Essays in Development and Urban Economics

Sandro Provenzano

Department of Geography and Environment

London School of Economics

A thesis submitted to the Department of Geography and Environment
of the London School of Economics for the Degree of Doctor of Philosophy

London, June 2022

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I declare that my thesis consists of approximately 50,000 words, excluding all appendices and the bibliography.

Sandro Provenzano (June 22, 2022)

Statement of co-authored work

I certify that Chapter 1 of this thesis is co-authored with Hannah Bull. I contributed 75% to the writing of this output in its current form. Chapter 4 of this thesis is co-authored with Felipe Carozzi and Sefi Roth, and I contributed 33% of the work. Chapter 5 of this thesis is co-authored with Sefi Roth and Lutz Sager, and I contributed 33% of the work.

Statement of inclusion of previous work

I confirm that Chapter 2 was the result of a fundamental revision of a study I undertook for my master's thesis at Stockholm School of Economics.

Acknowledgements

I would like to express my deep gratitude to my supervisors, Olmo Silva and Felipe Carozzi. During the past four years, they were utterly dedicated to providing the best possible guidance and teaching me the art of academic research. Olmo was always there to motivate and support my learning by sharing his vast experience, and always lifted the mood with his cheerful spirit. Felipe went out of his way to make sure to be available and provide feedback and perspective when it was needed. I very much enjoyed and have taken away a great deal from our countless conversations and joint research efforts. I am very grateful that I had the opportunity to have two such kind, gifted and bright scholars as my mentors.

A special mention goes to my coauthors Hannah Bull, Lutz Sager and Sefi Roth with whom I very much enjoyed working and learning from. I wish to thank Vernon Henderson for taking the time to reflect and share his insightful thoughts about my projects during several meetings, and serve as my mid-term reviewer. I am grateful to Gabriel Ahlfeld and Steve Gibbons who always provided me with critical feedback and constructive comments that helped me improve upon my work. A debt of thanks also goes to the other members of the faculty and fellow PhD students who made me enjoy coming to the office and whose comments during the departmental seminars and meetings enriched my work. A special mention to Louise with whom I could share and discuss the everyday challenges of working with large-scale spatial data, as well as Andreas, Lars, Manuel, Margarida, Matt, Sanchayan, Philippe, Tanner, Timo, Yohan and many others for many great memories. I also wish to thank Pete, Sam, Maxime and the PhD Academy for their outstanding administrative support over the years, as well as the UK Economic and Social Research Council for the financial support.

I am grateful to the members of the Paris School of Economics who welcomed me warmly for three months and for the many insightful seminars and discussions that contributed to my research. Furthermore, I wish to thank the discussants and seminar participants at conferences, as well as anonymous referees that provided me with valuable feedback.

I am indebted to my family and friends who mean the world to me and have gone out of their way to support me, visit me and keep close contact throughout the years. I am espe-

cially grateful to my parents, Irene and Antonio, for all their dedication and care and for teaching me the value of education. A special mention to my brother, Julian, for the countless evenings spent discussing research and always being available when I need him. To my sister, Philomena, who I can always fully count on. And, to Manfred for all the help and support over the years. I also wish to thank Maria and Ced for the good times we spent together, especially during the pandemic when our activities let me get a break and recharge.

Finally, I wish to thank my wife, Chaimae, for supporting me every day and sharing all the ups and downs throughout this lengthy process. Thank you for making life beautiful. I dedicate this work to you.

Abstract

This thesis studies how local economic, environmental and political conditions impact local prosperity of communities and individuals. Within the first block, three chapters focus on how the interplay between local geographical, institutional and political economy mechanisms shape development outcomes. In the second block, two chapters examine how urban and environmental conditions impact the outbreak of respiratory infectious diseases. In more detail, the first essay examines the economic consequences of mineral mining on local communities in Africa by analyzing changes in satellite images. While mining boosts the local economy in democracies, gains in autocracies are meagre and come at the expense of increases in localized conflict. Furthermore, mining does not automatically lead to self-sustained growth after the closure of mine sites. The second essay studies how isolation from the national capital city impacts economic performance in Sub-Saharan Africa. I show that remote areas are less economically developed due to information frictions that reduce the incentives of state executives to invest into isolated areas. The third essay adds to this topic by documenting that proximity to the regional capital city has a similar economic effect as national capital cities. The fourth essay revisits the debate about the role of population density for the spread of Covid-19 in the United States. While denser counties are hit earlier during the pandemic, overall they do not get hit harder. The fifth and final essay studies how exposure to air pollution is associated with the propagation of respiratory infectious diseases such as influenza and Covid-19 in the United States. The study finds no evidence that air pollution affects case numbers of respiratory infectious diseases in the short-run. The overriding contribution of this thesis is to provide detailed insights on the relevance and mechanisms through which local geographical, environmental and political frameworks impact economic and development outcomes.

Table of Contents

List of Figures

List of Tables

Introduction

This thesis consists of five chapters that study how local economic, environmental and political conditions impact local prosperity of communities and individuals. The motivation behind this focus is threefold. Firstly, by obtaining a more precise picture of the factors and mechanisms that hinder or contribute to economic growth at the local level, we reach a deeper understanding of the underlying processes that shape economic development. Secondly, studying the characteristics that determine economic outcomes at the local level fosters our understanding of the causes of the high levels of spatial heterogeneities across cities, regions and countries. And third, by gaining a better understanding of the drivers of economic development and spatial inequalities policy makers will be able to design more effective policy interventions that lead to more equitable and sustainable development.

Thesis Overview and Common Themes

The thesis is divided into two blocks. Within the first block, three chapters are centered around the question of how local economic, political and environmental framework conditions impact local prosperity of communities and individuals. While the main geographical focus of this work is on Africa, Chapter 3 further includes countries in the Americas and Asia. More specifically, these chapters study how the interplay of local geographical and institutional frameworks with political economy mechanisms shape development outcomes such as local economic performance, urbanization, public goods provision and conflict. In the second block, two chapters examine how urban and environmental conditions like urban density and air pollution impact the outbreak and severity of respiratory infectious diseases such as Covid-19 and influenza in the United States.

A common characteristic of all chapters is that economic and social phenomena are studied from a quantitative-empirical perspective, leveraging a variety of unconventional large-scale data sources to overcome data constraints and focus on new and innovative research questions. The empirical identification strategies are centered around obtaining quasi-random variation in treatment to estimate causal effects. Chapter 1 uses a stacked event study approach, Chapter 2 uses a boundary discontinuity design, Chapter 3 combines an event study with a boundary discontinuity design and Chapters 4 and 5 apply instrumental vari-

able approaches.

Further, the special consideration given to the spatial dimension of economic phenomena is an important feature throughout this thesis. One of the advantages of the spatial perspective is that it allows to obtain additional insights about research subjects. Based on the geolocation of individuals, households, cities and regions, it is possible to gather additional information from spatial data sources like geocoded surveys or satellite images. An additional advantage of such a 'spatial lens' on economic phenomena is that it opens up additional possibilities for the identification of causal effects. First of all, local economic, geographical, institutional and political conditions vary across space which facilitates the understanding of their relevance for development outcomes. Moreover, a spatial empirical approach makes it feasible to exploit discontinuous cutoffs in space, such as certain national or regional boundaries, to obtain quasi-random variation in variables of interest.

Lastly, all chapters focus on current policy-relevant questions. The first three chapters emphasize the context of developing countries and investigate societal challenges such as spatial inequalities, accountability or the use of natural resources. In contrast, motivated by the ongoing pandemic, the last two chapters focus on the United States and study the consequences of urban density and air pollution for the spread and severity of infectious diseases.

Contribution

The overriding contribution of this thesis is to provide detailed insights on the relevance and mechanisms through which local geographical, environmental and political frameworks impact economic and development outcomes.

More specifically, a central finding within the first block is that local geographical and environmental conditions impact and interact with political and institutional frameworks, and ultimately jointly determine local economic performance. In Chapter 2, I find that distance from the national centers of political decision making, the capital city, has a negative impact on economic development through reducing political accountability. I show that geographical isolation leads to isolation from information which ultimately impacts voting behavior and reduces the incentives of state executives to invest into isolated areas. These findings underline the importance of going beyond the national level and focus on small geographical units to better understand the determinants of economic development

and the causes of spatial heterogeneity. Furthermore, in Chapter 3, I document that the positive effect of reducing the distance to regional capitals depends upon the institutional context such as if the country is decentralized or a democracy. Similarly, in Chapter 1, we show that it is primarily democracies that benefit from mineral mining of natural resources. In contrast, mining in autocracies only leads to meagre economic gains that come at the expense of an increase in localized conflict. These findings emphasize the need to pay special attention to institutional frameworks and political economy mechanisms when investigating the role of geographical factors on local communities.

The implications of these findings go beyond the academic discourse and provide important insights for policy makers seeking to foster economic development. Most importantly, the findings of this thesis point at the importance of strengthening civic engagement and citizen participation within the political process. One way of achieving this goal would be to enable and encourage participation institutionally by reinforcing democratic rights, as well as developing governance processes that explicitly include citizen participation, for example when it comes to the management of natural resources. Another way of reaching this goal would be to promote the propagation of information about the decisions and actions taken by the political leadership, especially in remote areas. Ultimately, these actions could significantly improve political accountability and lead to more inclusive and effective policies, as well as reduce conflict. Other important policy implications include: i) paying more attention to the spatial distribution of public goods that is currently undersupplied in certain areas ii) taking into consideration that public policies that were a success in countries with strong institutions might be less effective or even net negative in a poor institutional setting and iii) a period of local economic expansion, for example induced by active mineral mining, does not automatically lead to self-sustained growth but requires public policy to support the process of economic transformation.

Within the second block, we provide novel findings on the role of population density and short-run air pollution for the propagation and severity of respiratory infectious diseases that are in contrast to previous research.[1] With regard to population density, we show that dense places feature both, characteristics that reinforce as well as characteristics that reduce the disease spread and severity. For example, while the circumstance that more people

[1] See for example Angel et al. (2020); Whittle and Diaz-Artiles (2020); Zhang and Schwartz (2020); Wheaton and Kinsella Thompson (2020) and Almagro and Orane-Hutchinson (2020) for evidence linking population density to a faster disease spread, as well as Clay et al. (2018); Isphording and Pestel (2021); Persico and Johnson (2021); Graff Zivin et al. (2021) for evidence linking air pollution to higher case numbers and fatalities of the Spanish flu, influenza and Covid-19.

live in close proximity in itself facilitates the spread of diseases, people in dense places reduce their activities more in response to disease outbreaks, have better access to healthcare and are on average younger and more resistant to diseases. These offsetting characteristics of density are often overlooked when considering risk factors for disease outbreaks, but need to be taken into consideration by public officials when designing measures to contain outbreaks. The contribution on the role of air pollution for influenza and Covid-19 outbreak is rather methodological and emphasizes the importance of using an identification strategy that is able to account for the cyclicality and seasonality of air pollution.

Last but not least, this thesis, and in particular the methodological work in Chapter 1, expands upon the current level and use cases of satellite images within economics and the social sciences. We show how the applicability of satellite imagery goes beyond serving a complementary role by adding one additional indicator to an existing dataset, but is well suited to be the main unit of observation providing a wide variety of information, such as urban, agricultural and mine land use or wealth index predictions. Satellite images allow researchers to obtain a new and much more granular perspective on economic and social processes that was previously out of reach. For example, it allows to precisely monitor urban growth and decline around the world. Moreover, since some satellite configurations such as Landsat have been taking similar images since the 1980s, we show that it is possible to construct panel datasets that are not only extensive in space but also in time. This especially opens up new opportunities for research seeking to better understand long-run effects. Given the technological advances in remote sensing and the increase in the number of orbiting satellites, remote sensing will continue to enable new perspectives on economic, social and environmental phenomena. This will allow us to gain entirely new insights on human organization, and will be particularly useful for studying developing countries where research is currently often constrained by a lack of reliable data. It is my hope that this thesis will contribute to further applications of earth observation data in the quantitative social sciences in the future.

Chapter Overview

Chapter 1

Mining projects often gain support from communities through promises of benefits to the local economy, though the empirical evidence is mixed: mineral assets are advantageous in some circumstances but lead to corruption and violence in others. To shed light on this

apparent discrepancy, we significantly extend the coverage of previous work in this area by gathering satellite data that spans several decades and encompasses several institutional environments. Our dataset consists of one million 30m-resolution Landsat images between 1984 and 2019 from a 40km radius around 1,658 mineral deposits across the continent. Using state-of-the-art techniques in computer vision, we translate these images into economically meaningful indicators measuring urban growth and material wealth. We then use stacked event studies and difference-in-difference models to estimate the impact of mine openings and closings. We find strong evidence of a political resource curse at the local level. Although mining boosts the local economy in democratic countries, these gains are meagre in autocracies and come at the expense of an increase in localized conflict. Furthermore, our results show that the growth acceleration in mining areas is only temporary. While former mining areas maintain their increased wealth and urban size with respect to non-mining areas, there is no sustained divergence.

Chapter 2

This chapter documents that in Sub-Saharan Africa areas isolated from the capital city are less economically developed and examines potential underlying mechanisms. We apply a boundary-discontinuity design using national borders that divide pre-colonial ethnic homelands to obtain quasi-experimental variation in distance to the national capital city. Based on nightlights and geocoded surveys, we find that a one percent increase in distance to the capital city causes a decrease in the probability of detecting nightlights by 3 percentage points and a reduction in household wealth corresponding to 3.5 percentiles of the national wealth distribution. Our results suggest that a lower provision of public goods in isolated areas is a key link between remoteness and economic performance. Despite receiving worse services, people who are isolated exhibit a higher level of trust in their political leaders. Further, isolated citizens consume the news less frequently and penalize their leaders less for misgovernance. We interpret these findings as pointing towards dysfunctional accountability mechanisms that reduce the incentives of vote-maximizing state executives to invest into isolated areas.

Chapter 3

Recent studies have emphasized the importance of proximity to national capital cities for economic performance. In this chapter, we investigate if regional capital cities are similarly relevant for local economic development. We focus on regional boundary reforms that split regions into smaller entities with new designated regional capitals and thereby induce

variation in distance from the regional capital city. Using satellite nightlights as a measure of economic activity and combining an event study design with a boundary discontinuity design, we document that proximity to regional capital cities has a strong positive causal impact on economic growth. On average, a one percent decrease in distance to the regional capital increases nightlight density by 0.16 percent within 8 years. Furthermore, we show that the economic benefits of proximity to regional capitals are mainly relevant in less developed, decentralized and democratic countries.

Chapter 4

This chapter revisits the debate around the link between population density and the severity of Covid-19 spread in the United States. We do so by conducting an empirical analysis based on graphical evidence, regression analysis and instrumental variable strategies borrowed from the agglomeration literature. Studying the period between the start of the epidemic and the beginning of the vaccination campaign at the end of 2020, we find that the cross-sectional relationship between density and Covid-19 deaths changed as the year evolved. Initially, denser counties experienced more Covid-19 deaths. Yet, by December, the relationship between Covid deaths and urban density was completely flat. This is consistent with evidence indicating density affected the timing of the outbreak – with denser locations more likely to have an early outbreak – yet had no influence on time-adjusted Covid-19 cases and deaths. Using data from Google, Facebook, the US Census and other sources, we investigate potential mechanisms behind these findings.

Chapter 5

Exposure to elevated levels of air pollution is linked with a wide range of adverse health outcomes and, recently, it has been suggested that air pollution might also be associated with the propagation of respiratory infectious diseases such as influenza-like illnesses (ILI) and Covid-19. In this chapter, we examine the relationship between air pollution and respiratory infectious diseases empirically by analyzing US administrative data on ambient air pollution and weekly cases of ILI and Covid-19. Assessing the link between pollution and infectious diseases is challenging due to the presence of correlated omitted variables and measurement error. We overcome these challenges by using an instrumental variable approach that relies on satellite-derived atmospheric temperature inversions as an instrument for air quality. Applying a variety of different specifications, we find no evidence that exposure to elevated levels of air pollution affect influenza and Covid-19 cases in the US – a finding that is in contrast to several recent papers in the economics literature.

Chapter 1

The Local Economic Impact of Mineral Mining in Africa: Evidence from four Decades of Satellite Imagery

1.1 Introduction

Africa's great wealth in natural resources attracts very large investments by international mining companies every year. Since the early 2000s, increased commodity prices have led to a surge in investor interest and mine openings on the continent. Mining projects often gain support from communities through promises of benefits to the local economy. Yet, we still lack conclusive evidence on how industrial mining shapes local economic development and the duration of this impact.

Using satellite data over multiple decades with wide-spanning geographic coverage, we can measure local economic changes at different stages of mining under various institutional contexts. We provide three key results on the local impact of mining. Firstly, mine openings, and in particular large mine openings have the potential to considerably increase urbanization, our main measure of economic development. Furthermore, mining also increases agricultural activities and material wealth in the local area. Secondly, we find strong support for the presence of a political resource curse at the local level. While mining boosts the local economy in democratic countries, these gains are meagre in autocracies and come at the expense of an increase in localized conflict. Thirdly, our results show that the growth boost in mining areas is only temporary. Although former mining areas remain more urbanized than comparable non-mining areas, there is no sustained divergence.

Previous research on the local impact of mineral mining has reached conflicting results: some found that mine openings boost the local economy (Aragón and Rud, 2013; Mamo et al., 2019), while others found no significant effects (Pokorny et al., 2019; Bazillier and

Girard, 2020)[1] and increases in corruption and conflict (Vicente, 2010; Berman et al., 2017; Knutsen et al., 2017). We argue that the reason for the mixed results is due to data limitations as these studies usually focus on a few mines or countries, and do not take into consideration the institutional context. Yet, research on the macro level suggests only countries with good quality institutions benefit from natural riches (Mehlum et al., 2006; Bhattacharyya and Hodler, 2010). Accordingly, institutions might play a similar role for the local impact of mining. Based on a recent field experiment in Northern Mozambique, Armand et al. (2020) provide tangible evidence in support of the local political resource curse hypothesis. The authors show that increased information and community participation helps to prevent civil conflict. In this study, we show that this mechanism applies more broadly and can explain cross-country heterogeneities. While mining communities in democracies benefit from mine openings, those in autocracies experience only meagre gains and instead increases in localized conflict.

Our approach is novel as we use almost one million satellite images to create an extensive panel covering 12% of the African landmass, tiling the area within 40km^2 of 1,658 mines across 47 African countries, over a period 35 years. We exploit the fact that the images are acquired from different stages in the mine life-cycle, including prior to the min-

FIGURE 1.1
AUTOMATIC ANNOTATION OF TWO LANDSAT IMAGES OVER TIME

(A) 1991: A LOCATION PRIOR TO THE ONSET OF MINING, WITH A SMALL URBAN AREA (MAGENTA)

(B) 2018: THE SAME LOCATION SHOWING DEVELOPMENT OF THE MINE (CYAN) AND URBAN GROWTH (MAGENTA)

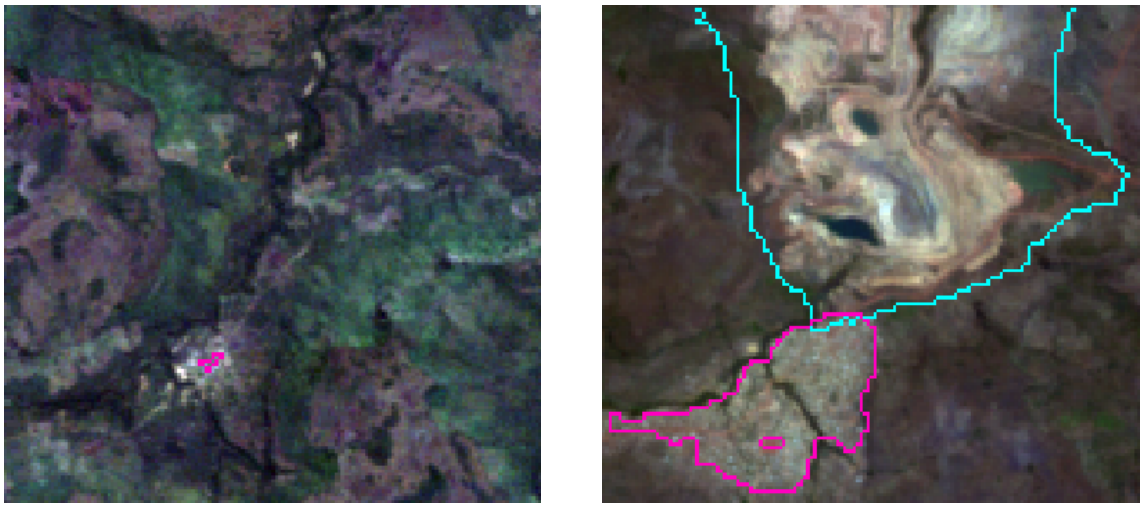

Note: Automatic annotation of satellite images allows long term observation of urban growth, land use and material wealth in mining areas. See Figure 1.A.1 for recent high-resolution images of this location.

[1]Note that these studies only found positive income and employment effects for small artisanal mines but not industrial mines which are the focus of this study.

eral discovery and after the mine's closure. This allows us to observe the local effects of mining over long periods of time and under strong and weak institutional contexts. We use state-of-the-art techniques in machine learning and computer vision to translate this vast collection of images into economically meaningful indicators including urban and agricultural land use and material wealth predictions[2]. Figure 1.1 illustrates how we can successfully detect landcover changes before and after the mine opening using satellite images and machine learning. We track these indicators across the one million images from 12 three-year periods. Using stacked event study designs and difference-in-difference (DiD) models, we estimate the impact of mine openings and closings on our satellite derived indicators.

To pinpoint how mine opening and closing events impact the evolution of the local economy we use a variety of counterfactuals. This includes comparing opening mines to areas where minerals were discovered but where production has not yet started. Using event-study graphs, we show that prior to the mine opening both groups evolve similarly in terms of urban growth, and that opening mines diverge and gain relatively only after the onset of mining. After 15 years, areas within 20km from opening mines gain on average around 80% in urban growth relative to not yet mined mine areas. In addition, we show that mining primarily impacts the area within 20km of the mine and only to a lesser extent areas further away, and that large mines tend to have a stronger effect as small mines. Furthermore, we find that a part of the economic gains from mining are indirect due to increased agricultural activities. As an alternative identification strategy, we use late mine openings prior to treatment as controls for early mine openings and obtain similar results.

In addition, we provide evidence for the presence of a political resource curse at the local level by showing that a democratic institutional setting is a decisive factor for making mine operations a success for local communities. Regarding the underlying mechanism, we demonstrate that one of the advantages of mining areas in democratic countries is that, as opposed to mining areas in autocratic countries, they can avoid a rise in conflict following the mine opening.

It is much harder to find a suitable control group for closing mines, as the mines experience two treatments, one of being an active mine as well as the subsequent closure. We

[2]Since wealth predictions are based on satellite images, 'wealth' does not explicitly refer to usual measures of wealth such as assets or savings. While the underlying models are trained using household level asset wealth indices as input, the wealth index ultimately reflects correlates and material manifestations of wealth in a local area such as urban shape and density, infrastructure, or roof reflectance.

can nevertheless benchmark the effects of mine closures by comparing their pre-treatment and post-treatment trends to continuously operating and not yet treated areas. Prior to the closure, communities around closing mines grow relative to both comparison groups. However, after the mine's closure, there is a trend break and sharp drop in the economic growth rate around closing mines. Closed mines continue to evolve no different than areas where mining had not yet started. These results indicates that the growth acceleration in mining areas is only temporary. While former mining areas maintain their increased urban size and wealth with respect to non-mining areas, there is no sustained divergence.

Related Work

The literature on the economic implications of the extractive resource sector in developing countries is vast. There is extensive research in macroeconomics that investigates potential underlying mechanisms regarding the 'resource curse' - the counter-intuitive finding that countries with a larger natural resource sector tend to be less developed (Auty, 1993). Researchers have made various attempts to solve the puzzle. In the 1980s and 90s, the literature has pointed towards macroeconomic mechanisms such as the 'dutch disease', i.e. a rise in the price level that make other exporting sectors less competitive (Corden and Neary, 1982; van Wijnbergen, 1984; Krugman, 1987; Sachs and Warner, 2001). Other explanations included the high volatility of commodity prices that make resource rich countries more vulnerable to macroeconomic shocks (Deaton, 1999). Subsequently, studies like Lane and Tornell (1996), Tornell and Lane (1999), Torvik (2002), Hodler (2006) and Caselli and Michaels (2013) saw the cause of resource misfortune in rent-seeking that prevent the redistribution of resource windfalls to the population and additionally shifts the economy towards less productive activities. Caselli and Michaels (2013) finds that only little of the increased oil revenues of Brazilian municipalities actually benefits the population, probably due to embezzlement. Finally, more recent research emphasizes the political dimension of the resource curse. These studies argue that natural resources are only advantageous in places with good institutions (Mehlum et al., 2006; Robinson et al., 2006; Collier and Hoeffler, 2009; Bhattacharyya and Hodler, 2010). In places with poor institutions, resource windfalls deteriorate institutions even further by increasing corruption and undermining the political process which results in a net negative impact of resources on growth.[3]

The literature on the local economic impact of mining is relatively recent. In their seminal paper, Aragón and Rud (2013) investigate a production expansion of a Peruvian Gold mine

[3]See van der Ploeg (2011) and Venables (2016) for extensive literature reviews and a discussion about the challenges of resource governance in developing countries.

over ten years. The authors find strong evidence for increases in real income in the mining city and neighboring areas, albeit with a decreasing positive impact with distance from the mining city. Moreover, service and agricultural workers also benefit from the mine expansion, and the authors argue that this is due to backward linkages. Similarly, Lippert (2014) studies the copper mining boom in Zambia and documents positive effects of mining beyond the mining sector and immediate mine area. Mamo et al. (2019) find positive effects of mining on the district level in Africa at the intensive (increased mining production) and especially at the extensive margin (new mineral discoveries and mine openings). Yet, the authors find little evidence for spillovers to neighboring districts. In regard to evidence from the developed world, Allcott and Keniston (2018) find positive effects of oil and gas production at the US county level on real wages and no negative effects on the productivity of the tradeable manufacturing sector.

In contrast, Hirschman (1958) and McMillan et al. (2014) maintain a more critical view on mining and consider it to be a highly productive and capital intensive 'enclave'. These authors argue that mining ventures are without linkages to the local economy, do not provide extensive employment opportunities for the local workforce and are therefore unable to induce structural change. Bazillier and Girard (2020) investigate the gold boom in Burkina Faso and find that artisanal mines as opposed to industrial mines are having a significantly positive impact on the local economy. Similarly, Pokorny et al. (2019) find that artisanal small-scale mining creates jobs and income whereas industrial mines fail to do so. Consequently, one of the reasons for different findings with regard to the effects of mining could be related to heterogeneities between different mine types (large industrial mines that are capital intensive that might provide relatively little job opportunities vs. artisanal mines that provide a source of income to local workers). There are other reasons why mining might not be beneficial for local communities. Aragón and Rud (2016) find that mining has considerable negative effects on local agricultural productivity via pollution that might even outweigh the gains from mining. The authors leave open whether this effect is mainly due to adverse effects on workers' health and hence labor productivity or due to soil and crop deterioration. The results of other studies that examine the impact of mining on health are ambiguous, mainly because there appear to be direct negative effects induced by pollution but positive indirect income effects. One the one hand, von der Goltz and Barnwal (2019) find that mining related pollution induces adverse health outcomes such as lead toxicity. On the other hand, Benshaul-Tolonen (2019) finds that gold mine openings reduce child mortality by 50% arguing that the indirect positive employment and

income effects outweigh direct negative health effects.

Lastly, recent studies point at the relevance of political economy mechanisms with regard to the local impact of resource wealth. Vicente (2010) and Knutsen et al. (2017) show that mining deteriorates local institutions through increasing corruption. Further, Dube and Vargas (2013) and Berman et al. (2017) document that resource windfalls result in political instability by increasing conflict.[4] Out of these studies, only Berman et al. (2017) examines if the effects are particularly relevant in countries with poor institutions as is suggested by the political resource curse hypothesis. The authors do not detect significant heterogenous effects of mining on conflict with regard to national institutions. However, the authors do find that putting in place more stringent mining specific anti-corruption measures and transparency initiatives helps to reduce conflict. Furthermore, based on a field experiment in areas of recent natural gas discoveries in Northern Mozambique, Armand et al. (2020) provide tangible evidence in support of the local political resource curse hypothesis. The authors show that information campaigns targeting communities and motivating them to participate in the decision-making led to decreases in conflict. In contrast, when the information only reaches the local leader there is no reduction in conflict, but instead increased elite capture and rent seeking.

The most important shortcoming in the existing literature on the local effects of mining is associated to the fact that previous studies usually only have a relatively small sample, often even focussing on only one or a handful of mines or countries. If the hypothesis of the political resource curse at the local level was true, we would expect there to be important heterogeneities between countries depending upon the institutional quality. Consequently, relying on a small sample might undermine external validity, which could explain why some studies reach conflicting conclusions about the impact of mining. A second limitation of the existing literature is that they usually only cover a relatively short time period and do not observe long-term effects of mine openings. Another drawback of short time periods is that it makes it hard to validate the identification assumption when using DiD. There are a few studies, such as Chuhan-Pole et al. (2016) and Mamo et al. (2019), that have a relatively long panel over around 20 years and cover a wider sample of mines and countries. Yet, these studies also face limitations as they rely on nightlights as a measure of economic activity. The use of nightlights is likely to result in distorted estimates in this context as

[4] Berman et al. (2017) show that one important channel through which mining increases conflict is by helping rebel groups in control of mine sites to finance their military capacity. Yet, the authors point out that there might be other important mechanisms that link mining and conflict, such as increasing the potential for rent-seeking.

the authors have difficulties distinguishing between a change in nightlights that reflects the mere activity of the mine and associated infrastructure, and the change that is related to wealth and income changes. With our approach of using archives of Landsat satellite images to track urban change and mask out the mines themselves, we extend previous research by obtaining a more extensive and granular view on how 1,658 mineral deposits across 47 African countries shape the surrounding area over a period of 35 years. One of our focuses will be to leverage our database to investigate potential differences in the local effects of mining between countries with good and poor institutions as is proposed by the local political resource curse hypothesis. Lastly, with our extensive panel we can evaluate how mining areas evolve after the closure of the mine, which is something that has not been covered in previous research.

The remainder of this paper is organized into five sections. Firstly, Section 1.2 will introduce the sample. Secondly, Section 1.3 will present the satellite data and machine learning techniques to process the images. Thirdly, Section 1.4 will present and discuss the identification strategy. In Section 1.5, we present our findings on the impact of mine openings and closings, investigate heterogeneities with respect to mine size and the institutional context and examine conflict as a potential mechanism. Finally, Section 1.6 summarizes our findings.

1.2 Sample

Our dataset on mining includes information on 1,658 mineral deposits and industrial mines in Africa, including the size and type of mine, the date of discovery, the dates of activity and the geographic location. Figure 1.2 summarizes the activity status of these mineral deposits during our study period between 1984-2019 which we divide into 12 equal-length periods of 3 years.[5] This categorization of mines is key to our identification strategy and defining treated and untreated units[6], and refer to these categories throughout the remainder of the article.

At the first level of grouping, we distinguish mines that were actively operating at any point during our study period. At the second level of grouping, we distinguish active mines into those mines that had been active in period 1 (1984-1986) and continued oper-

[5]The reason for using 3-year periods is to get a higher coverage of cloud-free satellite images in each period which is especially important for earlier images.
[6]We discuss the assignment of mines into treatment and control groups in more detail in Section 1.4.3.

ating throughout our study period and other mines that were only partially active. We also distinguish inactive mines into mines that had ceased operation at some point prior to 1984 and mineral discoveries with no active mine as of 2019. At the third level of grouping, we categorize mines that were partially active in our study period into three groups: areas where a mine began operation some time between period 2 and 12 and continued operation until period 12, mines that had been operating in period 1 but that closed at some point during our study period and remained inactive, and mines that opened and closed or closed and reopened during our study period.

FIGURE 1.2
MINE STATUS OVERVIEW IN THE PERIOD OF 1984-2019

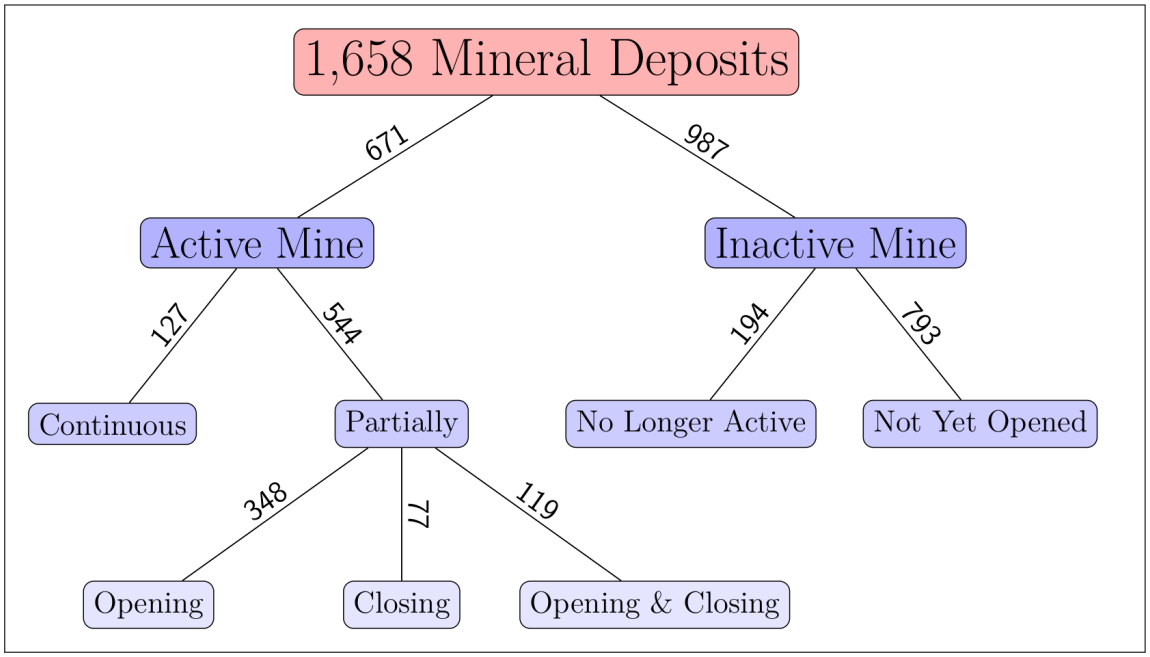

Note: Active mines are active during our period of interest. Continuous mines are active throughout the whole period while Partially active mines are not always active. Partially active mines may Open, Close or Open and Close within our period of interest. Our identification strategy exploits differences between these categories in order to measure the effect of active mining and mine closures on local communities.

Using the geolocalization of each mine, we assemble medium-resolution multi-spectral satellite images from a radius of 40km of each mine throughout our study period from 1984 to 2019. Using state-of-the-art techniques in machine learning and computer vision, we extract quantitatively meaningful economic indicators. In particular, we extract information on land use, including monitoring the extent of urban area and agricultural areas, and about the local material wealth level. Moreover, we are able to automatically segment the area of the mine from each image, in order to exclude the mine itself when analyzing the impact on the area surrounding the mine.

Our final dataset thus covers an area within a 40km radius of 1,658 mineral deposits across 47 African countries. This amounts to a total study area of 3.6 million km^2 (representing 12% of the total African landmass or 15 times the landmass of the United Kingdom). We divide this area into 84,207 tiles of 6.5×6.5km $\approx$ 43km^2, and observe the evolution of each tile over our 12 periods from 1984 until 2019. We associate this observed evolution to the respective neighboring mine activity. To the best of our knowledge, our dataset is the longest and most extensive panel database on mineral mining. The area studied is depicted in Figure 1.3. Further details on the assembling of this large panel dataset are provided in Section 1.3 and App. 1.A.1.

FIGURE 1.3
LOCATIONS OF MINERAL DEPOSITS AND TILE GRID

(A) 1,658 MINERAL DEPOSITS (B) 40KM RADIUS AROUND EACH MINERAL DEPOSIT

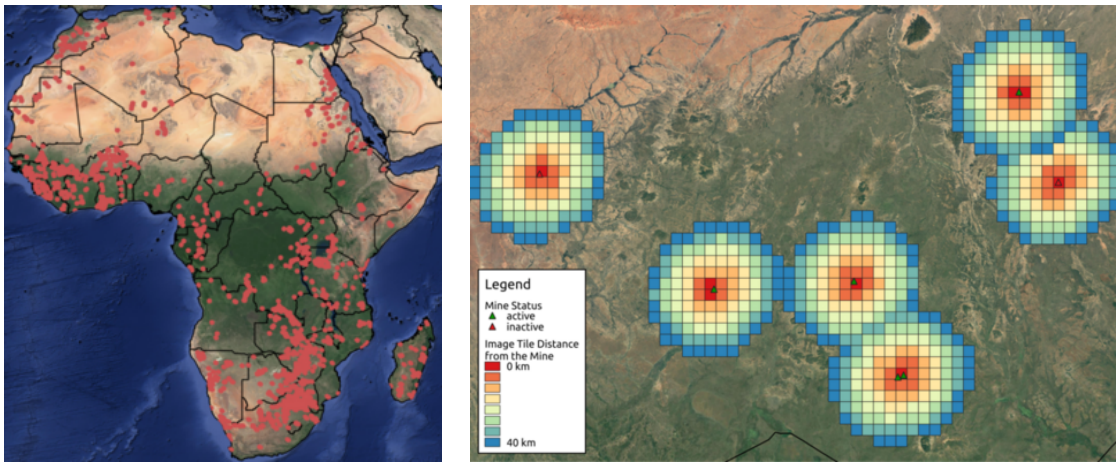

Note: The area within 40km surrounding a mineral deposit is divided into tiles of 6.5×6.5km $\approx$ 43km^2, providing a total of 84,207 images covering 12% of African landmass. Each tile corresponds to a Landsat image with 224×224 pixels at 30m resolution, with 7 multi-spectral bands. We acquire one image (3 year median) at each of the 12 three-year periods, corresponding to around one million tiles.

In order to assign tiles to mines, we compute the distance from each tile centroid to each active mine in each period using the 'Africa Sinusoidal'-projection that maps distances using the metric system.[7] In a next step, we classify tiles as either being within 20km of a mine *(close)*, between 20 and 40km from a mine *(far)* or being further than 40km from a mine *(non-mine area)*. We then classify these tiles as *Continuous*, *Opening*, *Closing* or *Inactive Mine* depending on their distance to the closest mine and group assignments over the periods. We remove a relatively small number of tiles that switch more than once between different assignments, as these groups receive multiple treatments that are hard

[7]For more information about the 'Africa Sinusoidal' projection visit: `http://spatialreference.org/ref/esri/africa-sinusoidal`.

to distinguish. Lastly, we assign each tile to the first active mine (< 40km) for active mine tiles and to the closest inactive deposit for inactive mine tiles to retrieve additional information about the mine such as if the mine was active in the past (*No Longer Active* vs. *Not Yet Opened*), mine size or which mineral is mined. The advantage of proceeding like this and using distance intervals to determine treatment is that we can categorize tiles into dichotomous treatment groups.[8]

1.3 Remote Sensing and Machine Learning

Our analysis requires using rich satellite imagery to extract variables on urban and agricultural land use and material wealth as well as the extent of the mine. We train neural networks to predict land use and mine areas from Landsat images, and use an existing model to predict local material wealth (Yeh et al., 2020).

The satellite imagery consists of the Blue, Green, Red, Near Infrared, Short-wave Infrared 1, Short-wave Infrared 2 and Thermal bands from the atmospherically corrected surface reflectance sensors of Landsat-5, 7 and 8. The Thermal band is at 120m resolution and has been resampled to 30m pixels, and the remaining bands are at 30m resolution. We take the median of all Landsat images over consecutive 3-year periods between 1984 and 2019, excluding any pixels with clouds, cloud shadows, or snow. We divide our area of interest (within 40km of a mine) into square tiles of 224×224 pixels, or approximately $6.5 \times 6.5\text{km}^2$. Thus, we obtain 84k 7-channel images of size 224×224 at each of the 12 three-year periods, or around one million satellite images. We download this data using Google Earth Engine.

1.3.1 Land Use and Mine Segmentation

Our land use model learns to segment our image tiles into 4 mutually exclusive classes: Urban Areas, Croplands, Water and Other. Our mine segmentation model predicts a binary outcome at each pixel, signaling the areas of mines. In order to train the land use model, we create a new land use segmentation dataset by combining annotations from multiple sources listed in App. 1.A.1. Annotations for the mine segmentation come from Maus et al. (2020).

[8] Other approaches like assigning tiles to the closest mine overall would lead to confusion about what the treatment is and distortions as a tile could be initially 30km and then 5 km away from the mine when another mine opens closeby.

To obtain land use indices, as well as the segmented area of each mine, we train two convolutional neural networks (CNNs) with a U-Net backbone architecture (Ronneberger et al., 2015) and with 16M parameters. The land use model and mine segmentation model are both treated as pixel classification problems (image segmentation) and are trained with cross-entropy loss. The encoder side of the U-Net model is a ResNet-50 model (He et al., 2016), but adapted to input 7 image channels instead of 3 RGB channels at the first convolutional layer. Moreover, the thermal band is at a resolution 4 times lower than the remaining 6 bands, and so we use 4-pixel dilated convolutions on this channel. For the land use model, the model weights are initialized by the ResNet-50 model pre-trained model on ImageNet (Deng et al., 2009). The model weights corresponding to the non-RGB bands in the first convolutional layer are initialized as the average weights corresponding to the RGB bands. For the mine segmentation model, the weights are initialized using the trained land use segmentation model.

TABLE 1.1
DATA FOR TRAINING AND EVALUATING THE LAND USE AND MINE SEGMENTATION

Task	**Land Use Segmentation**	**Mine Segmentation**
Landsat images	84k images (2014-6)	28k×2 images (2014-6, 2017-9)
Area	40km radius of mines (Africa)	40km radius of mines (Global)
Classes	Water Urban Areas Cropland Other	Mine
Annotations	ESA (Water, Urban) GHSL (Urban) Facebook (Urban) NASA (Croplands)	Maus et al. (2020)
Train-Val-Test	80%-5%-15%	80%-5%-15%

Note: Further details on land use annotation sources are listed in App. 1.A.1.

Our evaluation metric is the R^2 value between the predicted shares of a class and the ground truth shares of a class within each image tile. This is because in our analyses, our observation unit is an image tile. We are interested in computing the share of each land use category and the share of mines in each image tile and thus do not require pixel-level accuracy. We firstly train the land use model for 75 epochs (passes of the training set through the model) on a single NVIDIA Tesla K80 GPU (12 days) and choose epoch 72, as it has the best evaluation metric on the validation set. We then initialize the weights of the mine segmentation model with those of the land use model and train this model for a

further 2 days or 40 epochs on a single NVIDIA Quadro RTX 3000 GPU and chose epoch 29.

Performance of the land use model and the mine segmentation model on the test set is provided in Figure 1.4. The R^2 value is 0.96 for urban areas, 0.82 for cropland areas and 0.98 for water bodies, demonstrating that our model is a strong predictor of the share of land use categories in each image tile. The R^2 value for mine prediction is also high, at 0.78. Our model is thus capable of predicting image tiles with a high or low presence of mines. Qualitative examples of the land use and mine segmentation models are provided in Figure 1.5. Visually, our model is good at segmenting urban, cropland, water and mining areas.

FIGURE 1.4
TEST SET: PREDICTED VS. ACTUAL SHARES BY SEGMENTATION CLASS

(A) R^2 URBAN: 0.96 (B) R^2 AGRIC.: 0.82 (C) R^2 WATER: 0.98 (D) R^2 MINE: 0.78

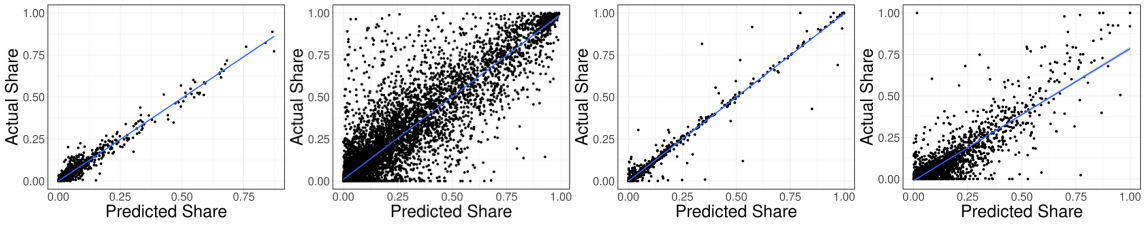

Note: The R^2 values show strong correlations for all classes, but particularly for the Urban and Water classes.

When applying our trained land use and mine segmentation models to all image tiles across all time periods, we firstly use the mine segmentation model to mask out the mines from all the tiles. We then compute our outcome measures on land use by computing for each image the log of the share of pixel that belongs to each land use class (urban, agriculture or water). Moreover, we exclude outliers by first flagging all observations that are more than 2 interquartile distances below the first quartile or above the third quartile. We then conduct a generalized extreme studentized deviate (ESD) test to sequentially test if the flagged observations are outliers at the 90% confidence level (Rosner, 1983).

1.3.2 Material Wealth Index

In Yeh et al. (2020), the authors train a CNN with ResNet-18 architecture (He et al., 2016) to learn a local material wealth index from multi-spectral Landsat images. The idea behind this approach is that there are material manifestations of local household wealth in satellite images such as the shape, density and roof reflectance of the urban area, or the length, size and color of road infrastructure. The CNN is trained using satellite images and

FIGURE 1.5
EXAMPLES: AUTOMATIC SEGMENTATION RESULTS (2011-3)

(A) URBAN (B) CROPLAND

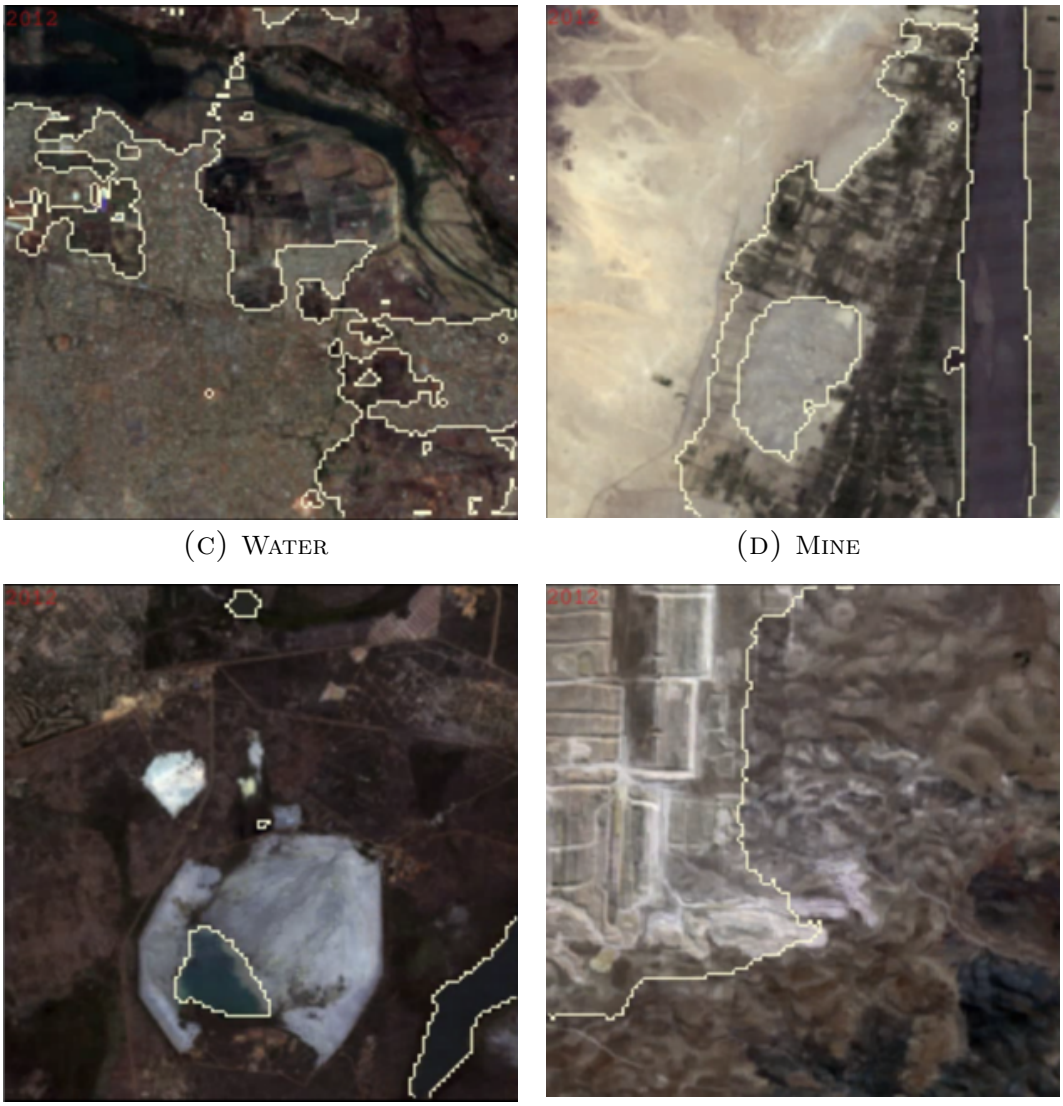

(C) WATER (D) MINE

Note: These examples show that our models are capable of identifying areas of interest.

corresponding asset wealth indices based on 43 Demographic and Health Surveys (DHS) conducted in 23 countries in Africa from 2009 to 2016. In the cross-section, Yeh et al. (2020) obtain an R^2 value on the relationship between the predicted and true (survey) values of the material wealth index on unseen data of around 0.65.

To assess the model performance in our context, we create a DHS-based asset wealth index for all tiles with DHS respondents in 2014-6 (N=1046) and compare it to the material wealth predictions based on Landsat imagery.[9] We obtain an R^2 value of 0.67, in coherence

[9]To protect the privacy of DHS respondents, their geo-coordinates are generally displaced by up to 2km in

with the results presented in Yeh et al. (2020). Figure 1.6 plots the wealth index values from the DHS survey and our predictions for tiles in our study area.

FIGURE 1.6
DHS-BASED WEALTH SCORE AND SATELLITE BASED PREDICTIONS

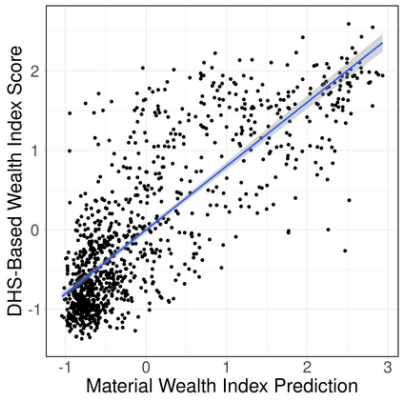

Note: The wealth index shows a strong correlation with DHS survey results on our study area of interest in 2014-6 (R^2 value of 0.67).

Moreover, the wealth index can be used to track changes in local material wealth over time. Living Standards Measurement Surveys (LSMS) surveys from the World Bank contain geolocalized household asset data for the same households at different points in time, providing an estimate of the change in wealth.[10]. There is a positive correlation (coefficient 0.35) between the predicted wealth change from the model and the estimated wealth change from LSMS panel data surveys, suggesting that the material wealth model is capable of detecting changes in wealth over time (Figure 1.A.4). This correlation is relatively weak for two main reasons. Firstly, the measurement error in the predictions of both, predicted wealth index changes and LSMS changes, weigh higher because the variation in changes over time is smaller (Yeh et al., 2020). Secondly, the LSMS wealth index relies on a different set of variables and is not representative of the local area. Qualitative examples show the coherence of the wealth index within a cross-section in Figure 1.A.3 and over time in Figure 1.A.2.

Saliency maps in Yeh et al. (2020) suggest that the model appears to weight urban areas, farmland, water bodies, and desert terrain when making predictions, in support of the view

urban areas and 5km in rural areas, but 1% of tiles are further displaced to up to 10km. We assign DHS clusters to tiles when their geocoordinates lie within the central region of tile at least 1km from the edges.

[10]For each pair of years, we compute the mean of the households that were surveyed in both years. We then run a PCA of these asset-differences across the 5 countries for which we have panel LSMS surveys (Ethiopia, Malawi, Nigeria, Tanzania, Uganda). The value of the first principle component is the household-level index of asset differences. Within each small geographical cluster, we then compute the mean household-level index of asset differences to get the cluster-level index of asset differences.

that the index actually reflects material wealth. However, the complexity of the model makes it difficult to guarantee that the image characteristics learned by the model are relevant explanatory variables for predicting material wealth. In our context, an additional complication is that the model may learn to detect the mine itself and for example learn that mines are associated with higher wealth. This would lead to spurious predictions of wealth and might lead to overestimating the impact of mining. In our analyses requiring the material wealth index, we exclude tiles with mines in order to alleviate this bias. Moreover, we exclude outliers using the same method as for land use and mine classes as described in Sec. 1.3.1.

1.4 Empirical Strategy

1.4.1 Descriptives

In the following, we provide descriptive evidence based on our main outcome, log urban landcover share. The graphs in Figure 1.7 plot the urban landcover share in periods 1 (1984-86) and 12 (2017-19) at distances from the mine in 5km intervals.[11] Mines with different activity statuses are displayed separately (see Figure 1.2 for an overview of different mine categories). There are three striking patterns. Firstly, there appears to be a gradient indicating that areas in closer proximity to the mine are more urban. Secondly, undeveloped deposits (in period 1:*Opening* and *Not Yet Opened*, in period 12: *Not Yet Opened*) tend to be at the bottom of the distribution. And thirdly, *Opening* and *Continuous* mines tend to gain between period 1 and 12, in particular relative to mines that are 'No Longer Active').

These patterns suggest that mining does have a positive economic impact on local communities in areas very close ($<$ 20km) to the mine. Further, the positive impact of mining does not seem to be only a one-off gain following the mine opening. Instead, mines appear to continuously induce growth during their activity. Lastly, after a mine closes the local area appears to stop growing at the same rate as before and loses relative to active mining areas.

Yet, it is hard to draw any definite causal conclusion based on such comparisons. If mine activity statuses were randomly distributed between mining areas, we would be able to infer

[11] The corresponding graphs for agricultural land use and material wealth index can be found in Figure 1.A.5.

FIGURE 1.7

CROSS-SECTIONAL COMPARISONS IN PERIODS 1 & 12

(A) URBAN LC (PERIOD 1) (B) URBAN LC (PERIOD 12)

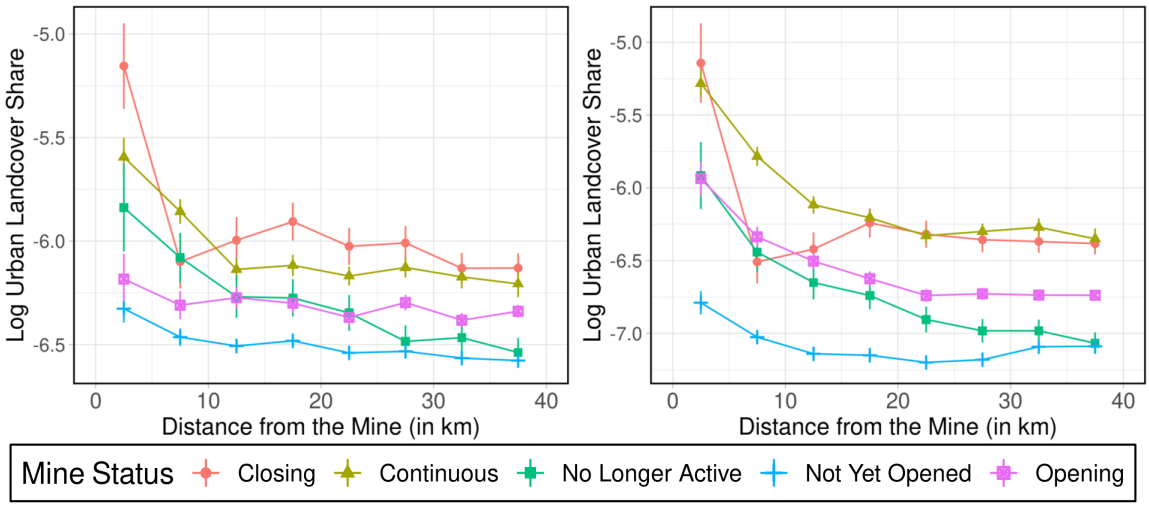

Note: Urban landcover tends to decrease with distance from the mine. Mines that were open in period 1 tend to have larger urban areas near the mines. Urban areas around Opening mines tend to grow between period 1 and 12, relative to mines that are No Longer Active or Not Yet Opened.

causal effects from such comparison. For example, we would be able to estimate the impact of mining in period 1 as the difference in the means between active mines (*Closing* and *Continuous*) and inactive deposits (*Not Yet Opened* and *Opening*). However, comparisons like these are likely distorted by underlying observable and unobservable characteristics of mining areas that determine both, the level of economic performance as well as the activity status itself. As an example, it could be that active mines are endogenously self-selected in countries and locations that are more developed and offer higher investor protection which would dramatically reduce the fixed costs and risks associated with opening a mine. Consequently, we probably overestimate the impact of mining in such comparisons because active mine areas would outperform inactive areas even in the absence of mining.

Similarly, while focussing on longitudinal variation and comparing one group of mines at different points in time (e.g. *Opening* areas before and after opening) would solve the problem of endogenous selection into treatment, such comparisons are distorted by time variant factors. Figure 1.8 displays the subset of *Opening* mines in period 1 (1984-86) when none of the mines had started operation, as well as in period 12 (2017-19) when all of the mines have started production. A naive interpretation of Figure 1.8a would suggest that mine openings increase urban growth close to the mine and reduce it a bit further away. Yet, such interpretations are fallacious, because they do not account for baseline changes over time that are unrelated to the treatment itself. There are many conditions in the country

or region that change simulataneously which could bias the estimate in either direction. As an example, country specific changes such as economic downturns, wars or increasing misgovernance would deteriorate the baseline economic condition and we would underestimate the impact of mining. Another potential bias might be the result of using Landsat 5 images in earlier periods and Landsat 7 and 8 images in later periods. The machine learning model might have higher or lower baseline probabilities for predicting urban areas with different satellite configurations due to slightly different technical configurations.

FIGURE 1.8

MINES IN THE CATEGORY *OPENING* IN PERIODS 1 & 12

(A) URBAN LC (UNADJUSTED) (B) URBAN LC (ADJUSTED)

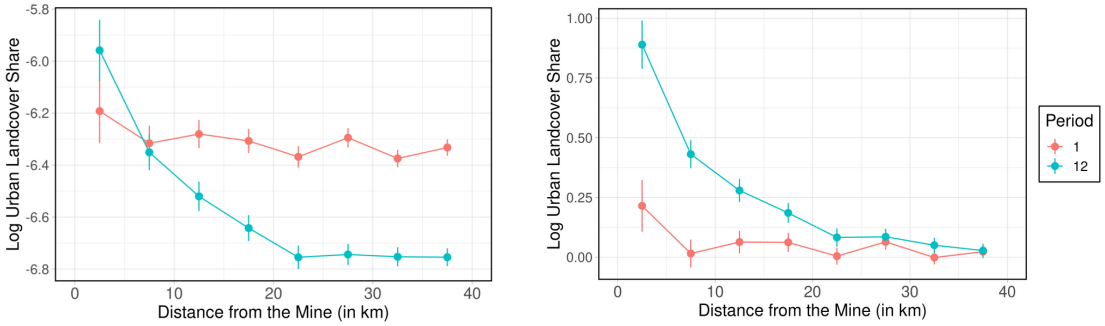

Note: At period 1, the mine is *Not Yet Opened* and there is little difference in wealth index between areas close to the mine and areas further from the mine. By period 12, these mines are active, and material wealth is much higher in areas close to the mine than during period 1.

As a result, the differences between both years are the sum of the treatment effect plus distortions. We can overcome this problem by using *Not Yet Opened* mine areas, that are likely equally affected by country specific trends and changing satellite configurations but not by mining, as a counterfactual to separate the changes induced by mining from other changes. When preprocessing the data of *Opening* mines by subtracting the average outcome in each country and period in the group of *Not Yet Opened* mines, we obtain Figure 1.8b. The figure indicates that while areas close ($<$ 20km) to the mine gained in terms of urban extent after the onset of mining, areas farther away remained unchanged.

This result supports the view that the onset of mining has indeed a positive economic impact on areas close to the mine, but there are still concerns about whether these are causal effects of mining. Most importantly, we cannot be sure that *Opening* and *Not Yet Opened* deposits would have evolved at the same rate in the absence of mine openings. Mine openings might self-select into areas where investors expect and economic upturn or where infrastructure investments occurred. For this reason, we require a more sophistic-

ated identification strategy and conduct additional tests to validate that *Not Yet Opened* deposits are valid counterfactuals.

1.4.2 Identification

In order to estimate the treatment effect of mine openings and closings, we elaborate on the adjustment procedure used in Figure 1.8b and exploit the panel dimension of our dataset using an event study approach with a control group. Over a time span of 15 years before until 15 years after treatment, we track the period-over-period development of the treatment group (e.g. *Opening* mines) against a control group (e.g. *Not Yet Opened* mine) to filter out unrelated shocks. While different mine categories are unlikely valid counterfactuals for each other at any given point in time, they might be similarly affected by spurious shocks. Hence, comparing one group to another within the same country and period can be useful in giving us an indication of how a treated group would have evolved in the absence of treatment and thereby facilitate the estimation of treatment effects. To test the underlying identifying assumption that both groups are equally affected by 'all other changes', we can test if both groups evolved similarly prior to the treatment onset i.e. test for parallel pre-trends.

$$Y_{i,t} = \sum_{\substack{t=-5,\\ t\neq 0}}^{5} \beta_t * D_t * Treat_i + b_{e,t} + b_{e,i} + \varepsilon_{i,t} \tag{1.1}$$

Equation 1.1 formulates the event study estimation equation with i referring to a tile, t to relative 3-year periods where period 1 corresponds to the first period of active mining, e to an event (all mine openings in a given country and period belong to one 'event'), D_t represents period dummies, $Treat_i$ is a dummy for the treatment group, $b_{e,t}$ are vectors of event × period fixed effects and $b_{e,i}$ are vectors of event × tile fixed effects. Standard errors are double-clustered at the tile and mine level. As the event study time is centered around the opening or closure event, we can keep only events for which we have observations over the entire interval of five periods (15 years) before and after the event.[12]

There is relatively new literature that focuses on the biases that might arise in DiD models, in particular in the presence of treatment heterogeneity and when treatment has a staggered onset over time (Chaisemartin, 2020; Sant'Anna and Zhao, 2020; Goodman-Bacon, 2021). For this reason, rather than relying on ordinary event studies and DiD estimations that

[12]This is necessary to ensure that in the presence of treatment heterogeneities between mines, time specific β-coefficients are not sensitive to distortions caused by a varying composition of the treatment group.

might be biased in either direction, we transform and stack our dataset such that there are no comparisons between the late treated units after treatment using earlier treated units that are also treated at the time as counterfactual (Cengiz et al., 2019). To do this, we define *events* that refer to all mines that open (or close) in one period within one country. In a second step, we assign each event as a control group all control units in the country. Finally, we stack all events together. As a result, each event has its own control group, which prevents implicit comparisons of late vs. early treated due to overlapping fixed effects. Since this also means that observations in the control group can be duplicated, in addition to clustering standard errors at the mine level, we also cluster standard errors on the tile level.

