11)-(12), the dependent variable is the log of mean light output in each cell. We study whether mines lead to an increase in luminosity in capital regions, in leaders' birth regions, and in generic regions. The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

D.5.3 Extractive Industries Transparency Initiative and non-mining regions

Following up on our analysis of the EITI's effects on mining regions, we explore how the EITI has affected non-mining regions. For this, we extend the specifications detailed in Equations C2, C3, and C4 by an interaction with a dummy for when a country has introduced the EITI (if at all).

We collect the results in Table D11 We find that the introduction of the EITI correlates with an increase in luminosity in capital and leaders' birth regions. As such, it does seem to not only have benefited the mining regions themselves, but also these specific non-mining regions. On the other hand, we find a significantly negative effect of the EITI for generic regions. This negative effect might indicate that the EITI has diminished the ability or the willingness of governments to redistribute resources towards other parts of the countries that are not of special importance for the country's leaders. As before, however, these results should not be interpreted as casual as the adoption of the EITI by a given country could be endogenous.

Table D11: Mineral resources and luminosity in non-mining regions – the effect of the EITI

	(1– Border)	(2– Capitals)	(3– Leaders)
Mines	-0.004		
	(0.004)		
Mines × EITI	-0.009***		
	(0.003)		
Capital × Mines		0.037***	
		(0.014)	
Capital × Mines × EITI		0.020**	
		(0.010)	
Leader × Mines			0.001
			(0.002)
Leader × Mines × EITI			0.017***
			(0.006)
Country FE	-	-	-
Cell FE	Yes	Yes	Yes
Year FE	Yes	-	-
Country-Year FE	No	Yes	Yes
Countries	48	48	48
Cells	5041	9056	9056
N	110902	199230	199230

This table reports an extension of the results in columns (1b), (2c), and (3c) of Table 4.4 where we explore heterogenous effects according to whether and when countries have adopted the Extractive Industries Transparency Initiative (EITI). We interact the mining variables with dummies for whether a country in a given year has adopted the EITI. In column (1), the dependent variable is the difference in log mean luminosity in cell i and log mean luminosity in neighboring cells located in other countries. The dependent variable in column (2)-(3) is the log of mean light output in each cell. The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

D.6 Sectoral specialization of regions

Regions in Africa specialize in different sectors. Using the *INDGEN* variable form the IPUMS, we calculate for each second-tier region (GEOLEV2) the share of respondents employed in four key sectors[36]: (i) mining and extraction, (ii) manufacturing, (iii) agriculture, and (iv) hotels and restaurants. The second regional tier is the lowest geographical identifier for respondents consistently indicated in the IPUMS. In particular, we are unable to link respondents to the standard grid used above.

As the census waves are available only every few years, we take the average over the sample period of the region-specific shares calculated from each wave. Figure D4 plots the sample-averaged shares for the countries with available data in the IPUMS at the regional-level.

[36]See `https://international.ipums.org/international-action/variables/INDGEN#codes_section` for details.

Figure D4: Regional industry specialization

(a) Mining

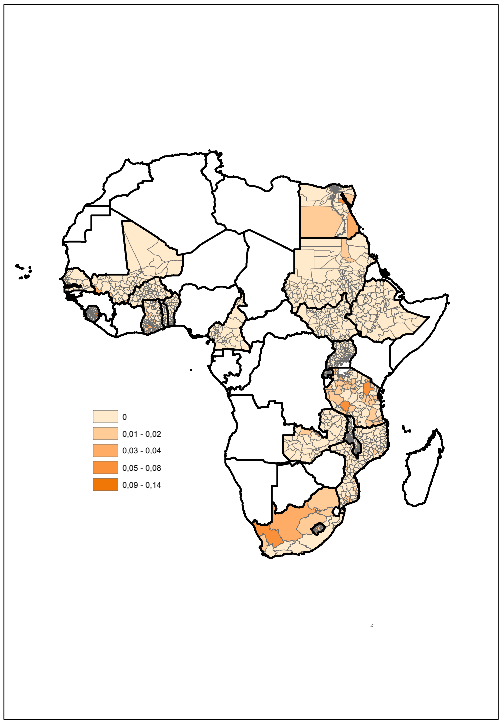

(b) Manufacturing

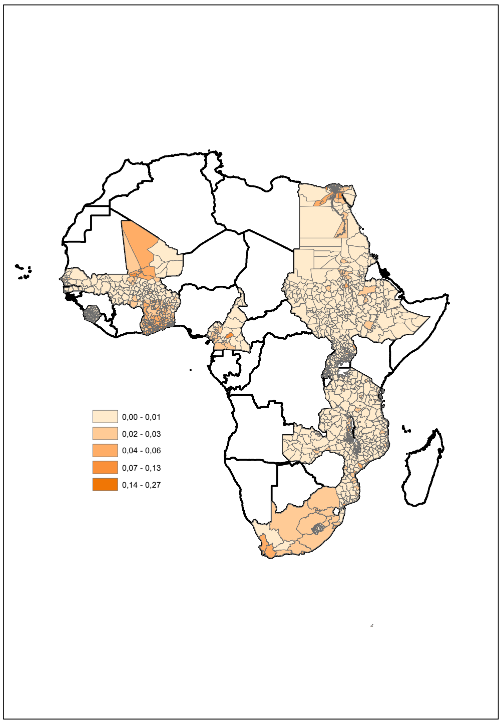

(c) Agriculture

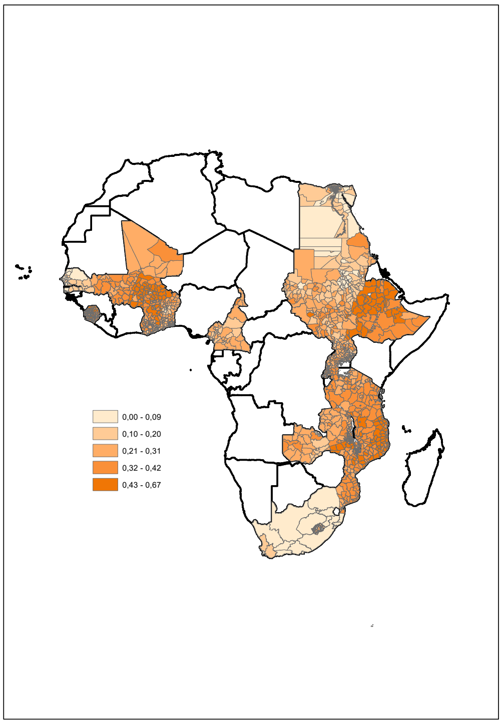

(d) Hospitality

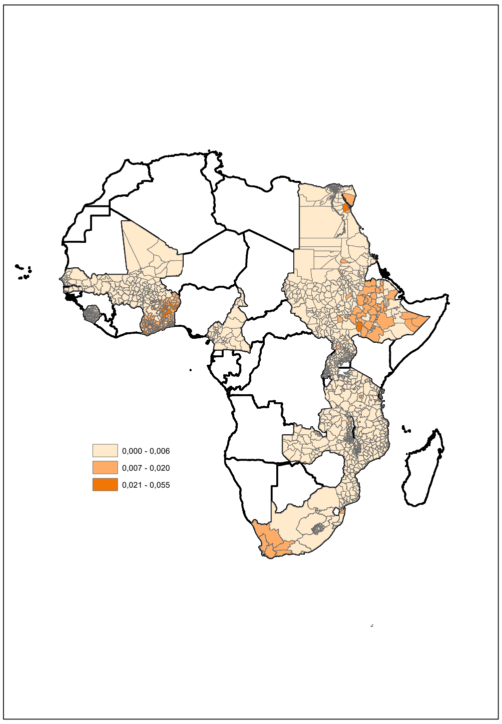

The figure shows four outcomes plotted to the level of countries' second regional tier. Subfigure (a) shows the share of census respondents employed in the mining sector. Subfigure (b) shows the share of census respondents employed in manufacturing. Subfigure (c) shows the share of census respondents employed in agriculture. Subfigure (d) shows the share of census respondents employed in the hospitality sector.

D.7 Additional figures and tables

Figure D5: Grid of Africa with capital city cells

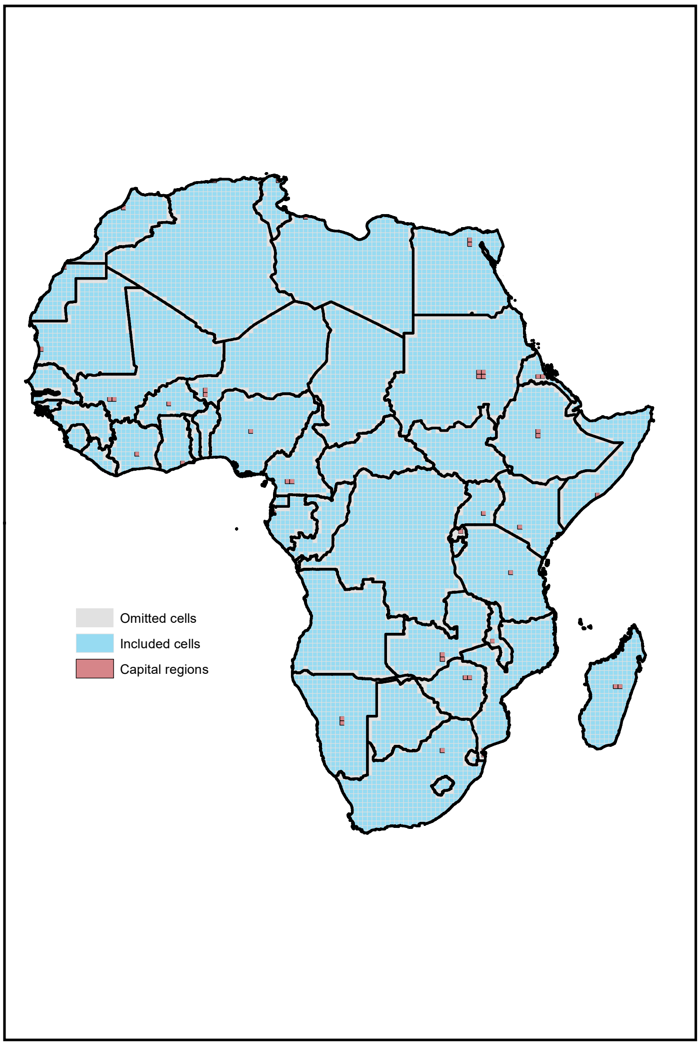

The figure shows a 0.5×0.5 degree grid over Africa. It indicates all cells that cross two or more countries, which are dropped from the analysis (gray), and all cells that are within a 10km buffer around the capital city of each country (red). Note that capital cities that are close to a country border are dropped from the sample and therefore not included (e. g., Kinshasa (DR Congo) or Bangui (Central African Republic)). Some capitals which are included in the sample might not be easily visible in this map due to, e. g., being located at the coast and thus the corresponding cell being cut-off and / or concealed by the solid line indicating country borders (e. g., Algiers, capital of Algeria).