1.4.3 Choice of Counterfactual Groups

Partially Active mines are most suitable for running event studies as we can track their evolution before and after the treatment (mine opening or closure). Consequently, we use *Opening* mines as the treatment group to assess the impact of mining onset, and *Closing* mines as the treatment group to study what happens to mining areas after the closure of the mine.[13]

For *Opening* mines, the most suitable counterfactual group are deposits that are in the same country and also without an active mine in the beginning of our study period, but remained undeveloped throughout the study period (*Not Yet Opened*). The advantage for this comparison is that both groups are in the same state (inactive) in the beginning of the study period with one group getting treated thereafter (becomes active) while the other remains inactive. A potential pitfall of using *Not Yet Opened* areas as controls is that they might not exhibit parallel trends due to being endogenously selected into their group i.e. investors might have avoided certain deposits because the local economies are underdeveloped. Yet, there are many potential reasons that determine whether a deposit is mined at a certain point in time or not (time since discovery, support of local leadership, legal obstacles or global commodity demand) and many of which are unrelated to local economic performance. Our approach of dealing with this potential selection bias is twofold. First, we will make both groups more comparable by restricting our sample to recent discoveries at or after 1984. This means that deposits discovered a long time ago without an active mine, which we interpret as a signal of endogenous obstacles to

[13] Mines that are both *Opening & Closing* are also of interest, but we have dropped these observations as they usually change their status for short time intervals which makes it hard to separate the effects of different treatments of being open or closed.

mining, are dropped from our sample. Conducting balancing tests reveal that *Opening* and *Not Yet Opened* deposits are indeed similar regarding geographical characteristics and pre-treatment outcomes (see Figure 1.A.6). The only significant difference between both groups is that *Opening* mines exhibit around 12% more agricultural land use prior to treatment. More importantly, we will show that both groups exhibit parallel pre-trends when conducting the event study analysis.

As an alternative to using *Not Yet Opened* mines as a counterfactual, we could use mines that only start operating at a later point in time, and then compare the evolution of economic development in areas with active mines to that of areas that will host mines in the future - early vs. late (or future) treated (Goodman-Bacon, 2021). Using future treated units as controls for current treated areas is a relatively common approach in studies of the causal effect of place-based policy interventions - see for example (Busso et al., 2013). In our setting, the applicability of such an approach is supported by the technical and institutional context: There are various site specific obstacles that are unrelated to the level of development before the treatment but determine the precise timing of the mine opening. To support this argument empirically, we conducted pre-period balancing tests that indicate that within the group of *Opening* mines the startup period is not related to pre-treatment outcomes (see Figure 1.A.7). Therefore, future mine sites are likely to constitute a valid counterfactual for current treated areas, and allow us to control for confounding factors when investigating the local economic and environmental impact of mining. A drawback of using future treated units as controls is that they might actually already experience gains in anticipation of treatment during the development stage of the mine. This would lead to underestimating the magnitude of the effects. For this reason, we interpret estimates using future treated units as controls as lower bound estimates.

In addition to understanding how mine openings shape local communities, we aim to estimate the effect of the closure of the mine site - which is mostly triggered by random geological factors such as the exhaustion of the ore body. Since closing mines are experiencing two treatments - being an active mine and closing - it is hard to find a valid counterfactual. For this reason, rather than seeking to construct a counterfactual group, we will compare the evolution of local economic indicators in closing mine areas to *Not Yet Opened* and *Continuously* operating mines.[14] These comparisons are unlikely to exhibit parallel pre-trends prior to the closure, but it will nevertheless be informative to compare their relative

[14]The corresponding balancing tests that show how *Closing* mines compare to *Not Yet Opened* and *Continuously* active mines are in Figure 1.A.8 and Figure 1.A.9 respectively.

evolution over time and have a close look at potential trend breaks induced by the mine closure e.g. *Closing* mines might initially outperform their comparison groups prior to the closure but not after. Furthermore, the comparison with *Not Yet Opened* deposits will allow us to assess if former mining areas continue to outperform non-mining areas, evolve at the same rate or potentially even lose relative to non-mining areas which might indicate that former mining areas converge back to their pre-mined equilibrium. Methodologically, we will proceed in the same way as for the *Opening* mines and conduct an event study.

1.5 Results

1.5.1 Mine Openings

In this section, we undertake the empirical analysis on the impact of mine openings as set out in Section 1.4.2.

We begin with the event study and plot the β_t-coefficients based on Equation 1.1 over relative time t, where $t = 1$ is the first period of active mining (see Figure 1.9). Period 0 represents the period before the actual mining started and serves as the baseline period. Therefore, the β-coefficients reflect to what extend the treatment group has gained (if $\beta > 0$) or lost (if $\beta < 0$) relative to $t = 0$. The treatment group is restricted to either areas close to the mine (left) or areas far from the mine (right).

FIGURE 1.9
EVENT STUDY: *OPENING* VS. *NOT YET OPENED* MINES

(A) CLOSE: LOG URBAN AREA SHARE (B) FAR: LOG URBAN AREA SHARE

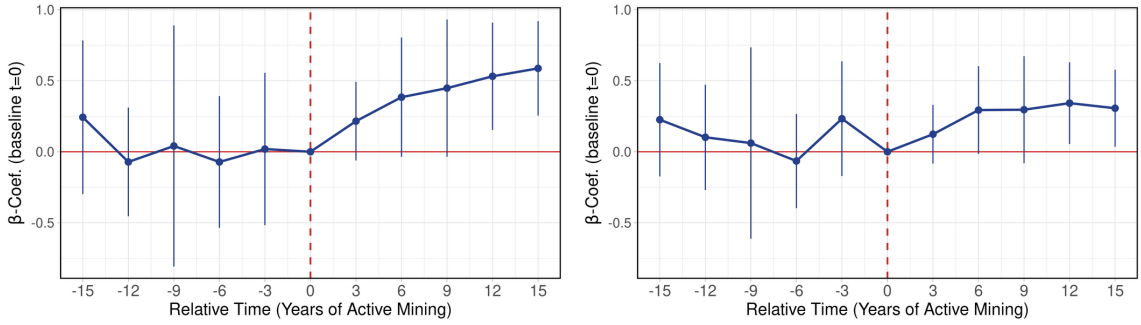

Note: In the years prior to the mine opening, areas near the mine (*Opening*) follow similar trends in terms of log of urban area share to areas with a mineral discovery but no active mine (*Not Yet Opened*). In the periods after the start of active mining, these two groups diverge, particularly in areas close to the mine (< 20km). A mine opening sets areas close to the mine on a high urban growth trajectory. A similar but less pronounced trend can be seen for areas further from the mine (20-40km).

One of the most important concerns of using *Not Yet Opened* (future/never-treated) areas

as control group for *Opening* mines is that they might already be diverging prior to treatment and therefore not constitute valid counterfactuals. However, the graphs in Figure 1.9 clearly indicate that prior to the onset of mining ($-5 \leq t \leq 0$) both groups evolve similarly. In contrast, as soon as active mining starts at $t = 1$, treated units in areas close to the mine significantly diverge from control areas. These gains are not just a one-off gain after the mine opening, but rather indicate that the local economy experiences continuous and persistent gains each period. Areas far away from the mine also diverge after the onset of mining but to a smaller extent. Furthermore, the fact that the divergence exactly coincides with the launch of the mine operation suggests that it is actually the mine activity that boosts the local economy rather than other structural changes. The magnitude of the effects is considerable. After 15 years, areas near *Opening* mines gain on average around 80%[15] in urban area extent relative to *Not Yet Opened* mine areas. Figure 1.A.10 show that the results are similar when conducting the event study with the sample based on the time frame from -6 years to 18 years.

$$Y_{i,p} = \beta * Treat_{i,p} + b_{e,p} + b_{e,i} + \varepsilon_{i,p} \tag{1.2}$$

Next, we estimate DiD regressions based on Equation 1.2. This is similar to Equation 1.1, but without centering time around the event or estimating multiple coefficients for relative time periods. Instead, we estimate the average treatment effect based on all units and time periods p after treatment using all observations instead of only those that are observed at least five periods before and after treatment. The results in Table 1.2 Panel A indicate that our results are robust with regards to using an ordinary DiD or stacked DiD model, as well as using *Not Yet Opened* deposits vs. *Late Treated* units (i.e. areas that get treated during our study period and serve as controls prior to their treatment onset). The results are in line with the event study results and indicate that mining does have a considerable positive impact on urban agglomeration and that these effects are primarily relevant for areas close to the mine. The results in Table 1.2 Panel B indicate that mining also increases the proportion of agricultural fields in areas close to the mine by around 19%. This finding is in line with previous studies that find that mining stimulates the local economy through backward linkages (Aragón and Rud, 2013).[16] Yet, these results are of a lower magnitude

[15] We compute the effect based on the coefficient in Figure 1.9 after 15 years (0.59) and retransforming it since the outcome is transformed using the natural logarithm: $exp(0.59) - 1 = 0.8$.

[16] Mining investments might also have a positive impact on other local firms via knowledge spillovers (see for example Ghebrihiwet (2019) for recent evidence on foreign direct investment (FDI) in South Africa's mining sector, or Abebe et al. (2022) for recent evidence on FDI in Ethiopia's manufacturing sector.) However, our dataset is not suitable to investigate such spillovers.

Table 1.2
DiD Regressions: Mine Openings

	Ordinary DiD		Stacked DiD			
	Opening vs. Not Yet Opened				Early vs. Late Treated	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Log Urban Landcover						
Treatment Dummy	0.093*	-	0.111*	-	0.068	-
	(0.051)		(0.063)		(0.064)	
Treatment × Close	-	0.223***	-	0.243***	-	0.173**
		(0.061)		(0.072)		(0.073)
Treatment × Far	-	0.048	-	0.066	-	0.031
		(0.050)		(0.062)		(0.063)
Panel B: Log Agriculture Landcover						
Treatment Dummy	0.127*	-	0.121	-	−0.004	-
	(0.067)		(0.084)		(0.083)	
Treatment × Close	-	0.182**	-	0.173*	-	0.011
		(0.078)		(0.094)		(0.091)
Treatment × Far	-	0.108	-	0.103	-	−0.009
		(0.066)		(0.082)		(0.083)
Panel C: Material Wealth Index (z-score)						
Treatment Dummy	0.028	-	0.028	-	0.026	-
	(0.031)		(0.036)		(0.048)	
Treatment × Close	-	0.105***	-	0.109***	-	0.097*
		(0.035)		(0.041)		(0.051)
Treatment × Far	-	0.001	-	0.001	-	0.001
		(0.030)		(0.035)		(0.048)
Country × Period FE	Yes	Yes	-	-	-	-
Tile FE	Yes	Yes	-	-	-	-
Event × Period FE	-	-	Yes	Yes	Yes	Yes
Tile × Event FE	-	-	Yes	Yes	Yes	Yes
Observations	400,239	400,239	1,297,985	1,297,985	626,035	626,035

Note: This table reports stacked DiD estimations based on Equation 1.2. The 'Treatment Dummy' (or 'Treatment') indicates if a tile's corresponding mine has started operating, it is always 0 for tiles in the control group. 'Close' and 'Far' in the interaction terms correspond to dummies indicating if a tile is within 20km from the mine or between 20km and 40km from the mine. Standard errors in parenthesis are clustered by mine in columns (1)-(2), and double-clustered by mine and tile in columns (3)-(6). *p<0.1; **p<0.05; ***p<0.01

and less robust when using early vs. late treated units. As discussed in Section 1.4.2, this is likely related to the fact that these estimates represent lower bounds as future treated control units might already experience some of the gains of mining during their development stage. In Table 1.2 Panel C, we document that mining increases the material wealth index by around 0.1 standard deviations in areas close to the mine.

1.5.2 Heterogeneities

In the next step, we investigate heterogeneities with regard to the impact of mining. Since previous research by Mamo et al. (2019), Pokorny et al. (2019) and Bazillier and Girard (2020) find that the mine size is important with regard to local economic outcomes, we investigate heterogeneities between small and large mines.[17] Further, since previous research at the macro level found evidence for the 'political resource curse' i.e. natural resources are only advantageous in countries with good institutions (Mehlum et al., 2006; Bhattacharyya and Hodler, 2010), we test if this relationship also applies on the micro level to mining areas. Based on the average Polity2 score[18] during our study period, we categorize countries as either being 'democratic' if they have a score greater than 0, or 'autocratic' for scores smaller than 0.

Table 1.3 presents the heterogeneity regression results based on Equation 1.2 for the sample of areas within 20km from the mine site, as this is where the effects are most relevant (see Table 1.2). The split between large vs. small and democratic vs. autocratic observations in the treatment group is relatively even, treatment group tiles around large mine represent 39% and tiles in democratic countries 43% of all treated tiles. As compared to small mines, large mines have a significantly stronger impact on urban growth and also lead to significant gains in areas further away from the mine. The characteristic of being a democracy appears even more important as the size and leads to significantly more urban area growth. Furthermore, in democracies there is no significant difference in the impact of different mine sizes. In contrast, in autocracies, small mines have a negative impact on urbanization in areas further away from the mine. The corresponding results for the agriculture land use and the material wealth index are similar and can be found in Table 1.A.1. In Table 1.A.2 we compare the characteristic of being a democracy to other measures of institutions as well as having a relatively high GDP per capita. We find

[17] The mine classification of 'small' or 'large' reflects the size of the mine operation and is based on a variety of indicators including the pre-mined resource, ore value, by-products and the type of mineral.

[18] More information about the Polity project by the Center for Systemic Peace can be found at: `www.systemicpeace.org/polityproject.html`.

TABLE 1.3
STACKED DID REGRESSIONS CLOSE TO THE MINE - HETEROGENEITIES

	Stacked DiD Regressions: *Log Urban Area Share*			
	Close		Far	
	(1)	(2)	(3)	(4)
Treatment Dummy	0.23***	−0.03	0.07	−0.21**
	(0.08)	(0.12)	(0.06)	(0.09)
Treat × Large Mine	-	0.29*	-	0.26**
		(0.17)		(0.12)
Treat × Democracy	-	0.53***	-	0.60***
		(0.16)		(0.13)
Treat × Large × Democracy	-	−0.46*	-	−0.40**
		(0.27)		(0.19)
Event x Tile FE	Yes	Yes	Yes	Yes
Event x Country x Period FE	Yes	Yes	Yes	Yes
Observations	1,172,561	1,172,561	1,254,498	1,254,498
Adj. R^2	0.76	0.76	0.76	0.77

Note: This table reports stacked DiD heterogeneity tests based on Equation 1.2. The 'Treatment Dummy' (or 'Treat') indicates if a tile's corresponding mine has started operating, it is always 0 for tiles in the control group. In Columns (1) and (2), we restrict the treatment group to tiles within 20km from the mine, and in Columns (3) and (4) to tiles between 20km and 40km from the mine. Standard errors in parenthesis are double-clustered by mine and tile. *p<0.1; **p<0.05; ***p<0.01

that democracy has a stronger effect than other institutional measures, and that GDP is irrelevant to the impact of mining.

1.5.3 Mechanism

In order to get a better understanding of why mine openings are mainly beneficial under democratic institutions, we now investigate a potential mechanism. Armand et al. (2020) conduct a field experiment showing that increased information and community participation helps to prevent conflict, and so it is plausible that democratic institutions have a similar effect. To test for the impact of mine openings on conflict, we combine our dataset with information about conflict events and locations from the Uppsala Conflict Data Program (UCDP) (Sundberg and Melander, 2013).[19] A big advantage of the UCDP is that it

[19]The main source of this database is global newswires reporting, but also other sources such as local news or NGO reports. The UCDP defines a conflict event as 'An incident where armed force was used by an organised actor against another organized actor, or against civilians, resulting in at least 1 direct death

provides a homogenous database of geolocalized conflicts from around the world and spans a relatively long time series starting in 1989 (which corresponds to period 2 of our dataset) until today. We combine the UCDP with our dataset by creating a tile-level dummy that indicates if a conflict event occurred in a tile during the respective period (or alternatively the total number of conflict events). We then use our stacked DiD model and estimate the impact of mine openings and closings separately for democracies and autocracies.[20]

In Table 1.4 column (1) and (3), we replicate the finding in the literature that mining fuels conflict (Berman et al., 2017).[21] When estimating regime specific effects for democracies and autocracies separately, we find that the onset of mining only leads to in autocracies, but not in democracies. The impact is also significant in areas further away but slightly smaller in magnitude. The value of the estimates indicate that the onset of mining increases the probability of conflict by 0.6% in areas close to the mine and 0.4% further away. In order to benchmark this coefficient, we compare it to the average probability of conflict prior to the onset of mining which is 0.16%. Hence, the increase in areas close to the mine corresponds to 3.75 times the baseline probability of conflict which signals a sizeable effect. This finding is novel in the literature. While Berman et al. (2017) undertake a similar exercise, they do not find significant effects associated with institutions.[22]

Our findings in Sections 1.5.1, 1.5.2 and 1.5.3 indicate that the onset of mining has a considerable positive impact on the local economy and that this is predominantly relevant for areas close to the mine. However, there are important heterogeneities when it comes to the gains from mine openings. Being a large mine and located in a democratic country are important factors for determining the economic gains for local communities. Further, mines that are small and in an autocratic setting might even have a negative impact on the local economy. One mechanism that explains this differential finding with regard to the institutional context is conflict. While mine openings lead to significant increases in conflict in autocracies, democratic institutions prevent a rise in conflict.

---

at a specific location and a specific date'. For more information visit: `https://ucdp.uu.se/downloads/ged/ged211.pdf`.

[20] The results are equivalent when alternatively using the model from Table 1.3 that also includes mine size as there is no significant relationship between mine size and conflict.

[21] All regressions in Table 1.4 are linear probability models. Using binomial regressions like logit instead is difficult to implement computationally given the high dimensional fixed effects and large number of observations.

[22] The reason for this is likely related to the fact that their dataset is much smaller, less granular and covers a much shorter time period between 1997-2010.

Table 1.4
Mining & Conflict

	Prob. of Conflict in Tile (baseline: 0.16%)			
	Close		Far	
	(1)	(2)	(3)	(4)
Treat	0.003**	-	0.002*	-
	(0.001)		(0.001)	
Treat × Democracy	-	−0.001	-	−0.001
		(0.001)		(0.001)
Treat × Autocracy	-	0.006**	-	0.004**
		(0.002)		(0.002)
Event x Tile FE	Yes	Yes	Yes	Yes
Event x Country x Period FE	Yes	Yes	Yes	Yes
Observations	1,078,351	1,078,351	1,153,489	1,153,489
Adj. R^2	0.12	0.12	0.12	0.12

Note: This table reports stacked DiD estimates for the impact of mine openings on conflict and is based on Equation 1.2. 'Treat' indicates if a tile's corresponding mine has started operating, it is always 0 for tiles in the control group. All models are linear probability models and the dependent variable indicates whether a tile experience any conflict during a given period. In columns (1) and (2), the treatment group is restricted to areas within 20km from the mine, and to areas between 20km and 40 km in Columns (3) and (4). Please note, that period 1 (1984-1986) is omitted from the sample as the Uppsala Conflict Data Program (UCDP) only starts in period 2 (1989). The baseline probability of conflict is 0.16% and refers to the average conflict probability in the treatment and control group prior to treatment onset. Standard errors in parenthesis are double-clustered by mine and tile.

*p<0.1; **p<0.05; ***p<0.01

1.5.4 Mine Closures

In this section, we investigate what happens to mine areas after the closure of the mine. In order to estimate the effects, we analyze the evolution of mines that had been active in period 1 followed by a closure during our study period. Since closing mines experience two treatments, firstly of being an active mine and secondly ceasing operation, it is much harder to find a suitable control group with similar pre-trends. For this reason, as discussed in Section 1.4.3, rather than using a control group indicating what would have happened if the mines had not closed, we use comparison groups to benchmark the performance of areas with closing mines.

We use an event study following Equation 1.1. We center relative time around the last period of active mining ($t = 0$), hence $t = 1$ is the first period without any mine activity. Since we are ultimately interested in whether the mine closure triggers a trend break, we need to compare the relative evolution between $t = -5$ and $t = 0$ (15 years prior to

shutdown vs. shutdown period) to the evolution between $t = 0$ and $t = 5$ (shutdown period vs. 15 years after shutdown). For this purpose, we add a dashed horizontal line with an intercept corresponding to the coefficient in period $t = 0$. We can then test if *Closing* mines gained relative prior to the closure by comparing the confidence interval of the estimate in period $t = 0$ to the horizontal line at $y = 0$. In the next step, we can test if *Closing* mines gained relative after the closure by comparing the confidence interval of the estimate at $t = 5$ to the horizontal line with the intercept of the coefficient at period $t = 0$.

FIGURE 1.10

EVENT STUDY: *CLOSING* VS. *NOT YET OPENED* MINES

(A) CLOSE: LOG URBAN AREA SHARE (B) FAR: LOG URBAN AREA SHARE

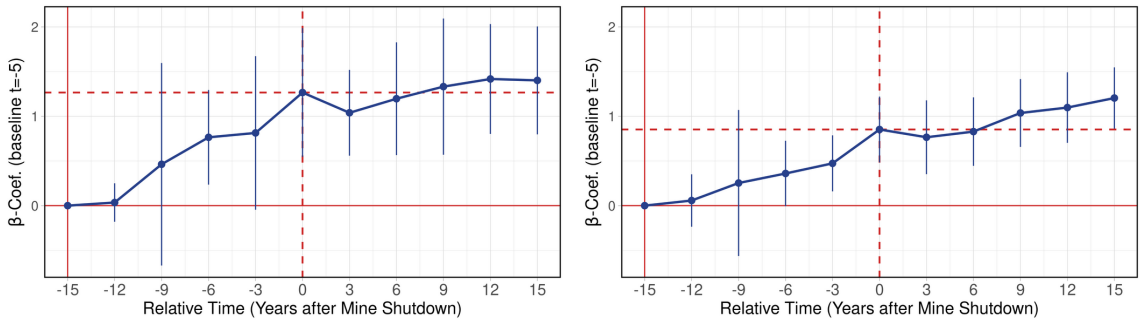

Note: Prior to closure, Closing mines have higher urban growth than Not Yet Opened mines. However, after closure, areas in proximity to Closing mines evolving at the same rate as non-mine areas, while areas further away only slightly outperform non-mine areas. This suggests that after the closure of the mine, mining areas grow no faster than non-mine areas.

The first comparison group are *Not Yet Opened* deposits. By comparing the relative evolution of *Closing* mines to mineral deposits that have not been mined yet, we can test if *Closing* mine areas continue outperforming non-mine areas, converge towards their pre-mined growth trajectory, or if they enter a state of economic decline back to the level they would have been at, had they never had an active mine in the first place.

Figure 1.10 plots the results when using *Not Yet Opened* areas as comparison group. Prior to the closure, *Closing* mines outperform *Not Yet Opened* mines in terms of urban growth. However, after the closure of the mines, we observe a strong trend break with areas in proximity to *Closing* mines evolving at the same rate as non-mine areas while areas further away only slightly outperform non-mine areas. The effects are very similar for agricultural land use and without any significant patterns for the material wealth index (see Figure 1.A.11). Overall, these results indicate that after the closure of the mine, mining areas grow no faster than non-mine areas.

In addition, we compare *Closing* mines to *Continuously* operating mines. The results in

Figure 1.A.12 show that *Closing* mines outperform *Continuous* mines before closure, but after closure, *Closing* mines have a slight but statistically insignificant downward trend.

These findings demonstrate that mining areas are unable to maintain elevated growth rates after the closure of the mine. Further, while areas close to the mines benefited more from the openings (see Section 1.5.1), they also suffer more when mines close.

1.6 Conclusion

In this study, we create a novel dataset based on large archives of satellite imagery covering 12% of the African continent over four decades to study the impact of mine openings on the development of local communities. Specifically, we follow the trajectory of urban and agricultural growth as well as material wealth in the villages and cities surrounding mines throughout the period of active mining as well as after the mines' closure.

Our results indicate that mineral mine operations, especially large operations, have the potential to give a considerable boost to local economic growth in areas surrounding the mine. A part of these gains are likely to be indirect gains due to backward linkages as is reflected by increased agricultural activities around the mine. Yet, our study also points out that there are important caveats with regard to the positive impact of mining.

Firstly, accelerated growth rates in mining areas are only temporary and are not sustained beyond the closure of the mine. Secondly, our analysis suggests that the benefits for local communities are mainly advantageous in democracies but not in autocracies. We show that one mechanism through which mining areas in democratic countries, as opposed to autocratic countries, avoid the resource curse is by avoiding conflict. There might be other relevant mechanisms that link mining and institutions on the local level. One example would be a fiscal channel through which the windfall gains from resource extraction might be less redistributed to the population when institutions are extractive. Another potential channel is that local corruption might rise under bad institutions with insufficient checks and balances. Furthermore, under poor institutional framework conditions there might be reduced incentives for local public officials to negotiate and enforce regulations such as local procurement rules, employment opportunities for local residents and other forms of resource governance that benefit local communities.

We conclude that not all mine operations benefit local communities, and mining could even

be net negative for economic growth, in particular in areas with poor institutions. For this reason, our study underlines the importance of paying special attention to the institutional framework conditions when considering mine openings. Furthermore, in order to achieve sustained growth in areas that benefit from mining, policy makers need to develop location-specific strategies for economic transformation during operation and after the closure of the mine.

1.A Appendix

1.A.1 Data Sources and Description

Mineral deposits

We purchase data on mining deposits from MinEx Consulting. This dataset includes the geolocalization, type, size and dates of discovery and activity of 1,658 mineral deposits in 47 African countries.

Land Use Label Data Sources

1. **Agricultural Areas**
   - Nasa Global Food Security-support Analysis Data (GFSAD), Croplands Africa 2015, 30m resolution, available at: `https://lpdaac.usgs.gov/products/gfsad30afcev001/`
2. **Urban Areas**
   - Global Human Settlement Layer (GHSL) 2015 by the European Commission, available at: `https://ghsl.jrc.ec.europa.eu/download.php?ds=buGHS_BUILT_LDSMT_GLOBE_R2018A_3857_30_V2_0`
   - CCI Land Cover S2 Prototype Africa 2016, resolution: 20m, available at: `https://2016africalandcover20m.esrin.esa.int/`)
   - Facebook (Meta) Population Map 2015, 30m resolution, available at: `https://data.humdata.org/dataset/highresolutionpopulationdensitymaps`, does not cover South Sudan, Sudan, Somalia and Ethiopia
3. **Water Bodies**
   - CCI Land Cover S2 Prototype Africa 2016, resolution: 20m, available at: `https://2016africalandcover20m.esrin.esa.int/`)

Geocovariates Sources

– **Crop Caloric Index**: A measure of agricultural suitability containing the potential agricultural caloric output per year and hectar (excluding zero yields) based on Galor and Özak (2016), available at: `https://ozak.github.io/Caloric-Suitability-Index/`.

– **Elevation**: NASA Shuttle Radar Topography Mission (SRTM) with 30m res-

olution, available at:https://www2.jpl.nasa.gov/srtm/.

- **African Cities**: Data from OECD Sahel and West Africa Club in collaboration with e-geopolis.org, available at: https://africapolis.org/data.

- **Country Borders & Coastline**: GADM (Version 3.6), available at: https://gadm.org/data.html.

- **Climate Data:** Annual precipitation (in mm), annual mean temperature, minimum temperature in the coldest month and maximum temperature in the warmest month (all in °C) based on Karger et al. (2017) were obtained from http://chelsa-climate.org/downloads/.

- **Ruggedness**: A measure for terrain ruggedness measured in degree of slope with an initial resolution of (20×20 arcseconds) based on Nunn and Puga (2012) was obtained from: https://diegopuga.org/data/rugged/.

1.A.2 Additional Figures

FIGURE 1.A.1
HIGH-RESOLUTION IMAGES CORRESPONDING TO FIGURE 1.1B

(A) 2020 HIGH-RESOLUTION GOOGLE EARTH IMAGE OF CITY CORRESPONDING TO FIGURE 1.1B

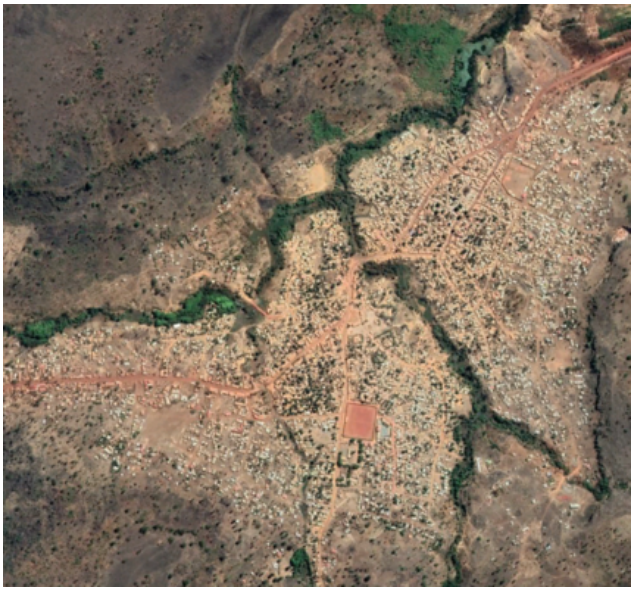

(B) 2020 HIGH-RESOLUTION GOOGLE EARTH IMAGE OF MINE CORRESPONDING TO FIGURE 1.1B

Note: Current high-resolution satellite images can be used for better visualization of the 30m Landsat images in Figure 1.1B. The fact that our model excluded the area that looks like a football field in image (A) from the urban extent (see Figure Figure 1.1B) illustrates the ability of the model to identify the urban extent at a very precise and granular level.

Figure 1.A.2

Examples: Low to High Wealth over Time

(A) 1991: Pre-mine, low wealth

(B) 2018: Post-mine, high wealth

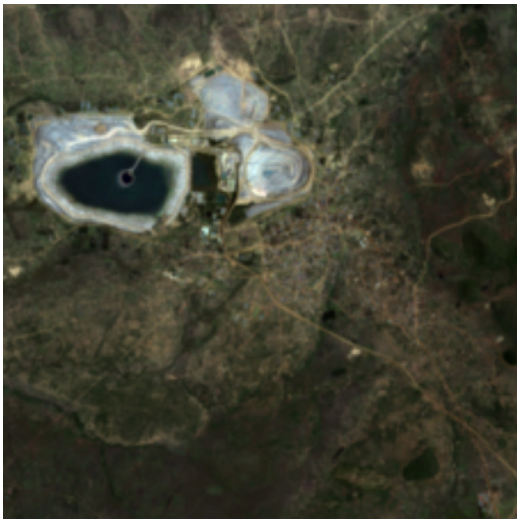

(C) 1988: Low wealth index

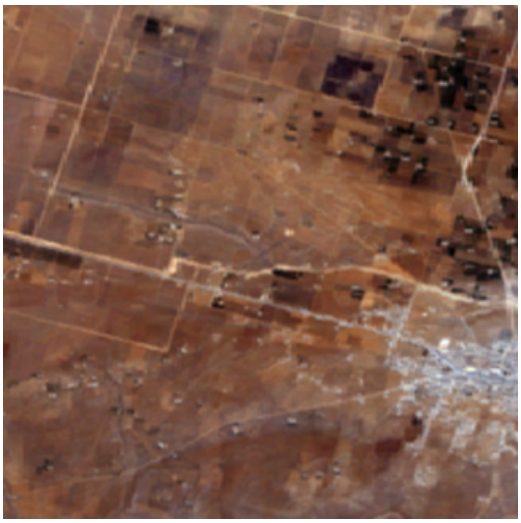

(D) 2018: High wealth index

Note: The growth in material wealth index over time seems to be coherent with our intuition for numerous examples from our dataset. We see the index increasing over time when towns and cities become more developed. In our analyses, we omit tiles with mines in order to remove any biases related to the presence of the mine in the image.

Figure 1.A.3

Landsat and Corresponding High-Resolution Images

(A) Landsat: Poor

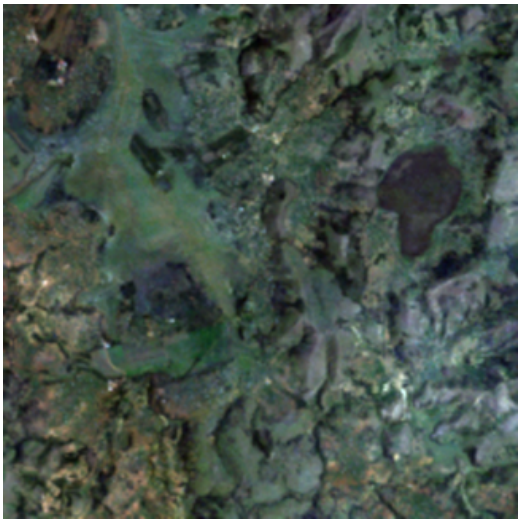

(B) Landsat: Rich

(C) Google Earth: Poor

(D) Google Earth: Rich

Note: High-resolution images from Google Earth provide a sanity check of the material wealth predictions using medium resolution images from Landsat.

Figure 1.A.4
Wealth Change: LSMS and Material Wealth Index

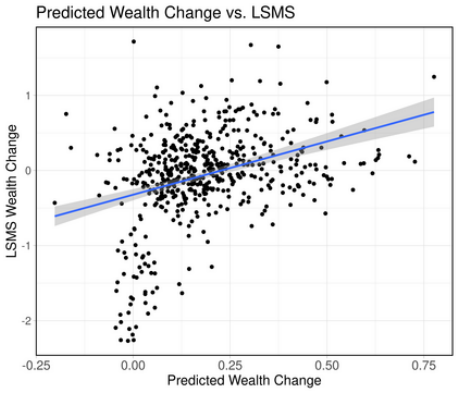

Note: There is a positive correlation (coefficient 0.35) between the predicted wealth change and the estimated wealth change from Living Standards Measurement Study (LSMS) panel data surveys by the World Bank, suggesting that the material wealth model is capable of detecting changes in wealth over time.

Figure 1.A.5

Further Cross-Sectional Comparisons in Periods 1 & 12

(a) Wealth Index (Period 1)

(b) Wealth Index (Period 12)

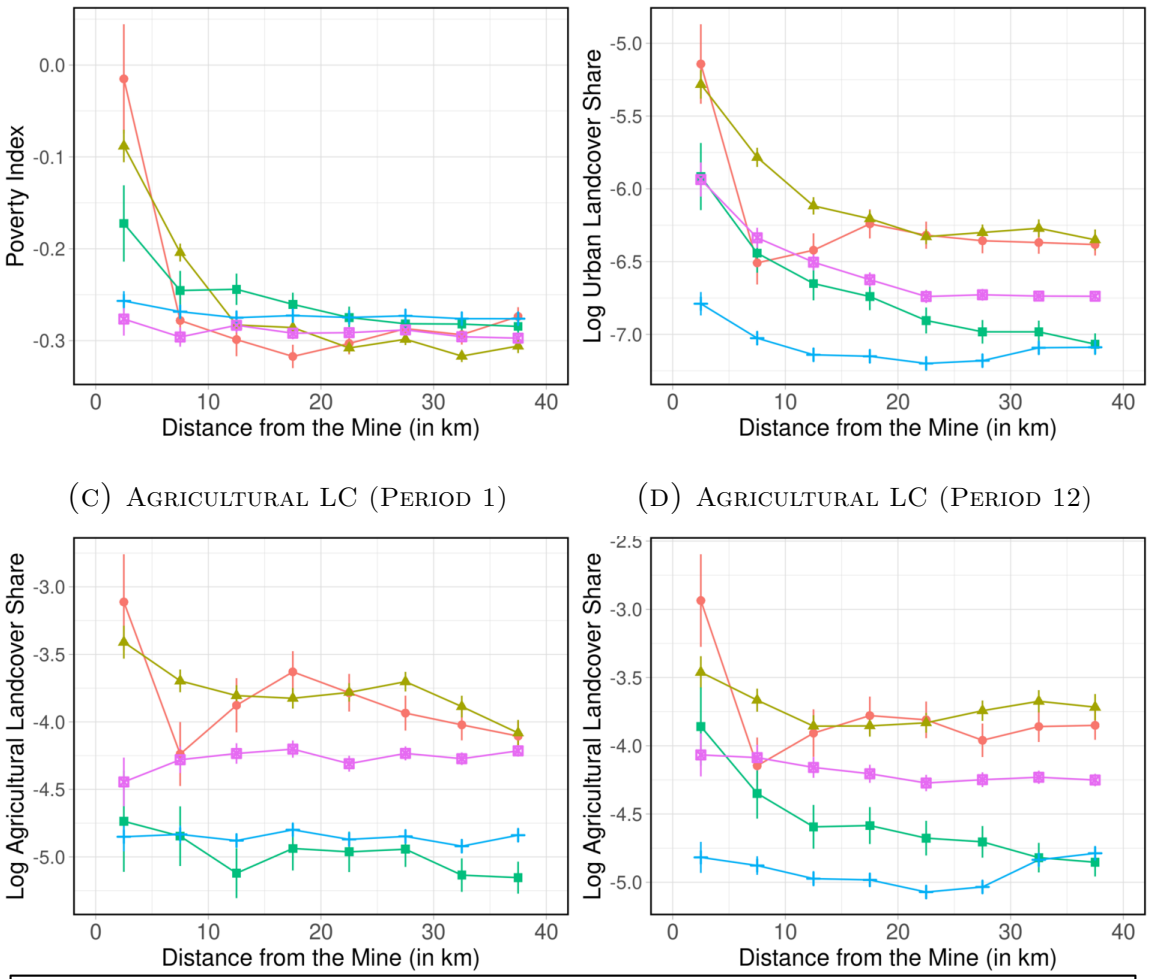

(c) Agricultural LC (Period 1)

(d) Agricultural LC (Period 12)

Note: The wealth index tends to decrease with distance from the mine. Mines that were open in period 1 tend to have a higher wealth index near the mines. The wealth index around Opening mines tend to grow between period 1 and 12, relative to mines that are No Longer Active or Not Yet Opened. The wealth index around Opening mines tends to grow between period 1 and 12, relative to mines that are No Longer Active or Not Yet Opened. The agricultural area around Opening mines also tends to grow between period 1 and 12, relative to mines that are No Longer Active or Not Yet Opened, but this trend is less pronounced.

1.A.3 Balancing Tests

FIGURE 1.A.6
BALANCING GRAPHS: OPENING VS. NOT YET OPENED

(A) TIME INVARIANT COVARIATES

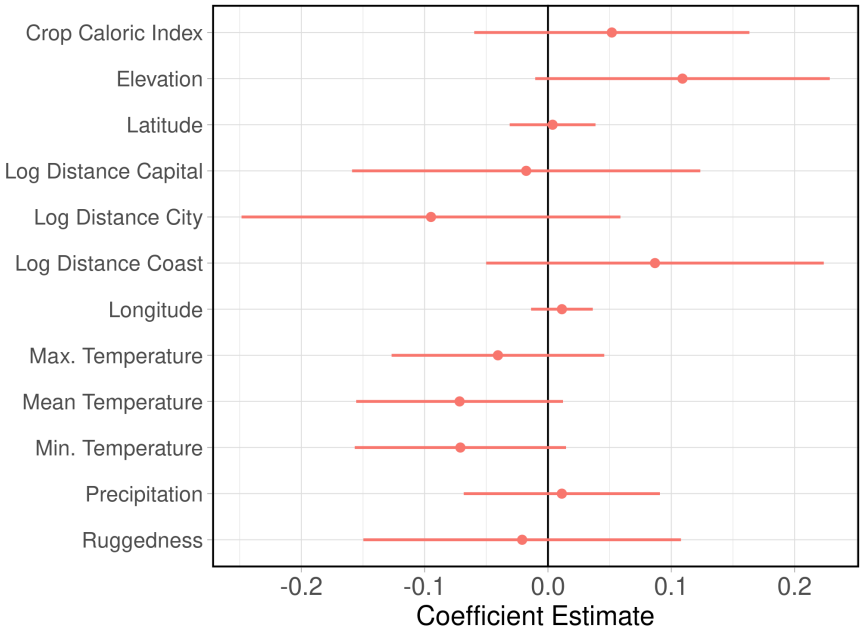

(B) OUTCOMES IN PERIOD 1

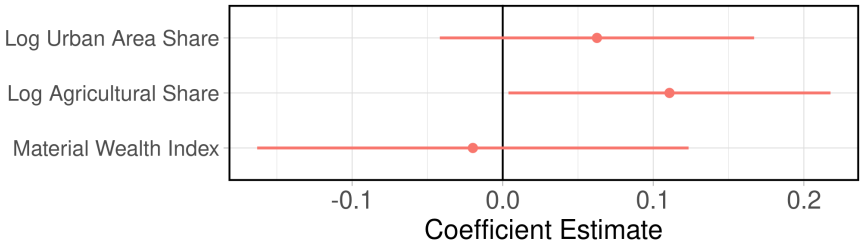

Note: These figures report balancing test between *Opening* and *Not Yet Opened* deposits based on estimating: $Y_i = \beta\ OpeningDummy_i + b_c + \varepsilon_i$, where b_c are country fixed effects and with SEs clustered at the mine level.

Figure 1.A.7
Balancing Graphs: Early vs. Late Opening

(a) Time Invariant Covariates

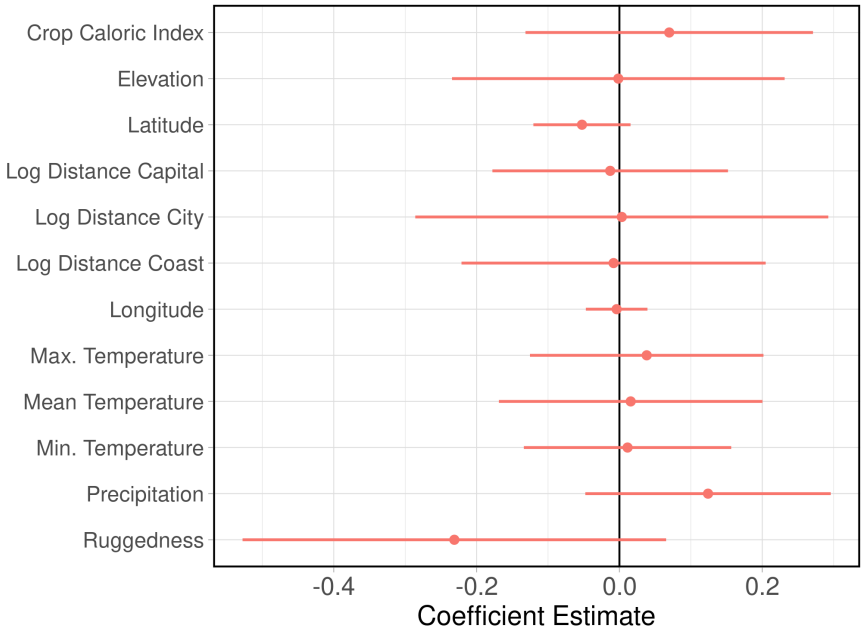

(b) Outcomes in Period 1

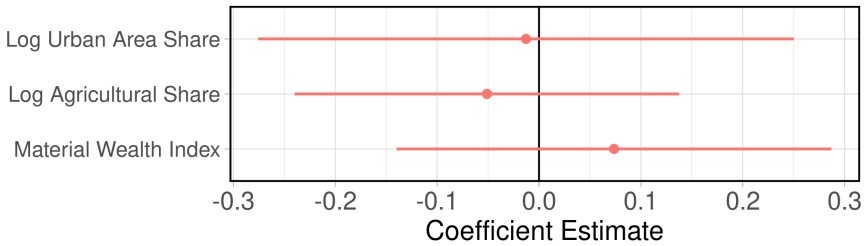

Note: These figures report balancing test between early and late *Opening* mines based on estimating: $Y_i = \beta\ log(OpeningYear)_i + b_c + \varepsilon_i$, where b_c are country fixed effects and with SEs clustered at the mine level.

Figure 1.A.8
Balancing Graphs: Closing vs. Not Yet Opened

(a) Balancing: Time Invariant Covariates

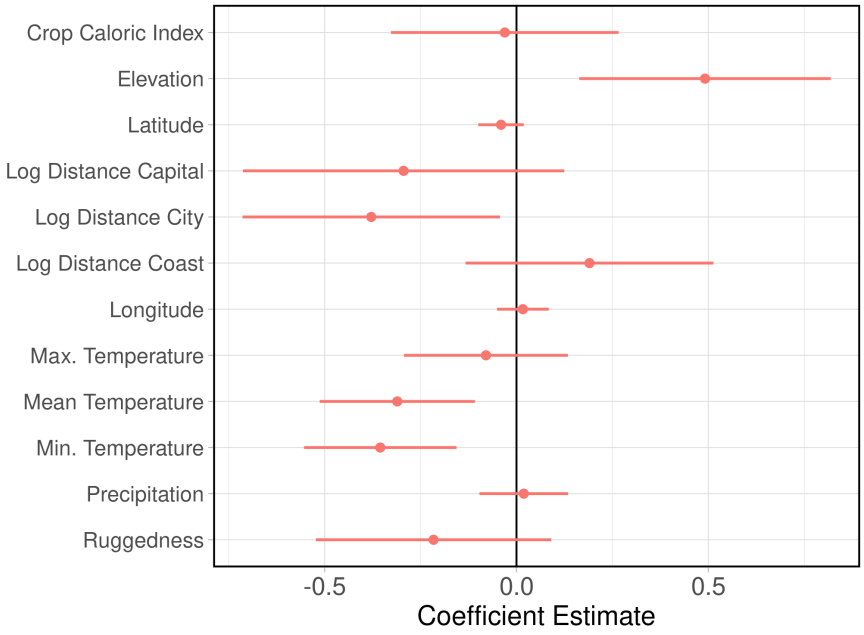

(b) Balancing: Outcomes in Period 1

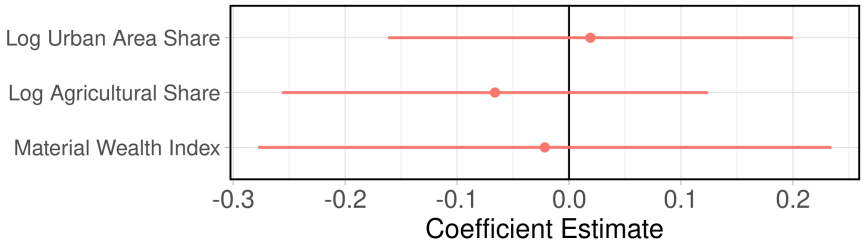

Note: These figures report balancing test between *Closing* and *Not Yet Opened* deposits based on estimating: $Y_i = \beta\ ClosingDummy_i + b_c + \varepsilon_i$, where b_c are country fixed effects and with SEs clustered at the mine level.

Figure 1.A.9
Balancing Graphs: Closing vs. Continuous

(a) Balancing: Time Invariant Covariates

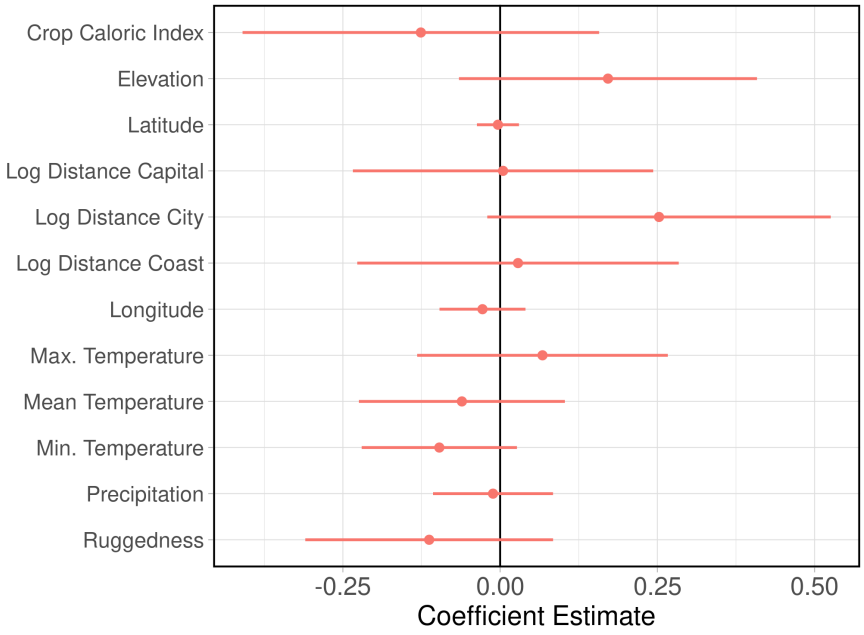

(b) Balancing: Outcomes in Period 1

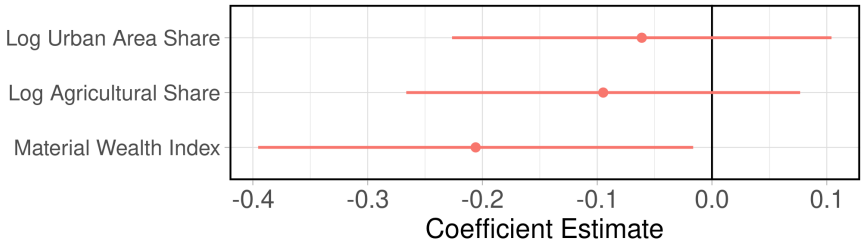

Note: These figures report balancing test between *Closing* and *Continuous* deposits based on estimating: $Y_i = \beta\ ClosingDummy_i + b_c + \varepsilon_i$, where b_c are country fixed effects and with SEs clustered at the mine level.

1.A.4 Additional Results

FIGURE 1.A.10
EVENT STUDY: OPENING VS. NOT YET OPENED (ROBUSTNESS)
(A) CLOSE: LOG URBAN AREA SHARE (B) FAR: LOG URBAN AREA SHARE

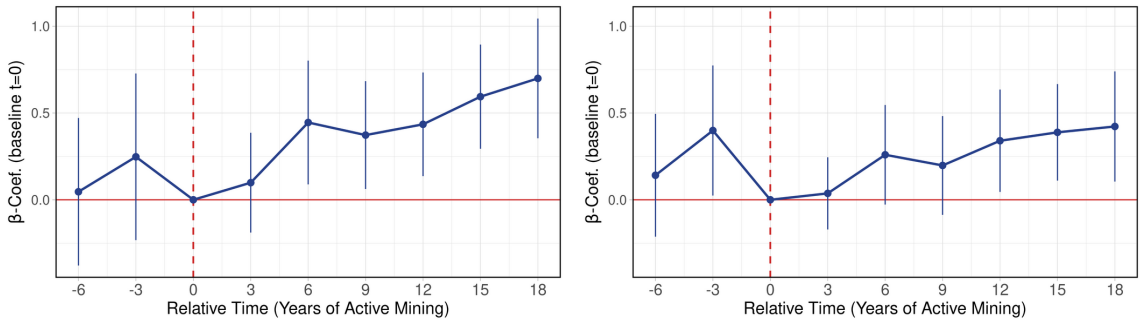

Note: See Figure 1.9 for similar results but using a shifted relative time scale.

TABLE 1.A.1
STACKED DID REGRESSIONS CLOSE TO THE MINE - HETEROGENEITIES (CONT.)

	Stacked DiD Regressions							
	Log Agricultural Area Share				Material Wealth Index (z-score)			
	Close		Far		Close		Far	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment Dummy	0.15	−0.16	0.10	−0.25**	0.11***	0.07	0.00	−0.01
	(0.10)	(0.15)	(0.08)	(0.13)	(0.04)	(0.07)	(0.04)	(0.06)
Treat × Large Mine	-	0.32*	-	0.29**	-	0.20**	-	0.08
		(0.18)		(0.14)		(0.10)		(0.07)
Treat × Democracy	-	0.60**	-	0.68***	-	−0.09	-	−0.06
		(0.23)		(0.20)		(0.10)		(0.08)
Treat × Large × Democracy	-	−0.36	-	−0.27	-	0.02	-	0.01
		(0.33)		(0.25)		(0.14)		(0.11)
Event x Tile FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Event x Country x Period FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,172,561	1,172,561	1,254,498	1,254,498	1,172,561	1,172,561	1,254,498	1,254,498
Adj. R^2	0.76	0.76	0.76	0.76	0.68	0.68	0.68	0.68

Note: This table reports stacked DiD heterogeneity tests based on Equation 1.2. The table corresponds to Table 1.3 and reports the results for the log of the agricultural landcover and the z-score of the material wealth index as dependent variables. The 'Treatment Dummy' (or 'Treat') indicates if a tile's corresponding mine has started operating, it is always 0 for tiles in the control group. In Columns (1), (2), (5) and (6), we restrict the treatment group to tiles within 20km from the mine, and in Columns (3), (4), (7) and (8) to tiles between 20km and 40km from the mine. Standard errors in parenthesis are double-clustered by mine and tile. *p<0.1; **p<0.05; ***p<0.01

Table 1.A.2
Heterogeneities: Country Characteristics

	Dependent Variable: Log Urban Landcover Share				
	Democracy	Voice & Account.	Rule of Law	Decentralization	GDP
Treat	0.08	0.04	0.09	0.20	0.24**
	(0.10)	(0.13)	(0.11)	(0.13)	(0.09)
Treat × Democracy	0.36**	-	-	-	-
	(0.15)				
Treat × High Participation	-	0.33**	-	-	-
		(0.16)			
Treat × High ROL	-	-	0.27*	-	-
			(0.15)		
Treat × Decentralized	-	-	-	0.07	-
				(0.16)	
Treat × High GDP	-	-	-	-	−0.03
					(0.16)
Event x Tile FE	Yes	Yes	Yes	Yes	Yes
Event x Country x Period FE	Yes	Yes	Yes	Yes	Yes
Observations	1,172,561	1,172,561	1,172,561	1,083,140	1,172,561
Adj. R^2	0.76	0.76	0.76	0.76	0.76

Note: This table reports stacked DiD heterogeneity tests for the sample of treatment tiles within 20km from the mine and is based on Equation 1.2. The 'Treatment Dummy' (or 'Treat') indicates if a tile's corresponding mine has started operating, it is always 0 for tiles in the control group. In each column, the treatment dummy is interacted with a group dummy based on country specific institutional characteristics. 'Democracy' is based on the average Polity2 score during the study period being larger than 0, 'Voice & Accountability' and 'Rule of Law' are based on the being above or below the median of the World Bank's Worldwide Governance Indicators (WGI), 'Decentralization' is based on the median decentralization index by Thomas Bijl and J. Vernon Henderson (LSE processed) and 'GDP' based on the median of the World Bank's GDP per capita estimate (in PPP) for the year 2020. Standard errors in parenthesis are double-clustered by mine and tile. *p<0.1; **p<0.05; ***p<0.01

Figure 1.A.11

Event Study: *Closing* vs. *Not Yet Opened* Mines

(a) Close: Log Agricultural Share

(b) Far: Log Agricultural Share

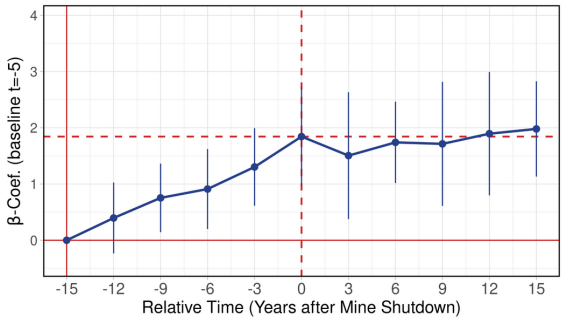

(c) Close: Material Wealth Index (z-score)

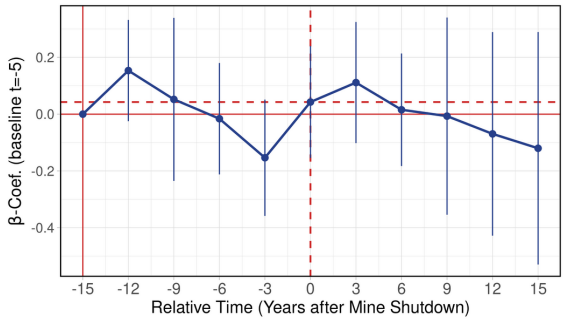

(d) Far: Material Wealth Index (z-score)

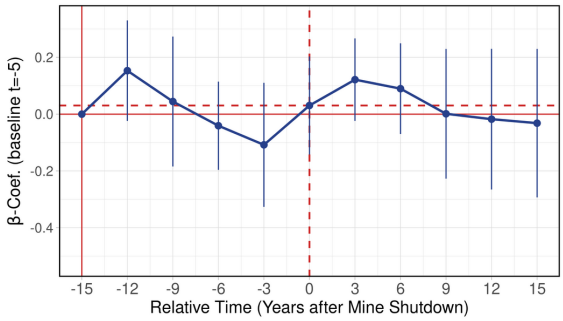

Note: Prior to closure, Closing mines have higher agricultural growth than Not Yet Opened mines. However, after closure, areas in proximity to Closing mines evolving at the same rate as non-mine areas, while areas further away only slightly outperform non-mine areas. This suggests that after the closure of the mine, agriculture around mining areas grow no faster than around non-mine areas. We do not observe any significant effects for the material wealth index.

Figure 1.A.12

Event Study: *Closing* vs. *Continuous* Mines

(a) Close: Log Urban Area Share

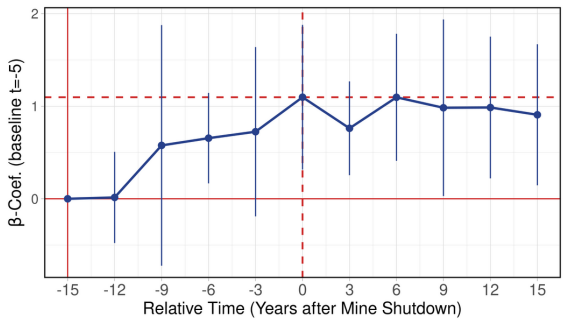

(b) Far: Log Urban Area Share

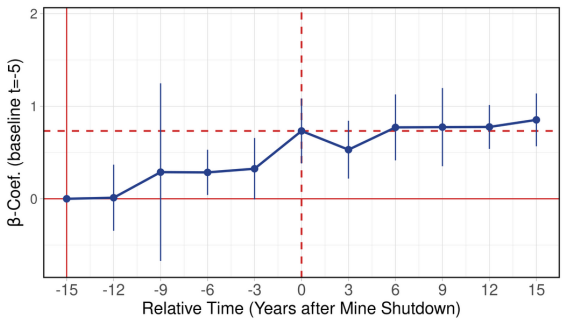

(c) Close: Log Agricultural Share

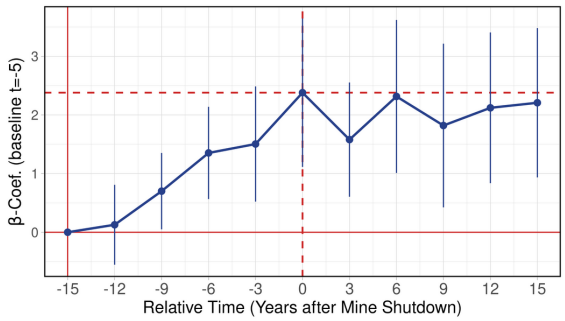

(d) Far: Log Agricultural Share

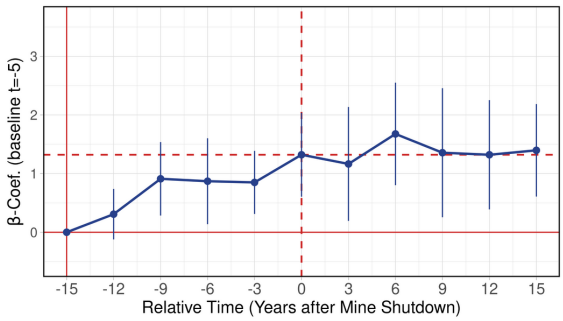

(e) Close: Material Wealth Index (z-score)

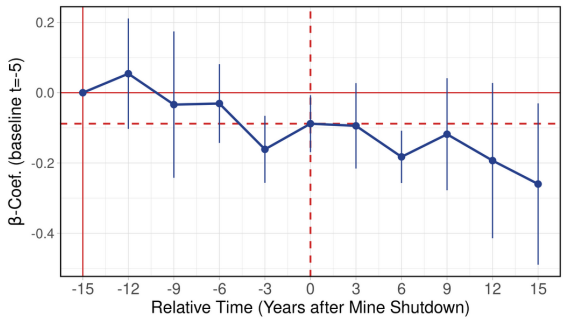

(f) Far: Material Wealth Index (z-score)

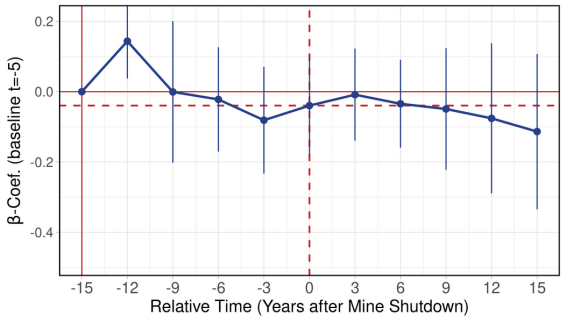

Note: Closing mines outperform Continuous mines before closure. There even seems to be a negative but not statistically significant trend in the urban area share and the material wealth index for areas close to a Closing mine after its closure.

Chapter 2

Accountability Failure in Isolated Areas: The Cost of Remoteness from the Capital City

The spatial distribution of economic activity in Sub-Saharan Africa is shaped by large differences in standards of living across regions (International Monetary Fund, 2015). Accordingly, gaining a better understanding of the underlying mechanisms that cause and maintain these spatial disparities is key for designing policies that could lift millions of individuals in the least developed regions out of poverty. Yet, the ongoing research on the subject has been largely descriptive (Odusola et al., 2017) rather than seeking to reveal underlying patterns and mechanisms. Only a limited number of scholars (see for example Kanbur and Venables (2005), Hodler and Raschky (2014) or Addison et al. (2017)) have examined spatial patterns and causes of inequality that go beyond the 'urban-rural bias' thematically (Lipton, 1977; Bates, 1981; Young, 2013; Lagakos, 2020).

For historical reasons, most African capital cities are located either at or close to the coast rather than in a central location which is why large parts of the population live far away from the capital city.[1] Previous research has pointed out that isolated capital cities impose important adverse effects on statewide outcomes such as aggregate corruption levels, conflict and quality of governance (Campante and Do, 2014; Campante et al., 2019). Yet, these studies look at capital city isolation as an aggregated state characteristics and do not investigate if there are spatial heterogeneities i.e. how locations farther from the capital perform economically relatively to their counterparts close to the capital.[2] Other research on the role of the capital city emphasizes that the ability of African states to broadcast power and impose institutions is restricted beyond the capital city (Herbst, 2000; Michalopoulos and Papaioannou, 2014). What is still unclear in this strand of research is what the

[1] For European colonizers who targeted the extraction of resources but had little access to the hinterlands, coastal trading points constituted suitable locations for colonial headquarters. Over the course of the colonial period, these administrative centers flourished and the majority of them subsequently persisted as post-colonial national capital cities in modern African states.

[2] Campante and Do (2014) include a few descriptives on the individual level and show that citizens further away from state capitals are less engaged with state politics.

reduced state presence means for the economic performance in areas far from the capital[3], as well as what the underlying mechanisms are that reduce the presence of the state in remote areas.

In this study, we suggest that the answer to this question is connected to the literature on accountability and public goods provision that emphasizes the importance of information about government actions to incentivize the attention of politicians (Besley and Burgess, 2002; Strömberg, 2004; Guriev et al., 2020). We argue that geographical isolation leads to information frictions which impair accountability mechanisms and leave the political leadership with reduced incentives to invest into remote areas. To support our claims empirically, we show that there is a significant and causal negative impact of distance from the capital city on economic performance. We further document a significant drop in public goods in remote areas, show that isolated citizens follow the news less frequently, have an overly positive view of the government and penalize their leaders less for misgovernance. Lastly, we show that alternative explanations such as market access or conflict are unlikely to be relevant with regard to the observed patterns.