Figure D6: Grid of Africa with leader region cells

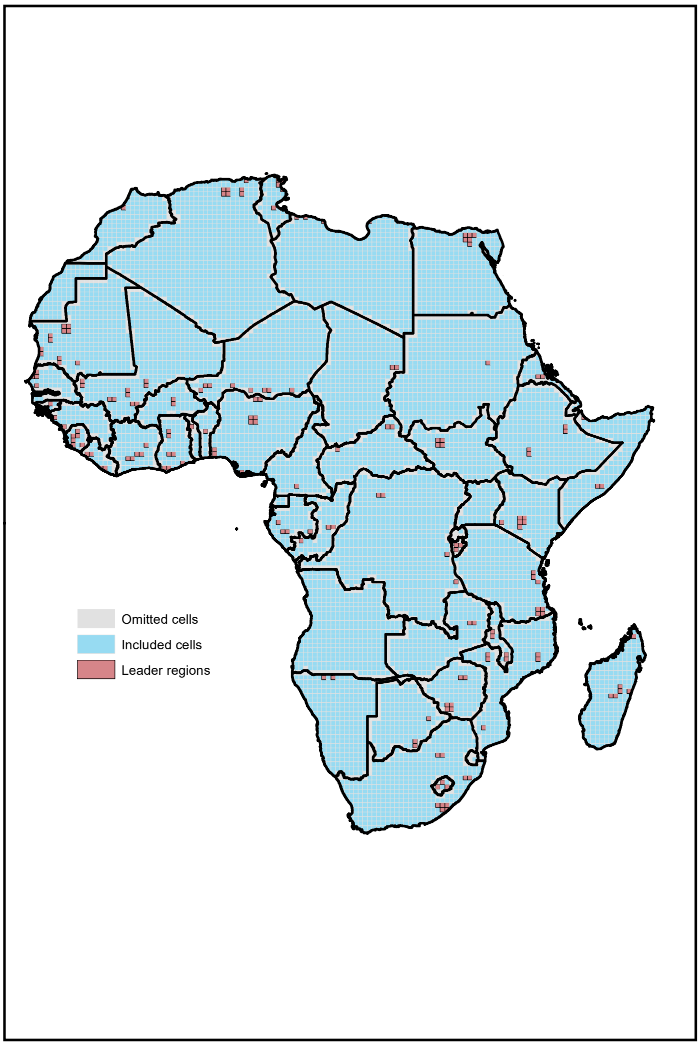

The figure displays a grid over Africa. It indicates all cells that cross two or more countries, which are dropped from the analysis (gray), and all cells that are within a 10km buffer around the coordinates of the birthplace of a country's leader (red). Note that birthplaces that are close to a country border are dropped from the sample and therefore not included.

Figure D7: Grid of Africa with cells included in the border regressions

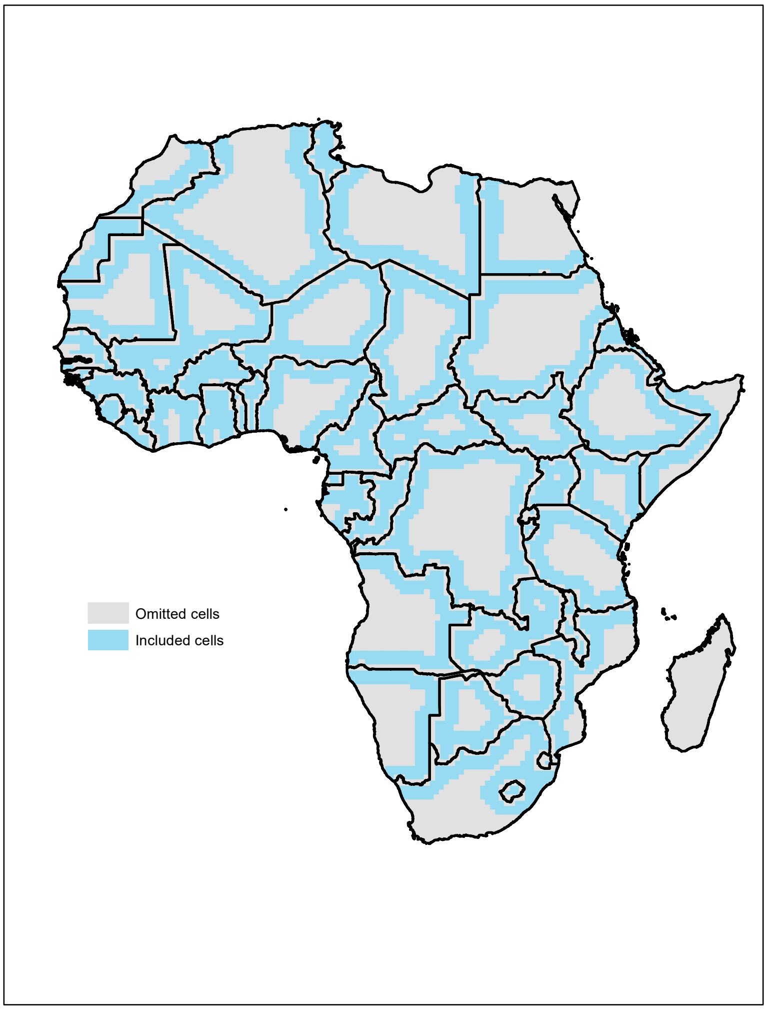

The figure shows a grid over Africa. It indicates all cells that are within 250km of any given country's border in red. These cells are included in the border cell regressions reported in models (1a)-(1b) of Table 4.4.

Figure D8: Evolution of mean luminosity in border cells and their neighboring cells during the sample period

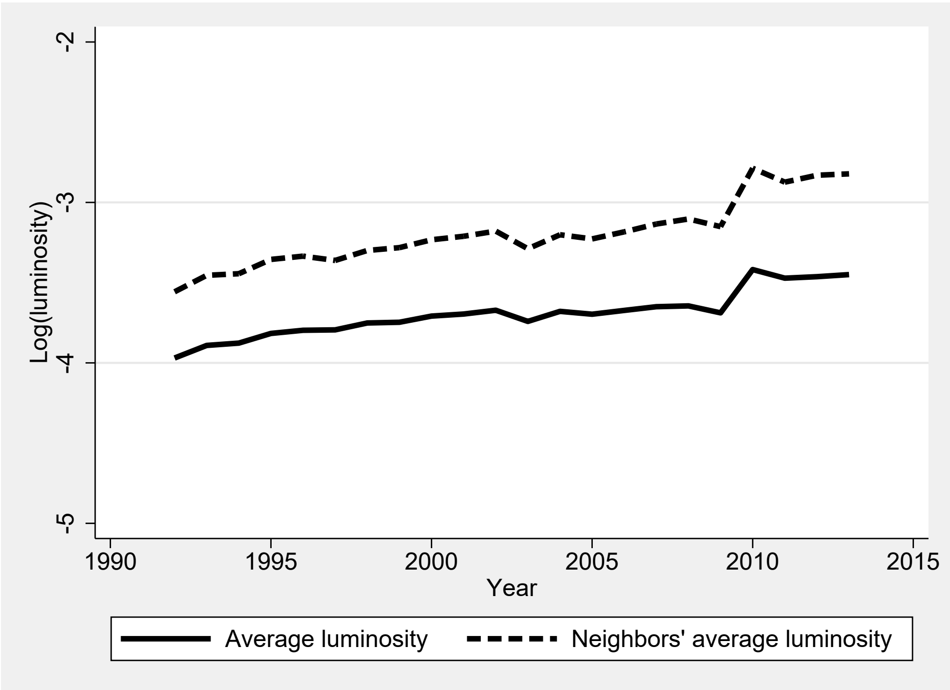

The figure shows the evolution of the average mean luminosity of all cells close to their country's border and the average mean luminosity of their neighbors. Specifically, for the solid line, we take the average of the luminosity values of each border cell. For the dashed line, we calculate for each border cell the average value of luminosity across all cells in neighboring countries within a distance of 250km and then calculate the average value of neighboring cells' luminosity for all border cells. The cells included in these calculations are colored blue in Figure D7.

Table D12: Definition and sources of main variables

Variable	Definition	Source
Log(mean luminosity)	Log of mean luminosity in cell i and year t (average of digital values of all pixels in the original DMSP-OLS data that fall within a 0.5×0.5 degree cell)	US Air Force (USAF) Defense Meteorological Satellite Program Operational Linesman System (DMSP-OLS)
Sum of conflict events	Sum of conflict events in a given cell i in year t	Armed Conflict Location and Event Data Project (ACLED)
Mine	Dummy = 1 for whether cell i has a mine in year t	MinEx Consulting
Major mine	Dummy = 1 for whehter cell i has a mine classifed as “major” by MinEx Consulting in year t	MinEx Consulting
Giant mine	Dummy = 1 for whehter cell i has a mine classifed as “giant” by MinEx Consulting in year t	MinEx Consulting
Mines	Total number of mines in country c in year t	MinEx Consulting
Discovered mine	Dummy = 1 after a mine has been discovered in cell i	MinEx Consulting
Discovered mine (not closed)	Dummy = 1 after a mine has been discovered in cell i until year $t + T$ where it has been closed	MinEx Consulting
Log(mineral resource revenues)	Log of ores and metal exports (in current US Dollars)	World Development Indicators, World Bank
Capital	Dummy = 1 for cell i if it the capital of the country is located in its area	CEPII's GeoDist database
Leader	Dummy = 1 if the birth place of the national leader in year t is located within a cell i	Archigos database
More democratic	Dummy = 1 if the country in which a cell is located is classifed as fully democratic in the Freedom House index	Freedom House, Dahlberg et al. (2020)
Less democratic	Dummy = 1 if the country in which a cell is located is not classifed as fully democratic in the Freedom House index	Freedom House, Dahlberg et al. (2020)
More corrupt	Dummy = 1 in cell i and year t if the country in which a cell is located has a value for the “control of corruption” index from the Worldwide Governance Indicators that is above the median value of the index in a given year	Worldwide Governance Indicators, Dahlberg et al. (2020)
Less corrupt	Dummy = 1 in cell i and year t if the country in which a cell is located has a value for the “control of corruption” index from the Worldwide Governance Indicators that is below the median value of the index in a given year	Worldwide Governance Indicators, Dahlberg et al. (2020)
Efficient government	Dummy = 1 in cell i and year t if the country in which a cell is located has a value for the “government effectiveness” index from the Worldwide Governance Indicators that is above the median value of the index in a given year	Worldwide Governance Indicators, Dahlberg et al. (2020)
Inefficient government	Dummy = 1 in cell i and year t if the country in which a cell is located has a value for the “government effectiveness” index from the Worldwide Governance Indicators that is below the median value of the index in a given year	Worldwide Governance Indicators, Dahlberg et al. (2020)
EITI membership	Dummy = 1 if the country in which a cell is located has joined the Extractive Industries Transparency Initiative	Own research
Regional sectoral specialization	Employment shares in the mining, manufacturing, agricultural, and hospitality industries across different regions	Own calculations based on IPUMS census data

This table lists the sources and definitions for the key variables used in the paper.