The core challenge when seeking to identify the impact of isolation from the capital city is that capital cities are not randomly located in space. There are numerous geographical characteristics, most notably isolation from the coast, that are simultaneously correlated with isolation from the capital city and economic performance and thus confound ordinary regressions. We overcome this obstacle by applying a boundary discontinuity design (BDD) across national boundaries and comparing places with otherwise similar geographical features but varying distances to their respective capital city. Moving across the boundary might not constitute a valid counterfactual if the switch between countries coincides with other variables that might themselves be linked with development such as ethnicity and culture. We therefore restrict our analysis to boundary segments that divide pre-colonial ethnic homelands. Figure 2.1 illustrates the intuition of our identification strategy using the example of the Nyanja ethnic homeland that is divided into two adjacent countries, Malawi and Mozambique. Our identification strategy exploits the jump in distance from the capital city at the border to explain differences in economic outcomes.

Using remote sensing data on nightlights as a proxy for economic activity, our results indicate that a one percentage point increase in distance from the capital city reduces the

[3]Michalopoulos and Papaioannou (2014) establish that there is a reduced state presence in isolated areas using cross-border comparisons for areas with similar distance to the capital city (close vs. close and far vs. far) between countries with good and bad institutions. In this paper, we focus on the relative performance of remote areas and compare locations close vs. far from the capital while keeping institutions constant.

Figure 2.1

Ethnic Homeland Partitioned between Malawi and Mozambique

(a) Lit vs. Not-Lit Pixels

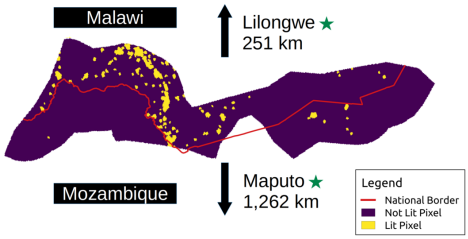

(b) Mean Light Density

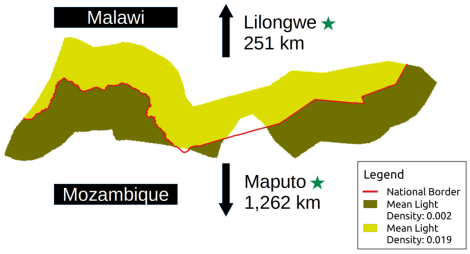

Note: This figure illustrates Nyanja ethnic homeland that is partitioned between Malawi and Mozambique. We observe that the Malawian side contains more lit pixels and features an overall higher average light density indicating that the Malawian side is more advanced economically. The Mozambiquean side is located 1,262 km from its capital city Maputo, while the Malawian side is situated only 251 km from its capital city Lilongwe. Our identification strategy aims at exploiting this jump in distance from the capital city to explain differences in economic outcomes.

probability of a pixel to be lit by three percentage points (the average probability to be lit is 2.5%). Alternatively, using DHS survey data we find that this corresponds to a reduction in household wealth of 3.5 percentiles of the national wealth distribution. To confirm that this effect is actually driven by isolation from the political center and not by remoteness from a major city, we show that the effects are unique to the capital city and do not apply to a set of placebo cities. Having established this reduced form relationship, we turn to the question of how isolation from the capital city impacts economic performance. We investigate three plausible mechanisms that are suggested by the literature: public goods provision, market access and conflict. Our analysis indicates that the latter two are unlikely to be relevant with regard to the effects under scrutiny. In contrast, we document a significant and strong causal impact of distance from the capital city on the level and quality of public goods provision suggesting that it plays and important role behind the observed patterns.

We then explore two potential explanations for reduced public goods in isolated areas: (i) isolated areas are less represented in national politics or (ii) isolated citizens are less able to hold political leaders accountable for providing government services i.e. dysfunctional accountability mechanisms. While political leaders are less likely to come from an isolated region, those regions are more likely to participate in the ruling government coalition. This finding contradicts the view that remote areas are systematically excluded from power. Yet, we find clear evidence that isolated citizens are not able to incentivize the incumbent government for the provision of public services to the same extent as those close to the capital. Despite being served with a lower level of public goods, people in isolated areas exhibit a higher level of trust in their national political leadership, evaluate their performance better and are less likely to believe that their leaders are corrupt. We argue that the positive

view of the government is associated to the fact that isolated citizens consume the news less frequently and therefore have less insights into government actions. To support this hypothesis, we show that isolated citizens are less responsive to changes in the quality of governance. In times of increased corruption and political misconduct, citizens close to the capital city lose trust and reduce their electoral support of the incumbent government to a significantly larger extent than citizens in isolated areas. This circumstance has important adverse consequences for political accountability in isolated areas. Since isolated citizens react less to government performance, vote-maximizing state executives are incentivized to invest more government resources into areas closer to the capital city - as this is where the marginal voters are.

Classical work on the origins of regional inequalities was conducted by Williamson (1965). Since then, research on the determinants of regional development has provided ample evidence that local geographical factors and endowments have a strong impact on the level of economic prosperity (see for example Diamond (1997); Nunn and Wantchekon (2011); Alesina et al. (2016); Jedwab and Moradi (2016); Bakker et al. (2018); Boxell (2019); Michalopoulos et al. (2019); Alesina et al. (2021)). Other research points out that historical institutional framework conditions such as pre-colonial ethnic institutions are a key factor for economic development (see for example Gennaioli and Rainer (2007), Michalopoulos and Papaioannou (2013) or Michalopoulos and Papaioannou (2020) for a recent extensive literature review of African historiography). Hodler and Raschky (2014) and Burgess et al. (2015), in turn, show that political factors such as the ethnic affiliation of the incumbent president plays an important role for regional economic growth. The authors document that under weak political institutions, public investments and economic growth are biased in favor of the president's home region. Thus, these studies clearly underline that research on comparative development needs to go beyond the national level and occupy itself with subnational patterns. Moreover, political mechanisms are a key determinant of comparative regional development. These insights are especially relevant in the African context where states are 'artificial', have not grown together as one over the centuries and feature very high levels of heterogeneity and ethnic fractionalization (Alesina et al., 2011; Michalopoulos and Papaioannou, 2020).

The remainder of this paper is organized into four sections. Firstly, Section 2.1 will introduce the dataset, and establish the empirical identification strategy. Secondly, Section 2.2 will present our empirical results. In Section 2.3, we examine potential mechanisms that

link isolation from the capital city and economic performance. Finally, Section 2.4 will summarize the findings and conclude the paper.

2.1 Data and Empirical Strategy

There are 38 Sub-Saharan African countries in our sample (see Figure 2.A.7).[4,5] Figure 2.1 presents the distribution of isolation from the capital city for our sample countries. As becomes clear from the graph, isolation from the capital city is not a characteristic that only applies to a small minority but with a median of 303 km (mean: 391 km) rather represents the common case.

FIGURE 2.1
ISOLATION FROM THE CAPITAL CITY - OVERVIEW

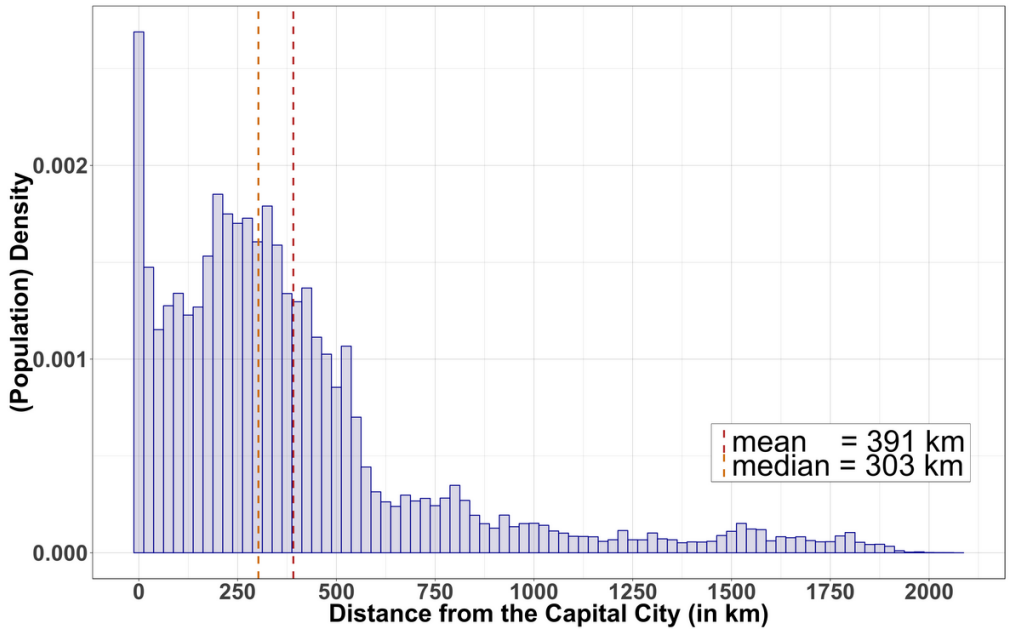

Note: This figure is based on own calculations using the UN-adjusted population density grid for the year 2015 by Worldpop. It displays the density distribution of isolation from the capital city for our sample countries. The population within a range of 20 km from the capital city that is omitted by default from all estimations (see Section 2.1.2). Each bar represents an interval of 25 km. The upper limit of isolation from the capital city is 2075 km. The average person lives 391 km (median: 303 km) away from the capital city.

[4]We exclude small island states, Sudan and South Sudan due to their recent separation, as well as Somalia and Somaliland due to the absence of a stable political power in Somalia and the special role of the government in Somaliland (Eubank, 2012). We further exclude South Africa as it has subdivided its three branches of government into three separate capital cities and Lesotho as it does not share a boundary with any remaining country in the sample.

[5]For the most part, the assignment of capital cities is uncontroversial as the majority have persisted as such since the colonial era. The exceptions are Ivory Coast and Nigeria where the capital city was ultimately shifted in 1983 from Abidjan to Yamoussoukro and in 1991 from Lagos to Abuja respectively. Tanzania has also been planning to move its capital from Dar es Salaam to Dodoma since 1973. However, since the Tanzanian parliament had not been relocated till 2019, we use Dar es Salaam as capital city of Tanzania.

2.1.1 Measuring Economic Performance Locally

To examine the impact of distance from the capital city, we require information on economic activity for small spatial units. Since reliable administrative data on the local level is not available in Sub-Saharan Africa, we use the 2016 VIIRS nighttime lights as our main measure for local economic performance. Additionally, in order to cross-validate our findings and address potential shortcomings of nightlights[6], we complement our analysis using the survey-based DHS wealth index.

Nighttime Lights (VIIRS):

The use of nighttime luminosity data as a proxy for economic activity has greatly increased in recent years. Several studies have investigated and validated the consistency of nighttime lights as a proxy for GDP (Henderson et al., 2012; Michalopoulos and Papaioannou, 2013; Donaldson and Storeygard, 2016). As is common in the literature, we use two measures of nightlights:

- Intensive approach (log nightlight intensity[7]):

$$Y_i = ln(Lights_i + 0.002000212) \tag{2.1}$$

- Extensive approach (the extent to which cells are lit or not lit):

$$Y_i = \begin{cases} 1 \;\; if \;\; Lights_i > 0 \\ 0 \;\; if \;\; Lights_i = 0 \end{cases} \tag{2.2}$$

Wealth Index (DHS):

The DHS (Demographic and Health Surveys Program) has collected nationally representative and geocoded[8] data on sociodemographic, economic and health characteristics and covers 30 out of the 38 countries in our sample. We use the most recent household recode survey available for each available country and illustrate the DHS sample in Figure 2.A.8a. Based on a household's ownership of selected assets (such as car, bicycle, refrigerator, computer, television) and household facilities (such as roof and floor material or type of toilet facility), the DHS estimates a household's 'wealth index' using principal component ana-

[6] Chen and Nordhaus (2011) and Cogneau and Dupraz (2014) point out that the predictive power of nightlights for economic activity is low and noisy for areas of low population and nightlight density.

[7] Since the vast majority of pixels has a light density of 0, we add the minimal observed light density that is greater than zero as a constant term before taking the natural logarithm. The results remain qualitatively equivalent when using alternative constant terms such as 0.0001 or 0.001.

[8] The longitude and latitude of each respondent is recorded using a GPS receiver. Note that to protect the privacy of respondents, the DHS displaces the GPS coordinates randomly up to 2 km for urban clusters and up to 5 km for rural clusters with 1% of rural clusters being displaced up to 10 km. This displacement is restricted such that respondents always stay within the same country and region.

lysis (PCA).[9] The PCA represents a composite measure of a household's cumulative living standard relative to other households within each country and year which makes it suitable for applications that seek to understand the relative distribution of living standards within countries. We normalize the PCA for each survey to make a household's wealth index more comparable to the relative position of households in other countries. As a second alternative measure, we rank and assign each household its relative position in the national wealth distribution. While the second measure loses valuable information regarding the absolute difference between two consecutive households, its interpretation is more intuitive as differences between households can be expressed in terms of percentile changes within the national wealth distribution.

2.1.2 Identification

Challenge:

The most intuitive way to assess the effect of isolation from the capital city[10] on local economic development is a simple univariate analysis. In Figure 2.2, we plot the share of lit pixels over distance from the capital city. The graph reveals that, on average, the probability of detecting nightlights in a pixel decreases exponentially with distance from the respective capital city. Yet, this correlation is hard to interpret as it is shaped by a variety of confounders. Most notably, isolation from the capital city is correlated with a range of location-specific geographical factors, that are themselves determinants of economic performance. For instance, African capitals tend to be located at the coast which means that proximity to capital cities concurrently translates into the advantages of proximity to ports and international trade (Henderson et al., 2017). Also, it is highly doubtful whether the relationship actually reverses for very high distance as is suggested by the slightly positive slope starting at around 1,250 km. It is more likely that these pixels just happen to be in economically more dynamic areas such as the mining areas in the South-Eastern part of DR Congo. Consequently, local economic framework conditions such as endowments with natural resources or local institutions and culture (Michalopoulos and Papaioannou, 2013)

[9] For more information about the construction of the DHS wealth index visit: `https://dhsprogram.com/topics/wealth-index/Wealth-Index-Construction.cfm`.

[10] We measure isolation from the capital city as the Euclidean distance between a location (a pixel) and the respective capital city. A drawback of this measure, as compared to more sophisticated travel time or travel cost estimates, is that it is less precise. However, measures that take into consideration the infrastructure development would induce reverse causality bias. This is due to the fact that places that are more dynamic economically tend to be better connected and are therefore closer to the capital city in terms of travel time. Combes and Lafourcade (2005, 346) underline the consistency of our metric by showing that simple distance measures "do a very good job in capturing transport costs in cross-section analysis". Nevertheless, in Section 2.A.6, we report estimates using travel time based on OpenStreetMap instead.

confound simple correlations.

FIGURE 2.2

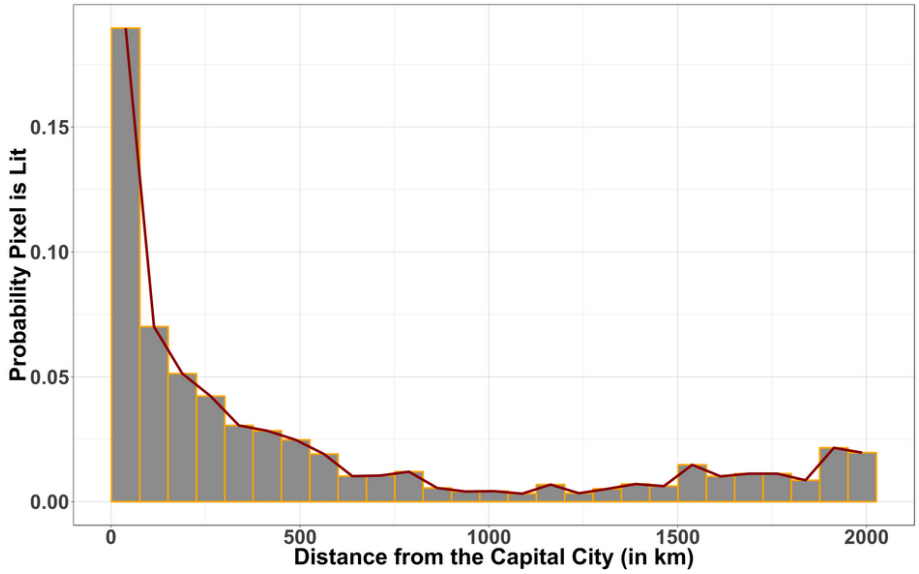

Note: This figure plots the share of lit pixels over distance from the capital city. The population within a range of 20 km from the capital city that is omitted by default from all estimations (see Section 2.1.2). Each bar represents an area of 75 km.

One way of addressing these shortcomings would be attempting to explicitly model all relevant relationships by including a wide set of geographical covariates, X_i, and country and ethnicity fixed effects, b_c and b_e. Equation 2.3 illustrates the respective OLS model equation where Y_i refers to our measure of *Nightlight Density* in pixel i and CAP_i to log distance from the capital city.[11] Yet, since we have to assume that we only control for a subset, $\hat{X}_i$, of all relevant location-specific factors ($X_i = \hat{X}_i + \widetilde{X_i}$), CAP_i is likely to remain endogenous with unobserved location-specific characteristics $\varepsilon_i = \widetilde{X_i} + u_i$ and $E(\varepsilon_i|CAP_i) \neq 0$. As a result, OLS-estimations based on Equation 2.3 are likely biased.

$$Y_i = \beta CAP_i + \gamma X_i + b_c + b_e + \varepsilon_i \tag{2.3}$$

Construction of the BDD Model:

In the African context, a tangible solution is to establish counterfactuals in a BDD model at national borders. Since African borders were arbitrarily drawn by the colonial powers and divide pre-colonial ethnic homelands with similar geographical, social and historical traits,

[11] We log-transform distance to the capital city using the natural logarithm to account for the exponential relationship suggested by Figure 2.2. The qualitative intuition behind the log-transformation is that the effects of isolation from the capital city are decreasing with distance from the capital city. Yet, given that the relationship in Figure 2.2 is strongly confounded and should therefore be taken with caution, we confirm the adequacy of the log-transformation statistically based on the Akaike information criteria applied to our more sophisticated BDD model.

the assignment of areas close to the boundary to a particular country and its respective capital city can consequently be interpreted as accidental (Asiwaju, 1985; Michalopoulos and Papaioannou, 2016, 2020). Hence, national borders constitute an arbitrary cutoff with a quasi-random jump in distance from the capital city. However, as opposed to Dell (2010), Basten and Betz (2013), Michalopoulos and Papaioannou (2014) or Dell et al. (2018), simply pooling areas around boundaries does not remove the heterogeneity in unobservables with respect to treatment intensity. As opposed to country-wide indicators, isolation from the capital city is autocorrelated i.e. evolves gradually along the boundary which induces a cross-border correlation and induces a spurious relationship between isolation from the capital city and other autocorrelated (un-)observable location-specific factors that concurrently gradually evolve along both sides of the boundary. For example, when considering the national boundary between Ghana and Togo (see Figure 2.3), distance from the capital city increases for both sides from South to North. The problem is that (un-)observable factors such as distance from the coast or climatic conditions simultaneously evolve along the boundary. As a consequence, a pooled BDD would still be confounded by unobservables.

FIGURE 2.3
THE GHANA-TOGO BORDER

(A) ACCURATE MAP - BOUNDARY AREA

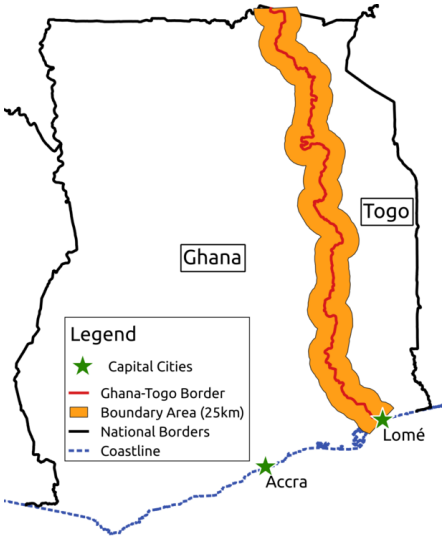

(B) SCHEMATIC MAP - BORDER SEGMENTS

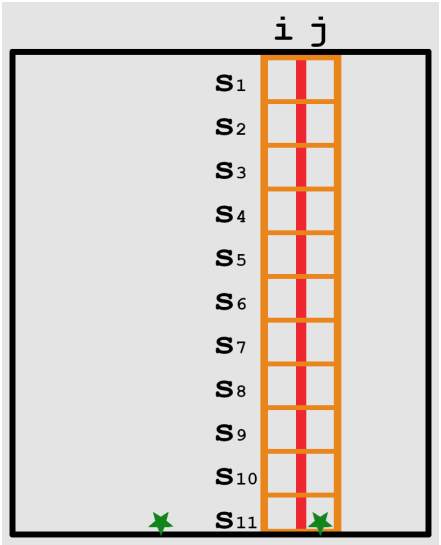

We overcome this problem by subdividing shared national boundary areas into smaller and more homogenous segments and conduct the BDD within these segments. In Figure 2.3, s_1-s_{11} sketch such segments for the example of the Ghana-Togo border schematically.

Technically, we implement this step by including segment fixed effects that essentially partial out anything that equally exists on either side of each boundary segment. Hence, under the assumption that segments are balanced with regard to geographic covariates, this procedure solves the problem of spurious geographic characteristics. We therefore reduce our sample to a small buffer of 25 km around those boundary segments and conduct balancing tests on observables. Yet, we still need to account for (i) the switch in the country and hence the institutional environment as well as (ii) national boundaries that coincide with ethnic boundaries. Firstly, since country characteristics are common across pixels within a country, we can resolve the former by using country fixed. Therefore, our BDD can be interpreted as a localized DID that compares the relative performance of different areas within a country relative to their respective counterfactual segment sides. Secondly, we can avoid that ethnic boundaries confound our estimation by nesting each segment within a partitioned ethnic homeland p. Thereby, this procedure does not only balance geographical covariates[12] but also historical, political and cultural framework conditions. Furthermore, we include polynomials of our running variable *Distance from the Boundary in km* (DFB) for each ethnic homeland in each country separately into our model. These polynomials pick up any potentially remaining heterogeneity within segments. Equation 2.4 presents our main BDD model equation.

$$Y_i = \beta CAP_i + \varphi X_i + b_c + b_s + \sum_{n=1}^{3} \lambda_{n,c,p} DFB_i^n + \xi_i \tag{2.4}$$

<u>*Placebo Tests:*</u>

A potential caveat might stem from the circumstance that isolation from the capital city always simultaneously means isolation from a major city and market. In order to affirm that it is in fact isolation from the political center that is driving the effects of isolation from the capital city, we run placebo tests. If hosting the political center is in fact the key characteristic of capital cities with respect to the effects under scrutiny, then the effects of isolation from other major cities within the country should be fundamentally different. For this purpose, we create PLC, representing the *log* of *distance from the placebo city (in km)*, as a new variable. A placebo city is defined as the largest non-capital city in the country.[13] We then include PLC into Equation 2.4 and compare the estimated coefficients

[12] Michalopoulos and Papaioannou (2014, 172) provide an in-depth discussion about the origin of Sub-Saharan African national boundaries and conclude that "differences in geography-ecology, location, and natural resources across the border within partitioned ethnic homelands are small and not systematically linked to differences in national institutions". Conducting a range of balancing tests, we validate that local (un-)obversable characteristics are also balanced with respect to isolation from the capital city.

[13] In order to avoid collinearity with the capital city, we further require the placebo city to be at least 50 km away from the capital city.

for capital and placebo city isolation. However, this placebo test might be confounded by the fact that capital cities tend to be the largest city within the country. To this end, we decompose the effect of isolation from a city into a city type (capital vs. placebo city) and city size effect (big vs. small city in terms of city population). We do so by partialling out the size effect by additionally including interactions between CAP and PLC with their respective population counts.

$$Y_i = \beta_1 CAP_i + \tau_1 CAP_i \times Pop_{cap,c} + \beta_2 PLC_i + \tau_2 PLC_i \times Pop_{plc,c} + \varphi X_i + b_c + b_s + \sum_{n=1}^{3} \lambda_{n,c,p} DFB_i^n + \epsilon_i \tag{2.5}$$

2.2 Empirical Analysis

2.2.1 OLS

The results of the OLS estimations based on Equation 2.3 for a variety of alternative specifications can be found in Table 2.1. As becomes clear from the table, irrespective of the precise model specification, isolation from the capital city is significantly negatively related to both the probability and the intensity with which a pixel is lit. A one percent increase in distance from the capital city, on average, decreases the probability that a pixel is lit by around 1.5 percentage points and the nightlight density by around 0.07 percent. Figure 2.A.13 presents the respective OLS estimates conducted for each country separately. Figure 2.A.13a compares the estimates of the full sample to those of the boundary sample using extensive nightlights corresponding to columns (1) and (2). Likewise, Figure 2.A.13b corresponds to the intensive margin corresponding to columns (5) and (6). The associated balancing tests, for the example of distance from the coast, can be found in Table 2.A.14.

2.2.2 Border Discontinuity Graphs

In the next step, the goal is to overcome the imbalance in location-specific characteristics with respect to the treatment intensity of isolation from the capital city. Yet, prior to moving to the most elaborate BDD model based on Equation 2.4, we undertake a simplified, yet more intuitive, graphical approach.

We begin by subdividing all national boundaries into segments of 50 km line length with a buffer of 50 km on either side (see 2.A.9 for an illustrative map).[14] Each segment side

[14] To obtain balanced subgroups, we drop segments where the minimum distance from the border on either side is greater than 5 km or the maximum distance less than 45 km (which occurs mostly around very uneven boundaries). This reduces our sample from 729,093 to 480,149 pixels.

Table 2.1
OLS Results

	Dependent variable:							
	Probability Pixel is Lit in 2016 (VIIRS)				Log Light Density in 2016 (VIIRS)			
	OLS		Ethnicity FE		OLS		Ethnicity FE	
	All	Border	All	Border	All	Border	All	Border
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Distance from the Capital City	-0.020*** (0.005)	-0.021*** (0.006)	-0.025*** (0.005)	-0.014** (0.005)	-0.068*** (0.021)	-0.083*** (0.026)	-0.101*** (0.023)	-0.065*** (0.024)
Geography Cov.	YES	YES	YES	YES	YES	YES	YES	YES
Country FE	37	37	37	37	37	37	37	37
Ethnicity FE	NO	NO	706	351	NO	NO	706	351
Observations	3,518,146	416,667	3,518,146	416,664	3,518,146	416,667	3,518,146	416,664
Adjusted R^2	0.080	0.082	0.130	0.133	0.069	0.069	0.126	0.120

Note: This table reports OLS and boundary area regression results based on Equation 2.3. In order to avoid capturing the break between the capital city and the hinterlands, we exclude 20 km around each capital city from our sample. To prevent misassignment of detected nightlights between countries due to blooming, we exclude 3 km on each side of the border. The boundary area regressions ('Border') are restricted to all pixels with centroids within the range of 25 km from shared national borders. The 'Geographical Cov.' include: distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km). The 'Ethnicity FE' are based on the ethnic homelands in the 'Tribal Map of Africa' (Murdock, 1959). Standard errors in parenthesis are clustered by ethnic homeland. *p<0.1; **p<0.05; ***p<0.01

belongs to a different country (and capital city). Next, we determine the average distance to the capital city for each segment side and assign them into the group 'close' if they are relatively closer to their capital city as their opposing boundary segment and into the group 'far' otherwise. Through this procedure, we obtain two groups that are balanced with respect to location-specific covariates but with systematically different distances from their respective capital cities. Thereby, this procedure enables us to assess the impact of crossing the boundary from 'far' to 'close' to the capital city on nightlight density while keeping geographical factors constant.

Figure 2.1 plots the border discontinuity graphs with 2.1a referring to the extensive and 2.1b to the intensive scale of nightlights. Areas on the left are on average 830 km and areas on the right 430 km away from the capital city. The graph indicates a large jump of around 25% in the probability with which pixels are lit when moving from relatively remote to areas near the capital city.

Figure 2.1
Boundary Discontinuity Graphs

(a) Extensive (b) Intensive

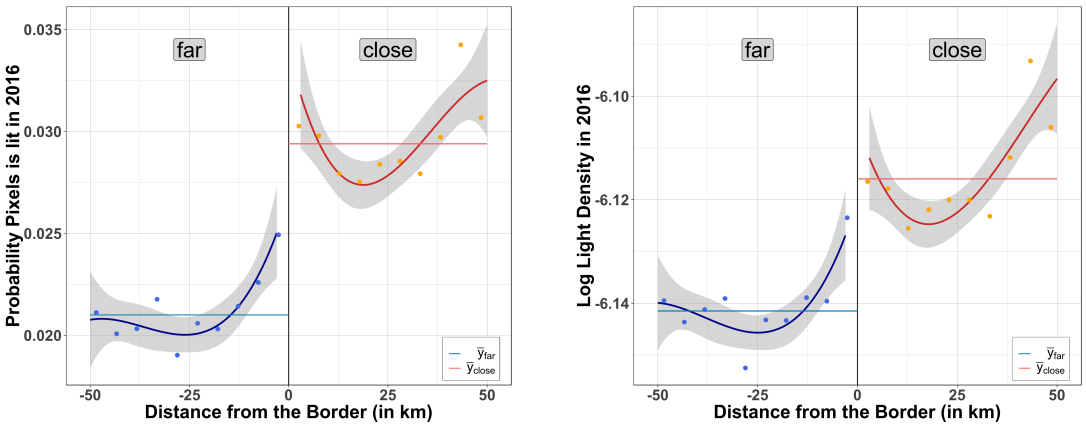

Note: The graphs illustrate the graphical BDD. The gray buffer around the lines represent the 95% confidence interval. The bins on the left-hand side are, with an average distance of 830 km, relatively far from the capital city and represent a total of 241,241 pixels. In contrast, pixels on the right-hand side are, with an average of around 430 km, relatively close to the capital city and represent 238,908 pixels.

Even if (un-)observable geographical factors should by construction be balanced between the two regimes ('far' and 'close'), this needs to be empirically confirmed. Figure 2.A.11 illustrates the respective graphs for a range of geographical indicators. All covariates, except for distance from the capital city, move smoothly across the cutoff and do not exhibit significant discontinuities. Based on the balancing tests, we can thus conclude that the jump in nightlight density stems from differences in isolation from the capital city.

A remaining concern is that the results might be confounded by country characteristics. For example, supposing that small countries perform better economically, and given that small countries tend to constitute the 'close' group, we would expect to see comparable patterns even in the absence of effects induced by isolation from the capital city. Moreover, it might be that the effects are a result of isolation from a major city rather than isolation from the political center. Yet, if either of the two concerns were valid, we should observe a similar result when using the location of other major cities instead. Figure 2.2 depicts the graphs when conducting the analogous analysis but for isolation from the placebo city.

As it turns out, the placebo graphs clearly indicate that there are no effects associated with isolation from other cities.[15] Yet, despite these highly encouraging results, at this stage, we cannot be entirely sure that the estimated effects are causal. While the graphical BDD, balancing and placebo tests give strong support to our hypothesis, in the next step,

[15]The respective balancing graphs, once again, indicate that (un-)observable geographical factors are balanced and move smoothly across the cutoff (see Figure 2.A.12).

FIGURE 2.2

BOUNDARY DISCONTINUITY GRAPHS - PLACEBO

(A) PLACEBO CITY ISOLATION - EXTENSIVE (B) PLACEBO CITY ISOLATION - INTENSIVE

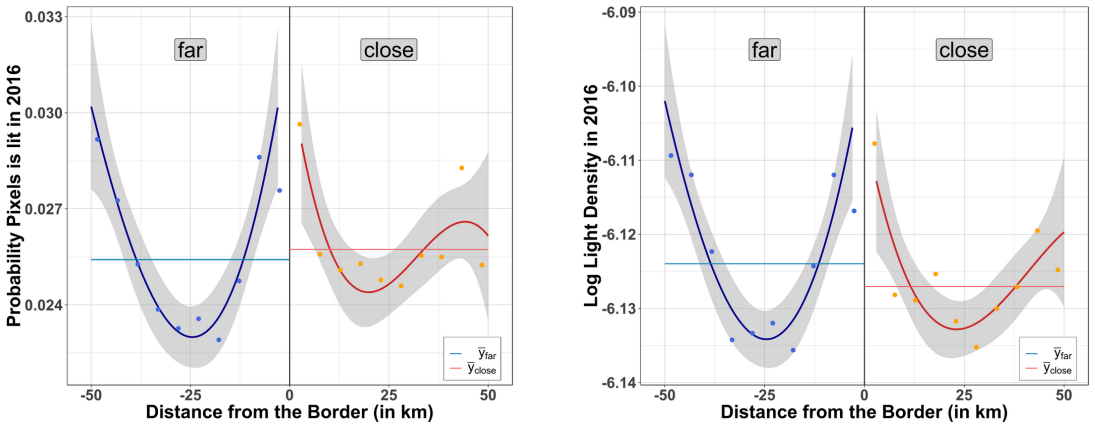

Note: The graphs illustrate the placebo tests regarding the graphical BDD. A placebo city is the largest non-capital city in the country. The gray buffer around the lines represent the 95% confidence interval. The bins on the left-hand side are, with an average distance of 820 km, relatively far from the placebo city and represent a total of 238,980 pixels. In contrast, pixels on the right-hand side are, with an average of around 440 km, relatively close to the placebo city and represent 237,153 pixels.

we need to properly account for the switch between countries and ethnic homelands at the cutoff.

2.2.3 Boundary Regression Discontinuity

In order to tackle the remaining shortcomings pointed out in Section 2.2.2, we move to a more sophisticated BDD regression model based on Equation 2.4. In this setting, the switch between countries at the boundary cutoff is accounted for using country fixed effects. Furthermore, we identify arbitrary borders by exclusively using border pieces that divide ethnic homelands. Additionally, we nest our boundary segments within the restricted partitioned ethnic homelands to prevent ethnic shifts within segments (see Section 2.1.2). Table 2.2 presents the BDD results for the extensive (columns (1)-(4)) and intensive (columns (5)-(8)) scales of nightlight density. Columns (1) and (5) exclude all geographical covariates and serve as a reference to assess the extend to which potentially omitted location-specific characteristics might confound our estimates. When comparing column (1) to (2) and column (5) to (6), it becomes clear that our identification strategy proves to be effective. Whether or not we include an extensive set of geographical covariates changes the magnitude of our estimates only by a small and statistically insignificant margin. With the exception of ruggedness, all balancing tests confirm that in our BDD model, (un-)observable factors are balanced with respect to treatment intensity - the coefficients are either insignificant and/or very close to zero and economically negligible (see Table

TABLE 2.2
BOUNDARY DISCONTINUITY ESTIMATION

	Dependent variable:							
	Probability Pixel is Lit in 2016 (VIIRS)				Log Light Density in 2016 (VIIRS)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Distance from the Capital City	-0.023**	-0.025**	-0.031**	-0.031**	-0.104**	-0.112**	-0.137**	-0.143**
	(0.010)	(0.010)	(0.012)	(0.013)	(0.048)	(0.048)	(0.058)	(0.061)
	Polynomials for: distance from the border × country × ethnicity (305 groups)							
2nd order	x	x	-	-	x	x	-	-
3rd order	-	-	x	-	-	-	x	-
4th order	-	-	-	x	-	-	-	x
Geography Cov.	NO	YES	YES	YES	NO	YES	YES	YES
Country FE	36	36	36	36	36	36	36	36
Segment FE	569	569	569	569	569	569	569	569
Observations	168,620	168,620	168,620	168,620	168,620	168,620	168,620	168,620
Adjusted R^2	0.175	0.177	0.182	0.185	0.155	0.156	0.161	0.164

Note: This table reports our main BDD regression results corresponding to Equation 2.4. In order to avoid capturing the break between the capital city and the hinterlands, we exclude 20 km around each capital city from our sample. To prevent misassignment of detected nightlights between countries due to blooming, we exclude 3 km on each side of the boundary. The 'Geographical Cov.' include: distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km). Boundary segments corresponds to a buffer of 25 km around border pieces of 50 km line length and are entirely nested within a restricted ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). The observations are weighted such that each segment side has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by boundary segment. *p<0.1; **p<0.05; ***p<0.01

2.A.15). The models in columns (2)-(4) as well as (6)-(8) feature an increasing number of polynomials of the running variable, distance from the border. In order to allow for a sufficient degree of flexibility regarding the dynamics of nightlights around the boundary, we choose third order polynomials as our default option. Consequently, applying our BDD estimation framework, we can verify that isolation from the capital city has a negative causal impact on economic development. A one percent increase in distance from the capital city decreases the probability of a pixel to be lit by 3 percentage points and the nightlight intensity by 0.14 percent. The causal estimates are therefore of a slightly higher absolute magnitude than the OLS and boundary area estimates (see Table 2.1). This differential is likely triggered by the fact that OLS estimates are confounded by local economic conditions. One such example are the economically high performing mining areas in the South-Eastern part of DR Congo far away from the capital city Kinshasa.

Regarding inference, it is important to account for spatial autocorrelation which is why we cluster standard errors at the boundary segment level by default. Alternatively, when double-clustering standard errors at the country and country-pair border level for our

baseline estimates in columns (3) and (7), we obtain -0.031* (0.017) and -0.137* (0.081) respectively, and when double-clustering at the country and broad ethnicity family level instead, we obtain -0.031* (0.018) and -0.137 (0.084) respectively (Cameron et al., 2011). Furthermore, we apply the methodology by Conley (1999), that accounts for arbitrary spatial dependence within a radius around each unit with and without Bartlett correction for various distance cutoffs (50, 100, 200, 500 or 1000 km; see Table 2.A.18). These alternative cluster specifications tend to yield slightly larger but overall similar standard errors.

2.2.4 Placebo Tests

In this section, we use our BDD framework to test whether the driving characteristic of isolation from the capital city is rooted in isolation from the political center or, alternatively, based on isolation from a major city within the country. In order to answer this question we compare the effects of isolation from the capital city to isolation from other major cities within the country, the placebo cities, using our estimation approach specified in Equation 2.5.

The results in Table 2.A.19 indicate very clearly that the effects of isolation from the capital city differ fundamentally from those of isolation from the placebo city. While, the impact of isolation from the capital city is significantly negative across all model specifications[16], the effects associated with placebo city isolation are insignificant and very close to zero. This result holds when decomposing isolation from a city into a city type and city size effect (see column (6) and (8) where we additionally control for the interactions of capital and placebo city isolation with their respective city population counts). These results imply that the type of a city (capital vs. other cities) is more important than the city size. We conclude that hosting the political center of the country is the driving force behind the effects of remoteness from the capital city.

2.2.5 Sensitivity Analysis Summary

This section contains a summary of the robustness and sensitivity tests (see Section 2.A.5 in the appendix for more detail). In Table 2.A.6, we show that the results are robust with regard to reducing the bandwidth of the segments around the boundary from 25 km to 15 km, only exclude pixels within 1.5 km instead of 3 km around the boundary, and applying a varying degree of zero up to fourth order polynomials of the running variable.

[16] Due to some collinearity between capital and placebo city isolation, the coefficient of isolation from the capital city is slightly lower as compared to Table 2.2.

We then show that the results are stable when increasing the sample restriction from excluding 20 km (default) around the capital city to 50, 75 and 100 km (Table 2.A.6). In Figure 2.A.5, we demonstrate that the results are not driven by an individual country or boundary. Further, in Table 2.A.8, we examine if there are heterogeneities between groups of countries. We find that the effects seem to be more relevant in democracies, as opposed to autocracies, and more relevant in relatively underdeveloped countries. In Table 2.2, we show that the pattern of newly lit pixels since 1992 is also negatively associated to remoteness. This finding supports the view that the adverse effects of isolation from the capital city are still actively shaping local economic growth as opposed to reflecting a persistent pattern from the past. We also show that the representation of ethnicities i.e. whether they are a minority or majority within the country are not driving the effects. In addition, we demonstrate that the implications of isolation from the capital city go beyond population agglomeration and hold when controlling for population density or using measures of light per capita as dependent variable (see Table 2.A.9). We also show that our results hold when controlling for distance to the closest river, waterbody, mine, city of varying thresholds and regional capital city[17] within the country (see Tables 2.A.10 and 2.A.11). Finally, in Table 2.A.12, we cross-validate our findings using the DHS wealth index as an alternative measure of economic performance.

Consequently, based on a wide range of sensitivity tests, we conclude that our results are highly robust to a wide range of alternative specifications and considerations. The robustness results thereby confirm our estimated negative causal impact of remoteness from the capital city on economic performance.

2.3 Mechanisms

So far, we have shown that isolation from the political center within the country, the capital city, has strong adverse causal net effects on local economic performance. Since isolation from the capital city itself is not an economic variable, there must be a more concrete economic link between remoteness from the capital city and economic performance.

The provision of public goods is a fundamental driver of economic development (Besley and Ghatak, 2006; Dittmar and Meisenzahl, 2019) and Campante and Do (2014) find that US states with isolated capital cities provide less public goods. This result might reflect

[17] We define the regional capital city as the first-level administrative capital city as defined in the global administrative unit layer (GAUL) by the Food and Agriculture Organization of the United Nations (FAO). We retrieved information about the name of the respective regional capital city from various sources including Wikipedia and geocoded the data using OpenStreetMaps. In a last step, we manually examined and corrected the shapefile of regional capital cities to avoid geocoding errors.

aggregations of the local economic patterns that are under scrutiny in this study and might be a result of low levels of public goods in isolated areas. Using a fixed effects OLS model, Michalopoulos and Papaioannou (2014) find that distance from the capital city is negatively correlated with law enforcement, which represents a public good in the wider sense. The authors' finding can therefore be interpreted as early suggestive evidence that points at reduced levels of public goods in African regions isolated from the capital city.

Using our BDD model, in this section, we show that a variety of local public goods are undersupplied in isolated areas which supports the view that the level of public goods provision constitutes an important mediator between isolation and economic development. In a subsequent step, we shed light into the causes of the low levels of public goods provision in remote areas and investigate political representation and accountability as potential mechanisms. While the results on political representation are ambiguous, the results on accountability clearly suggest that dysfunctional feedback and accountability mechanisms are relevant with regard to the effects. Lastly, we investigate two alternative potential channels that might constitute important mediators: Market access and conflict. We show that neither market access nor conflict are likely to be relevant with regard to the observed patterns.

2.3.1 Public Goods

Based on rounds 5, 6 and 7 of the Afrobarometer, we generate an index on the provision of public goods reflecting whether a cluster is provided with paved roads, electricity grid access, piped water and a sewage system. The public goods index is the average of binary responses about the availability of the respective provisions. Consequently, the index can take the values of 0, 0.25, 0.5, 0.75 and 1 corresponding to the respective share of public goods being present. Table 2.1 combines the results of OLS, boundary area and BDD regression models, similar to those in Section 2.2. We include a range of cluster and geography controls by default into all models. This includes a dummy variable corresponding to whether the cluster is located in an urban or rural area as this is directly related to public goods provision. However, since the degree of urbanization is an outcome of remoteness and thereby endogenous to the model, the estimated absolute magnitudes constitute lower bound estimates.[18] Furthermore, it should be noted that the Afrobarometer dataset has less observations than the nightlights sample. In order to avoid the sample size from

[18] When excluding the urban dummy, the estimated coefficient corresponding to Table 2.1 column (3) is -0.124** (0.057). Additionally or alternatively measuring agglomeration as the mean population density in a 5 or 10 km buffer around the cluster has no significant impact on the estimated coefficient.

getting too small, we do not restrict the ethnic homelands with a negative buffer before defining the boundary segments. Moreover, since Afrobarometer respondents are clustered in enumeration areas, with in some cases just one or two clusters per segment side, we do not account for distance from the boundary disaggregated but rather focus on capturing the general relationship.

TABLE 2.1
CHANNEL ANALYSIS: PUBLIC GOODS PROVISION (AFROBAROMETER)

	Dependent variable: Public Goods Index		
	OLS	Boundary Area	BDD
	(1)	(2)	(3)
Log Distance from the Capital City	-0.020**	-0.034***	-0.088**
	(0.009)	(0.012)	(0.038)
	Polynomials for: distance from the border		
3rd order	-	-	x
Cluster Cov.	YES	YES	YES
Geography Cov.	YES	YES	YES
Country × Round FE	73	71	63
Segment FE	-	-	138
Observations	11,959	1,868	1,069
Adjusted R^2	0.480	0.487	0.670

Note: This table reports the regressions on the impact of isolation from the capital city on the level of public goods provision using the Afrobarometer survey. The 'Cluster Cov.' include the average age, age squared and sex of all respondents in the cluster. The 'Geographical Cov.' include: distance from the coast (in km), longitude and latitude (projected in km) and whether the cluster is located in an urban or rural setting. Column (1) correspond to the full sample OLS regression, columns (2) to the 25 km boundary area regression and column (3) to the BDD regression. Boundary segments corresponds to a buffer of 25 km around border pieces of 50 km line length and are entirely nested within an ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). In column (1), the observations are weighted according to the Afrobarometer survey weights. In column (3), the observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by Afrobarometer cluster and ethnic homeland (columns (1) and (2)) and Afrobarometer cluster and boundary segment (column (3)). *p<0.1; **p<0.05; ***p<0.01

A one percent increase in isolation from the capital city decreases the probability index of public goods provision by around 10 percentage points.[19] This relationship also holds, albeit with a slight upward bias, for the OLS and boundary area regressions.

To cross-validate our findings on public goods we use alternative geocoded measures of public goods including (i) a large dataset on roads based on Michelin maps that were geocoded and digitized by Jedwab and Storeygard (2020) based on which we create dummies indicating whether a pixel intersects with a road or tarred road and (ii) an extensive dataset on the location of health facilities across the African continent assembled by Maina et al. (2019) based on which we compute distance to the closest health facility (including

[19]The corresponding placebo tests and associated F-tests indicate that the effects are indeed specific to isolation from the capital city and do not hold for placebo cities.

public and private lower tier health centers and hospitals), the closest hospital as well as the closest public health facility (public health center or hospital).
Column (2) in Table 2.A.20 indicates that a one percent increase in distance from the capital city reduces the probability of a pixel to be connected to a tarred road by 1.9 percentage points. Columns (4) and (5) report that a one percent increase in distance from the capital city increases the distance to the closest higher tier hospitals or publicly provided health facility by around 0.4 and 0.6 percent respectively. In contrast, for coarser measures like the availability of any type of road (column (1)) or any medical center (column (3)), there are no significant differences between areas close and far from the capital city. This finding underlines that it is important to use detailed data and take into consideration the quality when evaluating public goods provision.

We conclude that people in areas isolated from the capital city receive significantly less and lower quality public goods and services from their political leaders as compared to those closer to the capital city. Public goods are therefore likely to be an important mediator between isolation from the capital city and economic development.
However, it is not obvious why the provision of public goods in isolated areas is reduced. It might be that for economic reasons such as high transport costs, lack of specialized labor and intermediary goods the provision of public goods is simply more costly and therefore reduced in isolated areas. In context of India's national rural road program, Asher et al. (2018) document that the cost of road construction is the same in areas close and far from regional headquarters. While we do not have access to data to explicitly test for this potential channel in our context, these findings are indicative that the cost of public goods provision such as road infrastructure is generally similar across locations. Furthermore, if the cost of certain public goods would be a central mechanism, it should not primarily be determined by distance from the capital city but rather by remoteness from cities and markets in general. Yet, the fact that our placebo tests indicate that distance from other major cities has no comparable effect on public goods (Table 2.A.22), and that the effects of remoteness from the capital city are not primarily associated to market access or remoteness from the regional capital or closest city (see Tables 2.5 and 2.A.10) puts the cost argument as a key driver into question. We therefore consider it more likely that there are political mechanisms at play that lead to a reduced public goods provision in areas isolated from the capital city.

In the next step, we empirically investigate two potential political explanations. Firstly,

geographical isolation might translate into political isolation and thereby exclusion from government resources. For example, isolated areas might be less represented within the incumbent government and therefore benefit less from ethnic or regional favoritism (Hodler and Raschky, 2014; Dreher et al., 2019). Another explanation could be that public goods might be lower due to dysfunctional feedback and accountability mechanisms. Political agents might simply lack the incentives to provide those in isolated areas with public goods.

2.3.2 Public Goods and Political Representation

We start by investigating whether political representation and 'political favoritism' are relevant in this context and test whether segment sides farther away from the capital city exhibit a lower access to power. For this purpose, we examine whether the national leaders are more likely to come from the region of segment sides closer to the capital city using an updated version of the database by Dreher et al. (2016, 40-41). Since measuring an area's political representation and access to power solely based on the head of state might not sufficiently reflect the overall power access, we complement the analysis using the Ethnic Power Relations Core Dataset 2019 (EPR) (Vogt et al., 2015) and combine it with spatial information about the location of the respective politically relevant groups (Wucherpfennig, 2011). The results in Table 2.2 indicate that a one percent increase in distance from the capital city decreases the probability that the incumbent political leader comes from the region of the segment side by 9 percentage points. In contrast, it turns out that a one percent increase in isolation increases the number of years participated in the ruling government by 3.6 years (within a period of 18 years since 2000).[20] Since the results go in opposite directions, it remains ambiguous which of the two outweighs the other. Given that these results indicate that isolated areas are not systematically excluded from power, we conclude that political representation is unlikely to be the central mechanism.

2.3.3 Public Goods and Accountability

Next, we examine whether accountability is relevant to the mechanism between isolation and public goods provision. The idea behind this channel is that due to information frictions isolated citizens are less aware about government actions and, as compared to their counterparts close to the capital, penalize their leaders less for a low provision with public

[20] A simple explanation for this phenomenon might be that incumbent governments tend to include isolated groups into the government to counteract secessive aspirations. Another possible reason could be that isolated groups are striving to have an impact politically and overcompensate for their geographical isolation.

Table 2.2
Channel Analysis: Political Representation Balancing Tests

	Dependent variable: Political Representation							
	Region of Birth of Leader (1/0)			EPR Power Coalition Status (in years since 2000)				
	All	Since 2000	Ongoing	Irrelevant	Powerless	Junior	Senior	In Power
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Distance from the Capital City	-0.026	-0.152***	-0.091***	-2.569***	-1.096**	2.278***	1.386**	3.664***
	(0.039)	(0.032)	(0.027)	(0.838)	(0.530)	(0.785)	(0.667)	(0.909)
Population Share of Region	1.232***	1.188***	0.837***	-	-	-	-	-
	(0.308)	(0.247)	(0.224)					
Country FE	35	35	35	33	33	33	33	33
Segment FE	568	568	568	531	531	531	531	531
Observations	1,138	1,138	1,138	1,064	1,064	1,064	1,064	1,064
Adjusted R^2	0.438	0.347	0.248	0.468	0.704	0.442	0.612	0.564

Note: This table reports the balancing tests on political representation with regard to isolation from the capital city. The observational unit in these models are the BBD boundary segments. The dependent variables in column (1)-(3) are dummies indicating whether a head of state came from the same admin-1 region as the boundary segment side (since independence, since 2000 and only refering to incumbent state leaders). The dependent variables in columns (4)-(8) correspond the total number of years a segment side has spent under the respective power access status in the period between 2000 and 2018 (more information on these categories can be found in Section 2.A.1). Boundary segments corresponds to a buffer of 25 km around border pieces of 50 km line length and are entirely nested within a restricted ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). Standard errors in parenthesis are clustered by boundary segment. *p<0.1; **p<0.05; ***p<0.01

goods. As a consequence, the marginal benefit of politicians to provide local public goods is reduced in isolated areas.

We begin by analyzing how people in different geographical locations relate to their political leaders. Given the reduced provision of public goods, we might expect that isolated groups demonstrate mistrust towards their government. However, as it turns out, people farther isolated from the capital exhibit a significantly higher level of trust in their political leadership (see column (1) in Table 2.3). The results in Table 2.A.21 demonstrate that the increased trust in politicians is associated with the ruling party but not with the opposition party suggesting that the reason for the higher trust does not stem from a generally higher level of trust or credulity in isolated areas. Instead, it seems that the increased trust in state leaders is a result of a lower corruption perception and a higher performance evaluation of national leaders (see columns (1) and (2) in Table 2.3). These patterns support the view that isolated citizens are not aware of their disadvantaged position and therefore do not demand a higher provision of public services. Isolated citizens might trust their national leaders more because they know less about them. In line with this view, we find significant imbalances with regard to the frequency of news consumption which is

considerably smaller in isolated areas (see Table 2.3).[21] The limited exposure of media and information might also be the reason why isolated citizens are less aware of the importance of accountability mechanisms in politics and are less inclined to advocate for checks and balances on the government (Table 2.3).

Table 2.3
Channel Analysis: Perception of Political Leadership and Accountability

	BDD Model with dependent variable z-score of:				
	Trust in Political Leaders	Government Corruption Perception	Government Performance Evaluation	Frequency of News Consumption	Advocate Checks and and Balances
	(1)	(2)	(3)	(4)	(5)
Log Distance from the Capital City	0.191***	-0.078*	0.171***	-0.139**	-0.094**
	(0.061)	(0.043)	(0.064)	(0.066)	(0.041)
	Polynomials for: distance from the border				
3rd order	x	x	x	x	x
Household Cov.	YES	YES	YES	YES	YES
Geography Cov.	YES	YES	YES	YES	YES
Country × Round FE	70	71	71	71	71
Segment FE	140	140	140	140	140
Observations	7,812	6,705	6,593	8,113	7,763
Adjusted R^2	0.178	0.164	0.202	0.169	0.135

Note: This table reports the regressions on the impact of isolation from the capital city on the perception of political leaders and accountability. The 'Household Cov.' include: age, age squared and sex of respondent. The 'Geographical Cov.' include: distance from the coast (in km), longitude and latitude (projected in km) and whether the household is in an urban or rural setting. Columns (1)-(3) correspond to trust, corruption perception and performance evaluation of their national political leadership. Column (4) corresponds to the frequency of news consumption and column (5) to the extent to which respondents are advocating a system of checks and balances to monitor the actions of their political leaders. All models are BDD regressions using 'Segment FE' for boundary segments of 50 km length with a buffer of 25 km that are nested within an ethnic homelands based on the 'Tribal Map of Africa' (Murdock, 1959). All observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by Afrobarometer cluster and boundary segment.

*p<0.1; **p<0.05; ***p<0.01

To confirm that it is information frictions and not other differences between areas close and far from the capital city that lead to an asymmetric perception and trust in political leaders, we need to show that the difference in the perception of political leaders changes when their performance changes. Hence, if our hypothesis was true, we should observe that isolated citizens adjust their perception and evaluation of the political leadership less when the level of misgovernance and corruption changes. We can test this empirically by creating a panel of Afrobarometer rounds 5 (2011-2013), 6 (2014-2015) and 7 (2016-2018) and add information on government performance using the corruption perception index (CPI) by Transparency International as a proxy.[22] We define the 'curruption level' as $100 - CPI$,

[21] This result is in line with Campante and Do (2014) who find that people further away from US state capitals are less informed about state politics.

[22] Transparency International defines corruption as 'abuse of entrusted power for private gain' and produces

hence ranging from 0 to 100 with higher values indicating higher levels of corruption.

In Table 2.4, we analyze how citizens update their beliefs about their political leaders over time using a panel including all respondents[23] and a panel reduced to boundary segments. The estimate on the interaction of the corruption level and log distance from the capital city explicitly compares how people further from the capital city react relative to those closer to the capital city when the corruption level changes. The results indicate that when the corruption level rises, isolated citizens increase their corruption perception less than those closer to the capital city, vice versa. Similarly, isolated citizens lose less trust and reduce their marginal propensity to vote for the incumbent government less when the level of corruption and misgovernance is rising, vice versa.
Using the CPI as a proxy for government performance would be problematic when comparing levels across countries, as the level of the CPI could be related to endogenous characteristics such as institutions that could also be related to how citizens perceive the government perception of citizens. However, by including country fixed effects, we specifically focus on how changes in corruption relate to changes in perception and voting. One of the disadvantages of focussing on changes is that it introduces even larger measurement error in proxying government performance using the CPI. Since this measurement error is likely to be idiosyncratic and orthogonal to unobservables, the estimates suffer from attenuation bias - hence the absolute magnitude of the effects should be interpreted as a lower bound estimate.

These results have important repercussions on the functioning of effective accountability mechanisms that keep the actions of political agents aligned with the interest of the people. Since isolated citizens are less reactive to their political agents - including both penalties for bad and rewards for good government actions - national leaders are left with reduced incentives to allocate government resources to isolated areas. From the perspective of political leaders, the marginal increase in votes from improved government performance is lower in isolated areas. This can explain why we observe a reduced level and quality of public goods and services in areas isolated from the capital city. Moreover, the corresponding placebo tests indicate that these patterns and dynamics are specific to isolation

the CPI every year based on a variety of sources including assessments from several international institutions as well as a range of surveys with experts. Unfortunately, we cannot integrate earlier rounds of the Afrobarometer because the CPI uses a new methodology since 2012 and some of Afrobarometer questions that we use are not available for earlier rounds.

[23] Note that, as always, we exclude respondents within 20 km from the capital city by default.

TABLE 2.4
CHANNEL ANALYSIS: DYNAMIC POLITICAL SUPPORT

	Dependent variable:					
	Corruption Perception		Trust Pol. Leaders		Vote for Government	
	Panel	BDD-Panel	Panel	BDD-Panel	Panel	BDD-Panel
	(1)	(2)	(3)	(4)	(5)	(6)
Log Corruption Level × Log Dist CAP	-0.392***	-0.464	0.177*	1.318**	0.165***	0.963**
	(0.113)	(0.703)	(0.100)	(0.576)	(0.058)	(0.448)
	Polynomials for: distance from the border					
3rd order	-	x	-	x	-	x
Log Corruption Level	*absorbed*	Yes	*absorbed*	Yes	*absorbed*	Yes
Log Distance from the Capital City	Yes	Yes	Yes	Yes	Yes	Yes
Geography Controls	Yes	Yes	Yes	Yes	Yes	Yes
Household Controls	Yes	Yes	Yes	Yes	Yes	Yes
Cnt × Ethn × Round FE	Yes	No	Yes	No	Yes	No
Country FE	No	Yes	No	Yes	No	Yes
Ethnicity FE	No	Yes	No	Yes	No	Yes
Round FE	No	Yes	No	Yes	No	Yes
Segment × Round FE	-	161	-	161	-	149
Observations	77,712	4,089	86,678	4,755	63,137	3,195
Adjusted R^2	0.158	0.108	0.193	0.152	0.280	0.226

Note: This table reports the regressions on how Afrobarometer respondents adjust their beliefs about corruption among politicians, trust into the political leadership and whether they would vote for the incumbent government with regard to changes in degree of political misconduct as measured by the CPI as well as with regard to how far away they live from the capital city. The 'Household Cov.' include: age, age squared and sex of respondent. The 'Geographical Cov.' include: distance from the coast (in km), longitude and latitude (projected in km) and whether the household is in an urban or rural setting. The dependent variable in columns (1)-(2) is the normalized corruption perception among political leaders, in columns (3)-(4) it is the normalized level of trust into the political leadership and in columns (5)-(6) it is a dummy indicating whether the respondent would vote for the incumbent government if there was an election held. All observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by Afrobarometer cluster and ethnicity × round. *p<0.1; **p<0.05; ***p<0.01

from the capital city and do not apply to isolation from placebo cities (see Tables 2.A.22 and 2.A.23).

2.3.4 Market Access and Trade

As an alternative potential channel, distance to the capital city might affect economic growth through reduced market access (Redding and Sturm, 2008; Buys et al., 2010; Bosker and Garretsen, 2012; Storeygard, 2016; Donaldson and Hornbeck, 2016; Jedwab and Moradi, 2016; Gibbons and Wu, 2017; Jedwab et al., 2017). Places that are farther away from the capital city might be farther from markets in general and face higher costs when buying and selling intermediate and final goods and services. This, in turn, might have negative consequences on the opportunities for economies of scale and productivity growth.

For this reason, we test if segment sides closer to the capital city exhibit higher levels of market access, MA_s. Since in a BDD framework market access would by construction we balanced across the boundary under the assumption of open borders, for this exercise we restrict market access to within the country and keep in mind that this inevitably overestimates potential differences in market access across the border. We apply the conceptual framework of Harris (1954), Donaldson and Hornbeck (2016) and other recent applications in the African context by Chiovelli et al. (2018) and Jedwab and Storeygard (2020) and approximate market access as:

$$MA_s \approx \sum_{d=1}^{D} \tau_{s,d}^{-\theta} \, N_d \tag{2.6}$$

In Equation 2.6, $\tau_{s,d}$ refers to travel time between segment side s and all destination cities d within the country, θ represents the trade elasticity parameter that captures how trade declines with travel time, and N_d measures the market size that we approximate with the destination city population. We obtain a comprehensive database of geolocated African cities including their population size from Africapolis and compute travel times using OpenStreetMaps via the OSRM routing engine. Regarding the trade elasticity parameter, we follow Donaldson and Hornbeck (2016), Chiovelli et al. (2018) and Jedwab and Storeygard (2020) and use 3.8 as default but also report results when alternatively using 2.79 and 4.46 based on the estimations conducted by Simonovska and Waugh (2014). The results in Table 2.5 indicate that there is no significant different in market access in areas close or far from the capital city when we exclude the capital city itself irrespective of which trade elasticity parameter we chose. Moreover, the fact that market access is significantly reduced when including the capital city itself is not surprising as distance to the capital city mechanically decreases market access if the capital city is itself defined as a market. Nevertheless, since capital cities constitute important markets within the country, these effects might be significant. Yet, if the capital city actually mattered as a market rather than a capital city, we should observe similar effects for remoteness from other major markets within the country. The fact that, even when accounting for city and hence market size, the coefficients associated with placebo cities are insignificant (see Table 2.A.19) raises doubt on the relevance of market access in this context.

Another way of testing for this channel makes use of the fact that the shared national boundaries in our sample feature different degrees of border permeability with regard to trade. If borders did not constitute barriers for trade, in our BDD design, counterfactual pixels on both sides of the boundary would by construction have the same market access.

TABLE 2.5
CHANNEL ANALYSIS: MARKET ACCESS

	Dependent variable: Market Access Index					
	Trade Elasticity Parameter					
	$\tau = 2.79$		$\tau = 3.8$		$\tau = 4.46$	
	incl. CAP	excl. CAP	incl. CAP	excl. CAP	incl. CAP	excl. CAP
	(1)	(2)	(3)	(4)	(5)	(6)
Log Distance from the Capital City	-0.772*** (0.169)	-0.220 (0.193)	-0.748*** (0.275)	-0.208 (0.306)	-0.716** (0.348)	-0.205 (0.381)
Geography Cov.	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	35	35	35	35	35	35
Segment FE	551	551	551	551	551	551
Observations	1,104	1,104	1,104	1,104	1,104	1,104
Adjusted R^2	0.692	0.666	0.630	0.611	0.607	0.592

Note: This table reports the results on market access statistics. For each segment side we estimate a market index indicator based on Equation 2.6 using the OSRM routing engine and applying different trade elasticity parameters. Columns (2), (4) and (6) exclude the capital city itself as a destination market. The 'Geographical Cov.' include a segment side's average of distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km). Boundary segments corresponds to a buffer of 25 km around border pieces of 50 km line length and are entirely nested within a restricted ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). Standard errors in parenthesis are clustered by boundary segment. *p<0.1; **p<0.05; ***p<0.01

Therefore, if market access was indeed a relevant channel, we should observe that the effects of remoteness from the capital city are strongest at relatively closed and lower at relatively open national boundaries. In Table 2.A.24, we therefore test whether the impact of isolation from the capital city is lower at boundaries within trade blocs such as within free trade agreements (FTA), customs and monetary unions. For this purpose, we include interactions between *Log Distance from the Capital City* and dummies indicating low barrier boundaries into our main BDD model (see Equation 2.4). In columns (1) and (2) in Table 2.A.24, we focus on boundaries within FTAs. The results suggest that the effects are stronger within FTAs than at more restrictive boundaries which is in contrast to the idea that market access is driving the results. In columns (3) and (4), we tighten the criterion for relatively open borders and only consider customs unions and in columns (5) and (6) only those that also share a common currency. In both cases, we cannot reject the hypothesis that the effects at relatively open boundaries are any different than at restrictive ones.