Table D13: Summary statistics on key variables

Variable		Mean	Std. Dev.	Min.	Max.	Obs.
Log(mean luminosity)	overall	-3.583	1.816	-4.605	4.123	199274
	between	.	1.767	-4.605	3.990	9058
	within	.	0.421	-8.751	1.101	22
Sum of conflict events	overall	0.108	2.743	0.000	660.000	199496
	between	.	1.461	0.000	128.091	9068
	within	.	2.321	-127.982	532.018	22
Mine	overall	0.011	0.102	0.000	1.000	199496
	between	.	0.093	0.000	1.000	9068
	within	.	0.044	-0.944	0.965	22
Major mine	overall	0.008	0.089	0.000	1.000	199496
	between	.	0.081	0.000	1.000	9068
	within	.	0.036	-0.856	0.962	22
Giant mine	overall	0.005	0.069	0.000	1.000	199496
	between	.	0.065	0.000	1.000	9068
	within	.	0.023	-0.950	0.959	22
Mines (Total)	overall	5.536	13.934	0.000	73.000	199496
	between	.	13.852	0.000	65.864	9068
	within	.	1.518	-1.373	13.809	22
Discovered mine	overall	0.010	0.099	0.000	1.000	199496
	between	.	0.080	0.000	1.000	9068
	within	.	0.059	-0.945	0.964	22
Discovered mine (not closed)	overall	0.010	0.097	0.000	1.000	199496
	between	.	0.077	0.000	1.000	9068
	within	.	0.059	-0.945	0.964	22
Capital	overall	0.016	0.124	0.000	1.000	199496
	between	.	0.124	0.000	1.000	9068
	within	.	0.000	0.016	0.016	22
Leader	overall	0.017	0.130	0.000	1.000	199496
	between	.	0.103	0.000	1.000	9068
	within	.	0.080	-0.937	0.972	22
More democratic	overall	0.150	0.357	0.000	1.000	182683
	between	.	0.326	0.000	1.000	8970
	within	.	0.119	-0.759	1.104	20
Less democratic	overall	0.850	0.357	0.000	1.000	182683
	between	.	0.326	0.000	1.000	8970
	within	.	0.119	-0.104	1.759	20
More corrupt	overall	0.508	0.500	0.000	1.000	125309
	between	.	0.423	0.000	1.000	8970
	within	.	0.272	-0.425	1.441	14
Less corrupt	overall	0.492	0.500	0.000	1.000	125309
	between	.	0.423	0.000	1.000	8970
	within	.	0.272	-0.441	1.425	14
Efficient government	overall	0.512	0.500	0.000	1.000	125309
	between	.	0.413	0.000	1.000	8970
	within	.	0.291	-0.422	1.445	14
Inefficient government	overall	0.488	0.500	0.000	1.000	125309
	between	.	0.413	0.000	1.000	8970
	within	.	0.291	-0.445	1.422	14
EITI membership	overall	0.111	0.314	0.000	1.000	199496
	between	.	0.139	0.000	0.318	9068
	within	.	0.282	-0.207	1.066	22

The within number of observations is the average number of observations per cell.

Table D14: MINERAL PRICES: DEFINITION AND SOURCES

Variable	Definition	Source
Andalusite	–	–
Asbestos	US unit value $/t	USGS
Chromium	–	–
Cobalt	$/mt, nominal	IMF
Copper	$/mt, nominal	World Bank
Diamonds	US unit value $/t	USGS
Flourine	–	–
Gold	$/troy oz, nominal	World Bank
Lead	$/mt, nominal	World Bank
Manganese	US unit value $/t	USGS
Mineral Sands	US unit value $/t	USGS
Nickel	$/mt, nominal	World Bank
PGE	–	–
Platinum	$/troy oz	World Bank
Ruby	–	–
Sapphire	–	–
Tin	$/troy oz	World Bank
Tungsten	US unit value $/t	USGS
Uranium	–	–
Zinc	US Unit value $/t	USGS

This table lists the sources and definitions for the data on prices for the mineral resources. Minerals for which we were unable to obtain price data are indicated by "–". The price data from the USGS refer to US unit values rather than world prices.