2.3.5 Conflict

There is vast empirical evidence that conflict has negative implications for economic development (see for example Ray and Esteban (2017) for a general overview and Serneels and Verpoorten (2015) or Besley and Reynal-Querol (2014) for evidence from the African continent). Further, there are multiple ways in which isolation from the capital city might affect conflict, hence ultimately economic performance. Since a conflict or protest farther away from the capital represents a lower threat to a government (Johnson and Thyne, 2018; Campante et al., 2019), the state might be less inclined to prevent or resolve such isolated conflicts. Also, the capacity of a state to counter conflicts in isolated areas might simply be restricted and even attract conflict parties to target remote areas (Müller-Crepon et al., 2021). Moreover, a potentially lower state presence might impact the propensity for ethnic cleavages and conflict. Lastly, conflict might be a result of inequalities between areas close and far from the capital city and thereby reinforce the adverse implications of isolation from the capital city. On the other hand, conflict might be less prevalent in isolated areas because isolated citizens consume the news less frequently and potentially have less media and communication channels to get access to information or coordinate for protest (Manacorda and Tesei, 2020).

In order to assess whether conflict is actually relevant in this context, we test whether there is any pattern of increased or perhaps even decreased conflict in isolated areas. To this end, we use our ACLED dataset of conflicts by type (Violent Events, Demonstration Events and Non-Violent Action) (Raleigh et al., 2010). Out of the total 88,853 conflict events in our sample countries between 01.01.2000 and 27.11.2019, a total of 3,446 fall into the area of our boundary segments. We aggregate the frequency of conflict that fall within each segment side and run BDD regressions with segment sides as the observational unit (see Table 2.6). Since conflict frequency might directly increase with population density, in columns (5)-(8), we additionally account for the total population count in each segment side. Our results demonstrate that there is no significant relationship between distance from the capital city and the frequency of any type of conflict (or all types of conflicts aggregated). This result is in contrast to Campante et al. (2019) who find that conflict is more likely to emerge closer to the capital city. One reason for the different findings could be related to spatial spillovers of conflicts across the border along ethnic lines which is common in Africa (Bosker and de Ree, 2014; Michalopoulos and Papaioannou, 2016). As a results we might underestimate the true impact of remoteness on conflict. However, even if this was the case, as long as conflict has a similarly adverse effect on economic

performance irrespective of whether it originates or was spilled over to a segment side, it is very unlikely to be related to the difference in economic outcomes between segment sides close and far from the capital city.

TABLE 2.6
CHANNEL ANALYSIS: CONFLICT

	Dependent variable: Conflict frequency (ACLED)							
	Viol	Demo	Non-V	All	Viol	Demo	Non-V	All
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Distance from the Capital City	0.356	0.061	0.130	0.547	0.413	0.186	0.151	0.751
	(0.720)	(0.632)	(0.146)	(1.344)	(0.732)	(0.573)	(0.149)	(1.318)
Population Count in Segment Side	-	-	-	-	0.026	0.056**	0.010*	0.092*
					(0.021)	(0.026)	(0.006)	(0.050)
Country FE	35	35	35	35	35	35	35	35
Segment FE	568	568	568	568	568	568	568	568
Observations	1,138	1,138	1,138	1,138	1,138	1,138	1,138	1,138
Adjusted R^2	0.415	0.055	0.188	0.325	0.423	0.247	0.246	0.377

Note: This table reports the balancing tests on conflict with regard to isolation from the capital city. The observational unit in these models are the BBD boundary segments. The dependent variables are instances of conflict by type in the respective side of the boundary segment in the period between 01.01.2000 and 27.11.2019. Column (1) refers to 'violent conflicts', column (2) to 'demonstrations', column (3) to the number of 'non-violent actions' and column (4) aggregates all three kinds of conflict (for more information on these categories please refer to the ACLED homepage at: `https://www.acleddata.com/resources/general-guides/`). Boundary segments corresponds to a buffer of 25 km around border pieces of 50 km line length and are entirely nested within a restricted ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). Standard errors in parenthesis are clustered by boundary segment. *p<0.1; **p<0.05; ***p<0.01

To conclude, our analysis suggests that public goods provision is an important mediator between remoteness from the capital city and economic performance. Furthermore, our results support the view that dysfunctional accountability mechanisms are important regarding the reduced level of public goods provision. In contrast, based on our analysis, it is unlikely that market access, trade or conflict are relevant with regard to the implications of isolation from the capital city.

2.4 Conclusion

We investigate the impact of isolation from the capital city on economic development in Sub-Saharan Africa using extensive remote sensing data and large collections of geocoded and survey data. We obtain quasi-random variation in treatment at arbitrarily set national borders that divide ethnic homelands with similar geographical, social and historical characteristics. Conducting our analysis in a BDD regression framework, we deliver tangible

evidence that isolation from the capital city imposes strong adverse effects on the level of local economic performance. We perform a series of alternative specifications, balancing and robustness tests that underline that our estimates are in fact causal. Moreover, comparing the effects of isolation from the capital city to isolation from other major cities confirms that hosting the political center is the driving force behind the effects of isolation.

In addition, we investigate potential channels through which isolation from the capital city might affect economic performance: public goods provision, market access and conflict. We document that remoteness from the capital city, as opposed to other major cities, is linked to a significant drop in the level of public goods provision. In order to understand the imbalances in public goods, we explore two potential explanations: (i) geographical isolation translates into political isolation and thereby exclusion from government resources and (ii) due to limited accountability political leaders have lower incentives to allocate government resources to isolated areas. Our findings regarding political representation are ambiguous - the head of state is more likely to come from regions closer to the capital city but remote regions are more often part of the coalition in power. In contrast, our findings provide clear support for the accountability channel. Despite receiving less public goods, isolated citizens have more trust in their national political leaders, evaluate their performance better, believe less that they are corrupt or that their actions should be monitored. At the same time, people in isolated areas follow the news less frequently. We interpret this as reflecting an asymmetry in information about political affairs between areas far and close from the capital. We confirm this hypothesis empirically by showing that citizens in isolated areas are less sensitive to changes in the quality of governance - they are less likely to withdrawal government support in times of rising misgovernance and corruption but also reward them less for good governance. As a consequence, political leaders that seek to gain popular support with limited government resources are incentivized to allocate more public goods to areas closer to the capital where the expected political return to investment is higher. We therefore conclude that dysfunctional accountability mechanisms are likely to be a root cause behind the unequal spatial distribution of public goods and economic development in Sub-Saharan Africa.

Our findings are novel in the literature and provide new insights into the political economy of the location of the capital city (Campante and Do, 2014; Campante et al., 2019). Additionally, we provide new insights into the literature regarding the limited institutional outreach of the state beyond the capital city (Herbst, 2000; Michalopoulos and Papaioan-

nou, 2013) and the literature on power sharing and political representation in SSA (Francois et al., 2015). Furthermore, we add to the literature about information, accountability and public goods provision (Besley and Burgess, 2002; Strömberg, 2004; Guriev et al., 2020). Last but not least, by identifying proximity to the capital as an important dimension of spatial inequality, we contribute to the debate about the causes of the very high levels of regional economic disparity in Sub-Saharan Africa.

The results of this study underline the importance of considering political and, in particular, accountability mechanisms when seeking to understand the reasons for the large differences in living standards across African regions. Accordingly, policy makers should aim at strengthening information and feedback channels between citizens and their political leaders to improve and align public policy more with the interest of the people. Ensuring a good functioning of local media markets or increasing the awareness about the importance of active civic engagement in the political process represent two viable actions to this end. Lastly, targeted investments into public goods in isolated areas could provide remedy by boosting the local economy in these areas currently lagging behind.

2.A Appendix

2.A.1 Data Sources and Description

In the following, we present the various datasets that we collected from different sources and merged into the final dataset for the remote sensing analysis including an URL for download. Furthermore, we present in more detail our data sources for the mechanism analysis.

Remote Sensing Grids

Regarding nighttime luminosity data, there are currently two products: the 'new' VIIRS (Visible Infrared Imaging Radiometer Suite) by the Suomi National Polar Partnership between NOAA and NASA and the 'old' Version 4 DMSP-OLS (Defense Meteorological Satellite Program - Operational Linescan System) by the U.S. Air Force and the National Oceanic and Atmospheric Administration (NOAA). The VIIRS images are more recent and are superior to the DMSP-OLS with regard to their accuracy and resolution (Elvidge et al., 2013). We use the latest cleaned annual VIIRS product from 2016 that underwent extensive filtering including outlier (such as fires or ephemeral lights) and background (non-lights) removal as our main proxy for economic activity (see Elvidge et al. (2017) for details on the algorithms used to pre-process and filter the annual VIIRS images). This dataset is publicly available and can be downloaded from: `https://www.ngdc.noaa.gov/eog/viirs/download_dnb_composites.html`. The drawback of the VIIRS is that the earliest available grids (unfiltered and monthly) are available for April 2012 while the DMSP-OLS span from 1992-2013. Therefore, in order to obtain an 'early' disaggregated proxy for economic activity, we supplement our dataset with the annual DMSP-OLS nighttime lights from 1992 (available for download at: `https://www.ngdc.noaa.gov/eog/dmsp/downloadV4composites.html#AVSLCFC`).

We combine the nightlights data with a wide range of geographical covariates from various remote sensing data sources (see below). For the purpose of curbing measurement error and mismatching resulting from small inaccuracies between the datasets and to facilitate the computational intensity of the analysis, the data grids are aggregated to a resolution of 75×75 arcseconds which is equivalent to approximately 2.3×2.3 kilometers at the equator. The nightlight grids have an initial resolution of 15×15 arcseconds which means that 25 original pixels (5×5) constitute a new pixel and inherit the average value of its predecessors. This leaves us with around 3.5 million pixels for the 38 countries in our

sample. In order to be able to stack nighttime lights and other geographical covariates together, we interpolate all other grids bilinearly to match the nightlight grid resolution. We further transform our grids from the standard CRS 'WGS 84' using decimal degrees to 'Africa Sinusoidal' which properly maps distances in Sub-Saharan Africa using the metric system (see `http://spatialreference.org/ref/esri/africa-sinusoidal/` for more information).

The water surface grid (5×5 arcseconds) is aggregated to reflect the percentage of the surface that is covered with water and was obtained from the European Space Agency (ESA). It is publicly available for download at: `http://maps.elie.ucl.ac.be/CCI/viewer/index.php`.

The land elevation (in m) grid (30×30 arcseconds) by the NASA in context of the Shuttle Radar Topography Mission (SRTM) based on the work of Jarvis et al. (2008) is available for download at: `http://www.cgiar-csi.org/data/srtm-90m-digital-elevation-database-v4-1`.

A measure for terrain ruggedness measured in degree of slope with an initial resolution of (20×20 arcseconds) based on Nunn and Puga (2012) was obtained from: `https://diegopuga.org/data/rugged/`.

The 'Crop Caloric Index', a measure of agricultural suitability containing the potential agricultural caloric output per year and hectar (excluding zero yields) based on Galor and Özak (2016) was obtained from `https://ozak.github.io/Caloric-Suitability-Index/`.

Annual precipitation (in mm), annual mean temperature, minimum temperature in the coldest month and maximum temperature in the warmest month (all in °C) based on Karger et al. (2017) were obtained from `http://chelsa-climate.org/downloads/`.

The population grid (30×30 arcseconds) for 2015 was obtained from Worldpop (available for download at: `http://www.worldpop.org.uk/data/summary/?doi=10.5258/SOTON/WP00004`) and contains the (UN-adjusted) total number of inhabitants per pixel.

Vectorized Data

The list of capital and placebo cities including their population size were obtained from various sources including United Nations Department of Economic and Social Affairs Population Division (2018) (https://population.un.org/wup/Download/), CityPopulation (see here for further information: https://www.citypopulation.de/) and WorldPopulationReview (see here for further information: http://worldpopulationreview.com/). In a next step, these cities were geocoded in R using an OpenStreetMap (https://www.openstreetmap.org/) interface.

The country shapefiles were obtained from GADM (currently Version 3.6) and are available for download at: https://gadm.org/data.html.

Based on the above sources we have calculated the following indicators on the pixel level using the projected coordinate reference system 'Africa Sinusoidal' which properly maps distances in Sub-Saharan Africa using the metric system (in km): distance from the capital city, distance from the placebo city, distance from the coast, distance from shared national boundaries, latitude and longitude.

The ethnographic 'Tribal Map of Africa' based on Murdock (1959) was recently digitized by Nathan Nunn and is available for download at: https://worldmap.harvard.edu/data/geonode:murdock_ea_2010_3.

The African rivers data is provided by FAO and is derived from the World Wildlife Fund HydroSHEDS drainage direction layer and a stream network layer and can be downloaded at: http://www.fao.org/geonetwork/srv/en/main.home?uuid=b891ca64-4cd4-4efd-a7ca-b386e98d52e8.

We obtained the shapefile on African waterbodies that was create by the Regional Centre for Mapping of Resources for Development (RCMRD) from the World Bank https://datacatalog.worldbank.org/dataset/africa-water-bodies-2015.

The dataset of African cities including their population size for the market access analysis was created by the OECD Sahel and West Africa Club in collaboration with e-geopolis.org and can be downloaded from: https://africapolis.org/data.

The dataset on African roads and road quality was obtained from Jedwab and Storeygard (2020) who digitized an geocoded Michelin maps.

The data on medical centers based on an extensive data collection by Maina et al. (2019) can be obtained at: `https://www.who.int/malaria/areas/surveillance/public-sector-health-facilities-ss-africa/en/`

Afrobarometer Data

The Afrobarometer surveys comprise questions on household characteristics, public goods, public perception, as well as political attitudes and opinion. The surveys are conducted every few years in most African countries and provide the geolocation of their respondents which makes it highly suitable for studies in economic geography. We use rounds 5, 6 and 7 of the Afrobarometer covering 26 out of the 38 countries in our sample as is illustrated in Figure 2.A.8b. The data is available upon application at: `https://www.afrobarometer.org`.
The Afrobarometer is especially useful to investigate the spatial pattern of public goods provision as it contains information on whether the enumeration area (cluster) in which the respondent is located is supplied with certain public goods: paved road, electricity grid, piped water system and sewage system. Moreover, based on the Afrobarometer, we can examine how distance from the capital city affects corruption perception, trust into the political leadership, news readership and other characteristics reflecting how people think about their leaders and the political organization of their country. A complete list of all Afrobarometer survey questions and indicators can be found in Section 2.A.1.

In the following, we provide details about the construction of the indicators that are based on the Afrobarometer survey data including extracts from the underlying questions in the Afrobarometer codebook (Isbell, 2017). Note that respondents who answered 'Do not know', 'Refused to answer' or 'Missing' are excluded from the sample for the respective indicator.

Public Goods Index

Average of responses to each of the following questions. Note that the response was filled in by the interviewee prior to approaching the individual households within the cluster/ enumeration area:

Piped Water System: Are the following services present in the primary sampling unit/ enumeration area: Piped water system that most houses could access? Value Labels: 0=No, 1=Yes.

Electricity Grid: Are the following services present in the primary sampling unit/ enumeration area: Electricity grid that most houses could access? Value Labels: 0=No, 1=Yes.

Paved Road: Thinking of your journey here: Was the road at the start point in the primary sampling unit/ enumeration area paved/ tarred/ concrete? Value Labels: 0=No, 1=Yes.

Sewage System: Are the following services present in the primary sampling unit/ enumeration area: Sewage system that most houses could access? Values Labels: 0=No, 1=Yes.

Trust into the Political Elite

Average of responses to each of the following questions (see parenthesis): How much do you trust each of the following, or have you not heard enough about them to say: *The President (Parliament)*? Value Labels: 0 = Not at all, 1 = Just a little, 2 = Somewhat, 3 = A lot.

Corruption Perception of Political Leadership

Average of responses to each of the following questions (see parenthesis): How many of the following people do you think are involved in corruption, or haven't you heard enough about them to say: The President and Officials in his Office (Members of Parliament, Government Officials)? Value Labels: 0 = Not at all, 1 = Just a little, 2 = Somewhat, 3 = A lot.

Evaluation of Government Performance

Average of responses to each of the following questions (see parenthesis):

Now let's speak about the present government of this country. How well or badly would you say the current government is handling the following matters, or haven't you heard enough to say: Managing the economy (handling improving living standards of the poor, handling creating jobs, handling keeping prices down, handling narrowing income gaps, handling reducing crime, handling improving basic health services, handling addressing educational needs, handling providing water and sanitation services, handling ensuring enough to eat, handling fighting corruption, handling and maintaining roads and bridges, handling providing reliable electric supply)? Value Labels: 1 = Very badly, 2 = Fairly

badly, 3 = Fairly well, 4 = Very well.

Voter Turnout

Understanding that some people were unable to vote in the most recent national election in [20xx], which of the following statements is true for you? Value Labels: 0 = You were not registered to vote Or you decided not to vote Or you could not find the polling station Or you were prevented from voting Or you did not have time to vote Or you did not vote because you could not find your name in the voters' register Or Did not vote for some other reason, 1 = You voted in the elections

National vs. Ethnic Identity

Let us suppose that you had to choose between being a [NATIONALITY] and being a [Respondent's Ethnic Group]. Which of the following best expresses your feelings? Value Labels: 1 = I feel only (Respondent's ethnic group), 2 = I feel more (Respondent's ethnic group) than [NATIONALITY], 3 = I feel equally [NATIONALITY] and (Respondent's ethnic group), 4 = I feel more [NATIONALITY] than (Respondent's ethnic group), 5 = I feel only [NATIONALITY].

Vote for Government

Based on the response to the following question and the time of the interview we decode whether the respondent expresses to vote for the political party of the incumbent president (value of 1) or not (value of 0).

If a presidential election were held tomorrow, which party's candidate would you vote for? The variable labels correspond to the political parties in the country at the time of the interview.

Advocate Check and Balances

Average of support for statement 2 in the two respective questions:

Checks by citizens:
Statement 1: It is more important to have a government that can get things done, even if we have no influence over what it does.
Statement 2: It is more important for citizens to be able to hold government accountable,

even if that means it makes decisions more slowly.

Checks by parliament:
Statement 1: The President should be able to devote his full attention to developing the country rather than wasting time justifying his actions.
Statement 2: Parliament should ensure that the President explains to it on a regular basis how his government spends taxpayers' money.

Value Labels: 1 = Agree very strongly with Statement 1, 2 = Agree with Statement 1, 3 = Agree with Statement 2, 4 = Agree very strongly with Statement 2, 5 = Agree with neither.

News Consumption

How often do you get news from any of the following sources: Radio, Television, Newspaper or Internet? 0 = Never, 1 = Less than once a month, 2 = A few times a month, 3 = A few times a week, 4 = Every day.

Public Goods Data

To measure road infrastructure provision, we use a road network dataset created by Jedwab and Storeygard (2020) based on 2014 Michelin maps. The great advantage of this dataset is that it consistently maps a variety of roads across the African continent and provides information about the quality of the infrastructure such as whether or not a road is tarred.

We measure health care provision using an extensive database on the locations of 98,745 health facilities across 50 African countries sourced from a variety governmental and non-governmental sources (Maina et al., 2019). This dataset also includes information about whether the health facility is a hospital or lower tier facility and whether it is a public (e.g. run by the ministry of health or municipality) or non-public (e.g. NGOs, private or denominational) institution. Since the data for Angola and Guinea-Bissau is incomplete, we omitted those countries from our analysis.

Access to Power

In order to explore whether areas close and far from the capital city are, conditional on their respective population shares, equally represented in the government, we require a dataset on the origins of state executives. Francois et al. (2015) create an extensive dataset on the ethnicity of cabinet ministers in African countries since independence. Based on

their analysis, the authors are able to infer that ruling coalitions are large and ethnicities represented proportionally. However, the authors do not study if geographical locations are equally represented. Since their dataset is not geocoded and there exists no complementary dataset that corresponds to their categorization of ethnic groups, we are unable proceed with their database. For this reason, we limit our analysis to the heads of state using and supplementing the database by (Dreher et al., 2016, 40-41). We update this list regarding the period of office and region of origin up to the year of 2016 using publicly available information from various sources including Wikipedia.

Yet, solely mapping out the origins of heads of state might be imprecise and not give a holistic image about an area's access to power. Therefore, we complement the analysis with a second dataset comprising the degree of access to power for various politically relevant ethnic groups from 1946 to 2017: the Ethnic Power Relations Core Dataset 2019 (EPR) (Vogt et al., 2015). The advantage of the EPR dataset is that it can be combined with a complementary dataset on the geolocation (polygons) of the respective ethnic groups (Wucherpfennig, 2011). The EPR status indicators categorizes groups as *in power*: 'monopoly', 'dominance', 'senior partner', 'junior partner' as well as *excluded groups*: 'powerless', 'discrimination', 'self-exclusion' and areas that have are not been identified as being politically relevant which are coded as *irrelevant*. To obtain an image of contemporary patterns, we restrict the dataset to the period 2000-2017. We assign each pixel the total number of years in each respective status category as our variable of interest. Since some of the groups are overlapping geographically, we assign each area the status of the respective most powerful group.

Conflict Data (ACLED)

In order to investigate whether conflict is a relevant channel linking remoteness from the capital city and economic development, we use the Armed Conflict Location & Event Data (ACLED) containing the geolocation of conflicts between 1997 and today by type: 'violent events', 'demonstration events' and 'non-violent action' (Raleigh et al., 2010). For more information on the definition of different conflict types refer to the ACLED homepage at: `https://www.acleddata.com/resources/general-guides/`. Our conflict event sample comprises 88,853 instances between 01.01.2000 and 27.11.2019 distributed all over our sample countries.

2.A.2 Further Technical Details

In order to identify boundary segments that partition ethnic homelands, we use the ethnographic 'Tribal Map of Africa' on pre-colonial ethnic homelands by Murdock (1959) that is widely used in economics when marking out different ethnicities (Nunn, 2008; Michalopoulos and Papaioannou, 2013, 2014, 2016). Murdock subdivides the entire African continent into 843 distinct polygons that each relate to a local ethnic majority group. The map refers roughly to the era around the Berlin Conference and the establishment of artificial African boundaries. This circumstance makes this map well suited for our purpose as it is free from more recent and potentially endogenous migration movements. Moreover, "case study and anecdotal evidence suggest that in spite of population movements ethnic populations tend to reside in their respective historical homelands" (Michalopoulos and Papaioannou, 2014, 162). Since Murdock's map was originally printed in his book and only later digitized by Suzanne Blier and Nathan Nunn, there might be some inaccuracies with the precise delimination of ethnic homelands (Michalopoulos and Papaioannou, 2013, 143). To account for this potential shortcoming, we restrict each ethnic homeland by (a negative buffer of) -15 km prior to further processing the map.
In the next step, we subdivide all national borders that divide ethnic homelands into segments of 50 km length with a buffer of 25 km (15km, 50 km in sensitivity tests) on each side of the boundary. To verfy that segments are sufficiently small and homogenous, we conduct a range of balancing tests. In our econometric model, we implicitly assume that both segment sides are of equal size. However, in practice, since national borders are not always straight lines, opposing segment sides are not of equal size in all instances either. Therefore, we weight each pixel such that both segment sides have the same aggregated weight.[24]

For the purpose of preventing that inaccuracies and blooming confound the assignment of nightlights to a segment side (country), we exclude pixels whose centroid is within a range of 3 km (1.5 km in sensitivity tests) from the boundary from all estimations. Further, to account for uninhabitable areas, we exclude pixels that are entirely covered with water or are completely unpopulated. Additionally, in order to ensure that the estimates of isolation from the capital city are not simply capturing the break between the capital city and the hinterland, we exclude pixels within a radius of 20 km (but also 50 km, 75 km and 100 km in sensitivity tests) around the capital city (and placebo city) by default from all

[24] We maintain the overall weight of segments relative to each other (proportional to their total segment size). Further, in order to avoid overrepresenting very small segment sides, if a segment side contains less than 20 pixels, we drop the entire segment.

estimations.

Last but not least, our causal estimate of isolation from the capital city might be distorted by spillovers across the boundary. It might be that those living in the economically weak segment side migrate to the more dynamic segment side to seek employment (brain drain). Another possibility is that segment sides that are relatively disadvantaged benefit positively from the market access and resources on the other side (scale economies). Since negative spillovers would lead to overestimations and positive spillovers to underestimations of the effects, it remains unclear in what direction a potential bias would go. In addition, since potential spillovers at the boundary are likely a function of proximity, they are most probably being picked up by the polynomials terms of distance to the border.

2.A.3 Maps

FIGURE 2.A.1
VIIRS NIGHTLIGHTS SAMPLE MAP

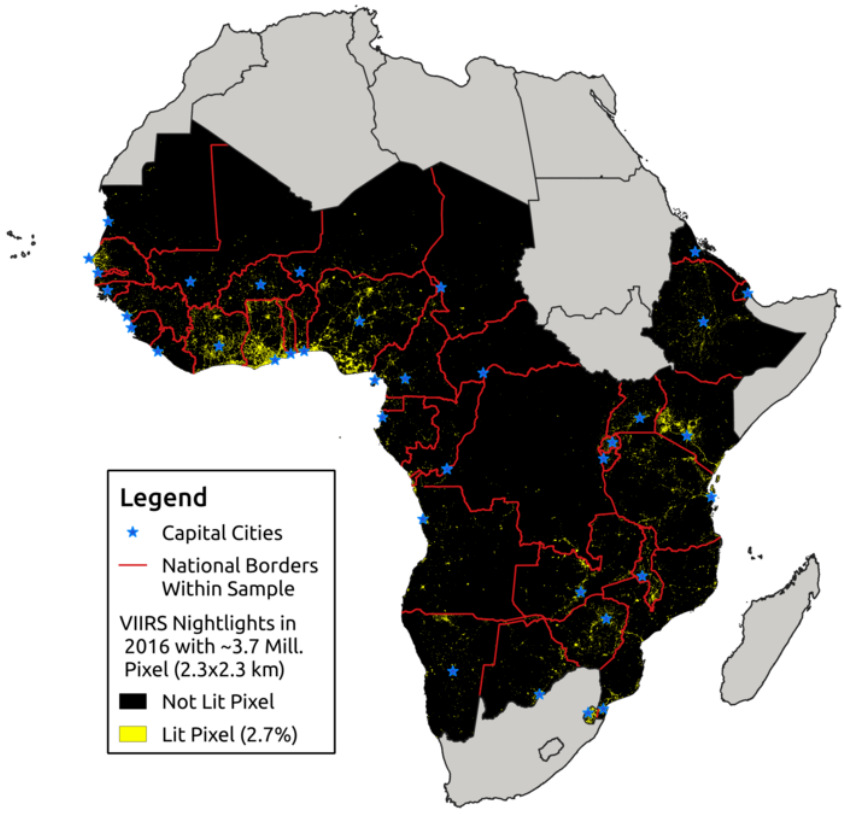

FIGURE 2.A.2
MAPPING SURVEY RESPONDENTS

(A) DHS SAMPLE

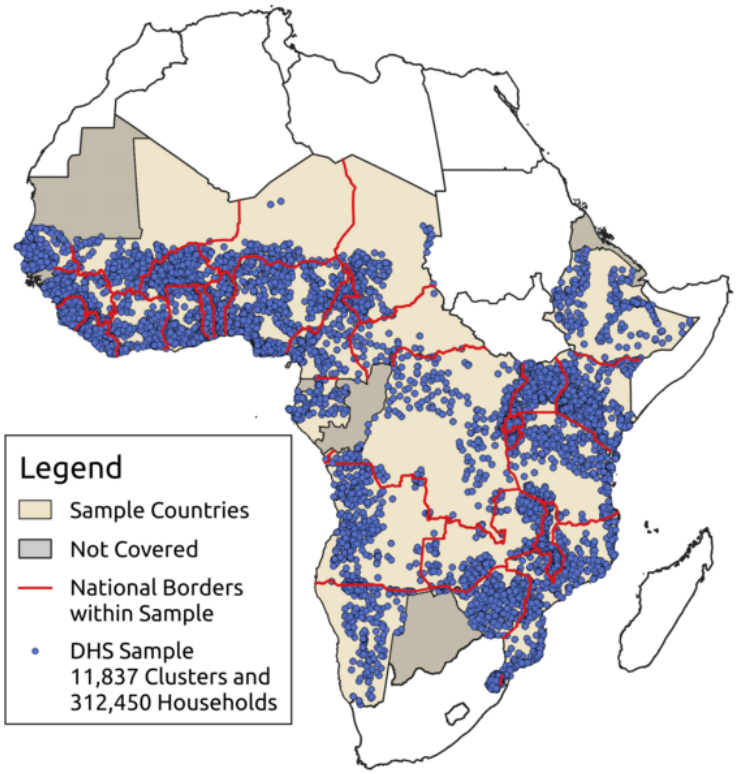

(B) AFROBAROMETER ROUND 5, 6 + 7 SAMPLE

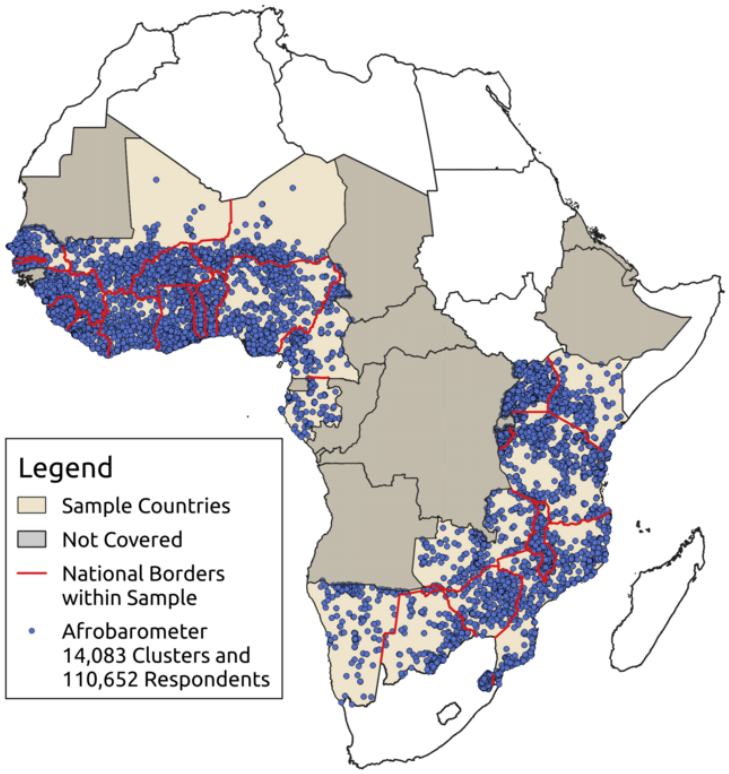

Figure 2.A.3
Unrestricted Segments

(a) Simple Border Segments

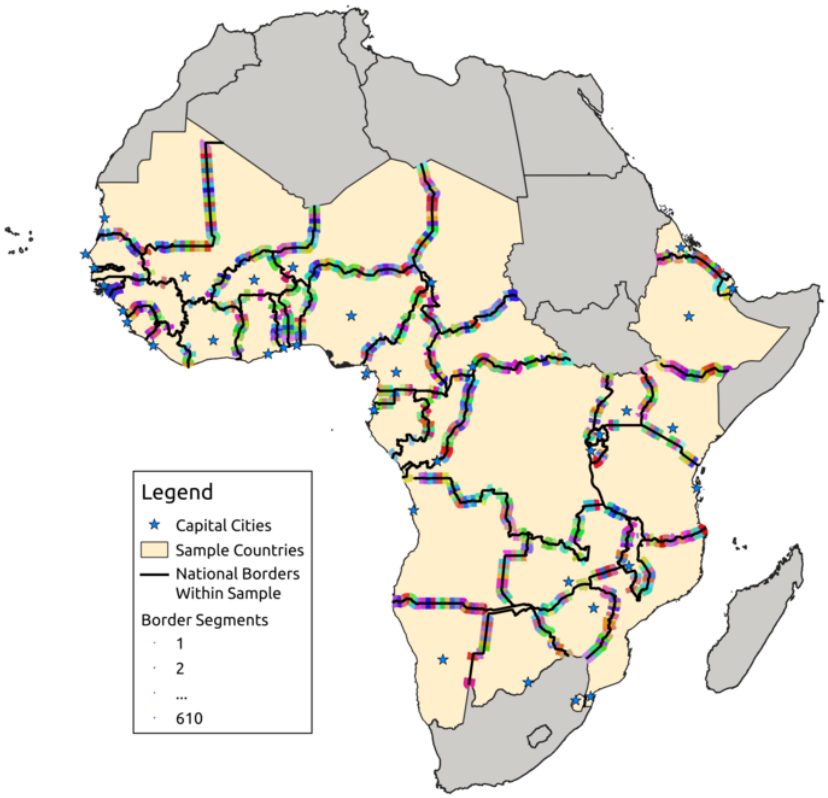

(b) Simple Border Segments - Magnified

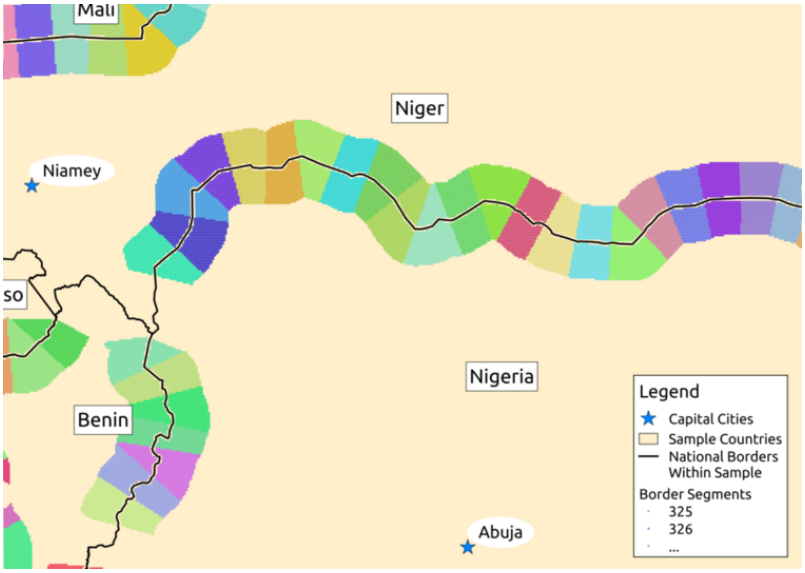

Figure 2.A.4
Restricted Segments

(a) Border Segments nested within Restricted Partitioned Ethnic Homelands

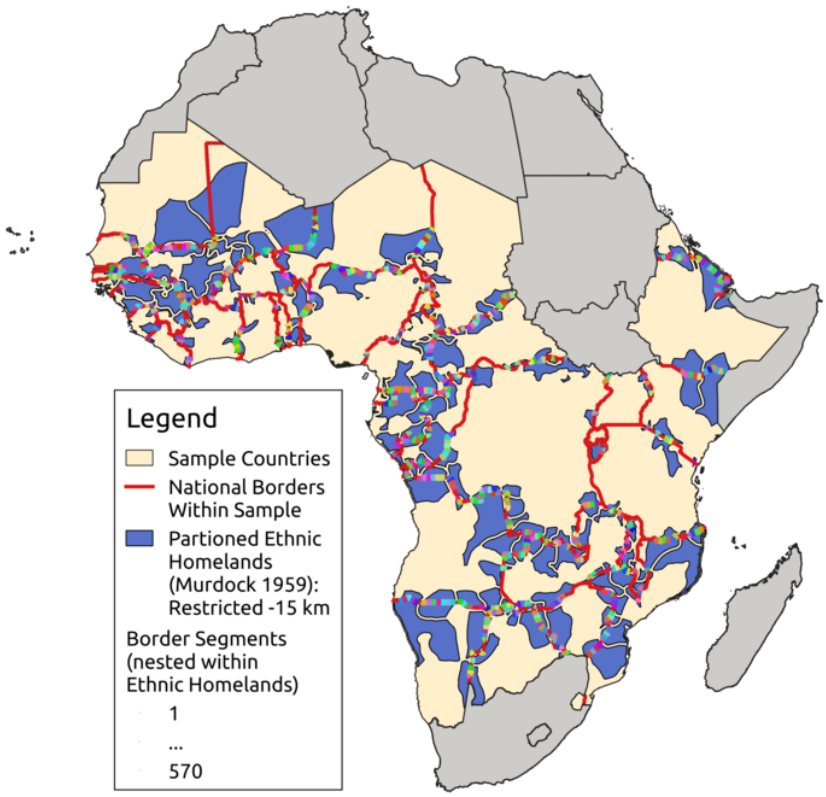

(b) Nested Border Segments - Magnified

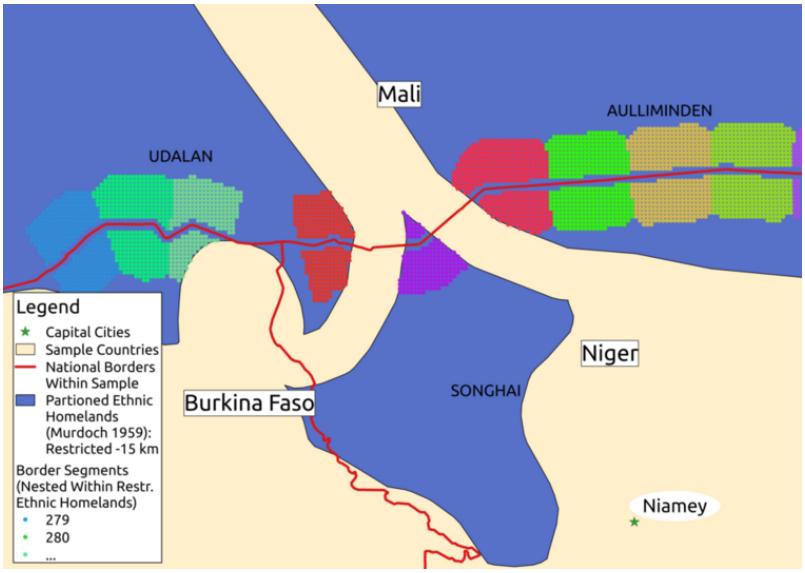

2.A.4 Summary Statistics

TABLE 2.A.1

SUMMARY STATISTICS - REMOTE SENSING DATA - ARITHMETIC MEAN BY COUNTRY (1/2)

Country Name	Light Density VIIRS 2016	Share Lit Pixels	Distance to the Capital	Distance to the Placebo	Distance to the Coast	Elevation (in m)	Ruggedness (in %)	% Surface covered in Water	Mean Temperature	Min Temperature	Max Temperature	Crop Caloric Index	Annual Precipitation	Longitude projected (in km)	Latitude projected (in km)
Angola	.024	1	700	477	476	1,071	2.3	.29	21	12	30	3,492	1,061	1,906	-1,360
Benin	.015	6	373	179	367	268	.38	.21	27	19	36	3,974	1,076	257	1,070
Botswana	.0095	1.9	422	479	784	1,035	.49	.082	22	6.8	32	1,549	390	2,453	-2,454
Burkina Faso	.0069	2.8	226	345	736	300	.64	.34	28	18	39	3,575	773	-189	1,358
Burundi	.0047	3.9	86	96	1,015	1,523	4.3	7.2	20	15	27	3,510	1,238	3,322	-375
Cameroon	.0045	2	353	447	433	650	1.4	.84	24	18	31	2,969	1,673	1,411	631
CAF	.0001	.075	453	457	1,221	600	.53	.2	25	18	33	3,482	1,366	2,265	728
Chad	.0035	.19	650	842	1,286	521	1.1	.28	27	13	39	1,432	312	2,000	1,700
Congo	.0079	.83	443	599	515	420	.41	.97	24	20	28	2,532	1,602	1,696	-91
Djibouti	.0036	1.7	91	91	37	441	6.5	2	28	19	39	11	180	4,638	1,302
DR Congo	.0037	.33	1,056	1,131	1,129	688	1.2	1.8	23	18	29	3,110	1,620	2,623	-318
Equat. Guinea	.034	6	299	107	84	494	1.3	1.2	23	20	27	2,003	2,519	1,155	183
Eritrea	.0018	.6	197	544	101	752	6.6	.85	27	18	36	831	283	4,169	1,702
Ethiopia	.0071	1.9	428	486	475	1,256	4.2	.68	23	15	31	3,035	861	4,357	954
Gabon	.032	1.5	330	370	219	379	.58	1.4	24	20	27	2,246	1,833	1,312	-69
Gambia	.0072	5.2	145	88	64	24	.91	3.9	28	19	37	3,488	797	-1,654	1,487
Ghana	.1	24	326	232	284	181	.59	2.8	27	21	35	3,625	1,242	-134	883
Guinea	.0078	2.8	370	403	287	443	2	.65	26	18	36	3,731	1,746	-1,198	1,156
Guinea Bissau	.0008	1.8	102	89	48	38	1.3	5.2	28	20	36	4,109	1,574	-1,631	1,328
Ivory Coast	.054	14	233	355	286	262	.59	.88	26	20	33	3,302	1,280	-613	844
Kenya	.016	6.1	364	584	451	785	1.7	2.1	25	19	32	3,007	630	4,210	62
Liberia	.0021	1.1	205	146	102	217	.7	.64	25	21	30	2,482	2,558	-1,028	712
Malawi	.015	8.7	214	310	508	854	2.2	21	22	13	30	4,062	1,113	3,716	-1,457
Mali	.0019	.43	808	822	1,038	320	.38	.28	29	14	42	1,086	309	-375	1,921
Mauritania	.0012	.26	724	776	534	263	.31	.089	28	13	42	233	101	-1,079	2,242
Mozambique	.014	2.5	1,110	634	209	354	1.6	2	24	15	31	3,767	1,020	3,775	-1,908
Namibia	.0088	1.9	391	646	330	1,088	2.5	.66	21	7.3	32	1,143	283	1,773	-2,450
Niger	.0017	.41	920	582	1,173	463	.48	.031	27	12	41	611	134	996	1,928
Nigeria	.13	8.1	382	675	499	330	.83	1.1	27	18	36	3,399	1,161	888	1,061
Rwanda	.031	13	68	117	1,056	1,688	8	6.4	19	14	25	3,327	1,173	3,332	-220
Senegal	.02	6.3	354	224	179	53	.64	1.8	29	18	38	2,861	658	-1,559	1,588
Sierra Leone	.0035	1.5	184	118	110	202	1.3	1.3	25	20	33	2,834	2,686	-1,297	948
Swaziland	.17	53	65	68	117	622	8.6	.41	19	8	27	3,934	820	3,138	-2,940
Tanzania	.0084	3.2	589	544	490	1,019	1.7	6.4	22	15	30	4,027	1,005	3,849	-693
Togo	.02	12	272	144	265	258	.74	.46	27	20	35	3,786	1,195	107	945
Uganda	.0067	4.1	217	242	977	1,155	2.1	15	22	18	28	3,873	1,214	3,604	141
Zambia	.019	2.9	443	411	942	1,120	1.4	1.8	21	11	32	3,960	1,015	3,009	-1,489
Zimbabwe	.013	6	292	295	517	973	3.2	1.2	21	8.9	30	3,330	680	3,144	-2,102
Overall	.016	2.5	615	627	716	671	1.4	1.5	24	15	34	2555	887	1,857	198

Note: This table reports the arithmetic mean of remote sensing variables by country excluding 20 km around the capital and placebo cities based on own calculations.

Table 2.A.2
Summary Statistics - Remote Sensing Data - Arithmetic Mean by Country (2/2)

Country Name	Capital Name	Capital Pop.	Placebo Name	Placebo Pop.	Decentralized In 2010	Decentralized 1990-2010	Democracy 2016	Democracy 1996-2016	GDP per capita 2016 (median)	GDP per capita 2016 (mean)	Number of Observations
Angola	Luanda	7,265	Huambo	620	0	0	0	0	1	1	238,832
Benin	Cotonou	683	Djougou	299	0	0	1	1	1	0	21,546
Botswana	Gaborone	269	Francistown	100	0	1	1	1	1	1	116,273
Burkina Faso	Ouagadougou	2,306	Bobo Dioulasso	794	1	1	1	0	0	0	51865
Burundi	Bujumbura	797	Muyinga	101	0	0	0	1	0	0	4,631
Cameroon	Yaounde	3,273	Douala	3,112	1	0	0	0	1	0	87,276
CAF	Bangui	815	Bimbo	268	0	0	1	1	0	0	116,642
Chad	N'Djamena	1,230	Moundou	138	0	1	0	0	0	0	246,434
Congo	Brazzaville	2,044	Pointe-Noire	1,047	1	1	0	0	1	1	63,746
Djibouti	Djibouti	541	Ali Sabieh	14	1	0	1	1	1	1	3,899
DR Congo	Kinshasa	12,100	Lubumbashi	2,101	0	0	0	1	0	0	437,353
Equat. Guinea	Malabo	297	Bata	332	0	0	0	0	1	1	4,781
Eritrea	Asmara	833	Assab	21	0	0	0	0	0	0	22,990
Ethiopia	Addis Ababa	4,040	Dire Dawa	344	1	1	0	0	0	0	213,500
Gabon	Libreville	769	Port-Gentil	148	0	1	1	0	1	1	49,380
Gambia	Banjul	426	Farafenni	103	1	1	0	0	0	0	1,790
Ghana	Accra	2,338	Kumasi	2,758	0	0	1	1	1	1	44,707
Guinea	Conakry	1,756	Nzerekore	195	1	1	1	1	0	0	46,305
Guinea Bissau	Bissau	516	Bafata	35	NA	NA	1	1	0	0	6,056
Ivory Coast	Yamoussoukro	231	Abidjan	4,659	0	0	1	1	1	0	60,325
Kenya	Nairobi	4,065	Mombasa	1,139	1	1	1	1	1	0	109,423
Liberia	Monrovia	1,318	Gbarnga	46	0	0	1	1	0	0	17,712
Malawi	Lilongwe	945	Blantyre	831	0	0	1	1	0	0	22,198
Mali	Bamako	2,292	Sikasso	320	1	1	1	1	0	0	245,165
Mauritania	Nouakchott	1,105	Nouadhibou	118	0	0	0	0	1	0	207,455
Mozambique	Maputo	1,100	Nampula	707	0	0	1	1	0	0	154,023
Namibia	Windhoek	380	Rundu	63	1	1	1	1	1	1	165,977
Niger	Niamey	1,143	Zinder	393	0	0	1	1	0	0	231,640
Nigeria	Abuja	2,591	Lagos	12,634	1	1	1	1	1	1	172,110
Rwanda	Kigali	986	Butare	90	1	0	0	0	0	0	4,323
Senegal	Dakar	2,830	Touba	753	1	1	1	1	1	0	37,753
Sierra Leone	Freetown	1,073	Bo	174	1	0	1	1	0	0	13,435
Swaziland	Mbabane	76	Big Bend	10	0	1	0	0	1	1	3,156
Tanzania	Dar es Salaam	5,409	Mwanza	890	1	1	1	0	1	0	176,892
Togo	Lome	1,671	Sokode	118	1	1	0	0	0	0	10,444
Uganda	Kampala	2,707	Gulu	150	1	1	0	0	0	0	44,715
Zambia	Lusaka	2,294	Kitwe	602	0	0	1	1	1	0	143,985
Zimbabwe	Harare	1,505	Bulawayo	645	0	1	1	0	1	0	76,778

Note: This table reports the arithmetic mean of remote sensing variables by country excluding 20 km around the capital and placebo cities based on own calculations. The population (Pop.) in the respective capital and placebo cities is in thousands of inhabitants.

Table 2.A.3
Summary Statistics - DHS - Arithmetic Mean by Country

Country Name	Z-Wealth Index	Distance to the Capital	Distance to the Coast	Urban Pop. Share	Longitude projected (in km)	Latitude projected) (in km)	Age of head of household	Number of household members	Mean Temperature	Min Temperature	Max Temperature	Precipitation Index	Elevation (in m)	Number of Observations
Angola	-0.33	482	242	0.39	1,686	-1,290	43	4.7	21	15	28	73	1,164	9,933
Benin	-0.27	214	204	0.33	243	904	45	5.3	28	23	33	78	187	10,161
Burkina Faso	-0.21	195	750	0.17	-216	1,365	45	5.6	29	22	35	66	311	10,643
Burundi	-0.27	87	1,017	0.045	3,326	-356	45	4.8	20	15	26	109	1,595	11,240
Cameroon	-0.069	412	398	0.48	1,306	742	46	5.1	26	21	31	117	513	8,824
CAF	-0.28	303	892	0.24	1951	603	42	4.6	25	19	31	123	605	2705
Chad	-0.19	287	1,028	0.13	1,761	1,192	43	5.8	28	21	36	71	414	7,797
DR Congo	-0.081	1,170	1,067	0.25	2,721	-251	43	5.2	23	18	29	122	864	8,655
Ethiopia	-0.28	335	521	0.19	4,231	982	44	4.8	21	14	28	93	1,863	8,967
Gabon	0.26	216	124	0.68	1,165	-12	47	4.2	26	22	30	124	210	4,849
Ghana	0.073	217	173	0.48	-120	766	44	3.5	28	23	32	84	184	8,416
Guinea	-0.32	308	211	0.18	-1,269	1,130	49	6.1	26	20	32	165	470	5,240
Ivory Coast	0.15	204	147	0.51	-562	708	44	5.1	27	22	32	104	167	6,829
Kenya	0.041	216	506	0.28	4,040	-46	45	4.1	21	15	27	115	1,539	30,365
Liberia	-0.15	176	106	0.35	-1,054	738	44	5.1	26	20	32	187	203	6,583
Malawi	-0.17	177	407	0.12	3,725	-1,605	43	4.5	23	17	29	73	907	19,326
Mali	-0.35	264	802	0.11	-692	1,442	47	5.6	29	22	36	65	314	6,018
Mozambique	-0.39	1,113	161	0.23	3,822	-1,931	43	4.4	25	19	30	72	375	9,906
Namibia	-0.25	486	445	0.36	1,812	-2,135	48	4.5	23	15	32	23	1,140	6,877
Niger	-0.31	487	885	0.12	691	1,533	45	5.9	29	22	36	38	356	4,694
Nigeria	-0.011	388	401	0.42	777	996	44	4.7	27	22	33	92	296	28,009
Rwanda	-0.26	68	1,062	0.09	3,323	-226	45	4.3	19	14	25	119	1,685	9,567
Senegal	0.075	211	88	0.36	-1,700	1,608	52	9.1	28	21	35	54	33	4,071
Sierra Leone	-0.31	181	99	0.19	-1,294	926	46	6	27	22	32	226	160	10,159
Swaziland	-0.18	60	121	0.28	3,135	-2,940	46	4.7	21	14	28	50	622	3,987
Tanzania	-0.19	586	483	0.23	3,863	-593	46	5	23	17	28	84	1,101	8,665
Togo	-0.31	219	210	0.23	114	889	46	4.8	28	23	33	83	223	6,611
Uganda	-0.14	198	962	0.16	3,604	116	43	4.7	23	17	29	121	1,211	14,324
Zambia	-0.34	452	830	0.18	3,173	-1,509	44	5.3	23	16	30	58	1,119	7,048
Zimbabwe	-0.27	243	457	0.26	3,205	-2,101	45	4.1	22	14	29	35	1,092	8,145
Overall	-0.18	336	511	0.26	1625	69	45	5.1	25	19	31	92	717	288,614

Note: This table reports the arithmetic mean of the Demographic and Health Survey (DHS) variables excluding 20 km around the capital cities based on own calculations.

Table 2.A.4

Summary Statistics - Afrobarometer - Arithmetic Mean by Country (1/2)

Country Name	Public Goods Index	Wealth Index	Trust into Leaders	Corruption Perception Index	Evaluation of Leadership Performance	Trust Ruling Party	Trust Oppos. Party	Voter Turnout	National vs. Ethnic Identity	Frequency of News Consumption	Advocate Monitoring of Leaders
Benin	.4	.35	1.7	1.5	2.2	1.5	1.2	.85	3.5	2.9	3
Botswana	.61	.38	2	1.1	2.6	1.9	1.1	.75	3.3	3	3
Burkina Faso	.2	.28	2	1.1	2.2	1.7	1.6	.73	3.7	2.8	2.9
Burundi	.2	.16	2.3	.82	2.3	2.1	1.1	.91	4.5	2.6	2.8
Cameroon	.67	.46	1.6	1.6	2.2	1.2	.87	.6	3.7	3.2	2.8
Gabon	.54	.42	1.1	1.9	1.6	.74	.95	.58	4.1	3.4	2.8
Gambia	.36	.43	2.1	1	2.5	1.7	1.4	.67	3.9	3.3	3.3
Ghana	.61	.44	1.6	1.5	2.2	1.5	1.5	.85	3.6	3.4	3.1
Guinea	.27	.28	1.7	1.3	1.9	1.5	1.3	.84	4.4	2.8	2.8
Ivory Coast	.53	.4	1.7	1.2	2.3	1.3	1.2	.62	3.7	3	2.8
Kenya	.41	.3	1.7	1.5	2.1	1.6	1.2	.82	3.8	3.4	2.8
Liberia	.15	.27	1.5	1.7	2.1	1.5	1.2	.83	3.2	2.9	2.8
Malawi	.21	.2	1.6	1.4	2.1	1.4	1.5	.85	3.4	2.4	2.8
Mali	.19	.29	1.7	1.4	2.1	1.6	1.2	.73	3.6	2.9	3
Mozambique	.26	.29	2.1	1.1	2.4	2	1.4	.79	3.7	2.9	2.6
Namibia	.43	.44	2.1	.98	2.4	2	1.2	.8	3.7	3.5	2.6
Niger	.24	.23	2.1	1.1	2.3	2	1.6	.82	4	2.5	2.8
Nigeria	.56	.49	1.1	1.7	2	1.1	.96	.71	3.2	3.4	2.6
Senegal	.52	.49	1.9	1.2	2.2	1.5	1.2	.7	3.7	3.4	2.9
Sierra Leone	.22	.26	1.8	1.4	2.2	1.7	1.5	.83	3.8	2.7	2.8
Swaziland	.47	.4	1.6	1.4	2.4	NA	NA	.69	3.6	3.4	3
Tanzania	.32	.33	2.2	.93	2.2	2	1.5	.82	4.1	3	2.7
Togo	.34	.35	1.5	1.6	2.1	1.2	1.1	.87	3.7	2.9	3
Uganda	.2	.29	2	1.4	2.2	1.9	1.2	.83	3.2	3.1	2.9
Zambia	.33	.33	1.8	1.3	2.2	1.7	1.2	.71	3.5	3	2.8
Zimbabwe	.37	.39	1.8	1.4	2.1	1.7	1	.75	3.8	2.6	2.8
Overall	.37	.34	1.8	1.3	2.2	1.6	1.3	.77	3.7	3	2.8

Note: This table reports the arithmetic mean of the Afrobarometer variables excluding 20 km around the capital cities based on own calculations.

Table 2.A.5

Summary Statistics - Afrobarometer - Arithmetic Mean by Country (2/2)

Country Name	Distance to the Capital	Distance to the Placebo	Distance to the Coast	Population Density	Age of Respondent	Share of Women	Rural Pop. Share	Longitude projected (in km)	Latitude projected (in km)	Number of Observations
Benin	222	241	213	291	39	0.5	0.6	250	914	2,928
Botswana	331	313	708	18	41	0.5	0.78	2,643	-2,487	1,567
Burkina Faso	192	309	739	96	43	0.5	0.85	-215	1,354	3,136
Burundi	81	85	1,020	400	38	0.5	0.9	3,321	-360	2,160
Cameroon	395	394	392	323	39	0.5	0.53	1,305	727	3,182
Gabon	315	338	199	55	35	0.5	0.35	1,286	-84	1,168
Gambia	119	93	51	196	38	0.48	0.63	-1,687	1,480	688
Ghana	234	177	183	432	38	0.5	0.52	-123	775	6,193
Guinea	313	418	226	150	43	0.5	0.75	-1,252	1,132	3,185
Ivory Coast	201	236	169	918	38	0.5	0.49	-584	729	3,527
Kenya	236	548	497	410	41	0.5	0.73	4,051	-51	5,378
Liberia	187	131	105	66	40	0.5	0.73	-1,044	729	2,510
Malawi	177	192	428	251	45	0.5	0.9	3,722	-1,581	4,993
Mali	326	348	813	62	41	0.5	0.88	-686	1,489	2,976
Mozambique	1,137	540	132	108	65	0.52	0.71	3,864	-1,925	6,384
Namibia	474	476	409	35	37	0.5	0.62	1,781	-2,164	2,896
Niger	453	365	889	105	42	0.5	0.88	651	1,545	3,295
Nigeria	379	518	346	1,099	35	0.5	0.57	776	934	6,303
Senegal	219	149	88	248	41	0.5	0.67	-1,697	1,598	2,759
Sierra Leone	166	105	93	171	41	0.51	0.71	-1,308	930	2,910
Swaziland	58	67	123	76	39	0.5	0.84	3,134	-2,940	2,936
Tanzania	586	487	481	171	41	0.5	0.76	3,868	-609	6,481
Togo	239	170	231	151	38	0.5	0.75	113	910	2,595
Uganda	205	258	979	255	41	0.5	0.86	3,583	113	5,408
Zambia	379	332	888	88	40	0.5	0.64	3,107	-1,486	2,890
Zimbabwe	277	275	480	95	40	0.5	0.73	3,171	-2,136	4,904
Overall	316	304	427	263	41	0.5	0.71	1,286	3.7	93,352

Note: This table reports the arithmetic mean of the Afrobarometer variables excluding 20 km around the capital cities based on own calculations.

2.A.5 Sensitivity Analysis

In this section, we examine whether the results from Section 2.2.3 are robust to variations in the precise model specification.

Firstly, we reduce the bandwidth of the segments around the boundary from 25 km to 15 km. Simultaneously, we only exclude pixels within 1.5 km, instead of 3 km, around the boundary. In addition, we run the regressions with a varying degree of up to fourth order polynomials of the running variable. As can be clearly seen in Table 2.A.6, the results are very similar to the previous findings in direction, magnitude and significance.

TABLE 2.A.6
BORDER DISCONTINUITY ESTIMATION - ROBUSTNESS BANDWIDTH

	Dependent variable:									
	Probability Pixel is Lit in 2016 (VIIRS)					Log Light Density in 2016 (VIIRS)				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Log Distance from the Capital City	-0.009 (0.006)	-0.019** (0.008)	-0.031*** (0.010)	-0.033*** (0.011)	-0.035*** (0.012)	-0.051* (0.029)	-0.078** (0.036)	-0.124*** (0.045)	-0.149*** (0.054)	-0.159*** (0.059)
Polynomials for: distance from the border × country × ethnicity (325 groups)										
1nd order	-	x	-	-	-	-	x	-	-	-
2nd order	-	-	x	-	-	-	-	x	-	-
3rd order	-	-	-	x	-	-	-	-	x	-
4th order	-	-	-	-	x	-	-	-	-	x
Geography Cov.	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Country FE	36	36	36	36	36	36	36	36	36	36
Segment FE	608	608	608	608	608	608	608	608	608	608
Observations	106,746	106,746	106,746	106,746	106,746	106,746	106,746	106,746	106,746	106,746
Adjusted R^2	0.171	0.198	0.207	0.212	0.215	0.161	0.187	0.194	0.200	0.203

Note: This table reports robustness tests on our main BDD regression results in Table 2.2 based on Equation 2.4 and feature variations on the boundary thickness, boundary buffer and the number of polynomials of the running variable. In order to avoid capturing the break between the capital city and the hinterlands, we exclude 20 km around each capital city from our sample. To prevent misassignment of detected nightlights between countries due to blooming, we exclude 1.5 km (instead of 3 km) on each side of the national boundary. The 'Geographical Cov.' include: distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km). Boundary segments corresponds to a buffer of 15 km (instead of 25 km) around border pieces of 50 km line length and are entirely nested within a restricted ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). The observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by boundary segment. *p<0.1; **p<0.05; ***p<0.01

We then investigate whether the effects of isolation from the capital city are limited to areas close to the capital city. Therefore, in the next step, we increase the sample restriction from excluding 20 km (default) around the capital city to 50, 75 and 100 km. Once again, the results in Table 2.A.7 reveal that the estimated effects are very stable with respect to this modification. This finding implies that isolation from the capital city is a more general phenomenon that is relevant for wide areas within the country.

Table 2.A.7
Border Discontinuity Estimation: Robustness - Capital City Buffer

	Dependent variable: VIIRS Nightlights 2016					
	Exclude all Pixels with a Distance from the Capital City					
	< 50 km		< 75 km		< 100 km	
	Prob	Log	Prob	Log	Prob	Log
	(1)	(2)	(3)	(4)	(5)	(6)
Log Distance from the Capital City	-0.030**	-0.135**	-0.034**	-0.151**	-0.037**	-0.164**
	(0.012)	(0.060)	(0.014)	(0.067)	(0.016)	(0.077)
	Polynomials for: distance from the border × country × ethnicity					
	305 groups		295 groups		290 groups	
3rd order	x	x	x	x	x	x
Geography Cov.	YES	YES	YES	YES	YES	YES
Country FE	36	36	36	36	36	36
Segment FE	563	563	548	548	531	531
Observations	166,741	166,741	163,014	163,014	158,033	158,033
Adjusted R^2	0.180	0.160	0.179	0.160	0.170	0.152

Note: This table reports robustness tests on our main BDD regression results in Table 2.2 based on Equation 2.4 and includes some variations on the exclusion around the capital city. In order to avoid capturing the break between the capital city and the hinterlands, we exclude 50, 75 and 100 km (instead of 20 km) around each capital city from our sample. To prevent misassignment of detected nightlights between countries due to blooming, we exclude 3 km on each side of the national boundary. The 'Geographical Cov.' include: distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km). Boundary segments corresponds to a buffer of 25 km around border pieces of 50 km line length and are entirely nested within a restricted ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). The observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by boundary segment.
*p<0.1; **p<0.05; ***p<0.01

In the subsequent step, we test whether the estimated effects of isolation from the capital city are driven by individual countries or boundaries. To this end, in Figure 2.A.5, we iteratively exclude both sides of boundary segments adjoining a particular country and compare the estimated coefficient to our unrestricted baseline coefficients (in Table 2.2 columns (3) and (7)). The figures clearly illustrate that the estimated coefficients are highly robust to excluding particular countries/boundaries. Even when excluding large amounts of border segments for large or centrally located countries, all confidence bands overlap clearly with that of the baseline estimate.