Table D15: SUMMARY STATISTICS ON MINERAL RESOURCE PRICES

Variable	**Mean**	**Std. Dev.**	**Min.**	**Max.**	**Obs.**
Asbestos	561.3	(426.9)	172	1570	22
Cobalt	28269.6	(14494.7)	6641.1	70839.3	22
Copper	4005.2	(2594)	1559.5	8828.2	22
Diamonds	2046454.5	(1140680.7)	752000	5250000	22
Gold	630.1	(451.3)	271	1669.5	22
Lead	1125	(766.7)	406.4	2580	22
Manganese	915.6	(477.1)	471	2380	22
Mineral Sands	13083.9	(14066.6)	3890	58100	22
Nickel	13001.4	(8321.2)	4629.5	37229.8	22
Platinum	851.9	(497.6)	359.7	1719.5	22
Tin	10260.6	(6956.6)	4060.5	26053.7	22
Tungsten	21100.5	(15675.9)	6820	56700	22
Zinc	1612.1	(748.3)	852	3500	22

Table D16: Mineral resources and economic activity in mining regions – geographical spillovers with mining cells dropped

	(1)	(2)	(3)	(4)	(5)
10 km	0.279***	0.285***	0.289***	0.296***	0.294***
	(0.090)	(0.090)	(0.090)	(0.090)	(0.091)
20-30 km		0.090**	0.094**	0.102**	0.100**
		(0.040)	(0.040)	(0.041)	(0.041)
30-50 km			0.029	0.037	0.035
			(0.034)	(0.035)	(0.036)
50-100 km				0.020	0.018
				(0.020)	(0.022)
100-200 km					-0.003
					(0.014)
Cell FE	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	Yes	Yes
Cells	8997	8997	8997	8997	8997
N	197135	197135	197135	197135	197135

This table collects difference-in-differences regressions that relate mineral resource activity (operating mines) to luminosity at the grid-level for all of Africa (0.5 × 0.5 decimal degree cells). As in Table 4.3, we study in this specification whether mines increase luminosity in neighboring cells. However, we drop all cells where mines are located (all cells that have no mines are included in the sample even if they are less than 10km away from a mine). The dependent variable is the log of mean light output in each cell. The independent variables are dummies that are one if a cell had an operating mine within the indicated distance. The sample includes operating mines with available information on startup and shutdown dates in the MinEx data. We drop the specific cells in which the mines are located. Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

Table D17: Mineral resources and economic activity in mining regions – different levels of clustering

	(1: All mines)	(2: ≥ Major mines)	(3: ≥ Giant)	(4: < Major)	(5: < Moderate)	(6: Prices)
Panel A: clustering at the first administrative level						
Mine	0.858***	1.078***	1.371***	0.410***	0.178*	
	(0.131)	(0.176)	(0.315)	(0.122)	(0.099)	
Mineral price						0.105***
						(0.020)
Panel B: clustering at the second administrative level						
Mine	0.858***	1.078***	1.371***	0.410***	0.178*	
	(0.122)	(0.151)	(0.254)	(0.122)	(0.097)	
Mineral price						0.105***
						(0.019)
Panel C: Conley SE						
Mine	0.860***	1.078***	1.373***	0.434***	0.076	
	(0.045)	(0.060)	(0.106)	(0.053)	(0.065)	
Mineral price						0.105***
						(0.006)

This table collects difference-in-differences regressions that relate mineral resource activity (operating mines) to luminosity at the grid-level for all of Africa (0.5 × 0.5 decimal degree cells). In this specification, we replicate the baseline regressions reported in Table 4.2 with different types of standard errors. In Panel A, we replace the unit of clustering from the cell to the first level administrative unit (e. g., states). In Panel B, we replace the unit of clustering to the second administrative unit (e. g., districts). In Panel C, we report results with Conley standard errors with a cutoff radius of 100km (Conley, 1999; Fetzer, 2014; Hsiang, 2010). The estimates in Panel C omit the country-year fixed effects due to computational limiations (i. e. we only include country and year fixed effects (but not country-year fixed effects) when reporting Conley standard errors). Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

Table D18: MINERAL RESOURCES AND ECONOMIC ACTIVITY IN NON-MINING REGIONS – DIFFERENT LEVELS OF CLUSTERING

	(1: First administrative level cluster)			(2: Second administrative level cluster)			(3: Conley standard errors)		
	(1– Border)	(2– Capitals)	(3– Leaders)	(1– Border)	(2– Capitals)	(3– Leaders)	(1– Border)	(2– Capitals)	(3– Leaders)
Mines	-0.011*			-0.011**			-0.011***		
	(0.007)			(0.006)			(0.002)		
Capital × Mines		0.051***			0.051***			0.049***	
		(0.019)			(0.015)			(0.007)	
Leader × Mines			0.002			0.002			0.001
			(0.002)			(0.002)			(0.001)

This table collects difference-in-difference regressions that relate mineral resource activity (operating mines) to luminosity at the grid-level for all of Africa (0.5 × 0.5 decimal degree cells). In this specification, we replicate the baseline regressions reported in Table 4.2 with different types of standard errors. In Panel A, we replace the unit of clustering from the cell to the first-level administrative unit (e. g., states). In Panel B, we replace the unit of clustering to the first administrative unit (e. g., district). Note that cells might cover several first or second-level administrative units. We randomly assign cells to one of the covered units for the purpose of clustering. In Panel C, we report results with Conley standard errors with a cutoff radius of 100km (Conley, 1999; Fetzer, 2014; Hsiang, 2010). The estimates in Panel C omit the country-year fixed effects due to computational limitations (i. e. we only include country and year fixed effects (but not country-year fixed effects) when reporting Conley standard errors). Stars indicate significance levels at 10%(*), 5%(**) and 1%(***). Heteroscedasticity- and cluster-robust standard errors are in parentheses. The unit of clustering is the cell.