Figure 2.A.5

Robustness - Dropping Countries Iteratively

(A) Extensive

(B) Intensive

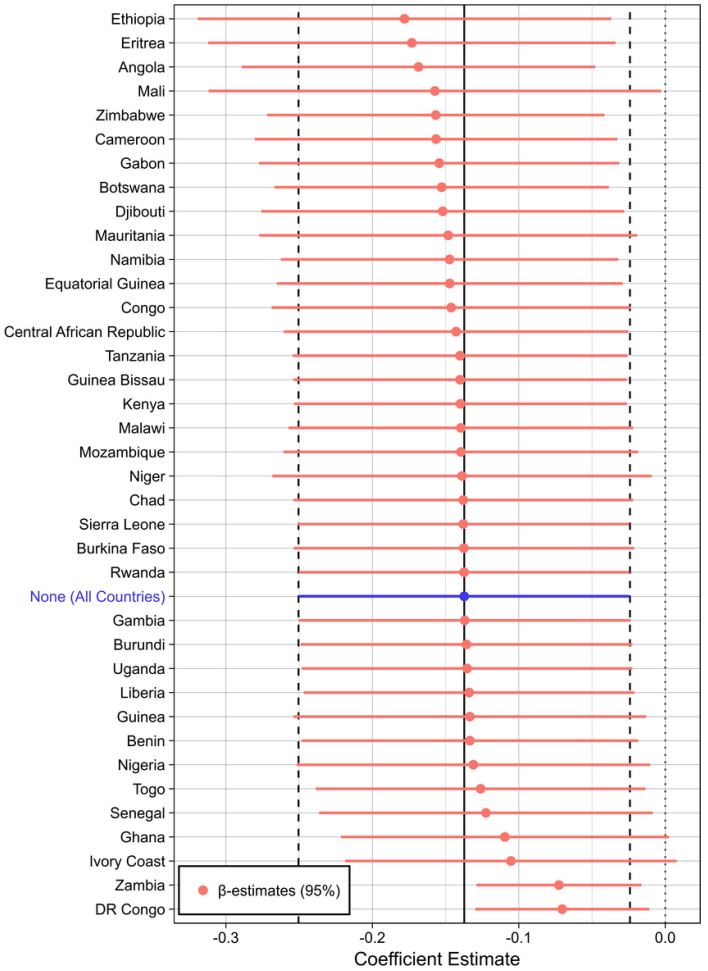

Note: The graphs illustrate the estimated coefficients and 95% confidence bands associated with isolation from the capital city based on Equation 2.4 when iteratively excluding all boundaries of individual countries. Since we are conducting a BDD, whenever we drop all segments of a country, we also drop a significant amount of pixels of adjacent countries. The change in the coefficient can therefore not be solely attributed to one particular country. The vertical solid black line corresponds to the baseline estimates of the full sample and the two vertical dashed black lines demarcate the respective 95% confidence interval. Standard errors are clustered by boundary segment.

While the coefficients prove themselves to be very stable to excluding even large parts of the sample, it might be that the deviations in the estimated coefficients are systematic. Therefore, we test whether there are heterogeneities with respect to the implications of isolation from the capital city depending upon a country's level of decentralization, democratization or overall level of development. For this purpose we categorize all boundaries as either delimiting two relative decentralized[25] (democratic[26], developed[27]) or two relatively non-decentralized (autocratic, underdeveloped) countries or one of each. We then run regressions allowing for heterogenous coefficients for the different boundary pair categories and compare the coefficients using F-tests (see Table 2.A.8).[28]

The results for decentralization in Table 2.A.8 in columns (1)-(4) are ambiguous and we cannot reject the hypothesis that the effects are the same for relatively decentralized and non-decentralized countries. In contrast, columns (5)-(8) reveal that the average effect in democracies is higher than in autocracies where they are consistently very close to zero. The difference is statistically significant for the 20-year average measure of democracy. Regarding GDP per capita, irrespective of whether we split countries at the median or mean, the results consistently indicate that the effect are significantly stronger in less developed countries. We conclude that the implications of isolation from the capital city seem to be more relevant under democratic, as opposed to autocratic, institutional framework conditions and more relevant in relatively underdeveloped countries. Yet, these results should be interpreted with caution as the patterns might be confounded by other characteristics these groups of countries have in common.

Next, we study whether the effects of isolation from the capital city had been relevant at some point in the past and have persisted until today or, alternatively, whether the effects are still ongoing and relevant today. For this exercise, we exclude all pixels that have already been lit in 1992[29] from the sample and thereby focus on the pattern of newly emerging nightlights (see Table 2.A.9). While the coefficients are, as expected, slightly

[25] We define countries as decentralized if they exceed the median decentralization value (based on a decentralization index by Thomas Bijl for J. Vernon Henderson (LSE processed)).

[26] Democratic countries are defined as having a Polity2 index score of greater than zero (Marshall et al., 2017).

[27] Since there is no obvious cutoff for distinguishing between relatively developed and underdeveloped countries, we categorize them as such depending on whether they exceed the sample median or mean GDP level based on estimates by the World Bank (2016).

[28] As policy indicators usually have a lot of fluctuations, we conduct the analysis based on the most recent snapshot of the indicators as well as a 20-year average.

[29] We use the earliest available DMSP-OLS nightlight grid from 1992 as VIIRS nightlights only date back to 2012.

smaller in magnitude as in Table 2.2, they are still strong and highly significant which underlines the present-day importance of the effects.

Table 2.A.8
Border Discontinuity Estimation - Heterogeneity

	Dependent variable: Nightlight Density in 2016 (Prob/Log)											
	Decentralization				Democracy (Polity2)				GDP per capita (World Bank 2016)			
	Mean 1990-2010		2010		Mean 1996-2016		2016		Split at median		Split at mean	
	Prob	Log	Prob	Log	Prob	Log	Prob	Log	Prob	Log	Prob	Log
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
CAP × Non-Decentr.	-0.045*	-0.229*	-0.041	-0.217	-	-	-	-	-	-	-	-
	(0.024)	(0.125)	(0.028)	(0.146)								
CAP × Decentralized	-0.031	-0.121*	-0.055***	-0.204***	-	-	-	-	-	-	-	-
	(0.021)	(0.070)	(0.017)	(0.060)								
CAP × Mixed	-0.026**	-0.107**	-0.025***	-0.101***	-	-	-	-	-	-	-	-
	(0.010)	(0.047)	(0.008)	(0.032)								
CAP × Autocracy	-	-	-	-	-0.001	0.007	-0.017	-0.041	-	-	-	-
					(0.010)	(0.042)	(0.013)	(0.049)				
CAP × Democracy	-	-	-	-	-0.064***	-0.292**	-0.027*	-0.117**	-	-	-	-
					(0.023)	(0.114)	(0.014)	(0.051)				
CAP × Mixed	-	-	-	-	-0.019	-0.089*	-0.041*	-0.206*	-	-	-	-
					(0.013)	(0.053)	(0.023)	(0.117)				
CAP × Underdeveloped	-	-	-	-	-	-	-	-	-0.033***	-0.140***	-0.041***	-0.177**
									0.011)	(0.049)	0.015	0.072
CAP × Developed	-	-	-	-	-	-	-	-	0.012	0.013	0.020	0.046
									(0.021)	(0.083)	0.022	0.096
CAP × Mixed	-	-	-	-	-	-	-	-	-0.049***	-0.210**	-0.027*	-0.114**
									(0.017)	(0.083)	0.014	0.057
Coefficient Tests	Non-Decentralized = Decentralized				Autocracy = Democracy				High GDP = Low GDP			
F-Statistic	0.23	0.68	0.21	0.01	6.66**	6.15**	0.27	1.18	4.67**	3.66**	5.62**	3.71*
	Polynomials for: distance from the border × country × ethnicity (302 groups columns (1)-(4) and 305 groups columns (5)-(12))											
3rd order	x	x	x	x	x	x	x	x	x	x	x	x
Geography Cov.	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Country FE	35	35	35	35	36	36	36	36	36	36	36	36
Segment FE	563	563	563	563	569	569	569	569	569	569	569	569
Observations	166,833	166,833	166,833	166,833	168,620	168,620	168,620	168,620	168,620	168,620	168,620	168,620
Adjusted R^2	0.182	0.162	0.182	0.162	0.182	0.161	0.182	0.161	0.182	0.161	0.182	0.161

Note: This table reports the heterogeneity test corresponding to our main BDD regression results based on Equation 2.4. For this purpose we categorize all boundaries as either delimiting two relatively decentralized, democratic or developed countries and interact the respective dummies with log distance from the capital city. In columns (1)-(4), we group boundaries based on the mean over 20 years and 2010 level of the decentralization index, in columns (5)-(8) based on the 20-year mean and 2016 level of the Polity2 score and in columns (9)-(12) based on the median and the mean of the 2016 GDP per capita value. We report the F-tests on the equality of the respective coefficients. In order to avoid capturing the break between the capital city and the hinterlands, we exclude 20 km around each capital city from our sample. To prevent misassignment of detected nightlights between countries due to blooming, we exclude 3 km on each side of the border. The 'Geographical Cov.' include: distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km). Boundary segments corresponds to a buffer of 25 km around border pieces of 50 km line length and are entirely nested within a restricted ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). The observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by boundary segment.
*p<0.1; **p<0.05; ***p<0.01

Further, we want to examine whether a potential shift in the role of the ethnicity within the country distorts our findings. One factor that we have not accounted for in our identification strategy is that a partitioned ethnic homeland might represent a minority group in one but a majority group in another country. This circumstance makes it necessary to

include a measure of ethnic representation within the country into our model. In columns (3) and (4) in Table 2.A.9, we include the share of pixels of an ethnic homeland in the overall number of pixels within a country. The results indicate that ethnic representation does not confound the effects of isolation from the capital city.

TABLE 2.A.9
BORDER DISCONTINUITY ESTIMATION - ROBUSTNESS

	Dependent variable: Nightlight Density in 2016 (Prob/Log/PC)						
	Exclude Pixels Lit in 1992		Ethnic Minority vs. Majority		Population Density		Log Lights Per Capita
	Prob	Log	Prob	Log	Prob	Log	Log pc
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log Distance from the Capital City	-0.022*** (0.008)	-0.085*** (0.031)	-0.032*** (0.012)	-0.142** (0.058)	-0.026** (0.011)	-0.107** (0.049)	-0.263** (0.115)
	Polynomials for: distance from the border × country × ethnicity (305 groups)						
3rd order	x	x	x	x	x	x	x
Population Density	No	No	No	No	YES	YES	No
Ethnicity Share	No	No	YES	YES	No	No	No
Geography Controls	YES	YES	YES	YES	YES	YES	YES
Country FE	36	36	36	36	36	36	36
Segment FE	569	569	569	569	569	569	569
Observations	167,841	167,841	168,620	168,620	168,620	168,620	168,620
Adjusted R^2	0.165	0.141	0.182	0.161	0.214	0.251	0.179

Note: This table reports robustness tests for our main BDD regression results in Table 2.2 and extends the model in Equation 2.4. In order to avoid capturing the break between the capital city and the hinterlands, we exclude 20 km around each capital city from our sample. To prevent misassignment of detected nightlights between countries due to blooming, we exclude 3 km on each side of the national boundary. The 'Geographical Cov.' include: distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km). Boundary segments corresponds to a buffer of 25 km around border pieces of 50 km line length and are entirely nested within a restricted ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). In columns (1)-(2), we exclude pixels that were already lit in 1992. In columns (3)-(4), we account for the share of the ethnicity with respect to the total area of the home country. In columns (5)-(6), we include population density as a covariate and in column (7), we use the log of lights per capita as our dependent variable. The observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by boundary segment.
*p<0.1; **p<0.05; ***p<0.01

Since increased economic performance generally boosts population agglomeration by attracting migration inflows as well as potentially increasing net fertility, we would expect that the gains in economic performance reinforce population agglomeration. Moreover, the two variables are likely in a mutually reinforcing relationship due to the reverse positive impact of agglomeration on economic performance (Marshall, 1920; Ahlfeldt et al., 2015). Moreover, population density maps do not exhibit a very high accuracy on the very fine pixel level and are themselves usually based on indicators of economic activity such as

schools, hospitals, roads or even nightlights. For this reason population agglomeration is endogenous to our outcome variable from a theoretical and technical point of view and there is no straightforward way to disentangle the two. In line with this hypothesis, we show that when we include population density as a control variable, the absolute magnitude of the estimated impact of distance from the capital city on economic performance decreases (see columns (5) and (6) in Table 2.A.9). While, we have to be cautious about interpreting this estimate due to endogeneity concerns, the fact that the coefficient remains significant and is not significantly different from the estimate in Table 2.2 suggests that the implications of isolation from the capital city on economic development go beyond population agglomeration. This is view is also supported by the fact that distance from the capital city also significantly reduces nightlight intensity per capita (column (7)).

Due to the design of our BDD model, distance controls are generally balanced by construction. Yet, for certain variables, it might matter whether the access point lies within the country or beyond the other side of the border. We therefore show that our results hold when we control for distance to the closest river, waterbody (lake, lagoon and reservoir) and mine within the country (see Table 2.A.10). We also report estimates when including distance to the closest city of a certain size and distance to regional capital city as controls because these variables might also be more relevant when within one's own border. Table 2.A.11 reveals that in these models the magnitude of the estimates associated with distance from the capital city are slightly smaller and insignificant for some of the models using the intensive margin of nightlights. Yet, due to the fact that population agglomeration is endogenous (see discussion in Section 2.A.5), these estimates should be interpreted with caution.

Table 2.A.10

Boundary Discontinuity Estimation: Additional Geographic Covariates

	Dependent variable:							
	Probability Pixel is Lit in 2016 (VIIRS)				Log Light Density in 2016 (VIIRS)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Distance from the Capital City	-0.031**	-0.026**	-0.031**	-0.027**	-0.137**	-0.122**	-0.132**	-0.118*
	(0.012)	(0.012)	(0.014)	(0.013)	(0.058)	(0.059)	(0.064)	(0.062)
Log Distance to the closest River	-0.001**	-	-	-	-0.003	-	-	-
	(0.000)				(0.002)			
Log Distance to the Closest Waterbody	-	-0.016***	-	-	-	-0.073***	-	-
		(0.006)				(0.028)		
Log Distance to the Closest Mine	-	-	-0.037***	-	-	-	-0.190***	-
			(0.011)				(0.063)	
Log Distance Closest Mine (opened > 1960)	-	-	-	-0.033**	-	-	-	-0.185**
				(0.014)				(0.078)
	Polynomials for: distance from the border × country × ethnicity							
3rd order	x	x	x	x	x	x	x	x
Geography Cov.	YES	YES	YES	YES	YES	YES	YES	YES
Country FE	36	35	29	28	36	35	29	28
Segment FE	569	539	475	457	569	539	475	457
Observations	168,620	161,545	142,231	138,252	168,620	161,545	142,231	138,252
Adjusted R^2	0.182	0.186	0.194	0.196	0.161	0.164	0.173	0.175

Note: This table reports regressions including additional geographic distance controls and corresponds to our baseline BDD estimates in columns (3) and (7) in Table 2.2 based on Equation 2.4. Waterbodies include lakes, lagoons and reservoirs. The 'Geographical Cov.' include: distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km). Columns (2) and (6) additionally include a dummy indicating whether the centroid of a pixel is inside a waterbody polygons. The observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by boundary segment.

*p<0.1; **p<0.05; ***p<0.01

Table 2.A.11
Boundary Discontinuity Estimation: City Distance Controls

	Dependent variable:							
	Probability Pixel is Lit in 2016 (VIIRS)				Log Light Density in 2016 (VIIRS)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Distance from the Capital City	-0.022* (0.013)	-0.026* (0.013)	-0.037** (0.015)	-0.018* (0.011)	-0.092 (0.065)	-0.103 (0.065)	-0.157** (0.074)	-0.063 (0.050)
Log Dist City $\geq$ 25	-0.076*** (0.009)	-	-	-	-0.387*** (0.054)	-	-	-
Log Dist City $\geq$ 50	-	-0.076*** (0.012)	-	-	-	-0.414*** (0.074)	-	-
Log Dist City $\geq$ 100	-	-	-0.081*** (0.015)	-	-	-	-0.473*** (0.099)	-
Log Dist Regional Capital	-	-	-	-0.053*** (0.011)	-	-	-	-0.291*** (0.068)
	Polynomials for: distance from the border × country × ethnicity							
3rd order	x	x	x	x	x	x	x	x
Geography Cov.	YES	YES	YES	YES	YES	YES	YES	YES
Country FE	36	34	29	36	36	34	29	36
Segment FE	569	556	425	560	569	556	425	560
Observations	168,620	164,616	125,570	165,004	168,620	164,616	125,570	165,004
Adjusted R^2	0.204	0.197	0.204	0.193	0.200	0.192	0.200	0.180

Note: This table reports regressions including endogenous city distance controls and corresponds to our baseline BDD estimates in columns (3) and (7) in Table 2.2 based on Equation 2.4. The city distance controls include the log of the distance to the closest city with at least 25,000 (50,000 and 100,000) inhabitants and the log of the distance to the corresponding first level administrative capital city. In each column, we exclude pixels where the closest respective city is the capital city itself. The 'Geographical Cov.' include: distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km). The observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by boundary segment. *p<0.1; **p<0.05; ***p<0.01

Finally, we cross-validate our results using the DHS wealth index as an alternative data source for measuring the spatial distribution of economic performance. For this purpose, we combine the latest available round for each country in our dataset. While the DHS comprises a geocoded wealth index for the vast majority of our sample countries, eight countries are not covered (see Figure 2.A.8a). Since the granularity of the data structure differs from nightlight grids, we have to implement some simplifications compared to the estimation framework underlying Table 2.2. Firstly, due to the lower sample size, we restrict the ethnic homelands by a negative buffer 5 km rather than 15 km. Secondly, the geographical accuracy of the DHS is lower than remote sensing sources. DHS households are clustered with an average of 26 households sharing the same coordinate pair leaving us with a minimum of one and a median of seven clusters per segment side. Additionally, the assigned geolocation is randomized by up to 10 km. As a consequence, there is too little variation and too much noise in the data to model the running variable disaggregated for each ethnicity in each country separately which is why we just focus on accounting for the general trend. Thirdly, as the observational unit are households, we include household level control variables: age of head of household, age of head of household squared and number of de jure household members. Lastly, since there are some concerns about the comparability of the DHS wealth index with regard to urban vs. rural households, we include a urban/rural dummy variable in our model. Yet, whether a location is urban or rural is endogenous as it is an outcome of economic development. As a consequence, the magnitude of our estimated coefficients in Table 2.A.12 constitute lower bounds estimates.[30]

The results in Table 2.A.12 confirm our finding about the adverse economic consequences of remoteness from the capital city. A one percent increase in distance from the capital city results in a drop of household wealth by around 0.12 standard deviations. This corresponds to a drop of 3.5 percentiles of the national wealth distribution. The respective balancing tests in columns (1)-(6) in Table 2.A.16 underline valid inference. The respective placebo tests in columns (7)-(8) in Table 2.A.16 once again validate that the effects are specific to isolation from the capital city and do not hold for other cities.

[30] When excluding the urban/rural dummy, the estimated coefficients corresponding to columns (2) and (4) in Table 2.A.12 are -0.183* (0.095) and -0.051* (0.029).

Table 2.A.12
DHS Wealth Index

	Dependent variable:			
	DHS Wealth Index (z-score)		DHS Wealth Rank (percentile)	
	(1)	(2)	(3)	(4)
Log Distance from the Capital City	-0.121***	-0.121**	-0.035*	-0.033*
	(0.045)	(0.052)	(0.018)	(0.019)
	Polynomials for: distance from the border			
3rd order	-	x	-	x
Household Cov.	YES	YES	YES	YES
Geography Cov.	YES	YES	YES	YES
Country FE	28	28	28	28
Segment FE	107	107	107	107
Observations	24,582	24,582	24,582	24,582
Adjusted R^2	0.426	0.381	0.453	0.455

Note: This table reports the BDD regressions results on household wealth using the DHS sample. The z-score of the wealth index constitutes the dependent variable in columns (1)-(2). The percentile rank within the country constitutes the dependent variable in columns (3)-(4). The 'Household Cov.' include: age of household head, age of houshold head squared, number of household members. The 'Geographical Cov.' include: distance from the coast (in km), mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), annual precipitation (in mm), longitude and latitude (projected in km) and whether the household is in an urban or rural setting. Boundary segments corresponds to a buffer of 25 km around border pieces of 50 km line length and are entirely nested within a restricted ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). The observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by DHS cluster and boundary segment. *p<0.1; **p<0.05; ***p<0.01

2.A.6 Travel Time Estimates

As an alternative to using the Euclidean distance to measure isolation from the capital city, in Table 2.A.13, we report estimates based on travel time and road distance based on OSRM which is a routing engine for OpenStreetMaps (`http://project-osrm.org`) instead. However, as discussed above, these estimates are likely to be endogenous and overestimate the true effect because underdeveloped areas are less likely to have equal access to good infrastructure. 'Travel Time' captures the estimated travel time (in minutes) from the pixel to the capital city. Yet, if a pixel is not adjacent to a road, the algorithm first searches for the nearest road to use as a starting point. Therefore, the variable 'Log adjusted Travel time' accounts for this additional distance by adding the time needed to get from the pixel to the closest road (assuming that there is a direct way and a reduced average speed of 10 km/h). Lastly, 'Road Distance' is computed as the sum of the Euclidean distance from the pixel to the nearest road (if applicable) and the distance from there to the capital city via roads (in km).[31]

The results in Table 2.A.13 demonstrate that our findings are similar irrespective of whether we use the Euclidean distance or other travel time estimates as our measure of isolation. Further, as expected, the estimated coefficients between isolation and economic performance are larger than in Table 2.2. This is most likely a result of the reverse causality bias that inflates the magnitude of the estimated effects.

[31] Note that Malabo, the capital city of Equatorial Guinea, is located on an island and therefore not in reach from the mainland via roads. Therefore, in this analysis, we replace Malabo with Bata, the main coastal city of Equatorial Guinea that serves Malabo via ferry or plane.

Table 2.A.13
Boundary Discontinuity Estimation - Travel Time

	Dependent variable:					
	Probability Pixel is Lit in 2016 (VIIRS)			Log Light Density in 2016 (VIIRS)		
	(1)	(2)	(3)	(4)	(5)	(6)
Log Travel Time	-0.051***	-	-	-0.209***	-	-
	(0.009)			(0.043)		
Log adjusted Travel Time	-	-0.073***	-	-	-0.291***	-
		(0.011)			(0.050)	
Log Road Distance	-	-	-0.039***	-	-	-0.165***
			(0.009)			(0.041)
	Polynomials: distance from the border × country × ethnicity (305 groups)					
3rd order	x	x	x	x	x	x
Geography Cov.	NO	YES	YES	YES	NO	YES
Country FE	36	36	36	36	36	36
Segment FE	569	569	569	569	569	569
Observations	168,620	168,620	168,620	168,620	168,620	168,620
Adjusted R^2	0.183	0.185	0.182	0.163	0.165	0.162

Note: This table reports estimation results using travel time and road distance instead of the Euclidean distance as explanatory variables and corresponds to our main BDD regressions in Table 2.2 based on Equation 2.4. In order to avoid capturing the break between the capital city and the hinterlands, we exclude 20 km around each capital city from our sample. To prevent misassignment of detected nightlights between countries due to blooming, we exclude 3 km on each side of the boundary. The 'Geographical Cov.' include: distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km). Boundary segments corresponds to a buffer of 25 km around border pieces of 50 km line length and are entirely nested within a restricted ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). The observations are weighted such that each segment side has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by boundary segment.

*p<0.1; **p<0.05; ***p<0.01

Figure 2.A.6
Euclidean Distance vs. Travel Time
(a) Travel Time

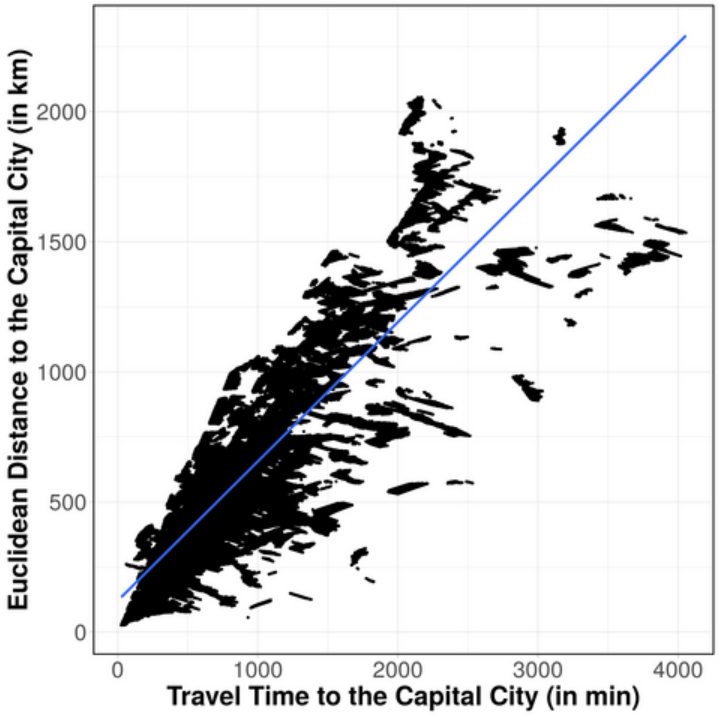

(b) Road Distance

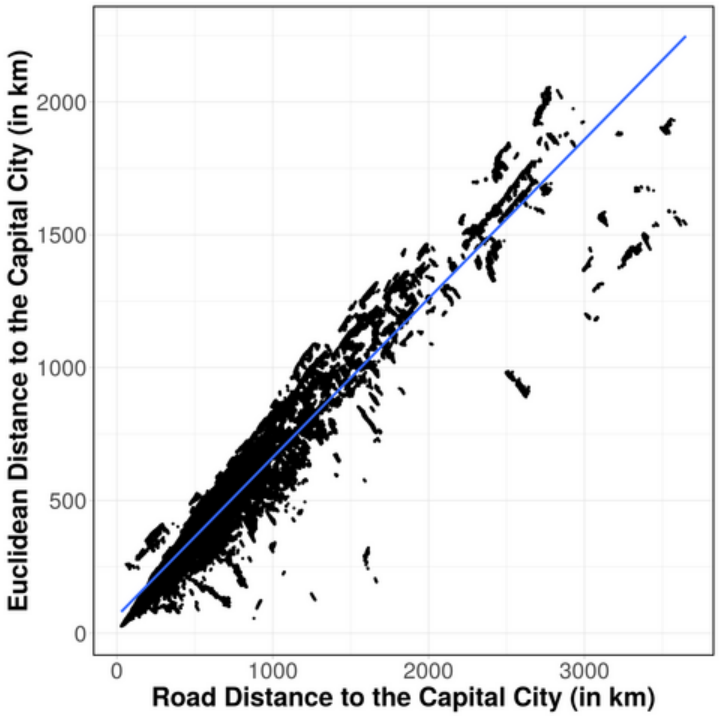

Note: The graphs illustrate the relationship between the Euclidean distance from the capital city and travel time or road distance as alternative distance measures. The blue lines represent the respective linear regression lines.

2.A.7 Maps

FIGURE 2.A.7
VIIRS NIGHTLIGHTS SAMPLE MAP

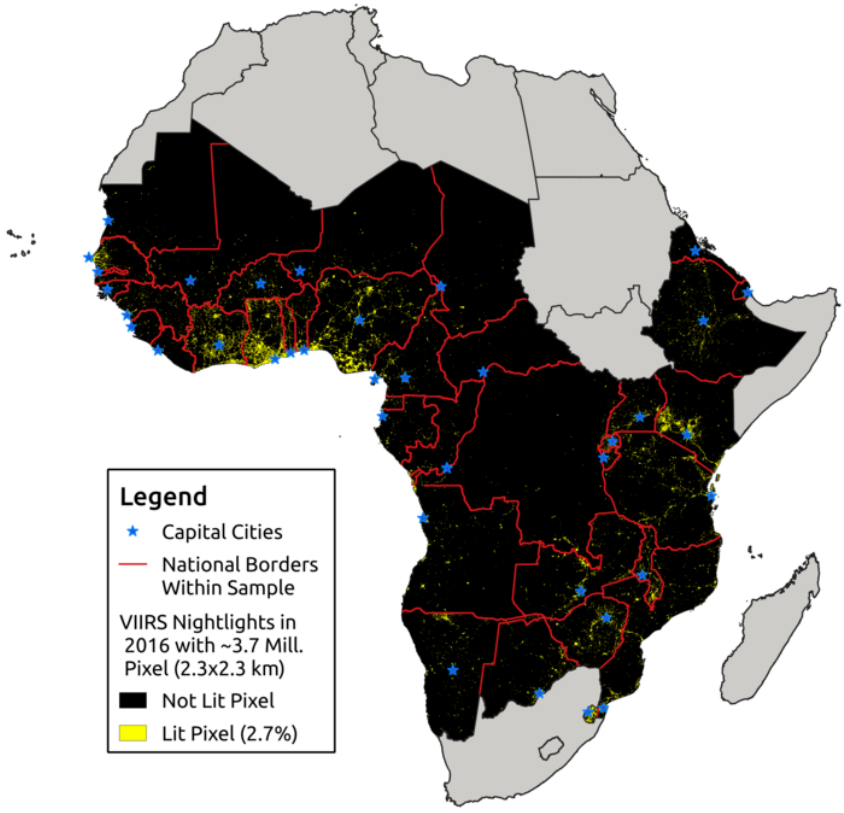

FIGURE 2.A.8
MAPPING SURVEY RESPONDENTS

(A) DHS SAMPLE

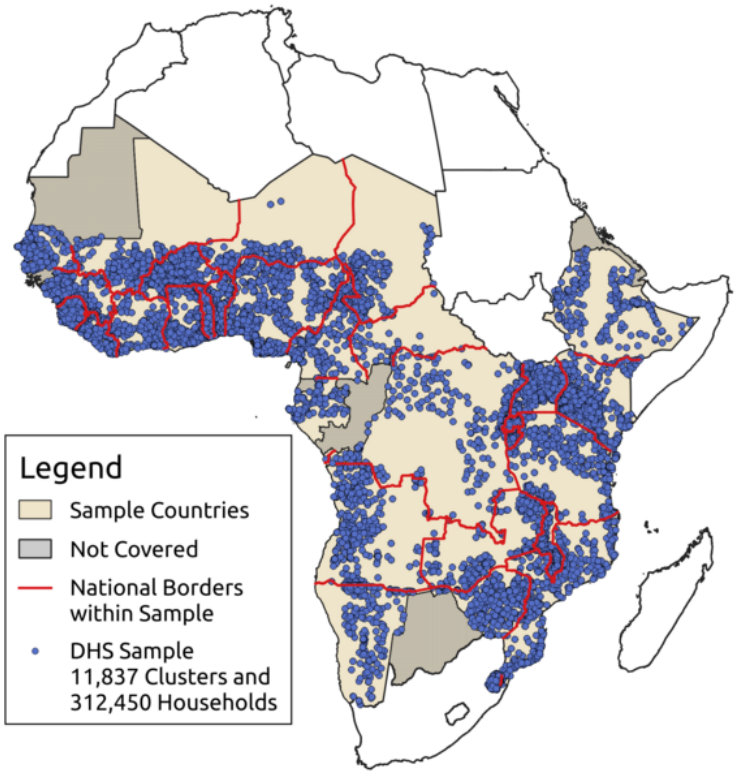

(B) AFROBAROMETER ROUND 5, 6 + 7 SAMPLE

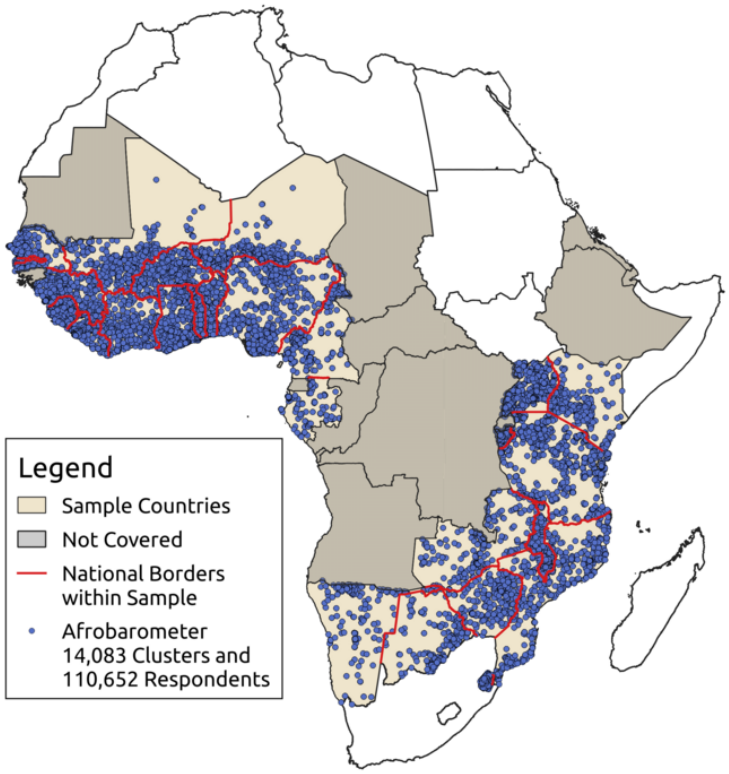

Figure 2.A.9
Unrestricted Segments

(a) Simple Border Segments

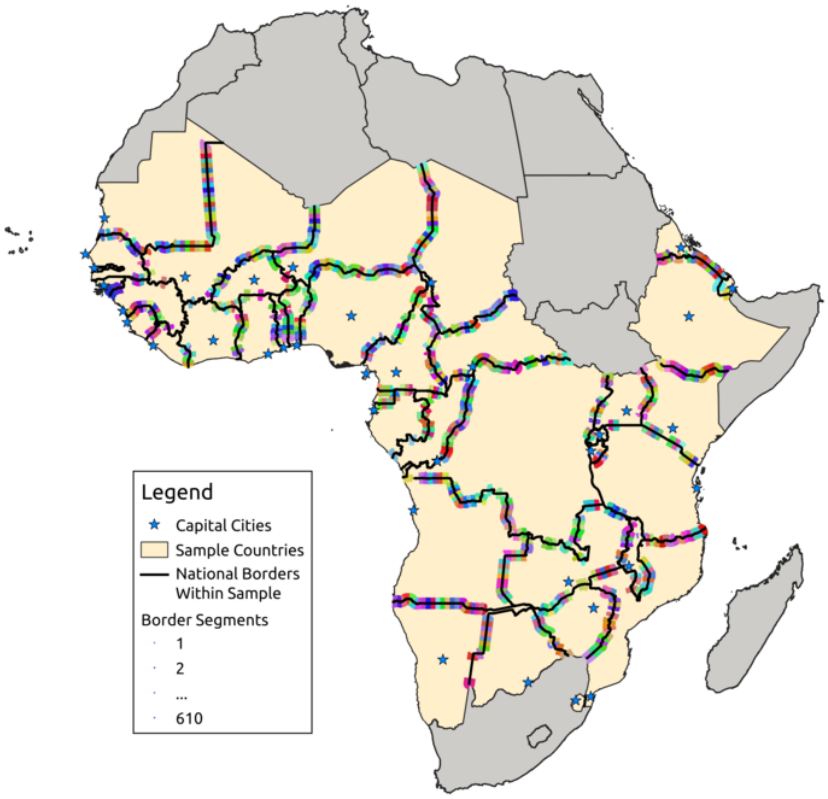

(b) Simple Border Segments - Magnified

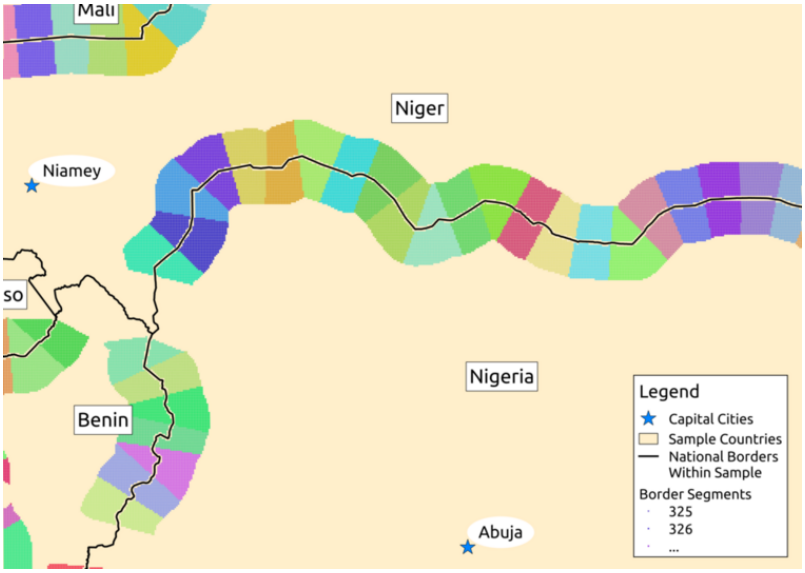

Figure 2.A.10
Restricted Segments

(a) Border Segments nested within Restricted Partitioned Ethnic Homelands

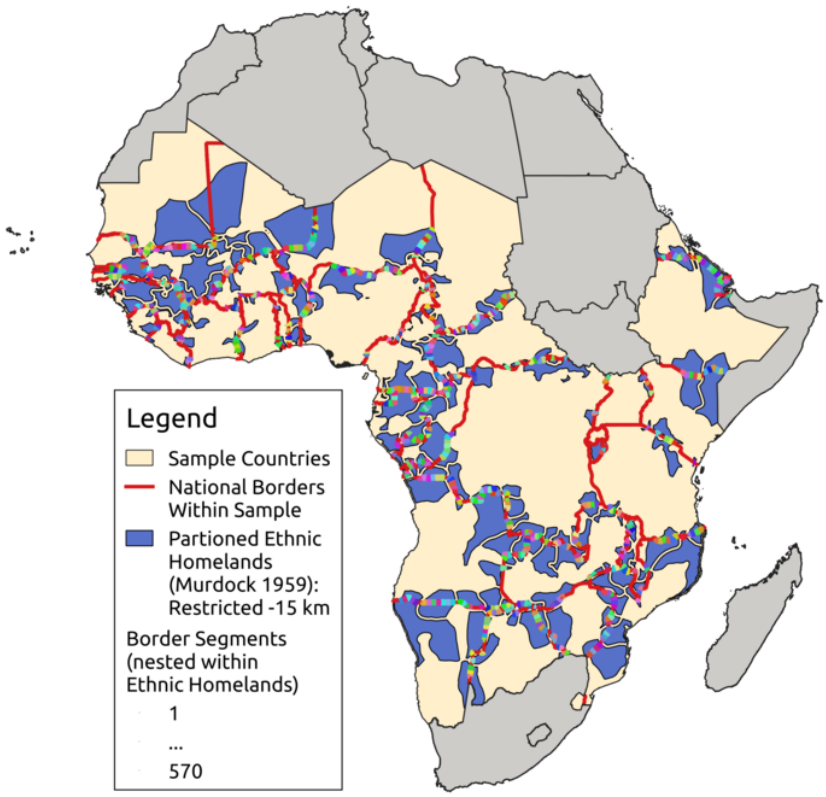

(b) Nested Border Segments - Magnified

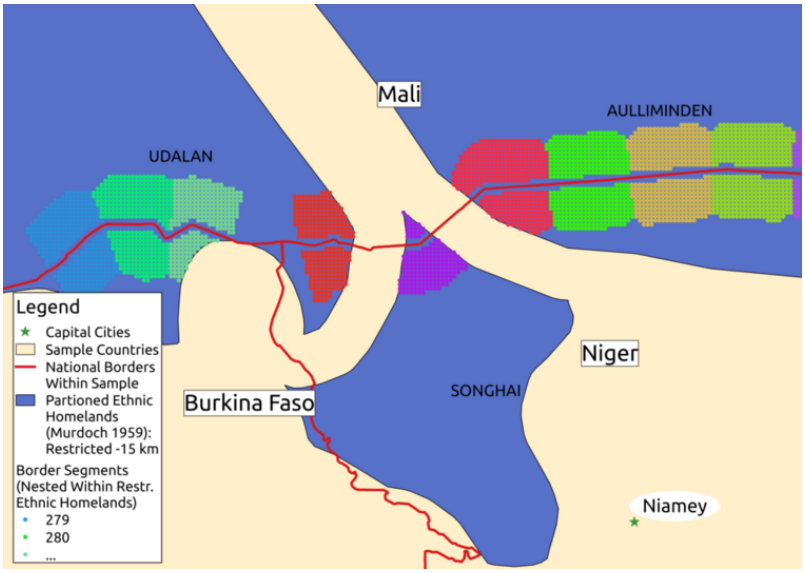

2.A.8 Balancing Tests

TABLE 2.A.14
OLS - BALANCING TESTS

	Dependent variable: Normalization of							
	Distance from the Coast				Crop Caloric Index			
	All Pixels		Border Area		All Pixels		Border Area	
	OLS	Ethnicity FE	OLS	Ethnicity FE	OLS	Ethnicity FE	OLS	Ethnicity FE
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Distance from	0.178***	0.077***	0.152***	0.044***	-0.137***	-0.042*	-0.046*	0.008
the Capital City	(0.022)	(0.016)	(0.029)	(0.013)	(0.022)	(0.022)	(0.024)	(0.014)
Geography Cov.	YES	YES	YES	YES	YES	YES	YES	YES
Country FE	37	37	37	37	37	37	37	37
Ethnicity FE	-	706	-	351	-	706	-	351
Observations	3,518,146	3,518,146	416,667	416,664	3,518,146	3,518,146	416,667	416,664
Adjusted R^2	0.838	0.974	0.792	0.980	0.785	0.935	0.670	0.930

Note: This table reports the balancedness tests for the examples of distance from the coast (in km) and crop caloric index corresponding to the OLS and boundary area regressions in Table 2.1 based on Equation 2.3. We normalize the dependent variables by dividing them by their sample mean. In order to avoid capturing the break between the capital city and the hinterlands, we exclude 20 km around each capital city from our sample. To prevent misassignment of detected nightlights between countries due to blooming, we exclude 3 km on each side of the border. The boundary area regressions ('Border') are restricted to all pixels with centroids within the range of 25 km from shared national borders. The 'Geographical Cov.' include: distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km) - except for when the respective variable is the dependent variable itself. The 'Ethnicity FE' are based on the ethnic homelands in the 'Tribal Map of Africa' (Murdock, 1959). Standard errors in parenthesis are clustered by ethnic homeland. *p<0.1; **p<0.05; ***p<0.01

Table 2.A.15
Border Discontinuity Estimation - Balancing Tests

	Dependent variable: Normalization of								
	Dist. Coast	Elevation	Water	Rugged.	∅ Temp.	Crop	Precip.	Lon.	Lat.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Log Distance from the Capital City	0.008***	0.006	-0.002	-0.367**	0.000	-0.005	0.010	0.006***	0.018**
	(0.003)	(0.004)	(0.004)	(0.156)	(0.000)	(0.006)	(0.017)	(0.001)	(0.009)
	Polynomials for: distance from the border × country × ethnicity (305 groups)								
3rd order	x	x	x	x	x	x	x	x	x
Geography Cov.	YES	YES	YES	YES	YES	YES	YES	YES	YES
Country FE	36	36	36	36	36	36	36	36	36
Segment FE	569	569	569	569	569	569	569	569	569
Observations	168,620	168,620	168,620	168,620	168,620	168,620	168,620	168,620	168,620
Adjusted R^2	1.000	0.999	0.311	0.680	1.000	0.990	0.971	1.000	1.000

Note: This table reports the balancedness tests corresponding to our main BDD regressions in Table 2.2 based on Equation 2.4. We normalize the dependent variables by dividing them by their sample mean. In order to avoid capturing the break between the capital city and the hinterlands, we exclude 20 km around each capital city from our sample. To prevent misassignment of detected nightlights between countries due to blooming, we exclude 3 km on each side of the national boundary. The 'Geographical Cov.' include: distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km) - except for when the respective variable is the dependent variable itself. Boundary segments corresponds to a buffer of 25 km around border pieces of 50 km line length and entirely nested within a restricted ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). The observations are weighted such that each segment side has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by boundary segment. *p<0.1; **p<0.05; ***p<0.01

Table 2.A.16
DHS - Balancing and Placebo Tests

	Dependent variable:							
	Balancing Tests: Normalization of						Placebo Tests	
	Distance to Coast	Longitude	Latitude	Elevation	Mean Temperature	Precip-itation	DHS wealth index (z-score)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Distance from the Capital City	0.000	0.000	0.029	-0.035	-0.000**	-0.008	-0.094*	-0.086
	(0.004)	(0.001)	(0.092)	(0.029)	(0.000)	(0.010)	(0.054)	(0.061)
Log Distance from the Placebo City	-	-	-	-	-	-	0.071	0.082
							(0.064)	(0.065)
Placebo Tests							$Coef_{CAP}$=$Coef_{PLC}$	
F-Statistic	-	-	-	-	-	-	8.76***	8.40***
	Polynomials for: distance from the border							
3rd order	x	x	x	x	x	x	-	x
Household Cov.	YES	YES	YES	YES	YES	YES	YES	YES
Geography Cov.	YES	YES	YES	YES	YES	YES	YES	YES
Country FE	36	36	36	36	36	36	36	36
Segment FE	569	569	569	569	569	569	569	569
Observations	24,582	24,582	24,582	24,582	24,582	24,582	23,671	23,671
Adjusted R^2	1.000	1.000	1.000	0.974	1.000	0.985	0.464	0.465

Note: This table reports the DHS balancing (columns (1)-(6)) and placebo tests (columns (7)-(8)) corresponding to Table 2.A.12. In columns (1)-(6), we normalize the dependent variables by dividing them by their sample mean. Regarding the placebo tests, we report the respective F-tests on the equality of the coefficients. The 'Geographical Cov.' include: age of household head, age of houshold head squared, number of household members. The 'Household Cov.' include: mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), annual precipitation (in mm), longitude and latitude (projected in km) and whether the household is in an urban or rural setting - except for when the respective variable is the dependent variable itself. Boundary segments corresponds to a buffer of 25 km around border pieces of 50 km line length and are entirely nested within a restricted ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). The observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by DHS cluster and boundary segment.

*p<0.1; **p<0.05; ***p<0.01

Table 2.A.17
Channel Analysis: Afrobarometer Balancing Tests

	Dependent variable: Normalization of								
	Distance from the Coast			Longitude			Latitude		
	OLS	Boundary	BDD	OLS	Boundary	BDD	OLS	Boundary	BDD
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Log Distance from the Capital City	0.154***	0.204***	0.015***	0.036***	0.061***	0.001	-0.185	-0.228	0.057***
	(0.023)	(0.050)	(0.005)	(0.010)	(0.013)	(0.001)	(0.128)	(0.229)	(0.020)
	Polynomials for: distance from the border								
3rd order	-	-	x	-	-	x	-	-	x
Household Cov.	YES	YES	YES	YES	YES	YES	YES	YES	YES
Geography Cov.	YES	YES	YES	YES	YES	YES	YES	YES	YES
Country × Round FE	73	72	71	73	72	71	73	72	71
Segment FE	-	-	140	-	-	140	-	-	140
Observations	93,242	15,419	8,347	93,242	15,419	8,347	93,242	15,419	8,347
Adjusted R^2	0.856	0.891	0.999	0.991	0.994	1.000	0.988	0.989	1.000

Note: This table reports the Afrobarometer balancing tests for distance from the coast (in km), latitude and longitude (in km) corresponding to Table 2.1, 2.3 and 2.A.21. We normalize the dependent variables by dividing them by their sample mean. The 'Household Cov.' include: age, age squared and sex of respondent. The 'Geographical Cov.' include: distance from the coast (in km), longitude and latitude (projected in km) and whether the household is in an urban or rural setting - except for when the respective variable is the dependent variable itself. The segments are entirely nested within an ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). In the full sample OLS regressions, the observations are weighted according to the Afrobarometer survey weights. In the BDD regressions, the observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by Afrobarometer cluster and ethnic homeland in the OLS and boundary area regressions and Afrobarometer cluster and boundary segment in the BDD regressions.

*p<0.1; **p<0.05; ***p<0.01

Figure 2.A.11

Discontinuity Balancing Graphs: Capital City

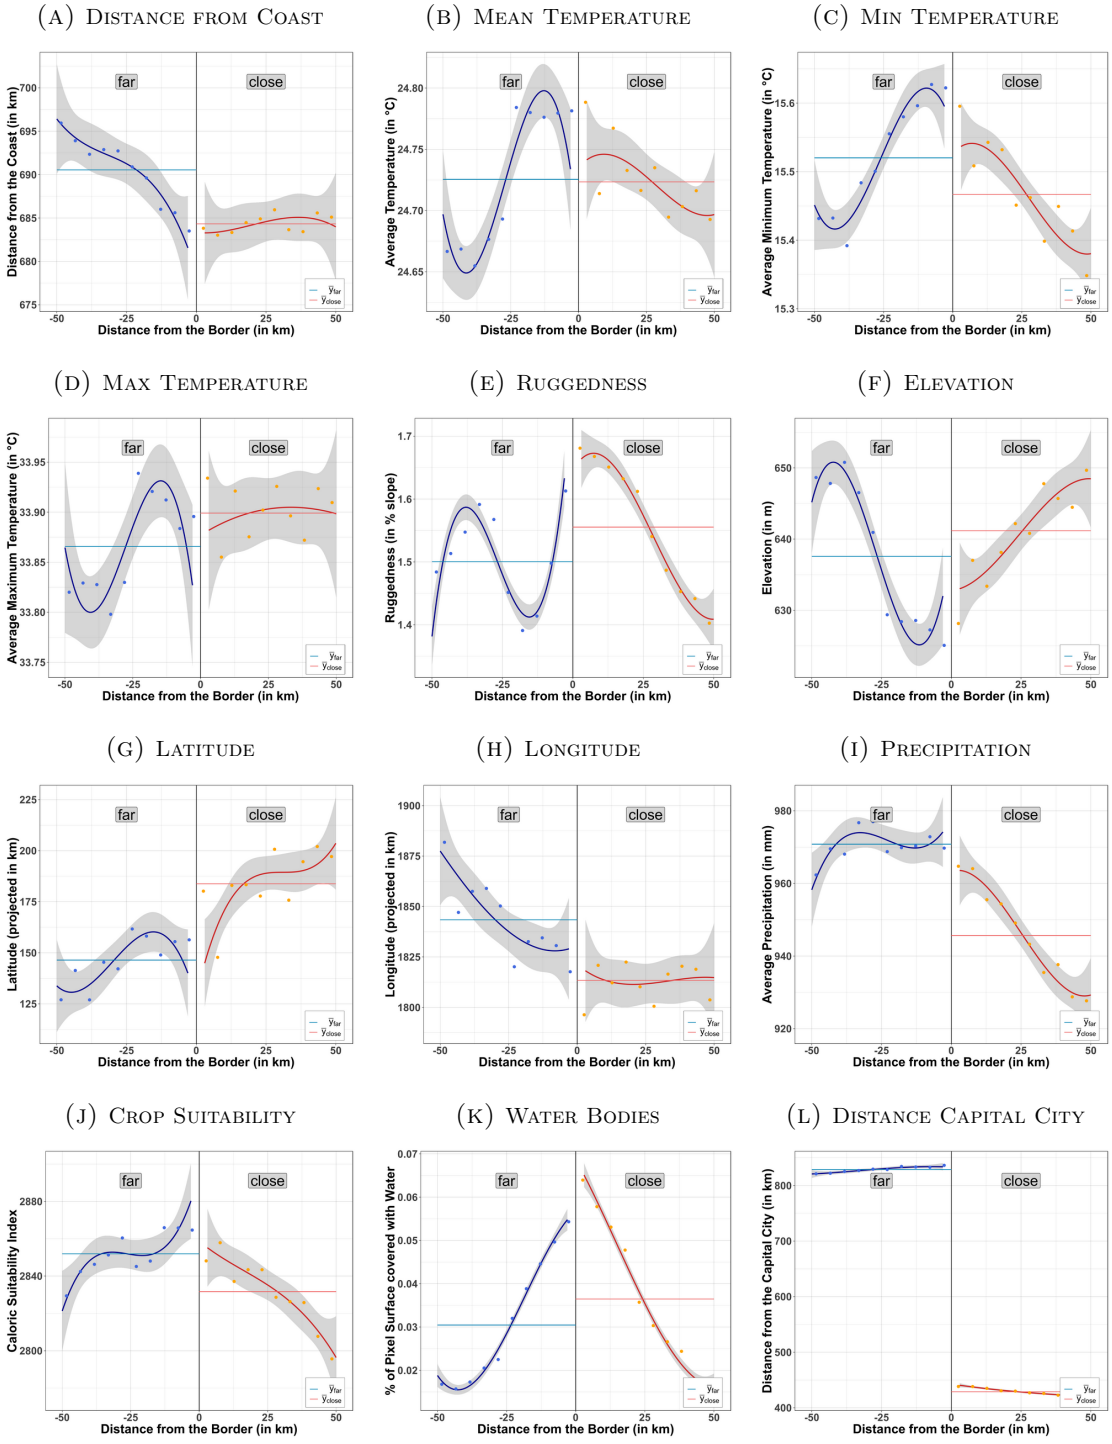

Note: The graphs illustrate the graphical balancing tests corresponding to Figure 2.1. The gray buffer around the lines represent the 95% confidence interval. The bins on the left-hand side are, with an average distance of 830 km, relatively far from the capital city and represent a total of 241,241 pixels. In contrast, pixels on the right-hand side are, with an average of around 430 km, relatively close to the capital city and represent 238,908 pixels.

Figure 2.A.12

Discontinuity Balancing Graphs: Placebo City

(a) Distance from Coast (b) Mean Temperature (c) Min Temperature

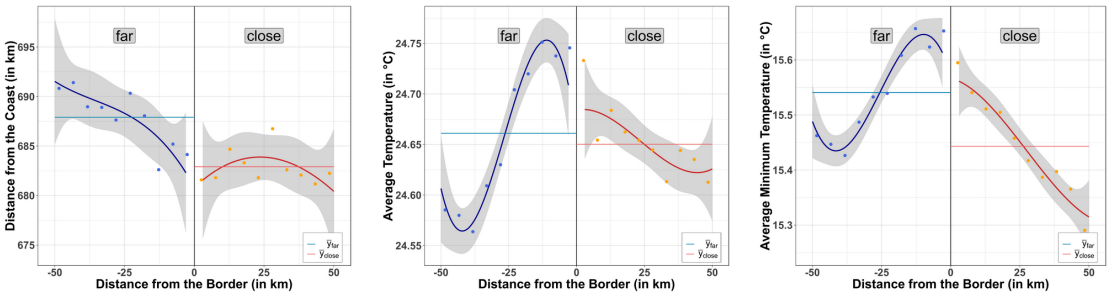

(d) Max Temperature (e) Ruggedness (f) Elevation

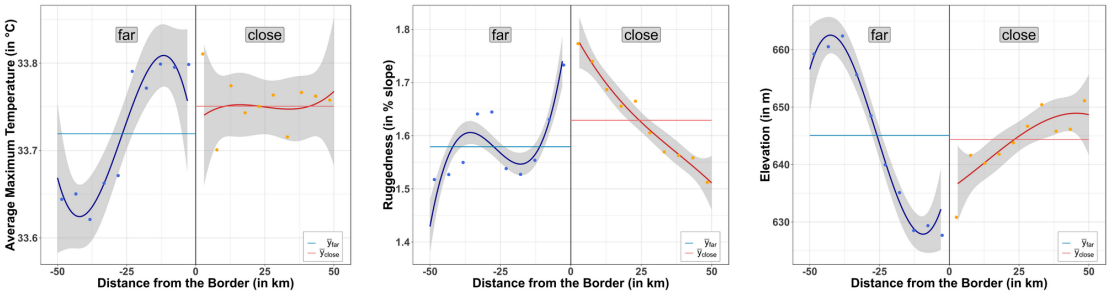

(g) Latitude (h) Longitude (i) Precipitation

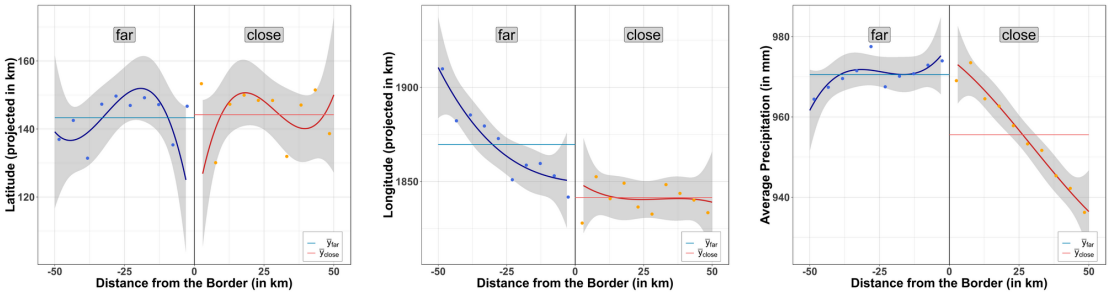

(j) Crop Suitability (k) Water Bodies (l) Distance Capital City

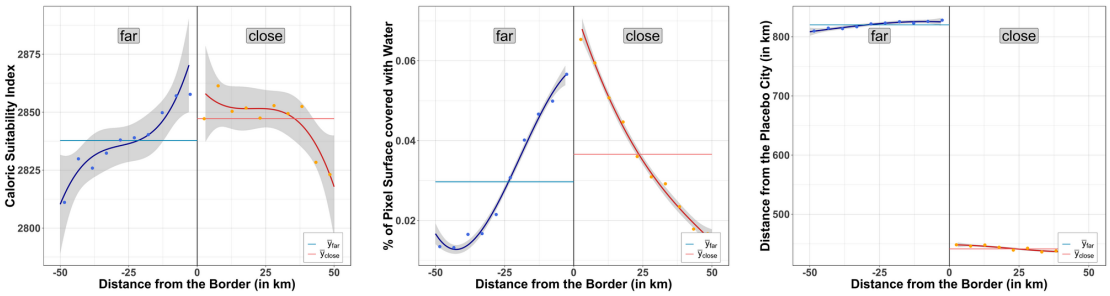

Note: The graphs illustrate the graphical placebo balancing tests corresponding to Figure 2.2. The gray buffer around the lines represent the 95% confidence interval. The bins on the left-hand side are, with an average distance of 820 km, relatively far from the capital city and represent a total of 238,980 pixels. In contrast, pixels on the right-hand side are, with an average of around 440 km, relatively close to the capital city and represent 237,153 pixels.

2.A.9 Additional Results

FIGURE 2.A.13
'NAIVE' OLS ESTIMATES BY COUNTRY

(A) OLS β - COEFFICIENTS (EXTENSIVE)

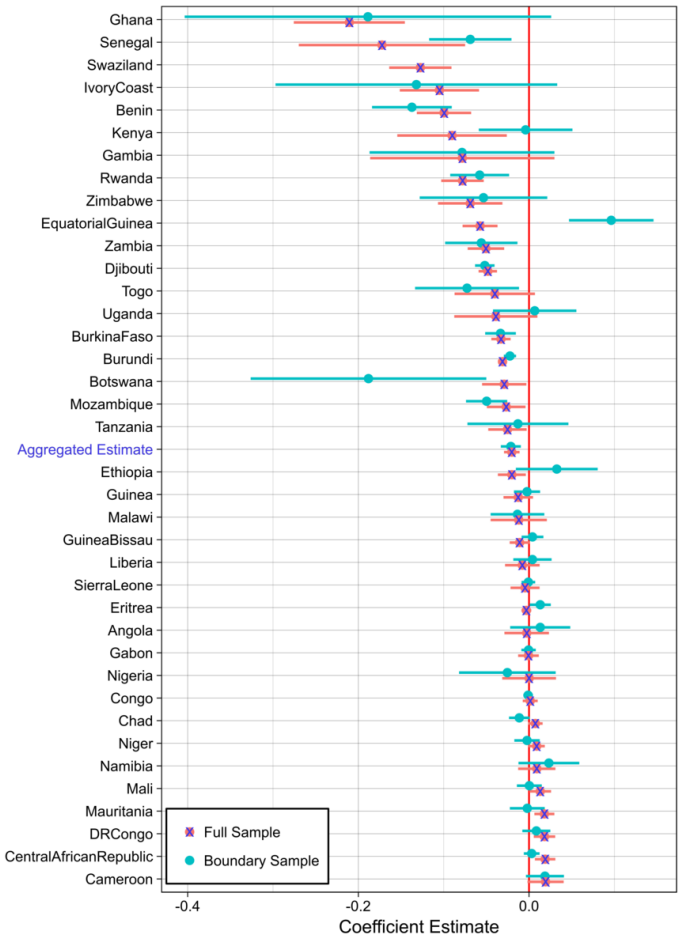

(B) OLS β - COEFFICIENTS (INTENSIVE)

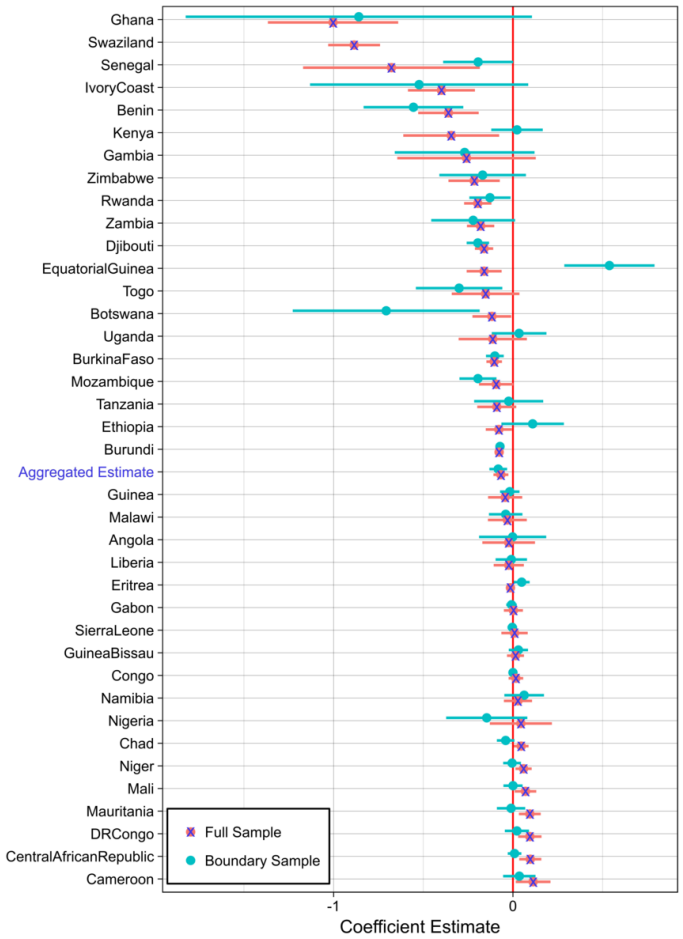

Note: On the left, Figure 2.A.13a present the OLS estimates by country based on columns (1) and (2) in Table 2.1. Likewise, Figure 2.A.13b corresponds to those of columns (5) and (6). The estimates corresponding to the boundary sample of Swaziland were omitted. The reason is that Swaziland only has relatively few observations and little variance in isolation from the capital city in the boundary area approach which inflates the standard errors and depreciates the readability of the graph. Standard errors are clustered by ethnic homeland.

Table 2.A.18
Border Discontinuity Estimation - Conley Standard Errors

	Dependent variable:									
	Probability Pixel is Lit in 2016 (VIIRS)					Log Light Density in 2016 (VIIRS)				
	Distance Cutoff (in km):									
	50	100	200	500	1,000	50	100	200	500	1,000
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Conley Standard Erros w/o Bartlett Correction									
Log Distance from the Capital City	-0.031**	-0.031**	-0.031*	-0.031*	-0.031**	-0.137*	-0.137**	-0.137	-0.137	-0.137*
	(0.015)	(0.015)	(0.018)	(0.017)	(0.015)	(0.075)	(0.070)	(0.085)	(0.084)	(0.073)
	Conley Standard Erros w/ Bartlett Correction									
Log Distance from the Capital City	-0.031***	-0.031**	-0.031**	-0.031*	-0.031*	-0.137**	-0.137**	-0.137**	-0.137*	-0.137*
	(0.012)	(0.014)	(0.014)	(0.017)	(0.016)	(0.058)	(0.067)	(0.068)	(0.080)	(0.079)
	Polynomials for: distance from the border × country × ethnicity (305 groups)									
3rd order	x	x	x	x	x	x	x	x	x	x
Geography Cov.	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Country FE	36	36	36	36	36	36	36	36	36	36
Segment FE	569	569	569	569	569	569	569	569	569	569
Observations	168,620	168,620	168,620	168,620	168,620	168,620	168,620	168,620	168,620	168,620
Adjusted R^2	0.182	0.182	0.182	0.182	0.182	0.161	0.161	0.161	0.161	0.161

Note: This table reports robustness tests for our baseline BDD estimates in columns (3) and (7) in Table 2.2 based on Equation 2.4 using Conley standard errors with and without Bartlett correction for various distance cutoffs (Conley, 1999). We estimate these SEs using the STATA package 'acreg' developed by Colella et al. (2019). The 'Geographical Cov.' include: distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km). The observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Conley standard errors are in parenthesis. *p<0.1; **p<0.05; ***p<0.01

Table 2.A.19
Placebo Tests

	Dependent variable:							
	Boundary Area Regression				Boundary Discontinuity Regression			
	Probability Pixel is Lit in 2016 (VIIRS)		Log Light Density in 2016 (VIIRS)		Probability Pixel is Lit in 2016 (VIIRS)		Log Light Density in 2016 (VIIRS)	
	OLS	Ethn FE	OLS	Ethn FE	BDD	City Size	BDD	City Size
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Distance from	-0.022***	-0.014**	-0.087***	-0.065***	-0.018**	-0.019**	-0.070*	-0.072*
the Capital City	(0.006)	(0.005)	(0.027)	(0.025)	(0.009)	(0.009)	(0.037)	(0.037)
Log Distance from	-0.008	-0.009	-0.035	-0.038	-0.005	-0.010	-0.041	-0.065
the Placebo City	(0.005)	(0.007)	(0.024)	(0.030)	(0.018)	(0.017)	(0.081)	(0.073)
CAP × SIZE_{CAP}	-	-	-	-	-	0.001	-	0.004
	-	-	-	-	-	(0.001)	-	(0.005)
PLC × SIZE_{PLC}	-	-	-	-	-	0.006	-	0.029
	-	-	-	-	-	(0.010)	-	(0.047)
	Polynomials for: distance from the border × country × ethnicity (299 groups)							
3rd order	-	-	-	-	x	x	x	x
Geography Controls	YES	YES	YES	YES	YES	YES	YES	YES
Country FE	37	37	37	37	35	35	35	35
Ethnicity FE	NO	351	NO	351	-	-	-	-
Segment FE	-	-	-	-	554	554	554	554
Observations	414,879	414,876	414,879	414,876	164,337	164,337	164,337	164,337
Adjusted R^2	0.083	0.133	0.071	0.122	0.179	0.179	0.156	0.156

Note: This table reports our boundary area and BDD placebo test based on Equation 2.5. In order to avoid capturing the break between the capital or placebo cities and the hinterlands, we exclude 20 km around each capital and placebo city from our sample. To prevent misassignment of detected nightlights between countries due to blooming, we exclude 3 km on each side around shared national boundaries. The 'Geographical Cov.' include: distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km). Columns (1)-(4) constitute boundary area regressions and columns (2) and (4) additionally include ethnicity fixed effects. Boundary segments corresponds to a buffer of 25 km around border pieces of 50 km line length and entirely nested within a restricted ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). Columns (6) and (8) include interactions between isolation from the capital and placebo city and their respective population counts. The observations in columns (5)-(8) are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by ethnic homeland in columns (1)-(4) and border segment in columns (5)-(8) *p<0.1; **p<0.05; ***p<0.01

Table 2.A.20
Boundary Discontinuity Estimation: Public Goods

	Dependent variable:				
	Road Infrastructure Provision		Health Care Provision		
	Prob. Road Exists	Prob. Tarred Road Exists	Log Dist to Health Center	Log Dist to Hospital	Log Dist Public Health Center
	(1)	(2)	(3)	(4)	(5)
Log Distance from the Capital City	0.004	-0.019**	0.123	0.435***	0.567**
	(0.018)	(0.007)	(0.150)	(0.123)	(0.281)
	Polynomials for: distance from the border × country × ethnicity				
3rd order	x	x	x	x	x
Population Density	YES	YES	YES	YES	YES
Geography Cov.	YES	YES	YES	YES	YES
Country FE	36	36	34	34	27
Segment FE	569	569	483	483	327
Observations	168,620	168,620	141,578	141,578	93,937
Adjusted R^2	0.088	0.203	0.715	0.872	0.751

Note: This table reports results on public goods as outcome variables using our BDD framework and is based on Equation 2.4. The road infrastructure depend variables are a dummy indicating if a road intersects with a pixel and if a tarred road intersects with a pixel. The health care provision variables include the log of the distance to the closest health center, the closest hospital and the closest public health center. The drop in the number of observations in column (5) is due to the fact that the ownership of the health facility is missing for Equatorial Guinea, Guinea, Liberia, Mauritania, Mozambique and Niger. The 'Geographical Cov.' include: distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km). In addition, in column (2) we include as dummy indicating if there exists a road. The observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by boundary segment.

*p<0.1; **p<0.05; ***p<0.01

TABLE 2.A.21

CHANNEL ANALYSIS: SUPPLEMENTARY RESULTS ON POLITICAL ATTITUDE

	BDD Model with dependent variable z-score of:			
	Trust in Ruling Party	Trust in Opposition Party	Voter Turnout	National (vs Ethnic) Identity
	(1)	(2)	(3)	(4)
Log Distance from the Capital City	0.169***	0.023	0.232***	0.039
	(0.060)	(0.052)	(0.051)	(0.049)
	Polynomials for: distance from the border			
3rd order	x	x	x	x
Household Cov.	YES	YES	YES	YES
Geography Cov.	YES	YES	YES	YES
Country × Round FE	67	67	70	68
Segment FE	140	140	140	140
Observations	7,794	7,809	7,702	8,007
Adjusted R^2	0.163	0.057	0.086	0.113

Note: This table reports the supplementary regressions on the impact of isolation from the capital city on the perception of political leaders and accountability. The 'Household Cov.' include: age, age squared and sex of respondent. The 'Geographical Cov.' include: distance from the coast (in km), longitude and latitude (projected in km) and whether the household is in an urban or rural setting. Columns (1) and (2) correspond to trust into the ruling and opposition party respectively. Column (3) to voter turnout and column (4) to the extent to which the respondent identifies with the nation rather than the ethnicity. All models are BDD regressions using 'Segment FE' for boundary segments of 50 km length with a buffer of 25 km that are nested within an ethnic homelands based on the 'Tribal Map of Africa' (Murdock, 1959). All observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by Afrobarometer cluster and boundary segment. *p<0.1; **p<0.05; ***p<0.01

Table 2.A.22
Channel Analysis: Afrobarometer Placebo Tests

	BDD Model with dependent variable z-score of:										
	Public Goods	Trust Leader	Corrupt Percept	Perform Eval	Trust Rule	Trust Oppos	Voter Turnout	National Identity	Educ Level	News Reader	Checks Balance
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Log Distance from the Capital City	-0.064*	0.173**	-0.070	0.226**	0.177**	0.064	0.122***	0.070	0.064	-0.092	-0.147**
	(0.038)	(0.071)	(0.068)	(0.105)	(0.081)	(0.070)	(0.027)	(0.068)	(0.090)	(0.083)	(0.059)
Log Distance from the Placebo City	0.059*	-0.024	0.023	0.080	0.012	0.062	0.041	0.037	0.003	0.097	-0.077
	(0.031)	(0.071)	(0.074)	(0.124)	(0.076)	(0.066)	(0.028)	(0.061)	(0.083)	(0.070)	(0.055)
Placebo Tests	$\text{Coef}_{CapitalCity} = \text{Coef}_{PlaceboCity}$										
F-Statistic	15.50***	8.49***	4.27**	3.51*	7.00***	0.0	9.82***	0.46	1.29	11.90***	2.16
	Polynomials for: distance from the border										
3rd order	x	x	x	x	x	x	x	x	x	x	x
Household Cov.	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Geography Cov.	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Country × Round FE	70	70	71	71	67	67	70	68	71	71	71
Segment FE	70	138	138	138	138	138	138	138	138	138	138
Observations	8,053	7,601	6,525	6,414	7,590	7,598	7,501	7,802	8,096	7,902	7,556
Adjusted R^2	0.614	0.183	0.162	0.202	0.168	0.058	0.087	0.110	0.296	0.173	0.137

Note: This table reports the BDD placebo tests corresponding to the Afrobarometer regressions in Table 2.1, 2.3 and 2.A.21. We report the respective F-tests on the equality of the coefficients. The 'Household Cov.' include: age, age squared and sex of respondent. The 'Geographical Cov.' include: distance from the coast (in km), longitude and latitude (projected in km) and whether the household is in an urban or rural setting. The BDD sample is restricted to a buffer of 25 km around the shared national boundaries. The 'Segment FE' are nested within an ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). The observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by Afrobarometer cluster and boundary segment. *p<0.1; **p<0.05; ***p<0.01

Table 2.A.23
Channel Analysis: Dynamic Political Support - Placebo Tests

	Dependent variable:					
	Corruption Perception		Trust Pol. Leaders		Vote for Government	
	Panel	BDD-Panel	Panel	BDD-Panel	Panel	BDD-Panel
	(1)	(2)	(3)	(4)	(5)	(6)
Log Corruption Level ×	-0.326***	-0.552	0.107	1.278**	0.144**	1.107**
Log Dist Cap	(0.093)	(0.730)	(0.095)	(0.579)	(0.056)	(0.469)
Log Corruption Level ×	0.277**	0.490	-0.366**	-0.373	-0.123*	-0.645
Log Dist PLC	(0.120)	(0.721)	(0.148)	(0.780)	(0.063)	(0.547)
	Polynomials for: distance from the border					
3rd order	-	x	-	x	-	x
Log Corruption Level	*absorbed*	Yes	*absorbed*	Yes	*absorbed*	Yes
Log Distance from the Capital City	Yes	Yes	Yes	Yes	Yes	Yes
Log Distance from the Placebo City	Yes	Yes	Yes	Yes	Yes	Yes
Geography Controls	Yes	Yes	Yes	Yes	Yes	Yes
Household Controls	Yes	Yes	Yes	Yes	Yes	Yes
Cnt × Ethn × Round FE	Yes	No	Yes	No	Yes	No
Country FE	No	Yes	No	Yes	No	Yes
Ethnicity FE	No	Yes	No	Yes	No	Yes
Round FE	No	Yes	No	Yes	No	Yes
Segment × Round FE	-	157	-	157	-	146
Observations	73,203	3,968	81,760	4,623	60,026	3,123
Adjusted R^2	0.158	0.111	0.195	0.156	0.284	0.232

Note: This table reports the placebo regressions corresponding to Table 2.4. The 'Household Cov.' include: age, age squared and sex of respondent. The 'Geographical Cov.' include: distance from the coast (in km), longitude and latitude (projected in km) and whether the household is in an urban or rural setting. The dependent variable in columns (1)-(2) is the normalized corruption perception among political leaders, in columns (3)-(4) it is the normalized level of trust into the political leadership and in columns (5)-(6) it is a dummy indicating whether the respondent would vote for the incumbent government if there was an election held. All observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by Afrobarometer cluster and ethnicity × round. *p<0.1; **p<0.05; ***p<0.01

Table 2.A.24
Channel Analysis: Market Access and Trade

	Dependent variable: VIIRS Nightlights in 2016 (Prob/Log)					
	Prob	Log	Prob	Log	Prob	Log
	(1)	(2)	(3)	(4)	(5)	(6)
Log Distance from the Capital City	-0.013	-0.067	-0.037***	-0.164**	-0.035**	-0.158**
	(0.011)	(0.044)	(0.014)	(0.068)	(0.014)	(0.066)
CAP × FTA	-0.028***	-0.113***	-	-	-	-
	(0.011)	(0.044)				
CAP × Customs Union	-	-	0.023	0.098	-	-
			(0.015)	(0.061)		
CAP × Customs Union and Monetary Union	-	-	-	-	0.018	0.083
					(0.015)	(0.060)
	Polynomials: dist. to border × country × ethnicity (305 groups)					
3rd order	x	x	x	x	x	x
Geography Cov.	YES	YES	YES	YES	YES	YES
Country FE	36	36	36	36	36	36
Segment FE	569	569	569	569	569	569
Observations	168,620	168,620	168,620	168,620	168,620	168,620
Adjusted R^2	0.182	0.161	0.182	0.161	0.182	0.161

Note: This table reports the tests for heterogeneities at boundaries that are less restrictive to trade and extend the model in Equation 2.4. In order to avoid capturing the break between the capital city and the hinterlands, we exclude 20 km around each capital city from our sample. To prevent misassignment of detected nightlights between countries due to blooming, we exclude 3 km on each side of the border. The 'Geographical Cov.' include: distance from the coast (in km), ruggedness (in % slope), % surface covered with water, mean annual temperature, minimum average temperature during the coldest month, maximum average temperature during the warmest month (in °C), crop caloric index, annual precipitation (in mm), longitude and latitude (projected in km). Boundary segments corresponds to a buffer of 25 km around border pieces of 50 km line length and are entirely nested within a restricted ethnic homeland based on the 'Tribal Map of Africa' (Murdock, 1959). We include interactions of CAP (*Log Distance to the Capital City*) and a dummy for boundaries between countries with FTAs (columns (1)-(2)), customs unions (columns (3)-(4)) and those that additionally share a common currency (columns (5)-(6)). The observations are weighted such that each side of a segment has the same aggregated weight as its counterfactual. Standard errors in parenthesis are clustered by boundary segment. *p<0.1; **p<0.05; ***p<0.01

Chapter 3

The Economic Impact of Proximity to Regional Capitals: Evidence from Boundary Reforms

3.1 Introduction

Capital cities are the centers of political decision making in a country and as such constitute a core building block of the political geography of states. Recent studies found that the location of national capitals has a strong impact on economic outcomes at the country and local level (Campante and Do, 2014; Michalopoulos and Papaioannou, 2014) (see also Chapter 2). In addition to national capital cities that serve as seat of the central government, countries usually have a set of multiple regional capital cities that serve as an extended arm of the national government and also have some autonomy to rule over certain regional and local matters.[1] Yet, only little is known about the economic relevance of regional capitals for the spatial distribution of economic activity within the region. In this paper, we aim to fill this gap by using a worldwide sample of region splits that led to the creation of new regional capital cities.

The literature on the role of capital cities has emphasized the challenge of states to impose rules and broadcast power beyond the capital cities into the hinterlands (Herbst, 2000; Michalopoulos and Papaioannou, 2014). Further, Campante and Do (2014) and Campante et al. (2019) show that states with isolated capital cities exhibit negative outcomes such as increased corruption, conflict and misgovernance. In Chapter 2, we investigate the micro-level impact of remoteness from the capital city on locations in Sub-Saharan African countries. We find that isolation from the capital city reduces public goods provision and ultimately economic performance by creating information frictions that disincentivizes the attention of the political leadership.

Regional capital cities might shape the spatial distribution of economic activity within regions in a similar way. On the one hand, regional capitals might serve as outposts of

[1] We use the Global Administrative Unit Layer (GAUL) by the Food and Agriculture Organization of the United Nations (FAO) and define regions as the first level of government subdivision.

national capital cities and reduce the distance between the central government and the hinterlands. On the other hand, regions can be viewed as mini-countries that, depending upon the level of decentralization, also have a regional government that is in charge of certain local affairs. In a related paper, Asher et al. (2018) provide first evidence from rural India that remoteness from subnational centers of political administration might have a similarly negative impact as the national capital city. The authors find that distance to district headquarters reduces income and public goods provision. Similarly, Henn (2021) shows that distance to administrative headquarters in Sub-Saharan Africa reduces economic performance.[2] However, both studies are facing a number of methodological caveats with regard to providing a final answer about the causal impact of distance to regional administrative headquarters that we seek to address in this study.[3] Another study by Richard Bluhm (2021) investigates the effects of designating new regional capital cities but mainly focuses on the impact on the new capital itself.[4] The authors document that gaining capital city status spurs city growth through public and private investments and migration, at least in places with favorable geographical conditions such as good market access. We extend this analysis by shifting the focus away from the new regional capital itself and give a more detailed picture of the effects on the wider region.

Another strand in the literature, in particular in political science, has studies the increased proliferation of administrative units across developing countries. This topic is very relevant in this context as the creation of new regional capital cities is a corollary of administrative unit proliferation. A potential advantage for breaking down regions into smaller ones could be a higher degree of ethnic homogeneity that reduces political violence and conflict (Cederman et al., 2015; Pierskalla, 2016). Another advantage of smaller constituencies could be related to more homogenous preferences that allow a more efficient provision of public goods (Oates, 1972). In line with this theory, Grossman et al. (2017) find that public goods provision increased following region splits in Sub-Saharan Africa between 1960 and 2012. In contrast, Billing (2019) find that newly created regions in Burkina Faso are significantly

[2]It should be mentioned that the methodological design of Henn (2021) does not aim at understanding the role of administrative headquarters. The author uses regional and local administrative headquarters to obtain variation in state capacity to study the constitutional framework conditions under which the traditional leadership serves as a substitute or as a complement of the state when state capacity is low.

[3]Most importantly, both studies use subnational boundaries in a boundary discontinuity design (BDD). However, as we will discuss in more detail below, subnational borders are not arbitrary and induce discontinuous variation in other covariates. In addition to a change in the identification strategy, we focus exclusively on regional rather than local headquarters. Moreover, we exclude the local area around the regional capitals to avoid picking up effects on the cities themselves.

[4]The authors document positive spillovers within the wider urban area around the new regional capital up to 75-100km away but not beyond.

weakened and lack the resources and administrative ability to provide public goods. Likewise, Lewis (2017) finds that region splits decreased public goods such as water and sanitation provision in Indonesia. A reason for this might be that region splits are sometimes aimed at reinforcing patronage and recentralizing power by weakening local constituencies (Green, 2010; Grossman and Lewis, 2014). Another intention behind region splits could be grouping constituencies in a specific way to win elections (Hassan, 2016). Baskaran and Blesse (2019) provide aggregate evidence from a large sample of African boundary splits and mergers between 1992 and 2013. This study finds that both splits and mergers have a positive impact on economic development. The authors argue that splits have a small positive effect by reducing conflict and mergers have a larger effect by improving administrative efficiency.

What becomes clear from this literature is that the motives behind region split might play an important role for the outcomes. Region splits might be advantageous if they aim to resolve ethnic cleavages or bring locations closer to their administrative center. However, any potential positive effects from getting closer to a regional capital city might be undermined when the entire administrative capacity of the new region is weakened. For this reason, we will attempt to shed light into how political framework conditions impact the effects of region splits. More specifically, we will show that being closer to the regional capital city only has positive effects under democratic but not under autocratic institutions.

The main challenge for identifying the causal effect of proximity to the regional capital city is that the location of capital cities is not random. There are a range of confounding geographical and geological factors like distance from the closest river or infrastructure that simultaneously change with distance from the regional capital. In Chapter 2, we induces quasi-random variation in distance from the national capital city in Sub-Saharan African countries by applying a boundary discontinuity design (BDD) at national borders. More specifically, we exploit the special circumstance that African national borders were arbitrarily set and imposed from the outside by the colonial powers.[5] As a result, certain boundary segments divide pre-colonial ethnic homelands with common geographical, social and historical characteristics.[6] Similarly, Asher et al. (2018) and Henn (2021) use a BDD at subnational administrative boundaries in an attempt to obtain quasi-random variation in distance from the administrative headquarter. However, using a BDD in this context is problematic as regional boundaries are not random but delineate ethnic and geographic

[5]See Michalopoulos and Papaioannou (2014, 172) for a detailed discussion of the origin of national borders in Sub-Saharan Africa.

[6]See Section 2.1.2 for a detailed discussion of the identification strategy used in Chapter 2.

shifts. For this reason, Asher et al. (2018) discovers other discontinuities such as that the boundary side that is closer to the headquarter is simultaneously closer to the next mid-sized city or river, and in case of Henn (2021) closer to colonial railroads and the national boundary. As a consequence, cross-sectional and BDD comparisons are not suitable to identify the causal impact of distance from the regional capital in this context.

To overcome the challenges associated with identifying the impact of distance to the regional capital city, we exploit the longitudinal variation induced by region splits. Whenever a region is split, there are two types of newly created regions. Firstly, there are the regions that contain the old regional capital which usually persists as such. However, those regions do not experiences a change in distance from the capital city which is why we omit them from the analysis. Secondly, the remaining 'splinter regions' require the assignment of a new regional capital city. Consequently, all areas in these new regions experience a change in distance from the capital city that can be exploited in an event study to analyze the effect of proximity to regional capitals. However, a simple before vs. after event study might still be biased by local economic shocks. To be able to effectively control for such confounders, we focus on the boundary area of these splinter regions and use neighboring regions which did not undergo a boundary reform themselves as control areas. To test if these controls are suitable counterfactuals, we run an event study and test for parallel pre-trends.
Our sample comprises of 17 countries in Africa, the Americas and Asia where in the period between 1996 and 2005 35 regions split and 37 splinter regions with new regional capitals were created. While areas that experienced a strong as compared to those that experienced only a small decline in distance to the regional capital city evolved similarly prior to the reform, they significantly gained relatively after the reform. On average, a 1 percent decrease in distance from the regional capital city increases the intensity of nightlights by 0.16 percent within 8 years. Furthermore, we find strong differences in the relevance of regional capital cities between countries. Regional capitals appear to be primarily relevant in democratic and decentralized countries, and less relevant in relatively developed countries. Finally, we do not document an impact of proximity to regional capitals on the propensity of conflict which makes it unlikely that conflict is a driver behind the economic gains.

The remainder of the paper is organized into 3 sections. Firstly, Section 3.2 will introduce the identification strategy and dataset. Secondly, Section 3.3 will present our empirical results. Finally, Section 3.4 summarizes the findings and concludes the paper.

3.2 Data and Empirical Strategy

3.2.1 Empirical Strategy

This study aims to examine if proximity to regional capital cities has a positive impact on a location's economic performance. The main challenge of estimating the causal effect is the endogenous selection of regional capitals. The choice for the capital usually falls on the largest and best performing city within the region. We therefore have to assume that distance from the capital is correlated with a variety of other factors that are themselves related to economic performance such as market access, proximity to infrastructure, lakes, rivers, the coastline or other geographic variables such as soil quality or climate. One potential solution might be to conduct a BDD at regional boundaries, however regional borders usually delineate geographic and ethnic shifts which is why areas on both sides of the boundary are unlikely to be valid counterfactuals in a cross-sectional comparison.

For this reason, we pursue a longitudinal identification strategy and focus on areas within regions that underwent boundary reforms and experienced a shift in their distance to the regional capital. More specifically, we will target region splits[7] i.e. the division of the original area of a region into two or more smaller regions. There are two types of new regions: i) regions containing the old regional capital city which usually persists as regional capital city - we never use these regions in our analysis as they do not experience a change in the capital city and ii) regions that do not contain the old regional capital city that get assigned a new designated regional capital city - which we refer to as 'splinter region'. As a consequence, after the region split, all locations within these newly created splinter regions experience a shift in distance to the new regional capital city. Our identification strategy seeks to exploit this shift in distance from the regional capital to study its impact on economic development.

In Figure 3.1, we illustrate a region split based on the Indian region 'Madhya Pradesh' which in November 2000 was split into the regions 'Madhya Pradesh' and 'Chhattisgarh'. Bhopal used to be the capital city of the original region and persisted as capital of the new Madhya Pradesh region after the split (the green stars in Figure 3.1 represent the capital cities). Raipur was selected as the new capital city of the newly formed region of

[7]There are two other less common types of regional boundary reforms: mergers and complete redrawings of boundaries. While mergers refer to aggregating the original area of two or more regions into one large region, there are a few rare cases where boundaries get redrawn entirely without following the original demarcation lines or only following them partially.

Figure 3.1

Region Split Example: Madhya Pradesh Region Nigeria

(a) Indian Regions in 2000

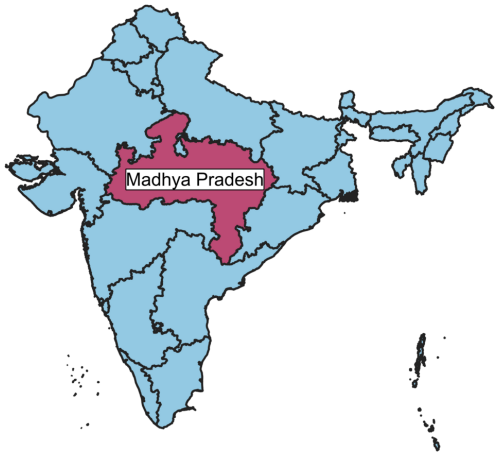

(b) Pre-split (2000)

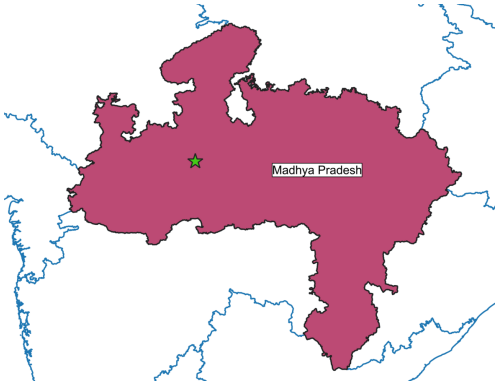

(c) Post-split (2001)

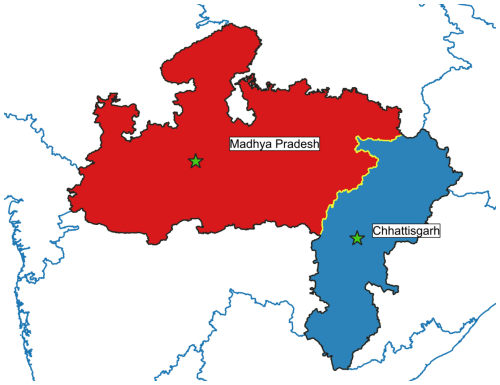

(d) Δ Proximity to Reg. Capital

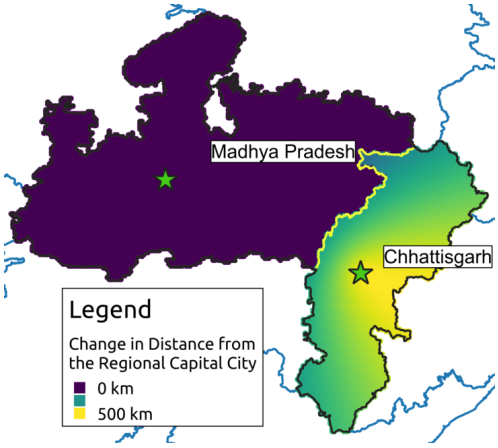

Note: These figures illustrate the split of the Indian state 'Madhya Pradesh' into 'Madhya Pradesh' and 'Chhattisgarh' in 2001. In this example, Chhattisgarh is the splinter region with 'Raipur' as the new state capital. 'Δ Proximity to Reg. Capital' is computed as the distance from the old regional capital minus distance from the new regional capital. Figure 3.1d illustrates that only the areas in Chhattisgarh experience a change in their distance to the capital.

'Chhattisgarh'. Consequently, after the split the location of the corresponding capital city remained the same for all areas inside the new Madhya Pradesh region and changed for all areas in the splinter region Chhattisgarh (see Figure 3.1d). For this reason, we only focus on the splinter regions in our sample to study the impact of proximity to the regional capital city, and exclude all areas in new regions that contain the old regional capital city.

3.2.2 Identification

Our identification strategy is centered around the idea of exploiting the shift in distance from the regional capital cities in splinter regions to study the economic impact of proximity to regional capitals. For this reason, our explanatory variable is a continuous treatment intensity that captures the change in log distance from the regional capital city. Figure 3.2a illustrates the distribution of distance from the regional capital city in our study sample of splinter regions before and after the split, and Figure 3.2b presents the distribution of the treatment intensity.

FIGURE 3.2
DISTANCE TO THE REGIONAL CAPITAL CITY
(A) DISTRIBUTION: BEFORE & AFTER (B) DISTRIBUTION: CHANGE

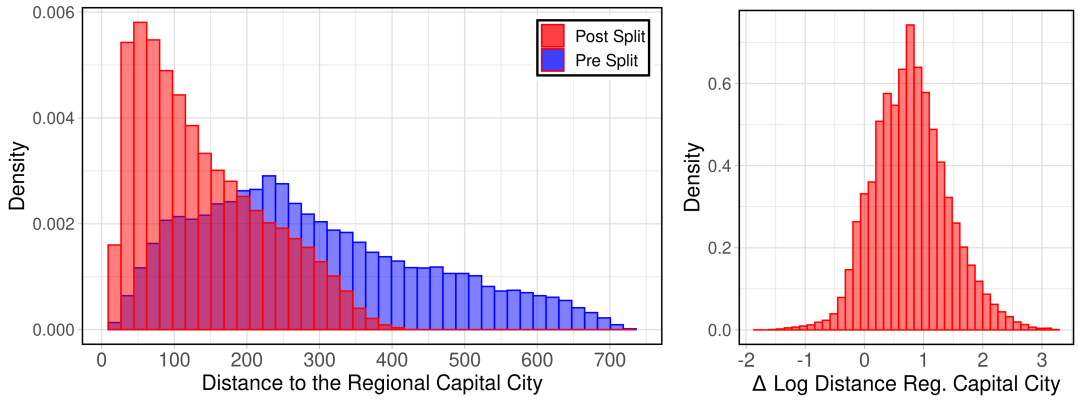

Note: These figures illustrate how the region split and establishment of a new regional capital city in the splinter region has altered how close locations are situated relative to the regional capital city. Figure 3.2a displays the histograms of distance to the regional capital city before and after the split. Figure 3.2b plots the histogram of the difference in log distance to the regional capital city, which is our measure of treatment intensity.

One way of estimating the relationship between distance from the regional capital city would be to use an event study that tracks locations in splinter regions a few years before until a few years after the region split. We can then examine if locations with a higher treatment intensity evolve similarly prior to the region split and diverge thereafter.

$$Y_{i,t} = \sum_{\substack{t=-4,\\ t\neq -1}}^{8} \beta_t \ \Delta CAP_i + b_i + b_{r,t} + \epsilon_{i,t} \tag{3.1}$$

Equation 3.1 presents the corresponding estimation equation where i refers to a pixel, t to relative time (e.g. $t = 1$ is one year after the split), $Y_{i,t}$ to our measure of economic performance in each pixel and relative time period and ΔCAP_i to the change in log distance from the regional capital city for pixel i. To account for time invariant pixel specific characteristics such as climate, ruggedness or distance from a river, we include pixel fixed effects b_i. In addition, we include region-year fixed effects, $b_{r,t}$, to control for all potential time-variant confounders that impact a specific region such as economic cycles or the effects of the split on the region. We cluster double-cluster standard errors by region-year and on the pixel level to account for the autocorrelation of residuals over space and time. Since we are interested in how proximity to the regional capital city impacts the distribution of economic activity in the wider region and not the capital cities themselves, we exclude the area within 20km of the capital city by default from all estimations. Furthermore, to avoid inaccuracies in the assignment of locations to the respective region as well as blooming of nightlights across borders, we exclude all pixels within 5km on either side of the boundary by default from all estimations.

The potential shortcomings of this approach are twofold. Firstly, it relies on the assumption that the change in distance from the capital city is unrelated to economic performance in the absence of the split. Yet, it could be that areas with a higher change in distance to the regional capital grow at a faster rate. One way of checking if this is a relevant caveat is to analyze if areas that subsequently gain more in terms of proximity to the regional capital city are already growing at a higher rate before the region split. Practically, this can be done by testing for parallel pre-trends. A second concern associated with this approach is that it does not control for local economic shocks that might distort the estimated effect. As an example, it could be that there happen to be mineral discoveries in areas with a higher or lower treatment intensity and thereby lead to over- or underestimations of the effects. Consequently, estimates based on this approach might be biased in either direction.

As a solution, we propose to combine the event study approach with a BDD approach. The idea is to focus on boundary areas and use locations in neighboring regions as controls for local economic shocks that are unrelated to the shift in distance from the regional

FIGURE 3.3
ILLUSTRATE IDENTIFICATION: CHHATTISGARH

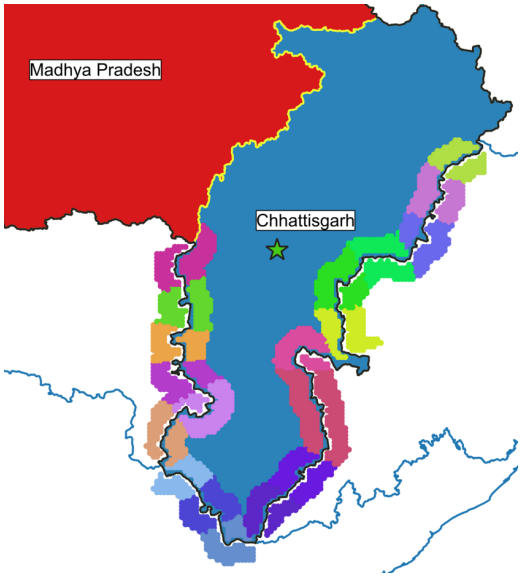

Note: This figure illustrates the boundary segments for the example of the splinter region 'Chhattisgarh'. The boundary segments play an important role for the identification strategy and allow us to absorb local economic shocks that equally affect both sides of the boundary for each period.

capital city. In practice, we divide the boundary of the splinter region into segments of 100 km line length with a buffer of 25 km on each side of the boundary.[8] Using segment-year fixed effects, we control for any events and local economic shock in each period that impact both sides of the boundary, as long as it spills over across the boundary. Figure 3.3 illustrates the boundary segments for the example of the Chhattisgarh region. In principle, we could use all areas in neighboring regions just across the boundary as controls. However, to avoid that the estimates are distorted, the ideal control group should not be affected by the treatment itself. In our example, areas in the new Madhya Pradesh region are not a suitable control as they were also affected by the region split (which is why there are no segments around the border between Chhattisgarh and the new Madhya Pradesh region). Moreover, neighboring regions that are themselves going through a boundary reform during the same period are also not a suitable control group. In our example the neighboring regions in the north also underwent splits with Jharkhand splitting from Bihar and Uttarakhand splitting from Uttar Pradesh in the year 2000.

$$Y_{i,t} = \sum_{\substack{t=-4,\\ t\neq -1}}^{8} \beta_t \ \Delta CAP_i + b_i + b_{r,t} + b_{s,t} + \varepsilon_{i,t} \tag{3.2}$$

[8] We only keep boundary segments that have at least 10 pixels on each segment side, to avoid that in some instances at very uneven boundaries only few pixels serve as counterfactuals for the respective other side.

Equation 3.2 presents the corresponding new estimation equation. It is very similar to Equation 3.1, but additionally includes segment-year fixed effects $b_{s,t}$. To account for the spatial and temporal autocorrelation of residuals, we double-cluster standard errors at the segment-year and at the pixel level. Since there is no overlap between split events i.e. each split event has its own control group without overlapping fixed effects, we are not facing the challenges of staggered onset such as the comparison between early vs. late treated units (Goodman-Bacon, 2021). As in the simple event study analysis, we exclude areas within 20km around the capital cities and 5km around the borders from the sample.

3.2.3 Data

For our empirical approach, we require three core inputs: (i) data on the annual spatial extent and demarcation of regions (which we define as first-level administrative regions), (ii) the respective new and old capital cities for splinter regions and (iii) an indicator of economic performance for small geographical units over time.

To map first-level administrative regions for all countries across the world, we use the global administrative unit layer (GAUL) by the Food and Agriculture Organization of the United Nations (FAO). This database contains annual shapefiles with polygons of regions and covers the period 1990-2014. This dataset allow us to correctly map administrative regions across the world, and also identify if and when regions were divided into two or more regions. Furthermore, we use these shapefiles to create the boundary segments.

The second data input is a geocoded list of regional capital cities. Since there exists no comprehensive off-the-shelf dataset on regional capital cities, we collected information on relevant old and new regional capital cities ourselves from various sources including Wikipedia.[9] We geocoded this list of cities using OpenStreetMaps.

The third input is a measure of economic performance for small spatial units across regional boundaries in different parts of the world. Despite some shortcomings related to the fact that nightlights are a noisy measure of economic activity (Chen and Nordhaus, 2011; Cogneau and Dupraz, 2014), they have become a widely used proxy for economic development at the local level (see for example Henderson et al. (2012); Michalopoulos and Papaioannou (2013); Donaldson and Storeygard (2016)). The big advantage of nightlights is that they provide a detailed and consistent global coverage over time with a resolution

[9]Note that since Wikipedia is not an official administrative data base and given that users can edit data entries, we cannot exclude the possibility of misassigning capitals.

of 30 arc seconds ($\sim$ 1km at the equator). We use the Version 4 Defense Meteorological Satellite Program - Operational Linescan System (DSMP-OLS) nightlighs by the U.S. Air Force and the National Oceanic and Atmospheric Administration (NOAA) as they are available annually from 1992-2013.[10] While there are newer and improved[11] satellite nightlights available from the VIIRS (Visible Infrared Imaging Radiometer Suite) by the Suomi National Polar Partnership between NOAA and NASA, they are not suitable for our purposes as they are only available starting April 2012 (Elvidge et al., 2013).

Since our sample covers a very large area, we aggregate the nightlights grid by a factor of 5 which means that each block of 5 $\times$ 5 original pixels get converted into a new pixel with the average nightlight density. Further, it is common practice to use two measures of nightlights: i) a measure at the intensive margin computed as[12]: $\ln(\text{Lights}_i + 0.08)$ and ii) a measure at the extensive margin which is a dummy that is 1 if there are nightlights detected in a pixel and zero otherwise. The advantage of the extensive measure of nightlights is that it deemphasizes the magnitude of nightlights which might be distorted by confounding factors such as local cultural preferences or the presence of waterbodies that increase the reflected light density. Yet, the disadvantage of an extensive measurement of nightlights is that it is unable to recognize continuous and gradual changes over time. This circumstance makes the use of nightlight dummies particularly problematic in longitudinal analyses as changes in nightlight density in pixels that are lit can no longer be detected. For this reason, we will focus on the intensive margin of nightlights in our analyses but report results based on the extensive margin in the appendix.

3.2.4 Sample

During our study period between 1992 and 2013, there were a total of around 190 region splits in 39 countries.[13] In our event study analysis, we only focus on a subsample consisting of 35 regions splits in 17 countries (6 in Africa, 2 in the Americas and 9 in Asia) that led to the creation of 37 splinter regions with new capitals. One reason for the large reduction in the sample is related to our event study approach in which we only consider splits between 1996 and 2005, to observe all areas 4 years prior and 8 years post treatment to be able to test for parallel pre-trends and leave enough time for the effects to material-

[10] The annual DMSP-OLS nightlights are publicly available for download at: `https://www.ngdc.noaa.gov/eog/dmsp/downloadV4composites.html#AVSLCFC`.

[11] The main advantages of the VIIRS nightlights is that they are not top-coded like the DMSP-OLS, have a higher resolution and suffer less from blooming.

[12] The constant term of 0.08 equals the minimum detected nightlight density greater than 0 and is added before taking the log to avoid losing all observations with a density of 0.

[13] We focus on this study period because the DSMP-OLS nightlights, that allow us to obtain a consistent measure of economic activity at the local level, are only available within this time frame.

ize. In addition, we exclude all regions that experienced more than one boundary reform during this period to avoid confusing the impact of different reforms. Furthermore, given our focus on boundary areas, we exclude a few regions with mostly unpopulated boundary areas that do not emit nightlights such as the *Northwest Territories* in Canada. Moreover, some newly created splinter regions such as the Indonesian region and island *Bangka Belitung* do not have land boundaries with neighboring regions, or have two capitals such as the Indian state *Uttarakhand*. Finally, a few additional region splits are omitted because of missing data on capitals. Table 3.A.1 provides and overview of all region splits in our sample, and Table 3.A.2 provides summary statistics on key variables for all region splits and the split sample respectively. Figure 3.4 illustrates all countries that experienced region splits between 1992 and 2013 (in red) as well as our study sample of region splits (in yellow).

FIGURE 3.4
OVERVIEW MAP: REGION SPLITS & SAMPLE

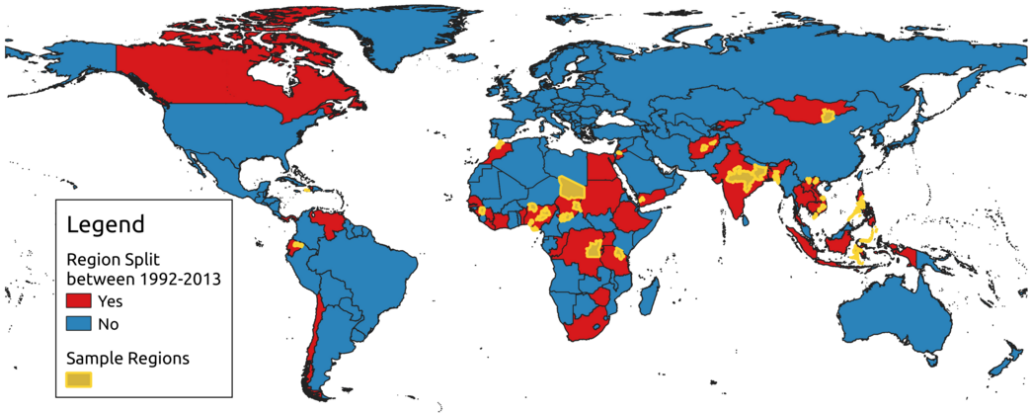

Note: This figure provides and overview of countries that experience region splits between 1992-2013 (in red), as well as our sample of region splits for our event study analysis (in yellow).

3.3 Results

In this section, we conduct the empirical analysis on how distance from the regional capital city impacts economic performance as set out in Section 3.2. We begin with a simpler event study before turning to a more complex estimation approach that combines an event study with BDD. After that, we will conduct a range of robustness test, and subsequently take a closer look at heterogeneities between continents and countries with different institutional settings.

3.3.1 Simple Event Study

In the simple event study analysis, we use all locations in the splinter region[14], but no areas in regions that contain the old regional capital as they experience no shift in the capital city, and track the association between the change in log distance from the regional capital and economic performance before and after the split. Equation 3.1 specifies the corresponding estimation equation. Figure 3.1 plots the event study results using the log of a pixel's nightlight density as dependent variable, the change in log distance from the old vs. the new regional capital city as the treatment intensity and the period prior to the split as the omitted baseline period.

FIGURE 3.1
SIMPLE EVENT STUDY (WITHOUT LOCAL CONTROL GROUP)

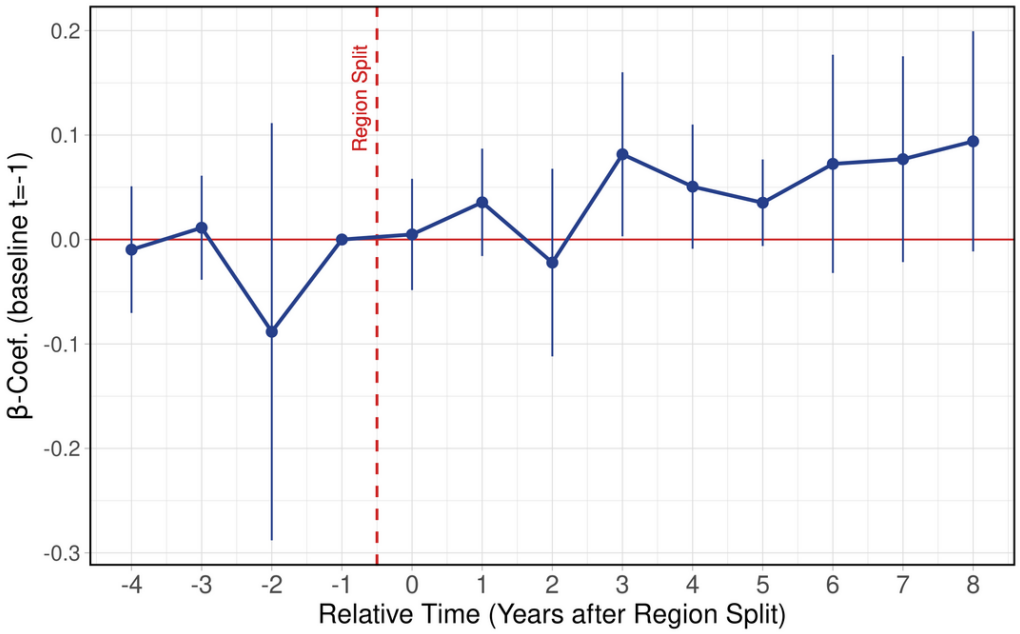

Note: This figure plots the estimated coefficients from our simple event study based on Equation 3.1. The dependent variable is the log of nightlight density in a pixel in a given period, and the β-coefficients correspond to an interaction between the relative period and the change in the log of the distance to the regional capital city. The relative period $t = -1$ represents the period before the region split occurred and serves as the baseline for the estimates. Standard errors are double-clustered at the splinter region-year and at the pixel level. The blue lines around each point estimate represent the 95%-confidence interval.

As can be seen in the figure, prior to the region split, areas with a high as compared to those with a low or even negative treatment intensity (i.e. those areas that ended up further away from the new as compared to the old capital city) evolved similarly. However,

[14]By default, we exclude the area within 20km from the regional capital city to avoid capturing the effect of the split on the regional capital cities themselves e.g. induced by the construction of government buildings. We also exclude the area within 5km from the regional boundary to avoid capturing nightlights from neighboring regions due to blooming.

after the split the curve exhibits a positive trend which means that, on average, areas with a higher treatment intensity gained relative to those with a lower one. The coefficient of 0.09 at t=8 implies that 8 years after the region split, areas that were located one percent closer to the new (as compared to the old capital city) on average experienced an increase in nightlight density of approximately 0.09 percent. The results are similar when using the extensive margin of nightlights instead (see Figure 3.A.1). Yet, we should be cautious about interpreting this estimate as indicating a positive causal impact of proximity to the regional capital city as i) it has a t-value of 1.75 which means that it is not statistically different from 0 at the 95%-confidence interval and ii) most importantly, in this estimation design, we are not accounting for local economic trends that could be confounding the estimates in either direction. As an example, it could be that areas with a lower (or higher) treatment intensity experienced a simultaneous economic upturn e.g. induced by mineral resource discoveries.

3.3.2 Event Study with Boundary Controls

In the next step, we conduct a similar event study but with a more restrictive design that combines an event study design with a boundary discontinuity design. For this reason, we reduce the sample to areas around borders with neighboring regions that are in the same country and did not undergo a boundary reform themselves. Further, we include areas within the boundary area from these neighboring regions as controls for changes in local economic conditions. Based on this design, we control for local economic shocks that are unrelated to the change in distance to regional capital city by partialling out anything that equally affects both sides of each boundary segment in a any given year.

As can be seen in Figure 3.2, the pre-trends with respect to treatment intensity are parallel. After the new capital city was assigned, locations that reduced their distance more as a result of the capital shift gained relative to the others. After 8 years, a one percent decrease in distance from the regional capital city caused an increase in nightlight density by around 0.16 percent. The results are similar to the previous ones without a local control group (see Figure 3.1), but the coefficient 8 years after the split is larger in magnitude and highly significant. This finding might reflect that there were local economic shocks that distorted the simple event study. Figure 3.A.2 presents qualitatively similar results using the extensive margin of nightlights. In this figure, it looks like areas with higher treatment intensity actually lost relatively during the pre-period, which could indicate a violation of the parallel pre-trends assumption. However, firstly, the pre-trends go into the other

Figure 3.2
Event Study with Boundary Controls

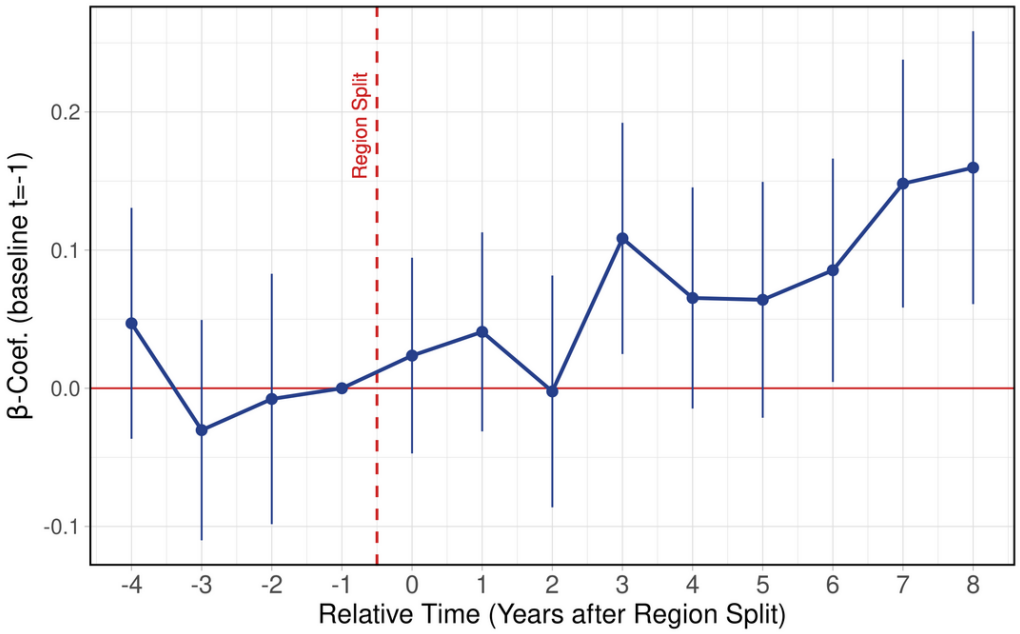

Note: This figure plots the estimated coefficients from our event study with boundary controls based on Equation 3.2. The dependent variable is the log of nightlight density in a pixel in a given period, and the β-coefficients correspond to an interaction between the relative period and the change in the log of the distance to the regional capital city. The relative period $t = -1$ represents the period before the region split occurred and serves as the baseline for the estimates. Standard errors are double-clustered at the segment-year and at the pixel level. The blue lines around each point estimate represent the 95%-confidence interval.

direction which suggest that we would underestimate the true treatment effect. Secondly, as discussed in Section 3.2.3, the extensive margin of nightlights is unable to capture all the variation in nightlights. Considering the fact that there are no pre-trends in Figure 3.2 which is based on the intensive margin that is more reliable in this context, we remain confident about our identification strategy and that our treatment and control group are valid counterfactuals.

Figure 3.A.3 reports the results using the simple event study methodology from Section 3.3.1 and using the boundary area sample in treated regions from Figure 3.2, but without the boundary controls in neighboring regions. This exercise is helpful to understand if the change in the estimates between Figures 3.1 and 3.2 is driven by a change in the sample[15], or by adding a local control group. The results in Figure 3.A.3 indicate that both aspects matter. Focussing on the boundary sample shifts the coefficients closer to those in Figure 3.2, which might reflect that reducing the sample to only boundary areas makes areas with

[15]Figure 3.2 focusses only on the boundary area while Figure 3.1 is based on the whole splinter region, except for the area within 20km from the capital which is excluded by default from all analyses.

different treatment intensities more comparable. However, without the local boundary controls that help with accounting for unrelated economic shocks, the estimates are less precise as is reflected in larger confidence intervals in Figure 3.A.3 as compared to Figure 3.2.

3.3.3 Robustness Tests

In the following, we conduct a range of robustness and sensitivity tests. First, we run difference-in-difference (DiD) regressions based on Equation 3.3 where $Post_t$ is a dummy that is 1 after the region split and 0 before. The model is similar to Equation 3.2 but with just one treatment estimate reflecting the average effect within the first 8 years after the split relative to the period within 4 years before the split.

$$Y_{i,t} = \beta \, \Delta CAP_i * Post_t + b_i + b_{r,t} + b_{s,t} + \xi_{i,t} \tag{3.3}$$

In Table 3.1 column (1), we estimate the default DiD model. In columns (2)-(4), we include a vector of polynomials of distance to the boundary for each boundary side and period separately as additional controls. In a classical cross-sectional boundary discontinuity design, it is vital to include these polynomials of the 'running variable' to account for the fact that by moving away from the discontinuous boundary cutoff, locations on either side become increasingly different from one another. In our case, it is important to account for the fact that local economic conditions and geographical differences become more pronounced the further we move away from the boundary. As an example, the new regional capital city shapes the spatial distribution of economic activity not only due to its role as administrative center for locations within the region, but also by providing market access to all surrounding areas (Donaldson and Hornbeck, 2016). While market access is continuously decreasing with distance, it equally affects locations closely around either side of the regional boundary (as these locations are at the same distance from the regional capital). However, the further we move away from the border, the more locations within the same region tend to gain in terms of market access if the new capital grows, while going further away on the side of the neighboring region decreases such gains. Consequently, by not accounting for these effects, we might obtain significant coefficients even if the new capital city was not relevant as a capital city in the first place. Another advantage of including these polynomials is to control for potential positive (or negative) spillovers from the treatment region to the neighboring region which might result in underestimating (overestimating) the true magnitude of the effects. Given that these spillovers would likely also be a function of distance, they should be absorbed by these polynomials.

As we would expect if market access was in fact a source of bias, the coefficients in column (2)-(4) are slightly smaller as compared to our default model. However, the coefficients are not significantly different from the default estimates and still significantly different from 0 which indicates that our results are not driven by this potential source of bias and underlines the robustness of our results.

TABLE 3.1
EVENT STUDY WITH BOUNDARY CONTROLS: ROBUSTNESS

	Dependent variable: Log Nightlight Density						
	Default	1st Poly (Dist Bound)	2nd Poly (Dist Bound)	3rd Poly (Dist Bound)	Dist Bound < 50km	Dist Cap > 50km	Dist Cap > 100km
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ Dist Reg. Cap	0.075***	0.067***	0.066***	0.066***	0.036**	0.086***	0.116***
$\times$ Post Dummy	(0.022)	(0.023)	(0.022)	(0.023)	(0.017)	(0.032)	(0.045)
	Polynomials for: distance from the border $\times$ boundary side $\times$ period (1,690 groups)						
1st order	-	x	-	-	-	-	-
2nd order	-	-	x	-	-	-	-
3rd order	-	-	-	x	-	-	-
Pixel FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region x Period FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Segment x Period FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	238,654	238,654	238,654	238,654	449,449	208,650	146,848
Adjusted R^2	0.91	0.91	0.91	0.91	0.91	0.92	0.92

Note: This table reports DiD estimations with boundary controls based on Equation 3.3. In order to avoid capturing the effect on the new capital cities themselves, we exclude 20 km around each regional capital from our sample. To prevent misassigning detected nightlights from a neighboring region due to blooming, we exclude 5 km on each side of the boundary. Boundary segments corresponds to a buffer of 25 km (in column (5) 50 km) around border pieces of 100 km line length. Standard errors in parenthesis are double-clustered at the segment-year and at the pixel level.
*p<0.1; **p<0.05; ***p<0.01

Next, we test the robustness of our results regarding changing the boundary buffer from 25km to 50km. The estimate in column (5) is a bit smaller which could be due to opposing sides being less comparable. Yet, the fact that it is still highly significantly positive indicates that our results are robust regarding this modification. In column (6) and (7), we exclude a larger area around the capital cities to make sure that our results are not driven by direct spillovers from cities on their immediate surrounding, but rather relevant for larger areas in the hinterlands. When removing 50km or 100km (rather 20km by default) around the capital cities, the estimated coefficients are still highly significant and even a bit larger in magnitude. However, we interpret this increase in the magnitude to be more likely a result of the shift in the sample and underlying characteristics as some smaller regions drop out of the sample - rather than reflecting that the effects get stronger in areas further away from the regional capitals.

As an additional robustness test, to make sure that our results are not driven by an individual segment or boundary, we iteratively exclude each boundary and re-estimate the coefficient. The results in Figure 3.A.4 show that the 95% confidence interval of all estimates largely overlaps with the aggregate estimate and that all estimated coefficients are significantly different from 0 which supports the robustness of our results.

3.3.4 Heterogeneity Tests

In Figure 3.A.5, we estimate continent specific effects and find that there are large differences in the effect of proximity to regional capitals on economic performance. The average effects are strongest in the Americas which is followed by Asia. In contrast, the effect is very close to zero and insignificant in Africa. This circumstance might reflect that there are important region or country characteristics that determine the impact of proximity to regional capital cities.

A potential reason for the large difference in the effects between countries could be related to political framework conditions. Firstly, decentralization might be important in this context. If regional governments have a low level of political decision making power, then regional capitals should be less relevant. Furthermore, African countries, for which the location of regional capitals appear to be less relevant (Figure 3.A.5), are among the least decentralized countries in the world (Ivanyna and Shah, 2014).

To estimate the relevance of political frameworks empirically, we run heterogeneity regressions based on Equation 3.3 and report the results in Table 3.2. Column (1) in Table 3.2 reveals that there are stark differences between centralized and decentralized countries[16], with proximity to regional capitals being significant only in the latter.
Another country characteristic that could be relevant in this context is democracy vs. autocracy[17] It might be that in autocracies region splits are more likely to be used as a political tool to weaken or reduce the bargaining power of lower tiers of government (Green, 2010; Grossman and Lewis, 2014; Billing, 2019) or win elections (Hassan, 2016). In such countries, there might not actually be a political will to create functioning new regions, which might undermine any potential gains in proximity to the regional capital city. The estimate in column (2) is in line with this suspicion by indicating that proximity to capital

[16] We use the decentralization index by (Ivanyna and Shah, 2014) and split countries into either centralized or decentralized depending on whether they are below or above the median score in our sample.
[17] We classify countries as democratic if their average polity2 score during our study period between 1992 and 2013 exceeds 0 (Marshall et al., 2017).

cities is only significantly positive in democracies, but very close to 0 and insignificant in autocracies.

Finally, in column (3), we find that capital cities appear to be more important in poor rather than rich countries.[18] This finding is similar to remoteness from national capital cities that also entails stronger effects in less developed countries (see Table 2.A.8 in Chapter 2). A potential explanation for this results could be related to distance being less relevant in more developed countries.

However, it is important to bear in mind that all of these estimates are based on different groupings of countries that share many other characteristics that might actually be driving the effects. For this reason, the results should be interpreted with caution. One way of trying to get more clarity about the relevance of these characteristics is to simultaneously estimate different group specific effects - but keeping in mind that there are always omitted relevant characteristics. For all paired combinations in column (4)-(6), decentralization is always significant and the other two once respectively. Finally, when including all three groupings at once, all coefficients remain significant. We interpret this as underlining that all three characteristics are relevant with regard to the effects.

TABLE 3.2
EVENT STUDY WITH BOUNDARY CONTROLS: HETEROGENEITY

	Dependent variable: Log Nightlight Density						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ Dist Reg. Cap × Post Dum.	−0.00	0.03	0.08***	−0.03	−0.00	0.03	−0.04**
	(0.01)	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)
Δ Dist Reg. Cap × Post Dum. × Decentralized	0.13***	-	-	0.11***	0.17***	-	0.15***
	(0.04)			(0.03)	(0.04)		(0.04)
Δ Dist Reg. Cap × Post Dum. × Democracy	-	0.08**	-	0.05	-	0.12**	0.10**
		(0.04)		(0.04)		(0.05)	(0.04)
Δ Dist Reg. Cap × Post Dum. × Developed	-	-	−0.05	-	−0.13*	−0.10	−0.17**
			(0.07)		(0.08)	(0.07)	(0.08)
Pixel FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region x Period FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Segment x Period FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Num. obs.	238,654	238,654	238,654	238,654	238,654	238,654	238,654
Adj. R^2	0.91	0.91	0.91	0.91	0.91	0.91	0.91

Note: This table reports DiD heterogeneity estimations with boundary controls based on Equation 3.3. In order to avoid capturing the effect on the new capital cities themselves, we exclude 20 km around each regional capital from our sample. To prevent misassigning detected nightlights from a neighboring region due to blooming, we exclude 5 km on each side of the boundary. Boundary segments corresponds to a buffer of 25 km around border pieces of 100 km line length. Standard errors in parenthesis are double-clustered at the segment-year and at the pixel level.

*p<0.1; **p<0.05; ***p<0.01

[18] We classify countries as poor or rich based on whether they exceed the average GDP per capital between 1992 and 2013 based on World Bank data.

Overall, these results indicate that proximity to regional capital cities might not always be advantageous for locations, but only boosts the local economy when the country is decentralized and democratic. Moreover, the spatial disparities associated with distance from the regional capital city appear to be less relevant in countries with a relatively high level of GDP per capita. Yet, as mentioned above, these group-specific estimates are rather descriptive than causal and should be interpreted with caution.

3.3.5 Regional Capitals and Conflict

Previous research found that region splits decrease conflict (Cederman et al., 2015; Pierskalla, 2016; Baskaran and Blesse, 2019). It could be that the pacifying nature of region splits actually stem from the presence of regional capital cities in closer proximity. In the following, we therefore explore if proximity to the regional capital city has a negative impact on the occurrence of conflict using our event study BDD estimation approach.

To obtain information on conflict occurrence, we extend our dataset with the Uppsala Conflict Data Program (UCDP) (Sundberg and Melander, 2013). This database features information about the location of conflicts across the world starting in the year 1989. The UCDP has a rather broad definition of conflict which it defines as: 'An incident where armed force was used by an organised actor against another organized actor, or against civilians, resulting in at least 1 direct death at a specific location and a specific date'.[19] We spatially match the UCDP with our database based on whether the geolocation of conflict events falls within a segment side, which will be the unit of observation in this analysis.[20] We create a dummy that equals 1 if a segment side has experienced conflict in a given year and is 0 otherwise.[21]

Figure 3.3 plots the results of the event study BDD estimation using conflict occurrence as dependent variable. As can be seen in the figure, there is no evidence that proximity to regional capital cities has an impact on the probability of conflict in a segment side. Similarly, when splitting the sample by decentralization, democracy or national GDP per capita, there are no significant changes in conflict after the region split (see Figure 3.A.6). It should be noted that conflict might spill over across the border to the neighboring region

[19] Visit: `https://ucdp.uu.se/downloads/ged/ged211.pdf` for more information about the UCDP database.

[20] To capture a larger number of conflict events in the local area, we increase the boundary buffer to 50km. But the results are equivalent to using the default boundary buffer of 25km. It should be noted that Mongolia, Vietnam and Yemen are not part of this analysis as we do not record any conflict events in the boundary segments in these countries.

[21] As an alternative we measure the conflict frequency in a segment side and obtain similar results.

which would bias our coefficients towards zero. Nevertheless, this potential caveat does not change the fact that we find a relative imbalance in economic performance across the boundary but no imbalance in conflict. This circumstance makes it unlikely that conflict is a key mechanism behind the observed patterns.