Table D19: Constitutional distribution of ownership rights for mineral resources to government units in African countries

Country	Assignment of ownership in the constitution	Regulation of management and sharing in the constitution or ordinary legislation
Angola	**1992-2010:** "All natural resources existing in the soil and subsoil, in internal and territorial waters, on the continental shelf and in the exclusive economic area, shall be the property of the State, which shall determine under what terms they are used, developed and exploited", art. 12 I.	**Since 2010:** The state shall "[...] determine the conditions for concessions, surveys and exploitation, under the terms of the Constitution, the law and international law", art. 16.
	Since 2010: "The solid, liquid and gaseous natural resources existing in the soil and subsoil, in territorial waters, in the exclusive economic zone and in the continental shelf under the jurisdiction of Angola shall be the property of the state", art. 16.	
Burkina Faso	**Since 1991:** "The natural wealth and resources belong to the people. They are utilized for the amelioration of their conditions of life and within the respect for sustainable development", art. 14.	**Since 2015:** A development mining fund is created and "[...] allocated to the financing of the regional and communal development plans. It is financed by the contribution [...] of the State at the level of 20% of the proportional royalties collected" and by "holders of mining exploitation permit and industrial exploitation quarries authorisation at the level of 1% of their monthly turnover [...]", art. 26 Mining Code 2015. "The Mining fund for geological and mining research and earth science support is financed by the allocation of 15%, proportional royalties, surface royalties, fixed fees and fees related to the licence to buy and sale collected gold", art. 29 Mining Code 2015.
DR Congo	**Since 2005** "The State exercises a permanent sovereignty notably over the soil, the subsoil, the waters and the forests, over the air, river, lakes and maritime spaces of the Congo as well as over the Congolese territorial sea and over the continental shelf", art. 9.	**2002-2018:** "The mining royalties are paid by the holder of the mining exploitation title to the Public Treasury. The latter is in charge of distributing the receipts of the mining royalties as follows: 60% remain in the hands of the Central Government, 25% is paid into an account designated by the Provincial Administration where the project is located and 15% into an account designated by the Town or the administrative territory in the area where the exploitation activities take place", art. 242 Mining Law 2002 **Since 2018:** "The mining royalty is paid by the holder of the mining title in the following proportions: 50% to the central government; 25% paid into an account designated by the Administration of the province where the project is located; 15% to an account designated by the decentralised territorial entity in whose area the operation is carried out; 10% to the Mining Fund for Future generations", art. 242, 2018 Mining Code.

Constitutional distribution of ownership rights for mineral resources to government units in African countries (cont.)

Country	Assignment of ownership in the constitution	Regulation of management and sharing in the constitution or ordinary legislation
Egypt	**Since 2012:** "The state's natural resources belong to the people, who have a right to their revenues. The state commits to preserving such resources, to their sound exploitation, and to take into consideration the rights of future generations", art. 18.	
Ethiopia	**1987:** "Natural resources, in particular land, mineral water and forest, are state property", art. 13 II S.2. **Since 1994/95:** "The right to ownership of rural and urban land, as well as of all natural resources, is exclusively vested in the State and in the peoples of Ethiopia", art. 40 III.	
Ghana	**Since 1996:** "Every mineral in its natural state in, under or upon any land in Ghana, rivers, streams, water courses throughout Ghana, the exclusive economic zone and any area covered by the territorial sea or continental shelf is the property of the Republic of Ghana and shall be vested in the President on behalf of, and in trust for the people of Ghana", art. 257 VI	The Mineral Development Fund established in 1993 receives 20% of the mining royalty payments. Half of the fund is distributed in the mining areas for projects to mitigate the effects of mining: 25% via the district assemblies and the rest to local communities. Brosio and Singh (2014); MDF Act (2016) mineral royalties: 91% to the central government 4.95% to municipalities and 4.05% to private landowners
Guinea		**1995-2011:** "The above duties, royalties and taxes are divided between the budgets of the State, local authorities and the Mining Promotion and Development Fund. The distribution rates are set by joint order of the Minister of Finance and the Minister of Mines", art. 142 Mining Code. **Since 2011:** Allocation of fixed fees and taxes: national budget 80%, direct support to local budgets of all Local Communities of the country 15%, Mining Investment Fund: 5%, art. 165 Mining Code.
Kenya	**1999 revised constitution:** "[...] provision may be made by or under an Act of Parliament enabling a person to be granted a right or interest to prospect for minerals or mineral oils on any area of Trust land, or to extract minerals or mineral oils from any such area", art. 115 (3). **since 2010:** "All land in Kenya belongs to the people of Kenya collectively as a nation, as communities and as individuals", art. 61 I. "Public land is - f) all minerals and mineral oils as defined by law", art. 62.	**since 2010:** "The State (a) ensures sustainable exploitation, utilization, management and conservation of the environment and natural resources, and ensures the equitable sharing of the accruing benefits", art. 69.

Constitutional distribution of ownership rights for mineral resources to government units in African countries (cont.)