FIGURE 3.3
REGIONAL CAPITALS AND CONFLICT

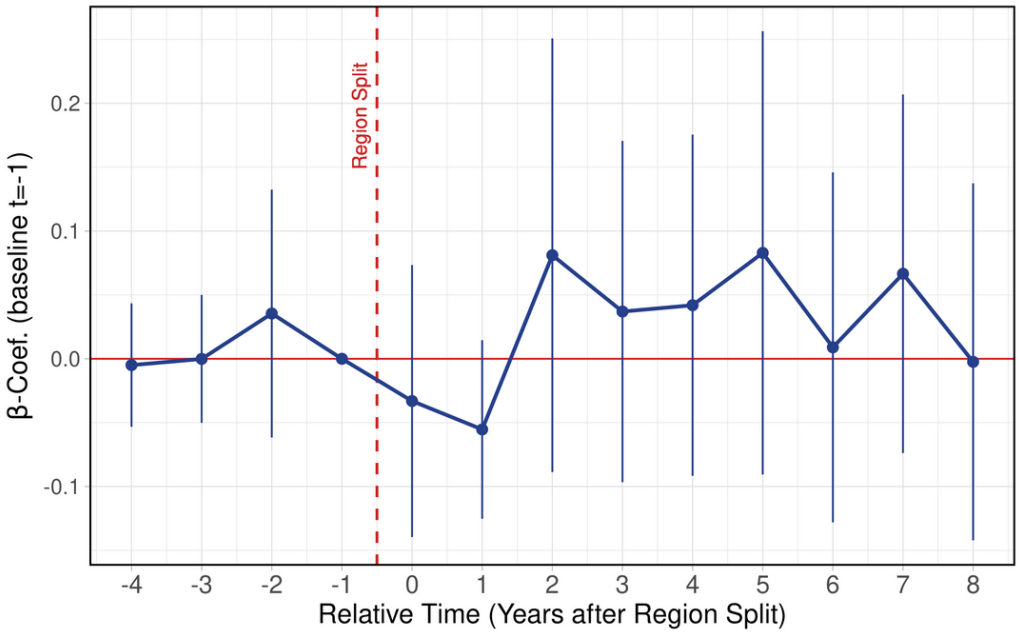

Note: This figure plots the estimated coefficients from our event study with boundary controls based on Equation 3.2. The dependent variable is a dummy indicating if a segment side experience conflict in a given period, and the β-coefficients correspond to an interaction between the relative period and the change in the log of the distance to the regional capital city. The relative period $t = -1$ represents the period before the region split occurred and serves as the baseline for the estimates. Standard errors are clustered by boundary. The blue lines around each point estimate represent the 95%-confidence interval.

3.4 Conclusion

This paper studies the impact of proximity to regional capital cities on local economic development. To obtain quasi-random variation in distance from the regional capital, we examine a worldwide sample of region splits that induce a shift in the location of the assigned regional capital city. Combining an event study with a BDD, we find that over a period of 8 years a decrease in distance from the regional capital city by 1 percent, on average increases nightlight density by 1.6 percent. Yet, there are stark differences between countries. Most importantly, the effects are primarily relevant in countries that are decentralized and feature democratic institutions. In contrast, the effects seem to be weaker

in relatively developed countries.

This paper contributes to at least two main streams of the literature. Firstly, by showing that regional capital cities have a similar effect on the spatial distribution of economic activity as national capitals, it contributes to the literature on the role of the location of capital cities (Campante and Do, 2014; Michalopoulos and Papaioannou, 2014; Campante et al., 2019) (see also Chapter 2). Secondly, it contributes to the literature on decentralization and administrative region proliferation by studying the impact of assigning new capital cities under different institutional settings.

Our findings highlight that regional capital cities are important determinants of the distribution of economic activity within a region. However, as the positive effects of proximity to regional capitals do not unfold in all settings, policy makers need to take the administrative organization and institutional context into consideration when considering assigning new capital cities. More specifically, creating new regional capitals appear ineffective when countries are too centralized, developed or autocratic. More research is required to study the precise mechanisms through which these framework conditions impact the effects. One possibility is that these contextual characteristics determine the political intentions behind the reforms, which might play a critical role for their outcomes (Green, 2010; Grossman and Lewis, 2014; Hassan, 2016). Similar to national capital cities, we have shown that conflict is unlikely to be a mediating factor. However, whether information frictions and accountability mechanisms play a similar role for regional as for national capitals is yet to be confirmed (see Chapter 2). Future research can build upon the methodology developed in this study and extend this research by gathering additional information about public goods provision, the performance of public officials and their perception and support from local communities over time. While there are still open questions, with this research we have shown that assigning regional capital cities strategically has the potential to promote economic development in remote hinterlands.

3.A Appendix

3.A.1 Descriptives

Table 3.A.1
Overview: Study Sample of Region Splits

	Split	Country	Original Region	New Region	Old Capital	New Capital
1	1996	Guinea	Labe	Mamou	Labe	Mamou
2	1996	Jordan	Amman	Madaba	Amman	Madaba
3	1996	Mongolia	Dornogovi	Govisu'mber	Sainshand	Choir
4	1997	DR Congo	Kivu	Maniema	Bukavu	Kindu
5	1997	DR Congo	Kivu	Nord-Kivu	Bukavu	Goma
6	1997	Morocco	Nord Ouest	Gharb-Chrarda-Beni Hssen	Rabat	Kenitra
7	1997	Nigeria	Rivers	Bayelsa	Port Harcourt	Yenagoa
8	1997	Nigeria	Ondo	Ekiti	Akure	Ado-Ekiti
9	1997	Nigeria	Bauchi	Gombe	Bauchi	Gombe
10	1997	Nigeria	Plateau	Nassarawa	Jos	Lafia
11	1997	Nigeria	Sokoto	Zamfara	Sokoto	Gusau
12	1997	Vietnam	Bac Thai	Bac Kan	Thai Nguyen	Bac Kan
13	1997	Vietnam	Song Be	Binh Phuoc	Thu Dau Mot City	Dong Xoai
14	1997	Vietnam	Minh Hai	Ca Mau	Bac Lieu	Ca Mau
15	1997	Vietnam	Vinh Phu	Phu Tho	Vinh Yen	Viet Tri
16	1997	Vietnam	Quang Nam-Da Nang	Quang Nam	Da Nang City	Tam Ky
17	1998	Bangladesh	Chittagong	Sylhet	Chittagong	Sylhet
18	1998	Ecuador	Napo	Orellana	Tena	P. Franc. Orellana
19	2000	Chad	Ouaddai	Assongha	Abeche	Adre
20	2000	Chad	Bourkou Ennedi Tibesti	Ennedi	Faya-Largeau	Fada
21	2000	Chad	Moyen Chari	Lac Iro	Sarh	Kyabe
22	2000	Chad	Lougoume Oriental	Mont De Lam	Doba	Baibokoum
23	2000	Chad	Ouaddai	Sila	Abeche	Goz Beida
24	2000	Chad	Tandjile	Tandjile Ouest	Lai	Kelo
25	2001	India	Madhya Pradesh	Chhattisgarh	Bhopal	Raipur
26	2001	India	Bihar	Jharkhand	Patna	Ranchi
27	2001	Indonesia	Sulawesi Utara	Gorontalo	Manado	Gorontalo
28	2002	Philippines	Southern Tagalog	Calabarzon	Quezon City	Calamba
29	2002	Tanzania	Arusha	Manyara	Arusha	Babati
30	2003	Haiti	Grande Anse	Nippes	Jeremie	Miragoane
31	2004	Afghanistan	Uruzgan	Daykundi	Tarin Kut	Nili
32	2004	Afghanistan	Parwan	Panjsher	Charikar	Bazarak
33	2004	Vietnam	Dak Lak	Dak Nong	Buon Ma Thuot	Gia Nghia
34	2004	Vietnam	Lai Chau	Dien Bien	Lai Chau	Dien Bien Phu
35	2004	Vietnam	Can Tho	Hau Giang	Can Tho	Vi Thanh
36	2004	Yemen	Sana'a	Raymah	Sanaa	Al Jabin
37	2005	Indonesia	Sulawesi Selatan	Sulawesi Barat	Makassar	Mamuju

Note: This table provides an overview of the splinter regions in our event study sample.

Table 3.A.2

Summary Statistics

	All Splits		Sample Splits	
	mean	sd	mean	sd
Nightlights 1992	0.47	3.15	0.61	3.56
Area 1992 (in km^2)	263,556	223,054	271,224	213,683
Share Africa	0.72	-	0.54	-
Share Asia	0.26	-	0.45	-
Share Americas	0.03	-	0.01	-

Note: This table reports the mean and the standard deviation of key indicators for the all areas that underwent a split between 1992 and 2013, as well as our sample of splits respectively. Please note that the region 'Northwest Territories' in Canada was omitted from the computation of statistics for 'All Splits'.

3.A.2 Additional Results

FIGURE 3.A.1
SIMPLE EVENT STUDY: EXTENSIVE MARGIN

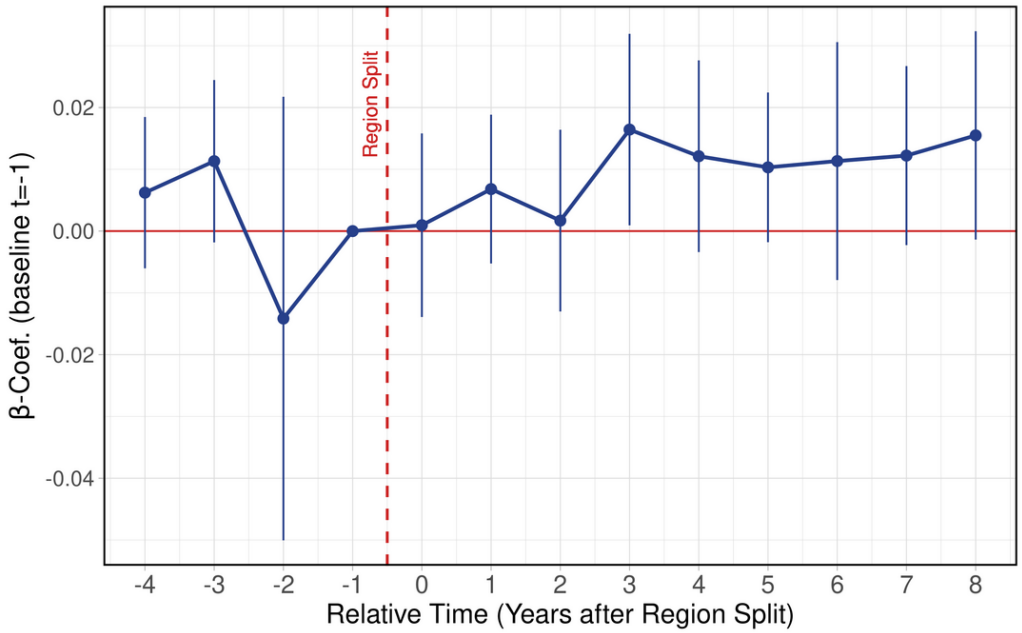

Note: This figure corresponds to Figure 3.1 using the extensive rather than the intensive margin of nightlights as dependent variable. The figure plots the estimated coefficients from our simple event study based on Equation 3.1. The dependent variable is a dummy variable indicating whether a pixel emits nightlights in a given period, and the β-coefficients correspond to an interaction between the relative period and the change in the log of the distance to the regional capital city. The relative period $t = -1$ represents the period before the region split occurred and serves as the baseline for the estimates. Standard errors are double-clustered at the splinter region-year and at the pixel level. The blue lines around each point estimate represent the 95%-confidence interval.

Figure 3.A.2

Event Study with Boundary Controls: Extensive Margin

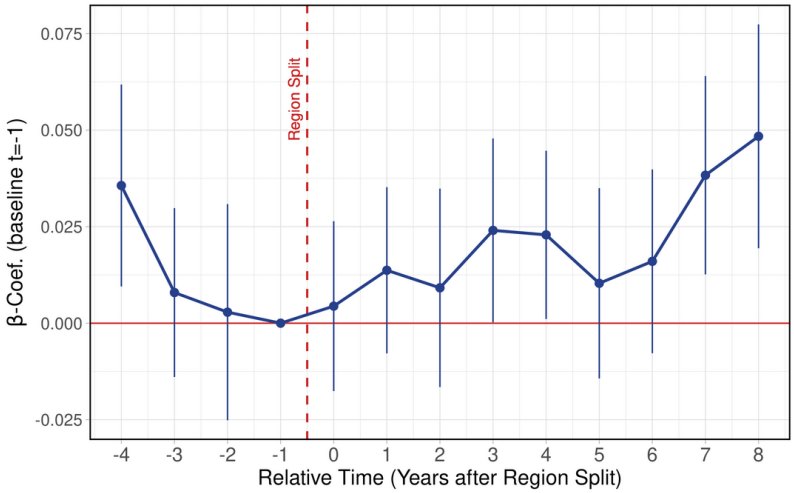

Note: This figure corresponds to Figure 3.2 using the extensive rather than the intensive margin of nightlights as dependent variable. The figure plots the estimated coefficients from our event study with boundary controls based on Equation 3.2. The β-coefficients correspond to an interaction between the relative period and the change in the log of the distance to the regional capital city. The relative period $t = -1$ represents the period before the region split occurred and serves as the baseline for the estimates. Standard errors are double-clustered at the segment-year and at the pixel level. The blue lines around each point estimate represent the 95%-confidence interval.

Figure 3.A.3

Event Study: Boundary Sample without Boundary Controls

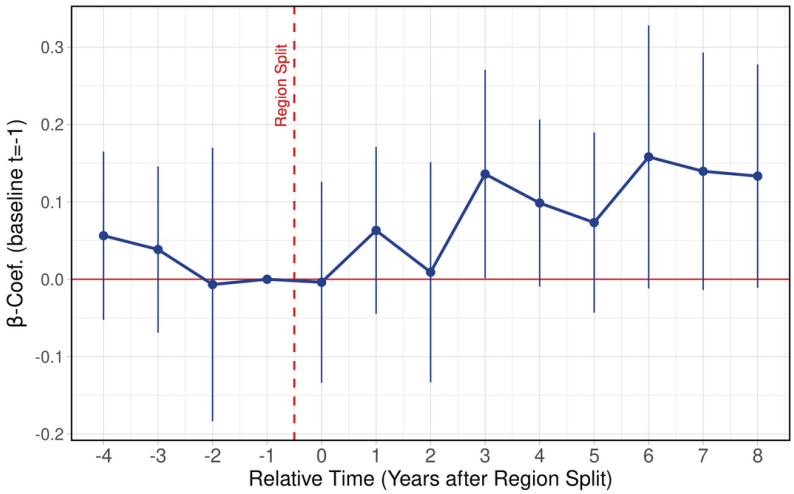

Note: This figure plots the estimated coefficients from our simple event study without boundary controls based on Equation 3.1, using the sample of boundary areas in splinter regions that serve as treatment group in Figure 3.2. The dependent variable is the log of nightlight density in a pixel in a given period, and the β-coefficients correspond to an interaction between the relative period and the change in the log of the distance to the regional capital city. The relative period $t = -1$ represents the period before the region split occurred and serves as the baseline for the estimates. Standard errors are double-clustered at the splinter region-year and at the pixel level. The blue lines around each point estimate represent the 95%-confidence interval.

Figure 3.A.4
Drop Boundaries Iteratively

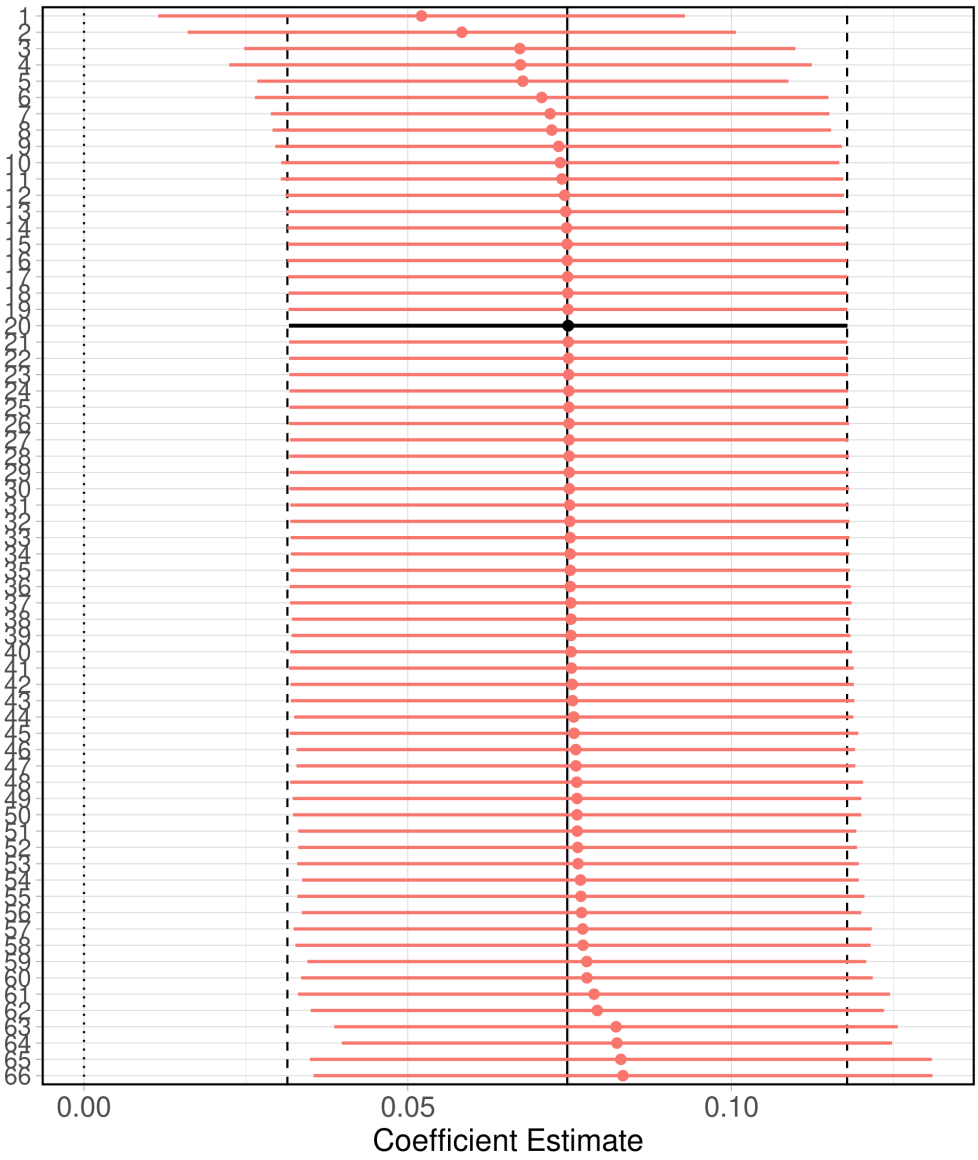

Note: This figure plots the estimated coefficients and 95%-confidence interval corresponding to Table 3.1 column (1) when iteratively dropping one boundary. The aggregate estimate based on the whole sample is number 20 (in black). The red lines around each point estimate represent the 95%-confidence interval. Standard errors are double-clustered at the segment-year and at the pixel level.

Figure 3.A.5

Heterogeneity: Continent

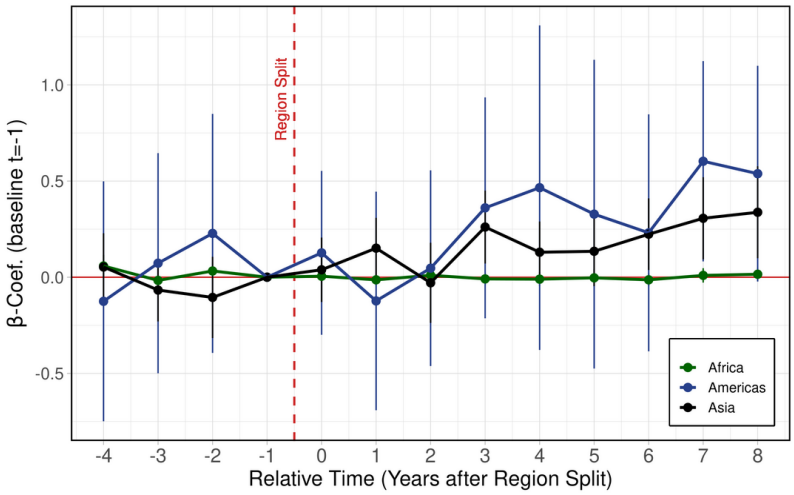

Note: This figure plots the estimated coefficients from our continent specific heterogeneity event study with boundary controls based on Equation 3.2 and corresponding to Figure 3.2. The dependent variable is the log of nightlight density in a pixel in a given period, and the β-coefficients correspond to an interaction between continent dummies, the relative period and the change in the log of the distance to the regional capital city. The relative period $t = -1$ represents the period before the region split occurred and serves as the baseline for the estimates. Standard errors are double-clustered at the segment-year and at the pixel level. The blue lines around each point estimate represent the 95%-confidence interval.

Figure 3.A.6
Conflict Heterogeneity
(a) Decentralized vs. Centralized

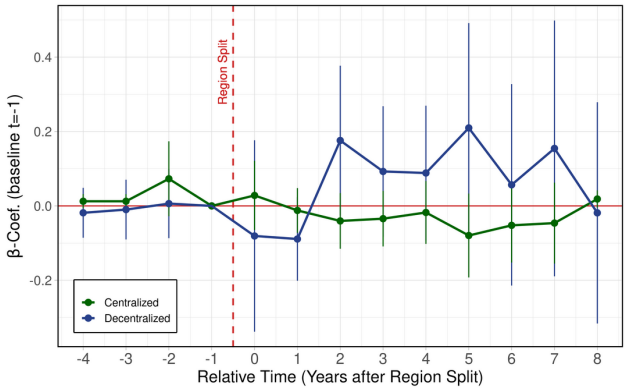

(b) Democracy vs. Autocracy

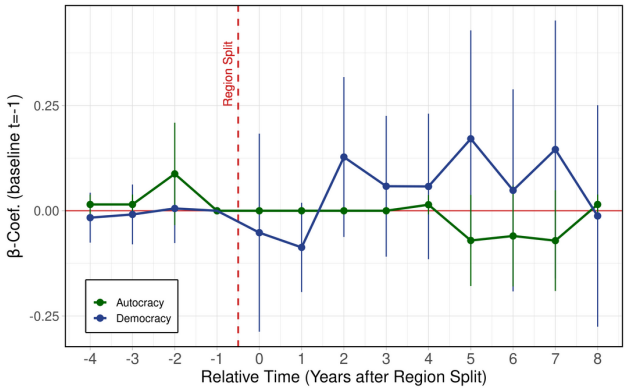

(c) High vs. Low GDP

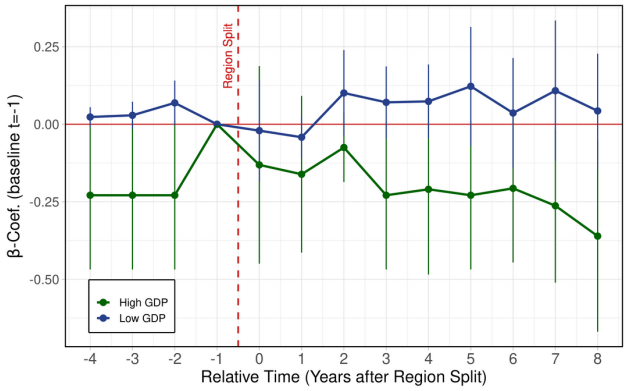

Note: This figure plots the estimated coefficients from our conflict heterogeneity event studies with boundary controls based on Equation 3.2 and corresponding to Figure 3.3. The dependent variable is a dummy indicating if a segment side experience conflict in a given period, and the β-coefficients correspond to an interaction between continent dummies, the relative period and the change in the log of the distance to the regional capital city. The relative period $t = -1$ represents the period before the region split occurred and serves as the baseline for the estimates. Standard errors are clustered by boundary. The blue lines around each point estimate represent the 95%-confidence interval.

Chapter 4

Urban Density and COVID-19: Understanding the US Experience

4.1 Introduction

Historically, cities have been associated with the propagation of infectious diseases.[1] It is therefore not surprising that the impact of density – the defining feature of cities – on the spread of COVID-19 was a frequent talking point from the very outset of the COVID-19 pandemic. As early as 22nd of March 2020, in the context of a critical outbreak in New York City, state governor Andrew Cuomo tweeted "There is a density level in NYC that is destructive. It has to stop and it has to stop now. NYC must develop an immediate plan to reduce density".[2]

The notion that dense cities would be hotbeds of virus transmission prompted a flurry of academic research on the topic. Initial empirical work – especially that looking at the United States' experience – suggested urban density fostered a faster spread of the disease.[3] Similar evidence was reported for other countries including India (Bhadra et al., 2021), Brazil (Pequeno et al., 2020) and Germany (Ehlert, 2021). However, subsequent research exploring a longer time series yielded mixed findings (see for example McFarlane 2021; Kim et al. 2021; Florida et al. 2021). This prompted a more nuanced approach to the question, and subsequent work on the roles of crowding, experienced density and other

[1] See Duranton and Puga (2020); Voigtländer and Voth (2013) for treatments of this relationship in economics.

[2] The attribution of detrimental effects of density for the evolution of the epidemic was not specific to the United States. On December the 9th 2020 Michael Gove (Chancellor of the Duchy of Lancaster and Minister for the UK Cabinet Office) said on ITV's Good Morning Britain that population density is one of the reasons why the UK has more COVID-19 related deaths in comparison to Germany.

[3] See for example Angel et al. (2020); Whittle and Diaz-Artiles (2020); Zhang and Schwartz (2020); Wheaton and Kinsella Thompson (2020) and Almagro and Orane-Hutchinson (2020). For a review of the empirical literature on the topic – covering papers in urban planning, economics and medical sciences – see Teller (2021).

more direct measures of social interactions.[4]

Now that massive vaccination campaigns have gradually reduced the threat of COVID-19 worldwide, can we draw any definitive conclusions about the mediating role of density in shaping the health impact of COVID-19 in cities? We turn to this question by looking at the evolution of the epidemic in the contiguous United States, in the period between the first registered cases in January 2020 and the beginning of the vaccination campaign in mid-December. By looking at the whole of 2020, we seek to understand how the results of initial studies indicating density was an important determinant of the impact of COVID-19 progressively led to more ambiguous findings as the pandemic evolved. Our empirical analysis combines descriptive evidence with an instrumental variable strategy borrowed from the agglomeration literature in economics. In doing so, our methodological approach avoids some of the pitfalls of conventional regression estimates and is close to methods that are familiar to both economists and economic geographers.

We find convincing evidence that density affected the timing of the outbreak in each county, with denser locations more likely to have an early outbreak. We show this leads to an initially positive and significant relationship between the impact of COVID-19 and population density at the county level, consistent with the results of early studies on the spread of the virus in the United States. However, after adjusting for the timing of the onset of the disease in each county, we find no evidence that population density is positively associated with the impact of COVID-19. Interestingly, we find a negative relationship between density and the spread of COVID-19 within a county at the very beginning of an outbreak, but this relationship fades completely within 2 months. We also show that, by the end of 2020, density could no longer explain the cross-sectional pattern of accumulated cases or deaths. Dense locations were hit first, but, as the pandemic evolved, they were not hit harder. Combined, these results help us frame other studies on this topic, and understand how the results in that literature changed as the pandemic developed.

The fact that – by the end of 2020 – density had no effect on the local impact of COVID-19 appears counter-intuitive. The virus spreads via human contact and denser areas provide more opportunities for human interaction. Yet, this is not the only way in which density can affect the spread of disease. Several mediating factors can make the direction of

[4]For example, there is evidence that a higher percentage of overcrowded households and poor housing conditions in US counties have both lead to higher mortality from COVID-19 (Ahmad et al. 2020; Krieger et al. 2020; Kamis et al. 2021).

this relationship theoretically ambiguous. We analyze social/behavioral factors that could explain our findings, bearing in mind that the spread of disease is a social as well as a biological phenomenon (Papageorge et al., 2020). To do so we use data from Google, Facebook, the US Census and The County Health Rankings and Roadmaps program. First, we show that density is positively associated with the reduction in work and leisure related activities throughout the pandemic, suggesting that compliance with social distancing measures was higher in denser locations. Second, we use our empirical strategy to illustrate the well-known fact that density is negatively associated with the share of Republican voters, which have been shown to be less engaged in social distancing and other efforts to reduce transmission (Allcott et al., 2020). Third, we show population density is positively associated with access to healthcare and income and negatively associated with inhabitants' age. Collectively, these results yield suggestive evidence of mechanisms generating offsetting negative effects of density on the spread and severity of the COVID-19 outbreak, and help us rationalize the estimates of the overall effects reported in our main analysis.

Estimating how population density shaped the spread and severity of the COVID-19 outbreak, as well as its effects on local behavioral responses and demographics is challenging for several reasons. First, population densities are not randomly assigned and they might be correlated with unobserved confounding factors. For example, population densities can be affected by locational productive advantages, whether natural or man-made (e.g. soil quality or transportation infrastructure), that may also simultaneously affect local economic conditions. Insofar as the COVID-19 outbreak is affected by economic factors, unobservable locational advantages can confound the effect of density on the spread and severity of the disease. Second, differences in the timing of the onset of the disease can generate cross-sectional differences in the severity of the outbreak at one point in time in the absence of true differences in the local reproduction rate. Finally, data on COVID-19 cases might be reported with error due to variation in local testing strategy and capacity.

We overcome the empirical challenges mentioned above in several ways. We use two Instrumental Variable (IV) strategies borrowed from the agglomeration literature in economics to induce plausibly exogenous variation in population density without affecting COVID-19 cases and deaths directly. More specifically, in our *geological* IV approach, we use the presence of aquifers, earthquake risk, and soil drainage capacity to as instruments for density (as in Duranton and Turner 2018). In our *historical* IV strategy, we use the traditional long-lag instrument, which measures urban population density in the 1880 US Census (as

in Ciccone and Hall 1996 and a large subsequent literature). We use these tools to study both how density affected the timing of the outbreak in each county and the time-adjusted number of deaths after that outbreak. We focus on the daily number of confirmed COVID-19 *deaths* rather than cases as our main outcome of interest since this is considered to be a more accurate indicator of local COVID-19 prevalence (Subbaraman 2020), and discuss COVID-reported cases as a robustness check. Finally, we cross-validate our COVID-19 figures with official data from the CDC to ensure reported deaths are consistent with other measures of COVID-19 mortality.

As discussed above, a number of papers have examined the link between density and COVID-19 incidence in the United States.[5] Alongside these studies, a vast number of papers in economics and economic geography have focused on other social determinants of differences in the spread of COVID-19 such as mobility Glaeser et al. (2020); Almagro et al. (2020), racial composition Benitez et al. (2020); Hamman (2021), social capital and institutions Ding et al. (2020); Rodríguez-Pose and Burlina (2021) as well as on the predicted long-run impact of the pandemic on cities (Florida et al., 2021; Nathan and Overman, 2020). We contribute to this literature by looking specifically at density – arguably one of the first explanatory factors that attracted the attention of the field in early 2020 – and its changing role throughout the US epidemic. Given that density is associated with many of the factors that were studied subsequently – mobility, race, urbanization – our findings also help interpret the results reported in the broader literature.

4.2 Data

Our dataset combines information on COVID-19 cases and deaths, population density, demographics, social connectedness, behavioral changes, voting behavior, healthcare provision, income and geological features at the US county level. We will use COVID data extending over the period between the 22nd of January, when the first US case was confirmed in King County, up until the 15th of December 2020, the day after the COVID vac-

[5]The literature on the relationship between the 1918 Influenza pandemic (the Spanish Flu) and population density is naturally more developed and can shed light on the link between pandemics and density more broadly. Interestingly, while it may seem intuitive that the influenza pandemic was positively associated with population density as the virus spread via human contact, a review of the literature produce mixed results. For example, Garrett (2007) finds a positive relationship between mortality rates and population density in the US. In contrast, Mills et al. (2004) find no statistical association between population density and the initial reproductive number (R) using data on 45 US cities. Chowell et al. (2008) also find no association between transmissibility, death rates and indicators of population density in England and Wales. Ferguson et al. (2006) studies the development of the 1918 pandemic and finds evidence for an early onset in dense urban cores before a more smooth development of the disease across space.

cination campaign began in the United States. We restrict our sample to urban counties[6] in the contiguous United States which leaves us with 1,759 counties comprising $\sim 93\%$ of the total US population. When analyzing the pace of the outbreak, we further restrict the sample further to those counties that had at least one confirmed COVID-19 related death 60 days before the end of our sample period. This *Outbreak Sub-sample* consists of 1,441 counties representing ~ 89 % of the total US population (see Figure 4.A.1). In the following, we describe the dataset and provide further information about the sources and URLs for download in Appendix 4.A.2 and descriptive statistics in Table 4.A.1.

COVID-19 Cases and Deaths

We obtain a panel of daily confirmed COVID-19 fatalities and cases for US counties from usafacts.org.[7] The most intuitive indicator to monitor the COVID-19 outbreak is the daily number of confirmed cases. However, this figure is likely to be distorted by varying local testing strategy and capacity. Furthermore, the ability of the virus to spread across asymptomatic people makes the task of recording the number of infections in the community extremely difficult (Subbaraman, 2020). Therefore, we mainly use the daily number of confirmed COVID-19 deaths as this is a more accurate indicator of the local COVID-19 prevalence.[8] In order to ensure that our COVID-19 data is reliable, we cross-validate our COVID-19 figures with official data from the Centers for Disease Control and Prevention (CDC). In the left panel of Figure 4.A.2, we compare our total COVID-19 fatality counts by county to the latest figures on officially confirmed deaths due to COVID-19. In the right panel, we compare total fatalities to CDC excess death estimates. Both graphs exhibit strong linear relationships and support the validity of our COVID-19 data.[9] The evolution of daily COVID-19 fatality numbers used in this paper is illustrated in Appendix Figure 4.A.3. In our analysis below, when we refer to deaths taking place in the *first-wave*, we refer to those taking place up to the 5th of July, which is the minimum in the moving average of deaths after April 2020.

[6] Urban counties are those that are classified as either 'metropolitan' or 'micropolitan' core-based statistical areas in the 2010 census.

[7] These are obtained from county-level reports by local health authorities across the United States. See Appendix 4.A.2 for further details.

[8] Recent work led by Diego Puga looks at the relationship between density and COVID-19 incidence in Spain using prevalence data obtained from randomized serological tests. Cross-sectional correlations using this information point to a flat (or weakly negative) relationship between the disease's spread and density.

[9] In contrast, the correlation between county level COVID-19 fatalities and USAFacts is -0.001 and insignificant indicating that COVID-19 mortality is not simply an amplification of fatalities occurring under normal circumstances but rather follows distinct patterns that are consistently capture by our database.

Population Density

Based on the US census for 2010, we compute two measures of population density. The first is simply the total population of a county over its total area. This will constitute the independent variable of interest throughout most of our analysis. The second variable takes the population density for all census-blocks within a county and computes the associated population-weighted mean. Population-weighted density is meant to measure average "experienced" density and was popularized in economics by Glaeser and Kahn (2004) and Rappaport (2008). It can be computed using spatially disaggregated data on the distribution of population and weighting each small unit of population density by its relative population in the county.

Instrumental Variables:

For our geological instrumental variable estimates we use three different instruments. More specifically, we use variables measuring earthquake risks and presence of aquifers from the United States Geological Survey (USGS) (also used in Duranton and Turner 2018), and data on soil drainage quality from NRCS State Soil Geographic Data Base. We match our grid cells to the geological data using grid cell centroids to spatially impute data on aquifers, earthquake risks and soil drainage quality. For our historical instrument, we use population density obtained from the 1880 United States census. We impute this data on the county level using spatial matching based on the assumption of uniform population distribution within 1880 counties.[10]

Behavioral Adjustment/Social Distancing:

To measure how much people in different counties adjusted their behavior as a response to the COVID-19 outbreak we use the 'COVID-19 Community Mobility Reports' by Google (Google CMR). This database aggregates extensive anonymized mobile device GPS user data and estimates the percentage change in activities (such as work, retail or transit) by county and day. The five week period from January 3rd to February 6th before the start of the COVID-19 outbreak in the US serves as the corresponding baseline period.

Other Variables:

We obtain data on county-level demographic characteristic estimates for 2018 from the

[10]Note that, while the assumption of uniform distribution is clearly a simplification which could lead to measurement error, this should not have a substantial impact on our main estimates. This is because measurement error in the instruments could affect the relevance of the instruments but should not generate bias in the coefficients of interest unless the measurement error itself is correlated with COVID-19 incidence.

US census. Social connectedness is measured with Facebook's Social Connectedness Index (Facebook SCI), which captures the intensity of the link between locations using the number of friend links in this social network (see Bailey et al. 2018 for further details on the SCI). Finally, data on access to healthcare and income comes from the County Health Rankings and Roadmaps program. Specifically, we use three indicators: (1) the ratio of population to primary care physicians (2) the percentage of adults under the age of 65 without health insurance and (3) median household income.

4.3 Empirical Analysis

Our empirical analysis proceeds in two ways. We first provide a series of figures that illustrate the main results, both in terms of the relationship between density and COVID-19 deaths, the evolution of that relationship over time and the explanations behind this evolution. We then provide formal quantitative estimates for these relationships using our OLS and IV strategies. The fact that by-and-large the quantitative findings are the same regardless of the methods employed in the analysis gives us confidence on the robustness of our results to methodological decisions made in the research process.

4.3.1 Graphical Evidence

The top-left panel of Figure 4.1 illustrates the positive cross-sectional correlation between a county's population density - calculated as the total population over the surface area - and the number of COVID-19 related deaths per capita by the end of the first wave on the 5th of July.[11] This is the basic fact that had been noticed in Wheaton and Kinsella Thompson (2020) and Dubner (2020) as early as April 2020. Similar graphs, again displaying positive relationships using population-weighted densities and number of cases, are reported in Appendix Figure 4.A.4.

Naturally, these cross-sectional patterns do not constitute conclusive evidence that urban density results in faster or more deadly COVID-19 spread. There are at least two problems that could arise in this context. First, the positive correlation in the top left panel of Figure 4.1 can be the result of differences in the timing of the onset of the disease across locations. Second, certain location characteristics which are correlated with both density and COVID-19 spread and severity could induce a correlation in the absence of any actual

[11] We define the first wave as the period between the onset of the disease in the United States in February 2020 and the minimal daily death rate before the second rise in COVID-19 fatalities. See Appendix Figure 4.A.3.

Figure 4.1
Population Density and COVID-19 in 2020

(a) Deaths by 5th of July

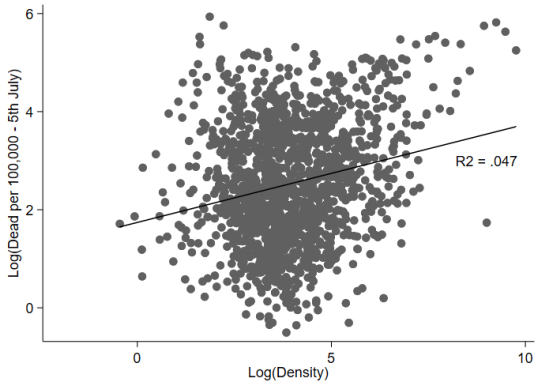

(b) Deaths by 15th of December

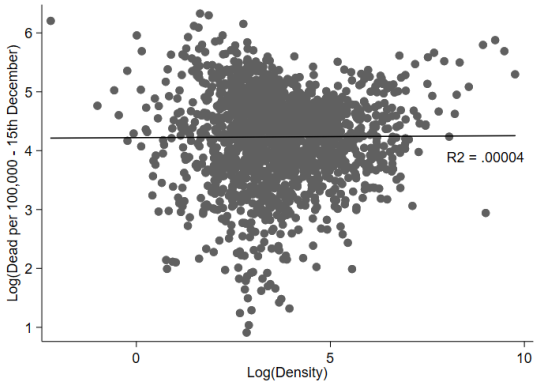

(c) Days to First Death

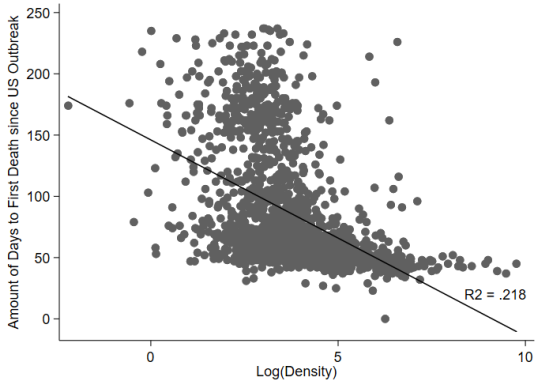

(d) Deaths 60 Days after 10th Case

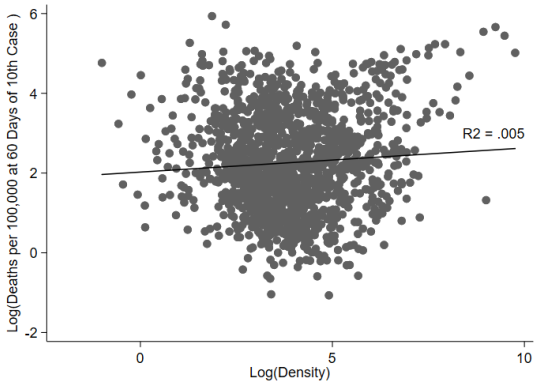

Notes: The horizontal axis represents the logarithm of the county's population density. *Top left panel* vertical axis represents the logarithm of the accumulated number of fatalities per hundred thousand inhabitants by the 5th of July 2020. *Top right panel* vertical axis represents the logarithm of the accumulated number of fatalities per hundred thousand inhabitants by the 1st of December 2020. *Bottom-left panel* vertical axis represents the number of days between the 22nd of January and the first fatality in each county. *Bottom-right panel* vertical axis represents the logarithm of the number of dead 60 days after the 10th case was reported in the county. Black markers correspond to counties forming part of a CBSA. Fitted lines estimated via Ordinary Least Squares. Univariate R-squared included in all Figures alongside fitted line.

causal link. We discuss this second issue in detail in the next section.

The top right panel of Figure 4.1 illustrates the point on differences in the timing of the onset of the disease across locations by showing that the positive correlation between population density and COVID-19 related deaths observed in the *first-wave* becomes almost flat when we use data extending to the 15th of December 2020. We investigate the timing dimension further in the bottom left panel of Figure 1 where we show the relationship between population density and the number of days between the 22nd of January and the

first fatality in each county. The figure exhibits a clear negative relationship, indicating that dense locations experienced COVID-19 fatalities earlier than more sparsely populated locations.

We can adjust for the differences in the timing of the onset of the disease by computing the number of deaths after a fixed number of days from that onset. This is what is typically shown in cross-country comparisons of the early evolution of the pandemic. In our case, we can compute the number of COVID-19 deaths at a specified time after the outbreak started in a county. We define the start of the outbreak as the first day with 10 reported cases and compute the number of deaths 60 days after this date for all counties.[12] The link between this time-adjusted variable and density is illustrated in the bottom-right panel of Figure 4.1. The relationship is almost flat after time-adjusting, suggesting that density does not simply translate into a higher rate of COVID-19 fatalities.

How is it possible that initial studies reported a clear positive influence of density on the impact of COVID-19, yet we report no relationship here? The answer is illustrated in Panel A of Figure 4.2, where we report how the slope of the relationship between population density and accumulated deaths evolved over 2020. These are simply the coefficients of a univariate regressions of the logarithm of total accumulated deaths – up to the period in the horizontal axis – on the logarithm of a county's population density.[13] Panel A of Figure 4.2 shows a positive relationship between deaths and density appeared at the beginning of the US epidemic, with the positive relationship peaking by May 15th 2020. Yet, in subsequent months the relationship progressively flattened, with the slopes of interest shrinking progressively until becoming statistically insignificant by November 15th. Thus, there was an apparently positive relationship at the beginning of the US epidemic, but this relationship became flat as the pandemic evolved.

Several factors could explain this result. We will turn to these in detail when we discuss mechanisms in Section 4.3.4, but consider as an illustration the role of changes in mobility across cities. Figure 4.3 shows the change in mobility relative to the January 2020 baseline

[12]The choice of 10 cases as marking the start of an outbreak from which we take the 60-day window is taken so as to ensure that there is some degree of within-county transmission at the time the window starts. We study how results change using different post-onset time windows in Section 3.2.

[13]Specifically, we estimate $Ln(\text{Acc. Deaths}_i^t + 1) = \alpha_0 + \alpha_t Ln(Pop.Dens_i) + \varepsilon_i$, where i is an index for counties and t indicates the end period, so that Acc. Deaths_i^t corresponds to accumulated deaths in county i from the start of the pandemic up to date t (e.g. the 15th of April).

Figure 4.2
Accumulated Covid Deaths-Density Elasticities over Time
(a) OLS Estimates

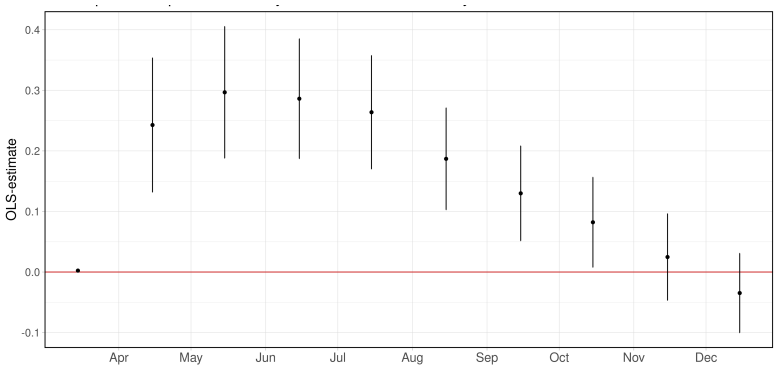

(b) IV Estimates

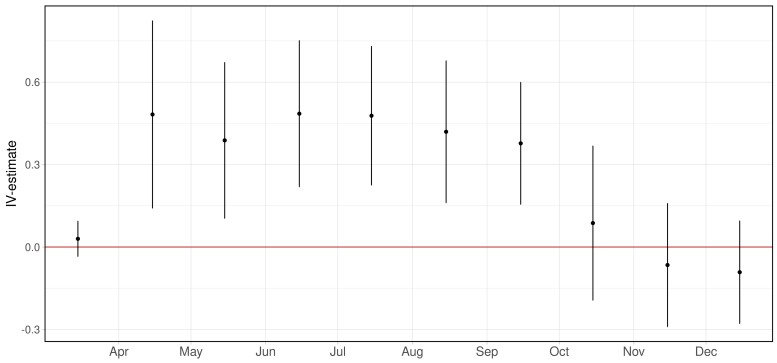

Notes: Both panels depict the cross-sectional relationship between the natural logarithms of accumulated deaths and population density for every monthly period ending in the 15th, from March through December 2020. Panel A: coefficients from univariate OLS regressions. Panel B: IV estimates obtained using both the geological and historical instruments for density. For each estimate we report the 95% confidence interval based on standard errors clustered at the CBSA level.

for sparse and dense counties, with the split based on median county density.[14] The left panel corresponds to changes in workplace-related mobility, the middle panel corresponds to changes in mobility for leisure activities and the right panel for transit. As expected, we observe a sharp reduction in mobility starting around mid-March. Importantly, in all cases we observe that this reduction is more acute in denser counties. Glaeser et al. (2020) show reductions in mobility had a substantial effect on the spread of COVID-19 over our sample period. Therefore, a sharper reduction in mobility in denser cities could contain the spread of the disease in these locations.

[14]The data is based on COVID-19 Community Mobility Reports released by Google and is based on data from portable device users in United States counties.

Figure 4.3

Changes in Mobility Relative to January Baseline (2020)

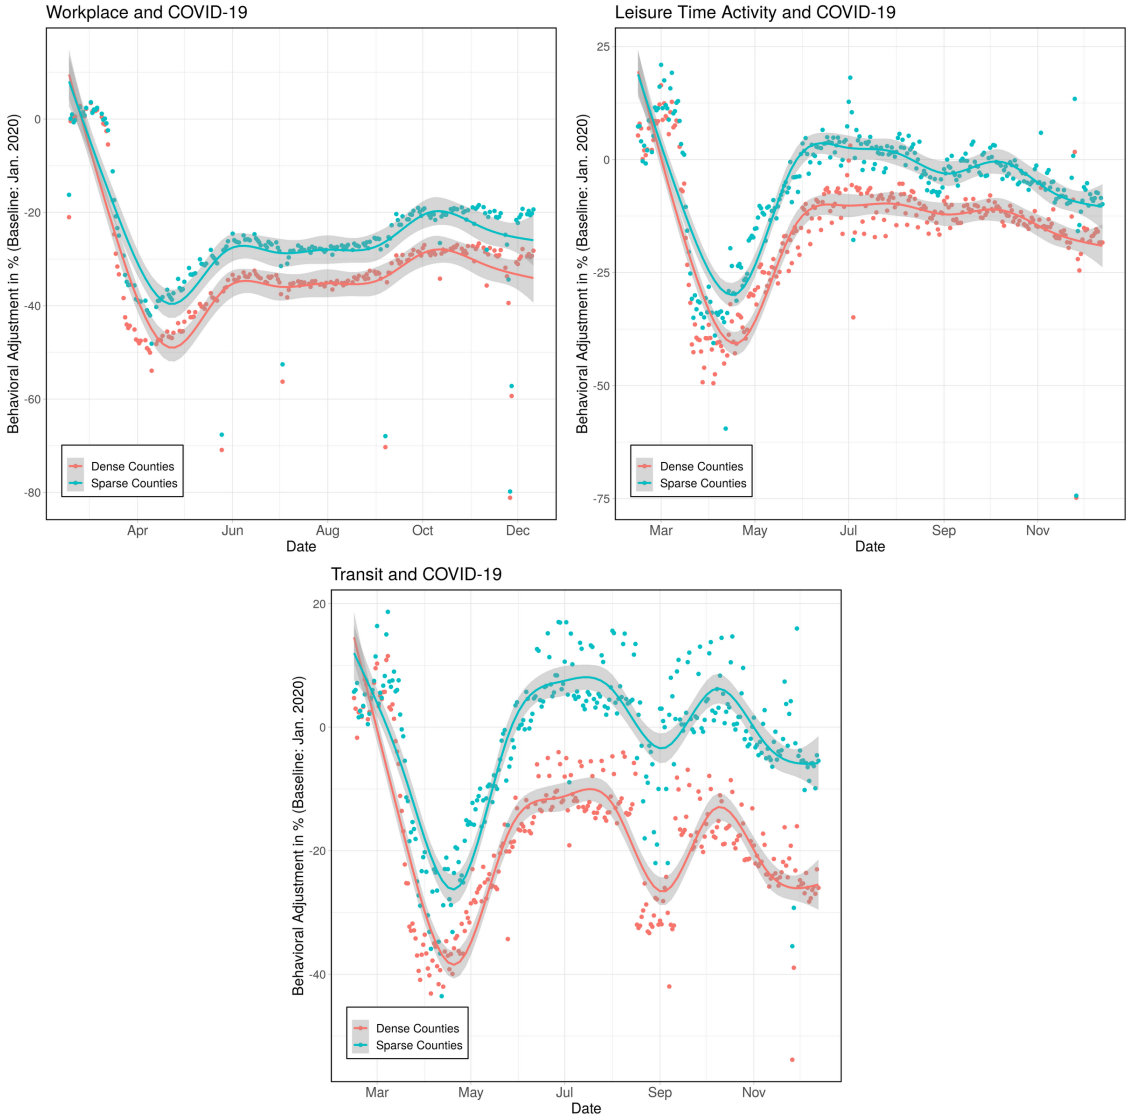

Notes: The figures plot the daily change and local regression curve (LOESS) over time in mobility relative to the January 2020 baseline for sparse counties and dense counties, with the split based on median weighted county density. The top left panel refers to adjustment of workplace-related activity. The top right panel refers to leisure time activities including restaurants, cafes, shopping centers, theme parks, museums, libraries, and movie theaters. The bottom panel refers to transit including public transport hubs such as subway, bus, and train stations.

4.3.2 Estimation

To obtain credible quantitative estimates of the relationship between time-adjusted COVID-19 related mortality and density, we also need to deal with potential confounders affecting both density and the prevalence and severity of the disease. Climate conditions, for example, can simultaneously influence household location decisions (see Glaeser et al. 2001)

and COVID-19 spread.[15] Local amenities such as waterfronts or low precipitation levels can themselves influence travel patterns – e.g. by increasing tourist arrivals – which could in turn affect COVID-19 rates. Insofar as some of these elements are observable, we can include them as controls in our regressions. Yet, some confounders may be unobservable due to their inherent nature or lack of accurate data. For instance, locational productive advantages can simultaneously affect local economic conditions and increase local densities.[16] Examples range from natural factors such as fertile or irrigable lands to man-made infrastructures such as ports or highways. Insofar as COVID-19 incidence and deaths are affected by economic conditions, unobservable locational advantages can confound the effect of density on the spread and severity of the disease.

To overcome the problem posed by potential unobservable confounding factors, we borrow canonical instruments for density from the agglomeration literature Combes et al. (2011) and our previous work on the relationship between density and air pollution Carozzi and Roth (2020). Specifically, we will instrument population density with either geological factors which can affect the costs of compact urban development or a long-lags in population density.

We use three geological instruments: the fraction of the urban footprint with aquifer presence, a measure of average earthquake risks and an estimate of soil drainage quality. The rationale for the aquifer instrument is that new dwellings in the periphery of urban areas need to either to pay for a costly connection with the municipal network or to directly connect with an underwater source. Given that the option of the underwater source is only available if there is an aquifer where the dwelling is located, cities with more land over aquifers can sprawl out further, contain more sparse development and lower densities. This instrument is motivated by the work in Burchfield et al. (2006) which reports that aquifers in the urban fringe are associated with urban sprawl. The rationale for our earthquake risk instrument is the expectation that the risk of an earthquake might influence building regulations, construction practices and the space between buildings, thus also affecting urban density. We also expect this instrument to satisfy the exogeneity condition, once we condition for distance to sea, average precipitation, latitude, longitude, and state fixed effects. Finally, the soil drainage quality variable is expected to affect land suitability for building at different densities. In fully urbanized land, a significant fraction of rainfall is

[15] A number of recent papers document a negative effect of temperature on COVID-19 incidence, at least in temperate weathers. See for example Prata et al. (2020); Tobías and Molina (2020).

[16] Locational advantages increase local densities because higher land prices in these areas trigger a substitution of land for capital in the production of structures (i.e. an increase in building heights).

drained through drainage networks and sewage systems (Konrad, 2003). However, at lower densities, soil drainage capacity is important to avoid stagnant water and, possibly, floods. In addition, high drainage soil is not ideal for laying down heavy infrastructure, making the task of building high density development more expensive.

We use a separate instrument for density based on historical population as recorded in the 1880 US census. Settlements in this period were in place before much of the technological revolutions in transportation that have affected location patterns in the last decades and also precede current patterns of industrial location. The use of historical population instruments for density was popularized by Ciccone and Hall (1996) and has been featured recurrently in the literature on agglomeration economies since (see Combes and Gobillon 2015 for a review).

Our main estimating equation will regress measures of COVID-19 presence on the logarithm of population density:

$$Y_i = \alpha_s + \beta Ln(Pop.Density)_i + \gamma' X_i + \varepsilon_i \tag{4.1}$$

where i indexes individual counties, α_s is a set of state effects and X_i is a set of controls. In all specifications, we control for average maximum and minimum temperatures, average yearly precipitation, latitude, longitude, distance between the county centroid and the closest sea front and distance to the closest waterfront. Our outcomes include different measures of COVID-19 presence. In most of our analysis, these are either variables capturing the time it took for the disease to arrive at a county or a time-adjusted measure of COVID-19 presence - the logarithm of the number of COVID-19 fatalities in the county 60 days after the 10th case was confirmed.

Before presenting our results, it is important to highlight that our estimates of parameter β from equation 4.1 will capture the overall effect of density on the outcome of interest. This includes the effect of geographic proximity facilitating transmission but also effects operating through the impact of density on agglomeration economies, personal behavior, local population compositions, healthcare systems, etc. After reporting estimates of the overall effect of density, we will turn to investigate the specific mediating factors behind it in Section 3.4 below.

4.3.3 Main Results

We first report baseline cross-sectional correlations between population density and COVID-19 cases and deaths during the *first-wave*. In Table 4.1, we estimate Equation 1 via Ordinary Least Squares (OLS) using the logarithm of the number of cases per 100,000 inhabitants and the logarithm of the number of deaths per 100,000 inhabitants as outcome variables. We find positive and statistically significant effects of population density on COVID-19 incidence, in line with the descriptive evidence reported in the top-left panel of Figure 4.1. Specifically, when using the conventional measure of population density we find elasticities of 22% and 13% for cases and deaths, respectively. This suggests that a 1% increase in population density increases cases and deaths per 100,000 people by 0.22% and 0.13%. When using our population-weighted measure of density, we also find very similar positive elasticities. The findings for COVID-19 cases are consistent with the evidence presented by Wheaton and Kinsella Thompson (2020) and Almagro and Orane-Hutchinson (2020). Yet, this should not be taken as conclusive evidence that density has a causal effect on the spread of COVID-19. As argued above, potential differences in the timing of the onset of the disease across locations or the presence of potential unobservable confounders can induce substantial bias in these coefficients.

Estimates reported in Table 4.2 deal with these empirical issues by looking explicitly at differences in the onset of the COVID-19 epidemic across locations and incorporating our instrumental variable strategy. In panels A and B, we report estimates for the effect of density on the number of days to the first case and the number of days to the first death. These numbers are measured relative to the date of the first reported case in the United States, so that small numbers correspond to an earlier onset of an outbreak. In column 1, we report OLS estimates obtained after controlling for state effects and covariates. In columns 2 and 3, we show IV estimates obtained using our Geological and Historical instruments, respectively. Note that the first-stage F-stats lie at 25 or above and the instruments explain between 5% and 10% of the variance in population density, indicating that they are not weak. Our second-stage estimates confirm that denser areas have indeed experienced earlier onsets of the disease whether we use days to the first case or days to the first death. A one log-point increase in density reduces the time to the first case by between 4 and 6 days depending on the specification. The effect on the time to the first deaths is even larger. These estimates demonstrate the importance of adjusting for differences in the timing of the onsets across locations when estimating the relationship between population

Table 4.1
Cases and Deaths in First COVID-19 Wave in 2020: Baseline OLS Estimates

A. Density				
	Log(Cases per 100,000)		Log(Deaths per 100,000)	
Log(Population Density)	0.218***	0.217***	0.130**	0.126***
	(0.027)	(0.021)	(0.051)	(0.034)
R2	0.20	0.42	0.11	0.36
Obs.	1,756	1,756	1,414	1,414
B. Weighted Density				
	Log(Cases per 100,000)		Log(Deaths per 100,000)	
Log(Weight. Density)	0.229***	0.219***	0.150**	0.110***
	(0.027)	(0.020)	(0.059)	(0.036)
State Effects	No	Yes	No	Yes
R2	0.21	0.42	0.11	0.35
Obs.	1,756	1,756	1,414	1,414

Notes: Baseline OLS estimates. Columns (1) and (2) use the log of cases per 100,000, columns (3) and (4) the log of deaths per 100,000 inhabitants as dependent variables, both taken as accumulated by the 5th of July. In Panel A, we report estimates for the effect of log of population density. In Panel B, we use the log of population-weighted density. In all models, we include controls for average maximum and minimum temperatures, average yearly precipitation, latitude, longitude, distance between the county centroid and the closest sea front and distance to the closest waterfront. The specifications in columns (2) and (4) add state effects. Standard errors in parenthesis are clustered at the CBSA level. ***p<0.01, **p<0.05, *p<0.1.

density and COVID-19 health outcomes.

In Panel C of Table 4.2, we examine our main outcome of interest; the effect of population density on time-adjusted COVID-19 related mortality. As mentioned previously, we focus on confirmed COVID-19 related deaths rather than cases as our main outcome of interest because it is considered to be a more accurate indicator of local COVID-19 prevalence. We provide a complementary analysis using reported cases in Section 4.3.5. In column 1, we find that the cross-sectional correlation observed in Table 4.1 becomes *negative* and statistically insignificant, suggesting that the positive link between population density and COVID-19 deaths might have been confounded by differences in the timing of the local outbreak. In columns 2 and 3, we use our instrumental variable approach to test this hypothesis more convincingly. Our second-stage results reveal a statistically insignificant relationship between population density and COVID-19 related deaths in both columns, portraying a similar picture as the OLS estimate presented in column 1. Our 2SLS results are unsurprisingly less precise, but the overall picture is clear. We find no evidence that

Table 4.2
Onset of the Disease and Deaths after 60 Days in 2020

	OLS	IV	
A. Days to First Case			
Log(Population Density)	-5.957***	-4.288***	-5.034***
	(0.447)	(0.864)	(0.870)
IV F-stat		25.9	111.0
R2	0.44	0.42	0.43
Obs.	1,759	1,759	1,733
B. Days to First Fatality			
Log(Population Density)	-15.935***	-6.673**	-13.384***
	(1.132)	(3.247)	(2.443)
IV F-stat		24.9	95.3
R2	0.34	0.29	0.33
Obs.	1,667	1,667	1,642
C. Log(Deaths per 100,000, 60 Days after 10th Case)			
Log(Population Density)	-0.046	-0.053	0.043
	(0.040)	(0.125)	(0.077)
IV F-stat		25.6	60.5
R2	0.30	0.30	0.29
Obs.	1,441	1,441	1,418

Notes: The main explanatory variable in all models is the natural logarithm of population density. Panels A and B report the estimates for the number of days to the first case and death respectively. Panel C reports the result for the log of the number of deaths per 100,000 residents in a county, 60 days after the 10th reported case. Column (1) corresponds to OLS estimates, column (2) and (3) presents 2SLS estimates using the Geological and Historical instruments respectively. In all models, we include controls for average maximum and minimum temperatures, average yearly precipitation, latitude, longitude, distance between the county centroid and the closest sea front and distance to the closest waterfront. The specifications in columns (2) and (3) add state effects. Standard errors in parenthesis are clustered at the CBSA level. ***p<0.01, **p<0.05, *p<0.1.

population density is positively linked with COVID-19 related deaths.

We can use our IV strategy to reproduce the findings illustrated in Panel A of Figure 4.2 showing the evolution of the cross-sectional relationship between COVID-19 deaths and population density over time. For this purpose, we estimate modified versions of equation 4.1 where the dependent variable is now the accumulated number of deaths up to the 15th day of each month in 2020 from March to December. Estimates of the different β_t slope coefficients obtained using 2SLS are reported in Panel B of Figure 4.2. In this case, we use both our geological and historical instruments as a source of exogenous variation. We observe that these results mimic those in Panel A, with an initially positive and signific-

ant relationship emerging by April 15th giving way to a progressive flatter relationship throughout 2020.

Finally, we test whether the relationship between density and time-adjusted COVID-19 deaths changes with the window used. To do this, we obtain estimates corresponding to 21, 30, and 45 day windows, all measured after the 10th case is reported in each county. The results are reported in Table 4.3 and show that the time-adjusted number of deaths is not positively affected by density, regardless of the window used. Interestingly, we find that at a beginning of an outbreak in a given county this relationship is in fact negative but becomes flat within two months.