Country	Assignment of ownership in the constitution	Regulation of management and sharing in the constitution or ordinary legislation
Madagascar		**Since 2005:** Fees from gold mining licences are shared between the region, commune and the Gold agency, art. 77 Mining Code amended by Law number 2005-021.
Namibia	**Since 1990:** "Land, water and natural resources below and above the surface of the land and in the continental shelf and within the territorial waters and the exclusive economic zone of Namibia shall belong to the State if they are not otherwise lawfully owned", art. 100.	
Niger	**Since 2010:** "The natural resources and the subsoil are the property of the Nigerien people. The law determines the conditions of their prospecting, their exploitation and their administration", art. 148.	**Since 2010:** "The receipts realized on the natural resources and on the subsoil are divided between the budget of the State and the budgets of the territorial collectivities according to the law", art. 152 of the constitution.
Nigeria	**Since 1979:** "The entire property in and control of all minerals, mineral oils and natural gas in, under or upon any land in Nigeria or in, under or upon the territorial waters and the Exclusive Economic Zone of Nigeria shall vest in the Government of the Federation and shall be managed in such manner as may be prescribed by the National Assembly", art. 40 III in the Constitution of 1979, art. 44 III in the constitution of 1999.	According to the Allocation of Revenue Act **1982** the federal government shall receive 56%, the state governments shall receive 24% and the remaining 20% shall go to local government councils. **Since 1999:** "The principle of derivation shall be constantly reflected in any approved formula, as being not less than 13 per cent of the revenue accruing to the Federation Account directly from any natural resources", art. 162 II.
Senegal	**Since 2001:** "The natural resources belong to the people", art. 25 I S.1.	**2009 - 2016:** "The share of annual resources from mining operations to be paid into the support and equalisation fund for local authorities is set at twenty percent", art. 1 decree number 2009-1334. **Since 2016:** "The proceeds of mining revenues are distributed between the general state budget, the Support and Equalisation Fund intended for local authorities [20% - art. 113] and the support fund for the mining sector mining sector [20% - art. 114]", art. 112 Mining Code.

Constitutional distribution of ownership rights for mineral resources to government units in African countries (cont.)

Country	Assignment of ownership in the constitution	Regulation of management and sharing in the constitution or ordinary legislation
South Africa	**Since 1996:** "Everyone has the right [...] to have the environment protected, for the benefit of present and future generations, through reasonable legislative and other measures that [...] secure ecologically sustainable development and use of natural resources while promoting justifiable economic and social development", art. 24. **Since 2002:** "Mineral and petroleum resources are the common heritage of all the people of South Africa and the State is the custodian thereof for the benefit of all South Africans. [...] As the custodian of the nation's mineral and petroleum resources, the State, acting through the Minister, may [...] grant, issue, refuse, control, administer and manage any reconnaissance permission, prospecting right, permission to remove, mining right, mining permit, retention permit, technical co-operation permit, reconnaissance permit, exploration right and production right", Mineral and Petroleum Act 2002 chapter 2 art. 3.	The **1999** Public Finance Management Act "establishes that the Minister of Minerals and Energy may determine that any community or local government may receive a payment from mining royalties. The payment goes to the Local Economic Development Fund managed by the national Department of Provincial and Local Government", Brosio and Singh (2014).
Tunisia	**Since 2014:** "Natural resources belong to the people of Tunisia. The state exercises sovereignty over them in the name of the people", art. 13.	**Since 2014:** "A portion of revenues coming from the exploitation of natural resources may be allocated to the promotion of regional development throughout the national territory", art. 136 of the constitution.
Uganda	**Before 2005:** No assignment of ownership in the constitution Article 244	**2003-2022:** "The central government is entitled to 80 percent of the mining royalties, the local government of the producing areas are entitled to 17 percent and the owner of the land gets 3 percent." Mining Act 2003
	Since 2005: Amendments to article 244 of the 1995 constitution: "the entire property in, and the control of, all minerals and petroleum in, on or under, any land or waters in Uganda are vested in the Government on behalf of the Republic of Uganda."	**Since 2022:** Federal Government receives 70% of the royalties from minerals, local government 15%, sub county or town council 10% and land owners 5%, Mining and Minerals Act 2022
Zambia	**Since 2015:** "All rights of ownership in, searching for, mining and disposing of, minerals wheresoever located in the Republic vest in the President on behalf of the Republic", Art. 3 I Mining Act.	**2008-2015:** "The Minister responsible for finance shall, in consultation with the Minister, establish a mineral royalty sharing mechanism for distributing royalty revenues", art. 136 2008 Mining Act. **Since 2015:** "The following principles shall apply to the mining and development of minerals – [...] development of local communities in areas surrounding the mining area based on prioritisation of community needs, health and safety" art. 4 (f) 2015 Mining Act.
Zimbabwe		**Since 2013:** "The State must ensure that local communities benefit from the resources in their areas", art. 13 IV of the constitution.

This table lists constitutional arrangements regarding mineral resources and notable changes during the sample period across countries in Africa. Further details can be found at, e. g., `https://www.constituteproject.org/countries/Africa` and `https://leap.unep.org/knowledge/legislation-and-case-law`.

Short CV

09/2020 - 02/2025	**PhD candidate** at the Graduate School of Economic and Social Sciences at the University of Mannheim.
04/2016 - 07/2019	**MSc Economics** at the University of Freiburg.
10/2011 - 09/2014	**BSc Business Informatics** at the Berlin School for Economics and Law.
2011	**A-levels** at the Evangelisches Gymnasium zum Grauen Kloster in Berlin.

Peer-reviewed publications

Birkholz, C., & Gomtsyan, D. (2023). **Immigrant religious practices and criminality: The case of Ramadan.** *Journal of Comparative Economics*, *51*(1), 90-104.

Asatryan, Z., Birkholz, C., & Heinemann, F. (2024). **Evidence-based policy or beauty contest? An LLM-based meta-analysis of EU cohesion policy evaluations.** *International Tax and Public Finance*, 1-31.

Asatryan, Z., Baskaran, T., Birkholz, C., & Hufschmidt, P. (2024). **The regional economics of mineral resource wealth in Africa.** *Economica.*