TABLE 4.3
DENSITY AND TIME-ADJUSTED DEATHS IN 2020 FOR DIFFERENT POST-ONSET WINDOWS

	OLS	IV	
A. Log(Deaths per 100,000, 21 Days after 10th Case)			
Log(Population Density)	-0.377***	-0.357***	-0.194**
	(0.044)	(0.116)	(0.076)
First stage F-stat		20.1	67.2
Obs.	1,159	1,159	1,141
B. Log(Deaths per 100,000, 30 Days after 10th Case)			
Log(Population Density)	-0.240***	-0.234*	-0.072
	(0.046)	(0.130)	(0.074)
First stage F-stat		20.4	72.8
Obs.	1,264	1,264	1,242
C. Log(Deaths per 100,000, 45 Days after 10th Case)			
Log(Population Density)	-0.126***	-0.162	-0.021
	(0.045)	(0.129)	(0.078)
First stage F-stat		24.0	77.7
Obs.	1,367	1,367	1,345
Instrument		Geological	Historical
State Effects	Yes	Yes	Yes

Notes: Estimates of the effect of the natural logarithm of population density on time-adjusted COVID-19 deaths per 100,000 population. Different panels correspond to different choices of the time-adjustment windows in the dependent variable. Standard errors clustered at the CBSA level. ***p<0.01, **p<0.05, *p<0.1

On first reflection, the null (or negative) results for COVID-19 spread in this section appear surprising given that the virus spreads via human contact and denser areas can

provide more opportunities for human interactions. Nevertheless, there are several mediating factors that might offset this intuitive mechanism. For example, density itself might attract younger residents who are less likely to develop significant symptoms. In addition, both behavioral and/or policy induced changes in behavior may be different in dense counties. In fact, studies on previous pandemics (e.g. the 1918 influenza pandemic) also show that population density is not necessarily linked with the spread and severity of a disease (Mills et al., 2004). In the next section, we explore potential mechanisms that can explain our reduced-form findings.

4.3.4 Mechanisms

Variation in density might lead to changes in several local conditions, which can themselves affect the spread and severity of the disease. These types of changes may provide mechanisms that reinforce or offset the hypothesized positive effects that have been suggested in the literature, both in terms of timing of the local onset of the pandemic and subsequent spread. We turn to study some of these mechanisms by estimating the effect of density on other determinants of COVID-19 spread and severity. To do so, we re-estimate Equation 1 using these hypothetical mediators as outcomes. The resulting estimates do not provide definite proof regarding the mechanisms explaining the effect of density on COVID-19 incidence and mortality, but should be interpreted as suggestive evidence in this regard.

We begin by looking at possible factors explaining the early onset of the disease in denser cities and show that density is associated with higher social connectedness with other US counties. Our proxy for this variable relies on Facebook's Social Connectedness Index (SCI).[17] This index is based on the relative frequency of friendship links between users of the social-network, with higher index values corresponding to a larger number of friendship links. To proxy for social connectedness with other counties we aggregate the SCI of each county with all *other* counties and normalize it by the own-county SCI. The resulting variable is large when inhabitants in a county are disproportionately connected to other counties. Coefficients resulting from estimating Equation 4.1 using the logarithm of this proxy as an outcome variable are provided in Panel A of Table 4.4. As above, we report both OLS estimates (column 1) and 2SLS estimates using our geological and historical instruments (columns 2 and 3). We observe consistently positive elasticities of roughly 0.4-0.5 across columns, indicating denser counties are more intensely related to other counties

[17] Kuchler et al. (2020) study how interpersonal networks provided a channel for the spread of the disease based on the SCI.

in the US.[18] These results provide a plausible explanation to our findings of early onsets of COVID-19 cases and deaths in denser counties illustrated in Figure 4.1 and Table 4.2.

Table 4.4
Suggested Mechanisms: Social Connectedness and Behavioral Responses

	OLS	IV	
A. Social Connectedness			
Log(Population Density)	0.619***	0.452***	0.372***
	(0.017)	(0.045)	(0.034)
IV F-stat		25.9	111.0
Obs.	1,758	1,758	1,732
B. Δ Workplace Related Activity			
Log(Population Density)	-3.789***	-4.860***	-3.796***
	(0.156)	(0.478)	(0.301)
IV F-stat		19.5	56.8
Obs.	1,355	1,355	1,336
C. Δ Retail Related Activity			
Log(Population Density)	-2.519***	-2.615**	-3.471***
	(0.325)	(1.022)	(0.641)
IV F-stat		19.7	50.4
Obs.	1,289	1,289	1,270
D. Republican Vote Share 2016			
Log(Population Density)	-0.050***	-0.009	-0.080***
	(0.003)	(0.011)	(0.008)
IV F-stat		25.9	111.0
Obs.	1,759	1,759	1,733
Instrument		Geological	Historical
State Effects	Yes	Yes	Yes

Notes: The main explanatory variable in all models is the natural logarithm of population density. In Panel A, we present the results for the social connectedness of a county based on Facebook's Social Connectedness Index. Panels B and C report the results on behavioral adjustment of workplace and retail activities relative to the January baseline respectively. Panel D features the results on votes for the Republican party in the 2016 presidential election. Column (1) corresponds to OLS estimates, column (2) and (3) presents 2SLS estimates using the Geological and Historical instruments respectively. In all models, we include controls for average maximum and minimum temperatures, average yearly precipitation, latitude, longitude, distance between the county centroid and the closest sea front and distance to the closest waterfront. The specifications in columns (2) and (3) add state effects. Standard errors in parenthesis are clustered at the CBSA level. ***p<0.01, **p<0.05, *p<0.1.

Next, we study how density affects behavioral responses to the pandemic (e.g. compliance with social distancing measures). We use data from the Google COVID-19 Community

[18] Dense counties are also candidates to have higher connectedness with locations outside of the United States.

Mobility Reports (CMR) to measure how mobility patterns in each county have changed relative to baseline levels measured in January 2020. In Panels B and C of Table 4.4, we show the relationship between county density and the change in mobility to workplaces and retail activity respectively. We find that population density is associated with a larger decline in mobility for both indicators. Doubling density reduces workplace-related mobility and retail related activity by approximately 2.6-3.4% and 1.7-2.4%, respectively. Given the significant variation in density across US counties, these estimates are large. Insofar as social distancing reduces the spread of the disease, these differences in behavior might explain why we find limited differences in spread by location after accounting for the timing of onset of the disease and confounding factors.

Several factors could explain this difference in behavior across dense and sparse counties. One candidate that could account for both policy responses and individual differences in behavior relates to ideological or political views. Allcott et al. (2020) show that the Republican county vote share has a positive and significant association with the number of weekly visits to points of interest during the peak of the social distancing measures in April. Anecdotal evidence also reveals substantial differences in the tone of the Democratic and Republican parties when discussing the pandemic and its consequences. If density is associated with reduced support for the Republican party, residents of denser areas may be more likely to comply with the social distancing advise. In Panel D of Table 4.4, we estimate this link using voting data from the 2016 presidential election as a proxy for Republican support. We find that population density has a negative association with the share of Republican voters, an observation that should come as no surprise for observers of US politics.[19] This difference in political preferences across locations could explain, at least in part, the observed differences in the behavioral response to the pandemic illustrated in Figure 4.3 and Table 4.4.

We can arrive at two conclusions from the results reported in Table 4.4. First, dense counties are more connected with other locations and this may account for earlier onset of the COVID-19 epidemic in these areas. Second, the behavioral response to the disease was larger in denser counties, with less mobility for work and leisure and reduced use of public transit in these locations.

[19]This relationship remains highly robust upon controlling for the share of black population as well as the population above 60 years of age. In fact, when adding these additional controls, the relationship remains between -0.04 and -0.05 and significant at the 99% confidence level for all three estimation approaches.

Finally, in Table 4.5, we examine the effect of density on access to healthcare and demographics, as these are likely to affect COVID-19 related mortality. In Panels A and B, we examine the effect of density on access to healthcare using the ratio of population to primary care physicians and the percentage of adults under the age of 65 without health insurance as proxies. We find that density is positively associated with the former and negatively associated with the latter, suggesting that denser locations benefit from better access to healthcare. In our context, this could be an important mediating factor for two main reasons. First, access to primary healthcare might affect the presence and management of underlying health conditions which consider being risk factors for COVID-19 mortality (Zhou et al., 2020). Second, access might also affect the probability of seeking and receiving medical treatment once infected with COVID-19. Relatedly, we also examine the link between population density and income in Panel C as it is likely to affect access to healthcare and also health status more broadly. As expected, we find that the density is positively associated with median household income, offering an additional explanation for our headline results. Finally, in Panel D, we examine the effect of density on the share of the population above 60 years of age. This is of particular importance given that older age considered to be a significant risk factor (Zhou et al., 2020) and that population density is likely to affect the age structure of local areas via its impact on employment opportunities Glaeser (1999). Indeed, we find some evidence that population density is linked with a smaller share of residents above 60 years of age. In other words, dense counties are "younger" than sparse counties and this could reduce the number of deaths in these areas.

Overall, our points relating to behavioral responses, healthcare provision and demographics provide probable explanations for the surprisingly flat relationship between density and COVID-19 related mortality reported in panel C of Tables 4.2.

4.3.5 Robustness Checks

In this section, we provide several tests to evaluate the robustness of our main findings. We first revisit our results for the time-adjusted COVID-19 deaths by controlling for time of onset. In Panel A of Appendix Table 4.A.2, we test whether the null effect of density is robust to flexibly controlling by week of onset in each state. This goes beyond simply time-adjusting the outcome variable of interest as it also incorporates differences in knowledge regarding the disease or country-wide behavioral adjustments. We find that our qualitative results remain unchanged, with coefficients being insignificantly different from 0 across specifications. In panel B, we test whether our results are affected by excluding the New

Table 4.5
Mechanisms: Healthcare Provision and Demographics

	OLS	IV	
A. Log Primary Care Physicians per Capita			
Log(Population Density)	0.244***	0.180***	0.148***
	(0.014)	(0.042)	(0.024)
IV F-stat		25.9	97.8
Obs.	1,714	1,714	1,688
B. Share of Pop. Uninsured			
Log(Population Density)	-0.003***	-0.005	-0.010***
	(0.001)	(0.003)	(0.002)
IV F-stat		25.9	111.0
Obs.	1,759	1,759	1,733
C. Median Houshold Income			
Log(Population Density)	3.975***	7.051***	2.167***
	(0.349)	(1.068)	(0.816)
IV F-stat		25.9	111.0
Obs.	1,759	1,759	1,733
D. Share of Pop. Above 60 Years			
Log(Population Density)	-0.019***	-0.002	-0.014***
	(0.001)	(0.005)	(0.003)
IV F-stat		25.9	111.0
Obs.	1,759	1,759	1,733
Instrument		Geological	Historical
State Effects	Yes	Yes	Yes

Notes: The main explanatory variable in all models is the natural logarithm of population density. In Panel A, we present the results for primary health care supply measured as the natural logarithm of the number of primary health care physicians in each county divided by population. Panels B refers to the share of adults without health insurance. Panel C reports the results on median household income in 1,000 USD. Panel D features the estimates for the share of population above 60 years of age. Column (1) corresponds to OLS estimates, column (2) and (3) presents 2SLS estimates using the Geological and Historical instruments respectively. In all models, we include controls for average maximum and minimum temperatures, average yearly precipitation, latitude, longitude, distance between the county centroid and the closest sea front and distance to the closest waterfront. The specifications in columns(2) and (3) add state effects. Standard errors in parenthesis are clustered at the CBSA level. ***p<0.01,**p<0.05, *p<0.1.

York metropolitan area.[20] In this case, we find a negative and statistically significant relationship between density and time-adjusted COVID-19 deaths in our OLS estimate but statistically insignificant effects when we use our IV methodologies. We interpret these

[20]We use the census 2010 definition corresponding to the New York-Northern New Jersey-Long Island CBSA.

results with caution, as we are imposing sample selection that simultaneously exclude the MSA with the largest initial outbreak and the highest density.

Much of the evidence featured in the discussion around the role of urban density in shaping the impact of COVID-19 has focused on the conventional, area-weighted definition of density (i.e., population divided by surface). In order to speak to that debate, this has been the object of our main analysis. But we can evaluate the robustness of our results to the definition of density by studying the effect of population-weighted densities. In Appendix Table 4.A.3 we reproduce our main results using this variable as our main independent variable of interest. Unfortunately, since our geological instruments do not provide a strong first stage for this variable, our IV analysis relies solely on our long lag instrument. Reassuringly, we find that the overall results are qualitatively similar to those obtained in Table 4.2. Panels A and B show denser counties had earlier onsets of the disease compared to sparse counties. In panel C, we find a negative association between weighted density and COVID-19 related deaths when using OLS. However, our IV estimates again show a statistically insignificant elasticity. We therefore conclude that variation in density did not result in more COVID-19 incidence and deaths in the United States beyond the effect on early onset of the disease despite prior descriptive evidence. We also check the robustness of our results regarding suggested mechanisms using population-weighted density as our main regressor of interest in Appendix Table 4.A.4. Reassuringly, we find that the overall results are qualitatively analogous to those reported in Table 4.2.

Finally, we test whether density affects the time-adjusted number of reported *cases* of COVID-19. As argued above, the number of cases is more likely to be affected by variation in testing resources and by the presence of asymptomatic cases. This motivates our focus on number of deaths in much of the main analysis. In Table 4.A.5, we report estimates of the relationship between density and the number of cases per 100,000 inhabitants measured 21, 30, 45 and 60 days after the 10th reported case in the county. IV estimates for the effect of density on time-adjusted cases are similar to estimates reported in Table 4.3. We conclude that the data does not yield any evidence indicating a positive effect of density on the spread of the disease.

4.4 Conclusions

Urban areas are often places of intense social interaction, crowded living and close contact. Whether Justinian's Constantinople, fourteenth century Florence or 1918 Philadelphia -

cities have historically been associated with the propagation of infectious disease. In the first three months of the global COVID-19 pandemic, large, dense urban areas around the world such as New York, Madrid and London were identified as disease hotspots. Increased awareness of the risks of present and future epidemics has understandably prompted a debate about the future of cities. Did density - the defining feature of cities - promote the spread of the disease?

Our analysis of the onset of the COVID-19 pandemic in the United States raises a series of important points regarding these questions. First, density is associated with an early arrival of COVID-19, so that urban cores and superstar cities get a head start on the spread of the disease. Second, the subsequent spread - once COVID-19 has arrived - is not faster or deadlier than in smaller towns or sparsely populated peripheries. Cities get hit first, but do not necessarily get hit harder. We argue this is one of the reasons why many of the early studies of the impact of density on the impact of COVID-19 reported positive findings. A wider look at the whole period before vaccination began yields a different overall view of this relationship.

Several mechanisms may explain these findings. Large cities are intensely inter-connected with other locations, which can explain early onset. Yet, in the case of within-city spread, different offsetting forces may be at play. Crowding may promote the spread of the disease but differences in precautionary measures, access to healthcare and demographics may contain it. As a result, our findings emphasize the importance of distinguishing between differences in spread between and within locations.

Our study contributes to the understanding of how a summary feature of urban structure – population density – shapes spread of disease and deaths. The way in which other elements or urban form, cities' transport infrastructure or housing conditions (e.g., overcrowding) shaped the impact of the COVID-19 pandemic is not addressed here and remains an active area of research (see e.g., Kamis et al. 2021, Borsati et al. 2022 and Brotherhood et al. 2022).

4.A Appendix

4.A.1 Additional Figures and Tables

Table 4.A.1
Descriptive Statistics

	Mean	Standard Deviation
A. Full Sample		
Population Density	147	696
Weighted Population Density	522	1,117
Population	173,406	432,333
COVID Deaths 60 Days after 10th case	49	272
COVID Cases 60 Days after 10th case	841	3,376
Share of Population above 60 Years	0.24	0.05
Δ Workplace Related Activity	-40.61	7.82
Δ Retail Related Activity	-35.62	11.98
Number of Counties	1,759	
Share of US population: 93%		
	Mean	Standard Deviation
B. COVID Outbreak Subsample		
Population Density	173	766
Weighted Population Density	585	1,220
Population	203,190	472,196
COVID Deaths 60 Days after 10th case	59	299
COVID Cases 60 Days after 10th case	1,008	3,704
Share of Population above 60 Years	0.24	0.05
Δ Workplace Related Activity	-41.17	7.92
Δ Retail Related Activity	-35.94	11.39
Number of Counties	1,441	
Share of US population: 89%		

Notes: Descriptive statistics presenting the mean and standard deviation for a set of key variables of interest. Panel A corresponds to the whole sample of urban counties (i.e. counties belonging to a CBSA). Panel B corresponds to the Outbreak sub-sample consisting of counties that had at least one confirmed COVID-19 death 60 days before the end of our sample period on the 15th of December 2020.

Figure 4.A.1
Sample Maps

(a) Sample Counties

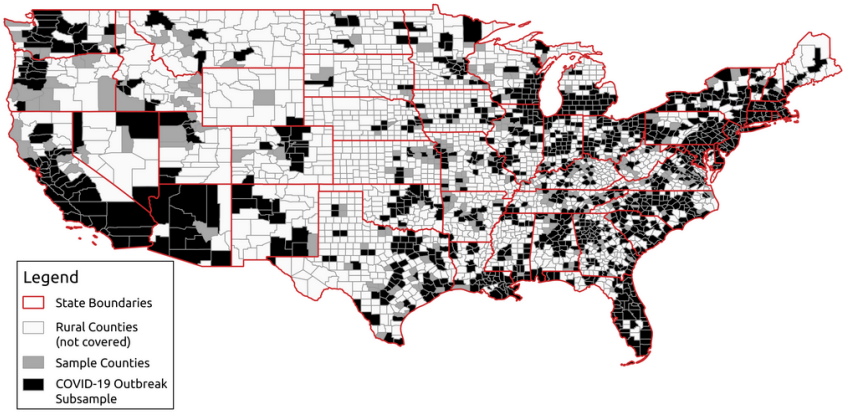

(b) COVID-19 Fatalities 60 Days after the First Case

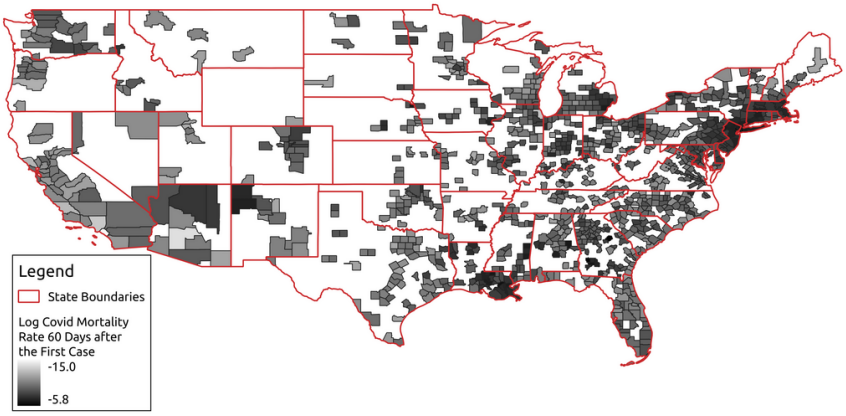

(c) Population Density

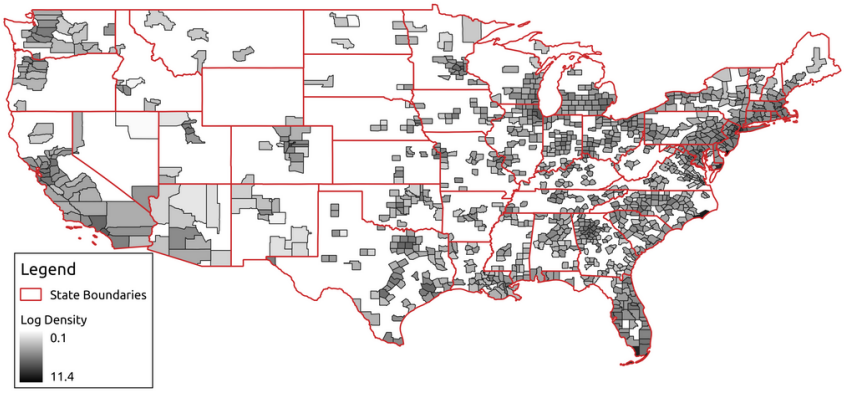

Notes: The map at the top shows which counties are part the urban county sample as well as the COVID-19 outbreak subsample. The map in the middle illustrates the distribution of the log COVID-19 mortality rate 60 days after the first case, and the map at the bottom the log population density for all counties that are part of the COVID-19 outbreak subsample.

Figure 4.A.2
Validating Covid-19 Figures

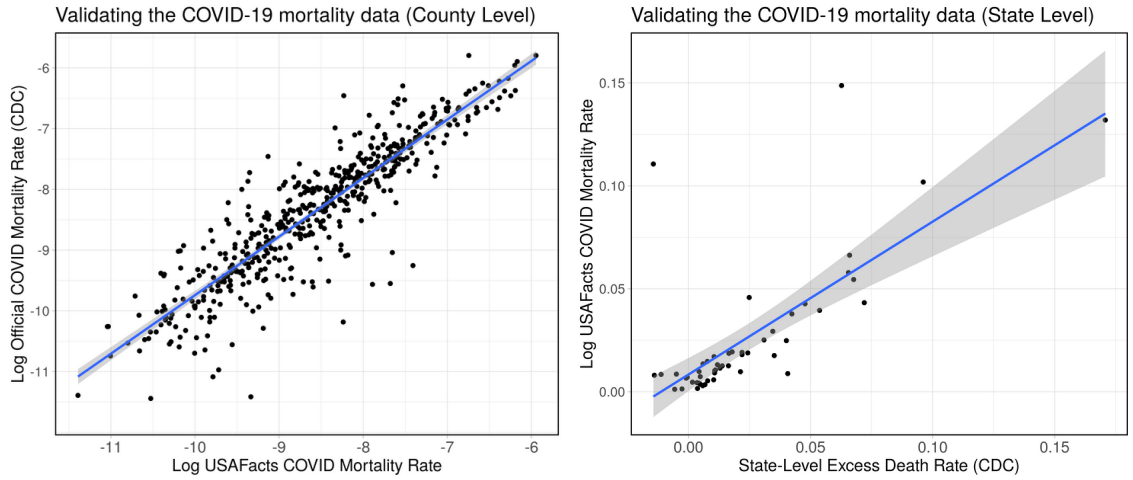

Notes: In the left panel, the vertical axis represents the log of the officially confirmed COVID-19 mortality rate per county by the CDC and the horizontal axis the COVID-19 mortality rate by USAFacts. The right panel plots the USAFacts state-level mortality rate (vertical axis) over the excess death estimates by the CDC (horizontal axis). Blue fit lines estimated via Ordinary Least Squares including the 95% confidence interval in gray.

Figure 4.A.3
Aggregate Daily Number of COVID-19 Deaths (2020)

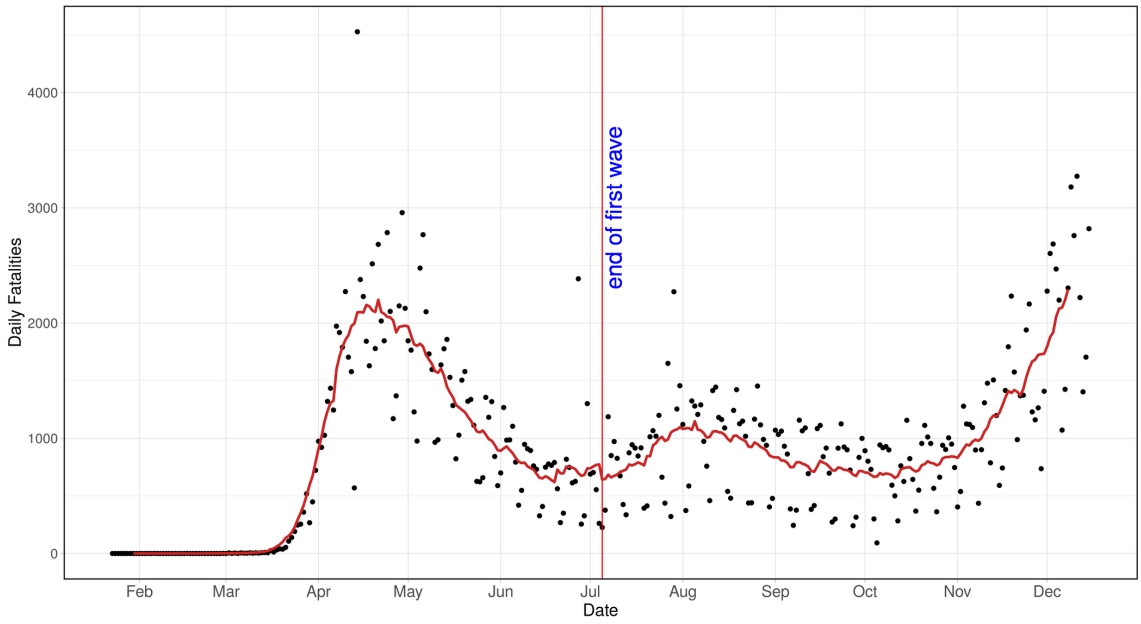

Notes: Daily COVID related deaths reported in the United States between February and the 15th of December 2020. Solid line represents moving average of daily deaths.

Figure 4.A.4
Cases and Deaths in First-Wave (Weighted Density)

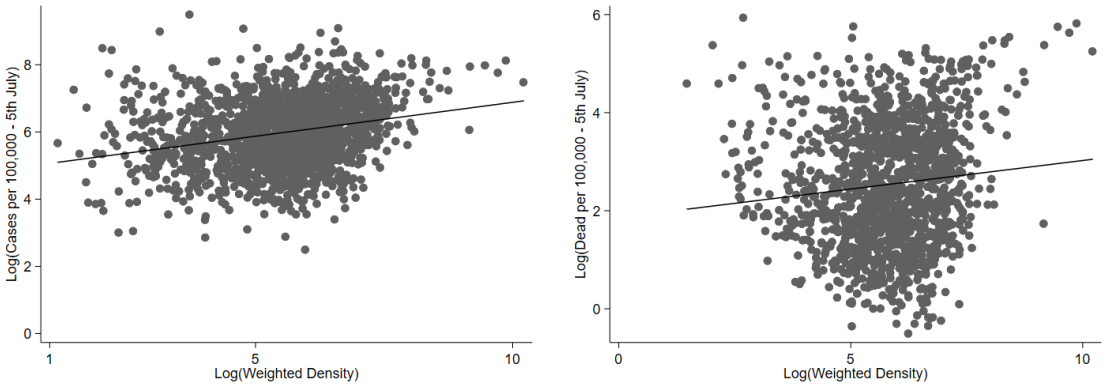

Notes: The horizontal axis represents the logarithm of the county's population-weighted density. In the left panel, the vertical axis represents the logarithm of the number of cases per 100,000 inhabitants. In the right panel, the vertical axis represents the logarithm of the number of fatalities per thousand inhabitants. Black markers correspond to counties forming part of a CBSA. Black fit lines estimated via Ordinary Least Squares.

Table 4.A.2
Robustness: Density and Deaths

	OLS	IV	
A. Controlling for Week of Onset Effects			
Log(Population Density)	-0.092**	-0.089	0.106
	(0.046)	(0.173)	(0.103)
First stage F-stat		23.8	60.6
Obs.	1,441	1,203	1,181
B. Excluding New York State			
Log(Population Density)	-0.128**	-0.053	0.043
	(0.053)	(0.125)	(0.077)
First stage F-stat		25.6	60.5
Obs.	1,203	1,441	1,418
Instrument		Geological	Historical
State Effects	Yes	Yes	Yes

Notes: Robustness tests corresponding to Table 4.2 Panel C, additionally controlling for the the week of the onset (Panel A) and excluding New York State (Panel B). The main explanatory variable in all models is the natural logarithm of population density. The dependent variable is the log of the number of deaths per 100,000 inhabitants in a county 45 days after the first case. Column (1) corresponds to OLS estimates, column (2) and (3) refer to 2SLS estimates using the Geological and Historical instruments respectively. In all models, we include controls for average maximum and minimum temperatures, average yearly precipitation, latitude, longitude, distance between the county centroid and the closest sea front and distance to the closest waterfront. The specifications in columns (2) and (3) add state effects. Standard errors in parenthesis are clustered at the CBSA level.

***p<0.01, **p<0.05, *p<0.1.

Table 4.A.3

Weighted Densities: Onset of the Disease and Deaths after 60 Days (2020)

	OLS	IV
A. Days to First Case		
Log(Weight. Density)	-5.867***	-9.789***
	(0.573)	(1.745)
IV F-stat		27.9
R2	0.44	0.34
Obs.	1,759	1,733
B. Days to First Fatality		
Log(Weight. Density)	-15.886***	-30.336***
	(1.205)	(7.102)
IV F-stat		20.8
R2	0.32	0.23
Obs.	1,667	1,642
C. Log(Deaths per 100,000, 60 Days after 10th Case)		
Log(Weight. Density)	-0.078**	0.090
	(0.039)	(0.160)
First stage F-stat		14.7
R2	0.30	0.28
Obs.	1,441	1,418
Instrument		Historical
State Effects	Yes	Yes

Notes: The main explanatory variable in all models is the natural logarithm of weighted density. Panels A and B report the estimates for the number of days to the first case and death respectively. Panel C reports the result for the log of the number of deaths per 100,000 inhabitants in a county, 60 days after the tenth case. Column (1) corresponds to OLS estimates and column (2) presents 2SLS estimates using the Historical instrument. In all models, we include controls for average maximum and minimum temperatures, average yearly precipitation, latitude, longitude, distance between the county centroid and the closest sea front and distance to the closest waterfront. The specifications in columns (2) and (3) add state effects. Standard errors in parenthesis are clustered at the CBSA level.

***p<0.01, **p<0.05, *p<0.1.

Table 4.A.4
Robustness: Suggested Mechanisms and Weighted Densities

	OLS	IV
A. Social Connectedness		
Log(Weight. Density)	0.532***	0.724***
	(0.021)	(0.091)
IV F-stat		27.7
Obs.	1,758	1,732
B. Δ Workplace Related Activity		
Log(Weight. Density)	-2.968***	-7.501***
	(0.200)	(1.248)
IV F-stat		15.6
Obs.	1,355	1,336
C. Δ Retail Related Activity		
Log(Weight. Density)	-1.893***	-7.293***
	(0.413)	(1.765)
IV F-stat		13.0
Obs.	1,289	1,270
D. Republican Vote Share 2016		
Log(Weight. Density)	-0.048***	-0.156***
	(0.004)	(0.023)
IV F-stat		27.9
Obs.	1,759	1,733
Instrument		Historical
State Effects	Yes	Yes

Notes: Corresponds to Table 4.4, using the log of weighted density as the main explanatory variable.

Table 4.A.5
Robustness: Cases

	OLS	IV	
A. Log(Cases per 100,000, 21 Days after 10th Case)			
Log(Population Density)	-0.249***	-0.334***	-0.148**
	(0.033)	(0.110)	(0.059)
IV F-stat		25.9	111.0
Obs.	1,759	1,759	1,733
B. Log(Cases per 100,000, 30 Days after 10th Case)			
Log(Population Density)	-0.187***	-0.269**	-0.076
	(0.035)	(0.117)	(0.063)
IV F-stat		25.4	109.8
Obs.	1,758	1,758	1,732
C. Log(Cases per 100,000, 45 Days after 10th Case)			
Log(Population Density)	-0.133***	-0.200*	-0.047
	(0.035)	(0.116)	(0.066)
IV F-stat		25.3	109.5
Obs.	1,757	1,757	1,731
D. Log(Cases per 100,000, 60 Days after 10th Case)			
Log(Population Density)	-0.107***	-0.145	-0.023
	(0.035)	(0.111)	(0.064)
IV F-stat		25.6	108.6
Obs.	1,754	1,754	1,728
Instrument		Geological	Historical
State Effects	Yes	Yes	Yes

Notes: The dependent variables are the log of the number of cases 60 days per 100,000 inhabitants after the 10th confirmed case. Column (1) corresponds to OLS estimates, column (2) and (3) refer to 2SLS estimates using the Geological and Historical instruments respectively. In all models, we include controls for average maximum and minimum temperatures, average yearly precipitation, latitude, longitude, and the distance between the county centroid and the closest sea front. The specifications in columns (2) and (3) add state effects. Standard errors in parenthesis are clustered at the CBSA level. ***p<0.01, **p<0.05, *p<0.1.

4.A.2 Data Sources

- **USAfacts.org COVID-19 Data**
  USAFacts is a non-profit civic initiative that provides data on the US population and government and works in partnership with the Penn Wharton Budget Model and the Stanford Institute for Economic Policy Research (SIEPR). The data can be retrieved at: `https://usafacts.org/visualizations/coronavirus-covid-19-spread-map/`. [Last visited: December 18th 2020]
- **CDC Official COVID-19 Mortality Rate** This database comprises confirmed or presumed COVID-19 fatalities and is limited to counties with at least 10 COVID-19 deaths. It should be noted, the dataset is incomplete because of the time lag between the death and the official certificate submitted to the National Center for Health Statistics (NCHS). For this reason, we this data corresponds only to 514 counties. The latest figures can be downloaded at: `https://data.cdc.gov/NCHS/Provisional-COVID-19-Death-Counts-in-the-United-St/kn79-hsxy`. [Last visited: December 18th 2020]
- **CDC Excess Mortality** Excess mortality corresponds to the deviation of total deaths to average expected deaths based on the experience in past years for each state. The latest estimates can be downloaded at: `https://www.cdc.gov/nchs/nvss/vsrr/covid19/excess_deaths.htm`. [Last visited: December 18th 2020]
- **US Census** contains information about demographics on the country level and can be accessed via: `https://www.census.gov/data/tables/time-series/demo/popest/2010s-counties-detail.html`. [Last visited: May 14th 2020]
- **'COVID-19 Community Mobility Reports' by Google**
  This report contains information about the behavioral activity change and social distancing in response to the COVID outbreak by county and day. For more detail on this database visit: `https://www.google.com/covid19/mobility/data_documentation.html?hl=en`. [Last visited: December 18th 2020]
- **Social Connectedness Data** Obtained after presenting a brief email application for the data based on this paper's outline to Mike Bailey and others at Facebook. April 6 2020 Release Version.
- **Healthcare and Income Data from The County Health Rankings and Roadmaps program** contains information on healthcare access and various social and economics indicators at the country level and can be accessed via: `https://www.countyhealthrankings.org`. [Last visited: July 3rd 2020]

Chapter 5

Air Pollution and Respiratory Infectious Diseases

5.1 Introduction

Exposure to elevated levels of air pollution is linked with a wide range of adverse health outcomes such as lower life expectancy, infant mortality, and emergency room visits (Dockery et al., 1993; Chay and Greenstone, 2003; Currie and Neidell, 2005; Schlenker and Walker, 2016). Recently, it has been suggested that air pollution might also be associated with the propagation of respiratory infectious diseases such as influenza ('the flu') and COVID-19. This potential link is of particular importance as respiratory infectious diseases lead to substantial disruptions and costs to health care systems and economies around the world. For example, it is estimated that the total annual economic burden of influenza on the US economy is $87.1 billion, and COVID-19 is projected to cost the US more than $16 trillion (Molinari et al., 2007; Cutler and Summers, 2020). As such, examining the potential link between air pollution and infectious diseases is important from a public health and economic standpoint.

In this paper, we study whether air pollution is linked with the most common and costly respiratory infectious diseases, namely influenza-like illnesses (ILI) and COVID-19. In theory, air pollution can affect respiratory infectious diseases in three main ways: First, exposure to air pollution can affect the body directly, either by making the respiratory system more vulnerable to such diseases or by inducing inflammatory reactions which impair the immune response to new infections (Ciencewicki and Jaspers, 2007). Second, the existence of pollution in the air might affect the airborne survival of respiratory viruses, allowing the virus to remain in the air for longer (Martelletti and Martelletti, 2020). Third, air pollution might also lead to changes in human behavior (e.g. staying indoors to avoid pollution exposure) that in turn can impact virus transmission via changes to the frequency and mode of social interactions. While the first two channels suggest that there might be a positive link between pollution and respiratory diseases, the last one is more ambiguous. We, therefore, aim to estimate the relationship between air pollution and respiratory infec-

tious diseases empirically by analyzing U.S administrative data on ambient air pollution and weekly cases of ILI and COVID-19.

Assessing the link between pollution and infectious disease is challenging due to the presence of correlated omitted variables and measurement error. More specifically, the volume of economic activity in a given region and time is just one of the many possible omitted variables that can lead to biased estimates as it is likely to affect both, air pollution and the propagation of infectious diseases. In terms of measurement, the assignment of air quality to the unit of analysis is bound to introduce some degree of measurement error due to variation within spatial units and across time. We overcome these challenges by using an Instrumental Variable (IV) approach that relies on deviations from long term atmospheric temperature inversion mean within a county/state and calendar week. The use of deviations of atmospheric temperature inversion rather than the popular straightforward inversion instrument is essential in this context due to the seasonal correlated pattern of inversions within a geography and infectious diseases that we find in the data.

Using our instrumental variable approach and considering a time frame of up to 6 weeks, we find no evidence that exposure to elevated levels of air pollution affect weekly influenza and COVID-19 cases in the U.S. Our results are precise, based on several time windows of exposure and are robust to different specifications. Importantly, our findings are in contrast to several recent papers in the economics literature which document positive links between pollution exposure and respiratory infectious diseases. Clay et al. (2018) documented a positive link between short-term exposure to air pollution and the number of deaths during the 1918 Spanish flu pandemic across U.S. cities, exploiting differential timing of the Spanish flu pandemic to overcome confounding factors. Using more recent data and exploiting random variation in wind direction as an instrument for air pollution, Graff Zivin et al. (2021) find that elevated levels of contemporaneous air pollution significantly increase influenza hospitalizations in the U.S. Austin et al. (2020) and Isphording and Pestel (2021) apply similar IV approaches to study the impact of particulate matter (PM) concentrations on COVID-19 cases and deaths in the U.S and Germany respectively. Both studies found significant positive effects. Finally, Persico and Johnson (2021) document that air pollution increases both the number of cases and case fatality of COVID-19, using the rollback of environmental regulations in the U.S as an instrument.

Our study provides two important contributions to the growing literature on the possible

link between air pollution and respiratory infectious diseases and to the literature on air pollution more broadly. First, our results are based on quasi-experimental methodology but yield very different results from previous studies. In contrast to other studies that document a positive association between pollution concentrations and respiratory infectious diseases, we find no link between the two looking at various lengths of exposure. We believe that given the precision of our estimates in conjunction with a sound methodology, it is vital to document such null results to foster further academic investigation on this matter. Second, we document that atmospheric inversion, which is widely used as an instrumental variable for air pollution in the economics literature (e.g. Arceo et al. (2016) or Bondy et al. (2020)), is subject to seasonal patterns (see Figure 5.A.1c). This is a serious concern to the validity of such instrument as it can also be correlated with other seasonal factors and the outcome itself. As such, we provide an alternative specification using deviations from long term inversion mean count within a county (or state) and calendar week, which overcomes this identification issue and still constitutes a strong predictor of pollution. We think that this is an important methodological contribution to the literature that examines the effect of air pollution on health but also on other aspects of human life including productivity, happiness, and human capital.

5.2 Data

To study the impact of ambient air pollution on the prevalence and severity of respiratory infectious diseases, we assemble two datasets. The first dataset is a weekly panel on influenza-like illnesses (ILI) at the US state level for the study period between 2010 to 2019, covering all 50 states as well as the District of Columbia. The second dataset is a weekly panel on COVID-19 covering 1,004 US counties representing 79.6% of the US population from early until late 2020 when the vaccination program rolled out. Summary statistics for our key variables in our data are presented in Table 1. In panel A, we show key statistics for our state level data on ILI and in Panel B for our county level COVID-19 sample. There are five main data inputs to create these datasets:

Influenza-like Illness Surveillance Network (ILINet): The first data source is provided by the Center for Disease Control (US CDC) in collaboration with the state and local health departments and health care providers. From this dataset, we obtain information about weekly counts of ILI patients across US states over 9 years/full flu seasons - from the 2010/11 flu season beginning in October 2010 until the 2018/19 flu season ending in Oc-

tober 2019.[1] We exclude more recent flu seasons to avoid an overlap with the COVID-19 pandemic. The CDC defines ILI patients as reporting symptoms of 'fever (temperature of 100°F [37.8°C] or greater) and a cough and/or a sore throat.[2] The ILINet data is based on information from over 3,000 healthcare providers across the US and allows to track the weekly ILI prevalence for all states over several years.

COVID-19: The second dataset covers weekly counts of COVID-19 cases and fatalities from usafacts.org, which is based on county-level reports by local health authorities across the US. This dataset covers the period between the start of the pandemic in the US in January 2020 until the launch of the vaccination program in December 2020. The dataset includes two measures of COVID-19 prevalence, cases and fatalities.[3] While cases might in principle be a better measure of disease prevalence, this indicator is hard to measure in practice. More specifically, the limited testing capacity, especially during the beginning of the pandemic, in conjunction with the virus' ability to spread in asymptomatic people limits the reliability of Covid cases as a main outcome of interest (Subbaraman, 2020). This type of reporting error is less likely to occur when using COVID-19 fatalities instead. Yet, fatalities are also a noisy measure of disease prevalence as the time between infection and death can take several weeks and vary a lot between cases. In addition, fatalities are not only a measure of disease prevalence but also of severity. Consequently, there are advantages and disadvantages from using either cases or fatalities. For this reason, we decided to focus on cases but we will show that we obtain similar results using deaths as an alternative measure.

Air Quality Index (AQI): We complement the health data with measures of air quality from the US Environmental Protection Agency (US EPA). The AQI is calculated on a daily basis using a variety of measures including carbon monoxide, nitrogen dioxide, ozone, sulfur dioxide as well as inhalable particulate matters (PM2.5 and PM10). For the period between 2010 and 2019, the EPA provides a daily AQI for 1,173 counties representing 83% of the US population. A higher AQI indicates a higher ambient air pollution level and ranges between 0 and 213.8 with a mean of 43.6 and standard deviation of 13.1 in our ILI sample, and between 0 and 848.6 with a mean of 36.9 and a standard deviation of 19.6 in our Covid sample (see Figure 5.A.1.[4] The AQI represents as noisy measure of air pollution since the

[1]Following the US CDC convention of an epidemiological week, we define weeks as starting on Sundays. A flu season is defined to run for one year starting in October.
[2]More information about the ILINet data can be found at: `https://gis.cdc.gov/grasp/fluview/fluportaldashboard.html`.
[3]More information about the COVID-19 data can be found at: `https://usafacts.org/visualizations/coronavirus-covid-19-spread-map/`.
[4]As a reference, an AQI between 0-50 is considered as 'Good'.

defining parameter varies across space and time, and measuring stations are capturing the air pollution at the station but not the average across the county. To combine the daily county-level AQI data with the ILI and COVID-19 data, we aggregate the AQI by week, and for ILI additionally by state weighing by county population.

Atmospheric Temperature Inversions: To obtain quasi-random variation in local air quality, we will use atmospheric temperature inversions as an instrument. Temperature inversions are short-term atmospheric episodes, usually occurring over a day or less, which lead to a reversal of temperature profiles that reduce atmospheric ventilation and consequently increase ground-level pollution levels. In other words, inversions might satisfy the condition of instrument relevance because they significantly impact air quality on the ground. Furthermore, while inversion episodes tend to be associated with other atmospheric and weather conditions, they are arguably independent of human behavior on the ground which makes them particularly popular as an instrument for air pollution at the daily (Jans et al., 2018; Sager, 2019) or weekly level (Arceo et al., 2016). Nevertheless, as we will show, inversions also follow cyclical patterns which is why it is important to either control for calendar-week fixed effects, or alternatively focus on deviations of inversions from long-run averages. We measure the occurrence of inversions based on satellite-derived three-dimensional temperature profiles of the atmosphere, which come from the MERRA-2 reanalysis project.[5] These provide 3-hourly mean temperatures by latitude, longitude and atmospheric pressure levels[6], which correspond to altitude. To match the inversions data with the other data inputs, we first spatially match inversion grid centroids with US counties to measure on how many days per week inversions occurred in a given county.[7] Specifically, we assign each grid point to the county it falls into and calculating mean temperature levels in a county on a day and each pressure level.[8] Whenever the daily mean temperature at the pressure level closest to the ground is lower than the temperature at the next higher-up level (25hPa less pressure which corresponds to roughly 200m in altitude), we define an inversion in that county on that day. Our final instrument measures the share of days within each week during which such an inversion occurred. To combine the inversions data with the ILINet data, we aggregate inversions to the state level using the county population share as weight.

[5]Further information about the dataset can be found at: `https://disc.gsfc.nasa.gov/datasets/M2I3NVASM_5.12.4/summary`.

[6]Location by latitude and longitude is divided into grid cells of size $0.625° \times 0.5°$. Altitude is divided into 42 atmospheric pressure levels with 25hPa intervals, which corresponds to approximately 200 meters.

[7]Note that we lose observation at this step as not all counties overlap with grid centroids.

[8]For those few counties which do not contain a grid point, we assign readings from the grid point closest to the county centroid.

Additional Weather Controls: While inversions may well be independent of human behavior, they are known to covary with certain weather conditions—including precipitation patterns, frozen rain, and fog formation—that may themselves affect human behavior. Consequently, we add control variables measuring surface air temperature, precipitation as well as relative humidity from the North American Regional Reanalysis (NARR) project by the National Oceanic and Atmospheric Administration (NOAA).

5.3 Methodology

Our aim is to estimate the hypothesized effect of contemporaneous air pollution exposure on the prevalence of respiratory infections. To do so, we use two samples with two different measures of the spread of respiratory infections:

(1) Weekly cases of influenza-like illness (ILI) at the state-level (2010-2019), and

(2) Weekly cases of COVID-19 at the county-level

We use similar methods for both of these. First, consider the case of ILI where we model the expected number of ILI cases in state i and during week t, denoted $Cases_{i,t}$, as the following exponential function:

$$E(Cases_{i,t}) = exp[\beta\ AQI_{i,t} + f(Weather_{i,t}) + \mu_t + \gamma_i] \tag{5.1}$$

The expected number of ILI cases exponentially depends on air quality, weather and additional invariant factors. Specifically, $AQI_{i,t}$ is the average air quality index (AQI) in state i and during week t, $Temp_{i,t}$ is average temperature, $RH_{i,t}$ is relative humidity, and $Rain_{i,t}$ is cumulative rainfall. We flexibly account for weather conditions in $f(Weather_{i,t})$ by including 20 temperature bins[9], a measure of relative humidity, an interaction between relative humidity and temperature, as well as a measure of cumulative rainfall and its square. In addition, we account for unobserved time-invariant heterogeneity using state fixed effects γ_i, and we account for common variation across time with year-week fixed effects, μ_t. We will show that this choice of fixed effects influences the results due to the likely strong degrees of seasonality and periodicity in both respiratory outcomes and air quality.

For our second sample, $Cases_{i,t}$ denotes the number of COVID-19 cases in county i and during week t, and all other variables are also measured at the county-level. In both cases, our coefficient of interest is β, which describes the relationship between air pollution and

[9]We include dummies for average temperature bins defined by the following cutoff values: -30, -10, -5, 0, 2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24, 26, 28, 30, 45 (all in °C.).

(exponential) cases of respiratory disease.

We estimate Equation 5.1 using the Poisson pseudo-maximum likelihood (PPML) regression as proposed by Silva and Tenreyro (2006) and implemented using the computationally efficient routine in the presence of high-dimensional fixed effects as developed by Correia et al. (2020). However, estimation of Equation 5.1 might yield biased results for two reasons - identification and measurement. In terms of the former, any observed association between air quality and respiratory outcomes as measured by $\hat{\beta}$ could be biased when certain variables are omitted from Equation 5.1 that affect both air quality and respiratory outcomes. The volume of economic activity in a given region and during a given week is just one of the many possible candidates for such an omitted variable. Regarding the latter, the assignment of air quality is bound to be imprecise due to variation within spatial units, be it counties or states, and throughout the week. Such measurement error may well also bias estimates $\hat{\beta}$, generally towards 0.

To address these concerns, we turn to a second identification strategy that relies on atmospheric temperature inversions as an instrument to induce plausibly exogenous variation in the levels of air quality. Temperature inversions are short-term atmospheric episodes, usually occurring over a day or less, which lead to a reversal of temperature profiles that lower atmospheric ventilation and thus temporarily increase ground-level pollution levels. They are thus best suited as instruments for short-term fluctuations in air quality at the daily (Jans et al., 2018; Sager, 2019) or weekly level (Arceo et al., 2016).

Specifically, we estimate the following linear first-stage relationship:

$$AQI_{i,t} = \rho\, INV_{i,t} + \delta(Weather_{i,t}) + \eta_t + \theta_i + v_{i,t} \tag{5.2}$$

Air quality in a given county or state i and during week t, $AQI_{i,t}$, is a linear function of the share of days in that week during which inversions occurred, $INV_{i,t}$, as well as the same covariates for weather and fixed effects as in Equation 5.1. As we will show, inversions are systematically associated with higher levels of air pollution throughout all specifications and both samples. To leverage the inversion instrument to estimate the exponential relationship between pollution and respiratory diseases as stipulated in Equation 5.1, we employ a control function approach as proposed by Wooldridge (2015). In a first step, we estimate Equation 5.2 using Ordinary Least Squares (OLS) estimation. We then add the residuals from that regression, $\hat{v}_{i,t}$, to the PPML estimation of Equation 5.1.

5.4 Results

5.4.1 PPML Estimates

We now turn to the results, beginning with the ILI state-year sample. Results from the non-instrumented PPML regression of Equation 5.1 are shown in Table 5.1. In column (1), we include weather controls but not fixed effects. Like in several previous studies, we find a positive association between pollution and ILI cases. Specifically, our estimate suggests that each additional 1-point increase in AQI[10] in a given state and week is associated with an increase in the expected number of ILI cases by about 1.2%. However, once we include fixed effects for time-invariant differences between states and for common shocks at the year-week level, results change. As shown in column (2), we find a precisely estimated zero effect of AQI on ILI cases. In column (3), we add one-week lags of ILI cases and AQI to account for any potential autocorrelation across time that may lead to either biased estimates or exaggerated precision. The coefficient of interest hardly moves and remains statistically and economically insignificant.

Another key concern is that both respiratory disease and air quality are highly seasonal (see Figure 5.A.1a and 5.A.1e). The year-week fixed effects may factor out some seasonal elements, but those that are common across all spatial units. However, if seasonal patterns of disease and pollution differ across states and counties - as indeed it seems quite plausible that they would - we may obtain biased estimates. Simply put, we may see more pollution and more flu cases during late January in some states that routinely experience severe winters, but less so in other states that have milder temperatures. In our view, this introduces a substantial risk of bias when trying to identify a hypothesized relationship between air quality and respiratory disease without accounting for region-specific seasonality trends. We take two approaches to region-specific seasonality.

First, we include state-calendar week fixed effects in column (4) of Table 5.1. Simply put, we take out all variation that is repeated across years in a given state and calendar week, such as the second week of each year in Texas. The coefficient of interest falls somewhat, but remains small and negative. Second, we take a different approach to seasonality-adjustments in column (5) of Table 5.1, by including the AQI only as deviations from the long-run average in each state and each calendar week. Perhaps unsurprisingly, given the similarity to the approach with calendar week fixed effects, the coefficient is also zero.

[10] As a reference, the standard deviation in our ILI Sample is 13.1 (see Figure 5.A.1).

Table 5.1
The Association between AQI and ILI Cases (PPML)

	Dependent variable: ILI Cases				
	(1)	(2)	(3)	(4)	(5)
AQI	0.012***	0.000	0.001	0.000	0.000
	(0.004)	(0.001)	(0.001)	(0.001)	(0.001)
Weather Controls	Yes	Yes	Yes	Yes	Yes
State FE	No	Yes	Yes	N/A	Yes
Week FE	No	Yes	Yes	Yes	Yes
Flu/AQI Lags	No	No	Yes	No	No
State-calendar week FE	No	No	No	Yes	No
AQI Deviations	No	No	No	No	Yes
Observations	21,519	21,519	21,418	21,519	21,519
Pseudo R2	0.09	0.87	0.88	0.89	0.87

Note: This table reports Poisson pseudo-maximum likelihood (PPML) estimates based on Equation 5.1 for the influenza-like illnesses (ILI) sample. The dependent variable are weekly ILI cases at the US state level provided by the Center for Disease Control (US CDC), and the main explanatory variable is the air quality index (AQI) (in column (5) deviations) by the US Environmental Protection Agency (US EPA), with higher AQI values indicating higher air pollution. Standard errors in parentheses are cluster-robust to autocorrelation within each flu season by state. *p<0.1; **p<0.05; ***p<0.01

However, using deviations of AQI from long-run averages in each state-calendar week, is our preferred specification as it will enable us to apply it to the COVID-19 sample where the calendar week fixed effects approach is infeasible.

Table 5.2 shows equivalent results for COVID-19 cases reported by each county in each week. Accordingly, the state fixed effects are replaced with county fixed effects. Due to the recency of the COVID-19 outbreak in the United States, we do not yet have data spanning multiple years and are thus not able to estimate a specification with county-calendar week fixed effects. While the simple specification without fixed effects in column (1) suggest a positive association between higher levels of AQI and reported COVID-19 cases, we detect no such positive relationship after appropriately accounting for time-invariant factors, common time-varying shocks, and region-specific seasonality. In fact, our results are small but the sign of the coefficient actually reverse, suggesting that higher level of pollution reduces the number of COVID-19 cases. As discussed in the data section, we also examine the effect of air pollution on COVID-19 fatalities - with a time lag of two weeks to allow for more time for the potential effects to materialize - as an alternative outcome measure,

given the limited testing capacity during the beginning of the pandemic and the virus' ability to spread in asymptomatic people which limits the reliability of the data on Covid cases. The results are presented in Table 5.A.2 and are qualitatively similar to those obtained in Table 3 in the sense that once we use AQI deviations, we find no evidence for a link between pollution and COVID-19 related mortality.

TABLE 5.2
THE ASSOCIATION BETWEEN AQI AND COVID-19 CASES (PPML)

	Dependent variable: Covid-19 Cases				
	(1)	(2)	(3)	(4)	(5)
AQI	0.008***	−0.002***	−0.001	-	−0.003***
	(0.0012)	(0.0008)	(0.0008)		(0.0007)
Weather Controls	Yes	Yes	Yes	-	Yes
County FE	No	Yes	Yes	-	Yes
Week FE	No	Yes	Yes	-	Yes
Flu/AQI Lags	No	No	Yes	-	No
AQI Deviations	No	No	No	-	Yes
Observations	47,431	47,430	46,325	-	47,430
Pseudo R2	0.11	0.90	0.90	-	0.90

Note: This table reports Poisson pseudo-maximum likelihood (PPML) estimates based on Equation 5.1 for the Covid sample. The dependent variable are weekly COVID-19 cases at the US county level provided by usafacts.org, and the main explanatory variable is the air quality index (AQI) (in column (5) deviations) by the US Environmental Protection Agency (US EPA), with higher AQI values indicating higher air pollution. Standard errors in parentheses are cluster-robust to autocorrelation at the level of counties. *p<0.1; **p<0.05; ***p<0.01

5.4.2 Control Function Estimates with Instruments

Next, we turn to the control function approach that uses inversions (columns (1)-(4)) and inversions deviations (column (5)) as an instrument for air quality. Table 5.3 shows OLS estimates for the linear first-stage relationship between inversions and AQI according to Equation 5.2. Reassuringly, we consistently find a positive and statistically significant (at the 1% level) relationship between more frequent inversions and higher AQI levels at the state and county levels. Specifically, we estimate that increasing the share of days with inversions from 0 to 1, i.e. going from a week with 0 inversions to one with 7, is associated with an increase in average AQI of between 12 and 18 points in our state level sample. Importantly, columns (4) and (5) show that the inversion instrument is robust to the two ways of accounting for region-specific seasonality. In column (4), this is done via state-

calendar week fixed effects. In column (5), we use the deviation of inversions from the long-run mean in each state and calendar week to predict the deviation of AQI from the long-run mean.

Table 5.3
The Association between Inversions and AQI (First stage - OLS)

	Dependent variable: AQI				
	(1)	(2)	(3)	(4)	(5)
Inversions	12.3***	15.9***	16.4***	17.7***	17.8***
	(1.6)	(0.8)	(0.7)	(0.9)	(0.9)
Weather Controls	Yes	Yes	Yes	Yes	Yes
State FE	No	Yes	Yes	N/A	Yes
Week FE	No	Yes	Yes	Yes	Yes
Flu/AQI Lags	No	No	Yes	No	No
State-calendar week FE	No	No	No	Yes	No
AQI Deviations	No	No	No	No	Yes
IV F-stat	59.9	378.4	506.5	432.3	414.4
Observations	21,519	21,519	21,418	21,519	21,519
Adj. R^2	0.36	0.68	0.74	0.77	0.39

Note: This table reports OLS first-stage regression estimates based on Equation 5.2 for the influenza-like illnesses (ILI) sample. The dependent variable is the air quality index (AQI) by the US Environmental Protection Agency (US EPA), with higher AQI values indicating higher air pollution. The instruments are the share of days in a week with inversions (columns (1)-(4)) and inversions deviations from their long-run average in the respective calendar week (column (5)). Standard errors in parentheses are cluster-robust to autocorrelation within each flu season by state. *p<0.1; **p<0.05; ***p<0.01

After confirming that inversions are indeed associated with higher levels of pollution, we now turn to our instrumental variable estimates of the relationship between air quality and respiratory disease. Our identification assumption is that the frequency of inversion episodes in a given week is not, after controlling for weather conditions and fixed effects, associated with any change in respiratory health other than through changes in air quality. We are not aware of any mechanism that would lead to such confounding, though we cannot be certain. In our view, the consistently strong first-stage estimates shown in Table A1, where coefficient hardly move between specifications, is very reassuring.

We begin by investigating the effect of air pollution on ILI. These results are shown in Table 5.4, with each column again showing equivalent specifications to those in Tables 2.

Table 5.4
The Association between AQI and ILI Cases (CF/PPML)

	Dependent variable: ILI Cases				
	(1)	(2)	(3)	(4)	(5)
AQI	0.063***	0.022***	0.017***	0.001	0.005
	(0.012)	(0.003)	(0.002)	(0.004)	(0.003)
Weather Controls	Yes	Yes	Yes	Yes	Yes
State FE	No	Yes	Yes	N/A	Yes
Week FE	No	Yes	Yes	Yes	Yes
Flu/AQI Lags	No	No	Yes	No	No
State-calendar week FE	No	No	No	Yes	No
AQI Deviations	No	No	No	No	Yes
Observations	21,519	21,519	21,418	21,519	21,519
Pseudo R^2	0.09	0.87	0.89	0.89	0.87

Note: This table reports Poisson pseudo-maximum likelihood (PPML) estimates based on the control function approach as proposed by Wooldridge (2015) that uses inversions (column (1)-(4)) and inversions deviations (column (5)) as instruments for air quality (columns (1)-(4)) and air quality deviations (column (5)) for the influenza-like illnesses (ILI) sample. The dependent variable are weekly ILI cases at the US state level provided by the Center for Disease Control (US CDC), and the main explanatory variable is the air quality index (AQI) (in column (5) deviations) by the US Environmental Protection Agency (US EPA), with higher AQI values indicating higher air pollution. The corresponding first-stage regressions can be found in Table 5.3. Standard errors in parentheses are bootstrapped using cluster-wise resampling at the level of flu seasons by state. *p<0.1; **p<0.05; ***p<0.01

As shown in column (1), we estimate a positive relationship between AQI and ILI cases when we do not account for any fixed effects. After including state and week fixed effects (column 2) and controlling for one-week lags for ILI cases, AQI and inversions (column 3), the coefficient of interest falls by more than half, but remains positive and statistically significant. These results may support the positive relationship between pollution and influenza cases found in other contributions. However, as we already discussed above, we believe that it is crucial to account for region-specific seasonality in such analyses. When we do so, by including state-calendar week fixed effects (column 4) or by taking deviations from the long-run mean for ILI cases, AQI and inversions (column 5), the coefficients of interest fall significantly and are no longer significantly different from 0.

The equivalent control function results for the COVID-19 sample are shown in Tables 5.A.4. Again, we fail to detect any systematic relationship between air quality and COVID-19 cases after appropriately accounting for region-specific seasonality patterns through demeaning our instrument in column (5) of Table 5. Table 5.A.5, shows similar results, using

Figure 5.1

Association between Leads/Lags of AQI Deviations and Disease

(a) ILI Cases: PPML

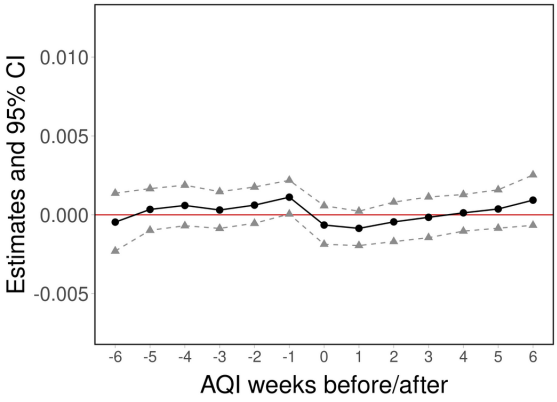

(b) ILI Cases: CF/PPML

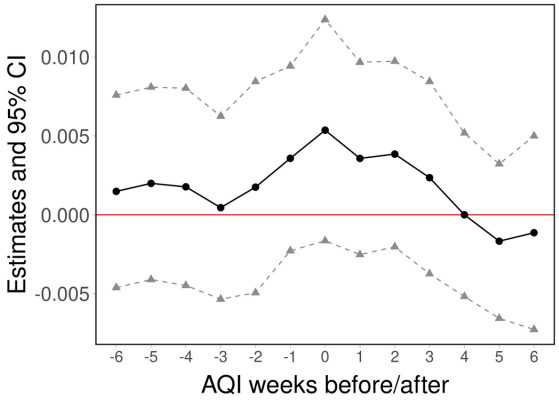

(c) COVID-19 Cases: PPML

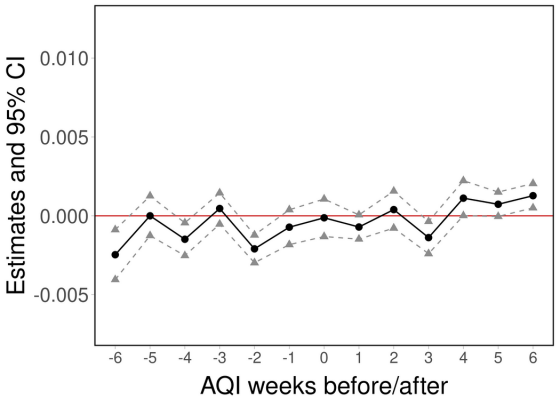

(d) COVID-19 Cases: CF/PPML

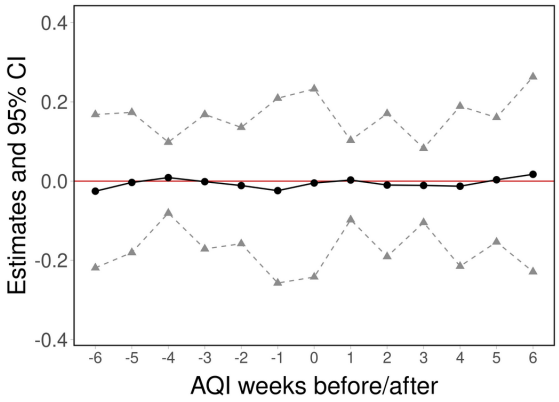

(e) COVID-19 Fatalities: PPML

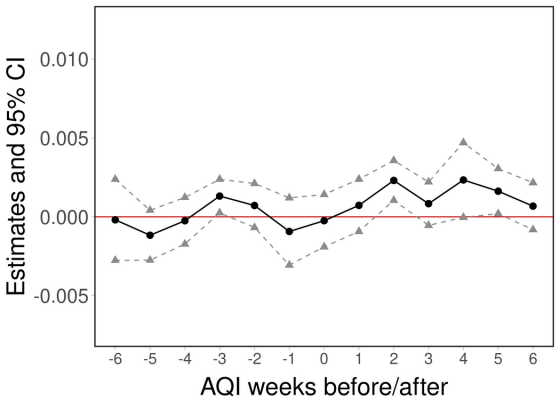

(f) COVID-19 Fatalities: CF/PPML

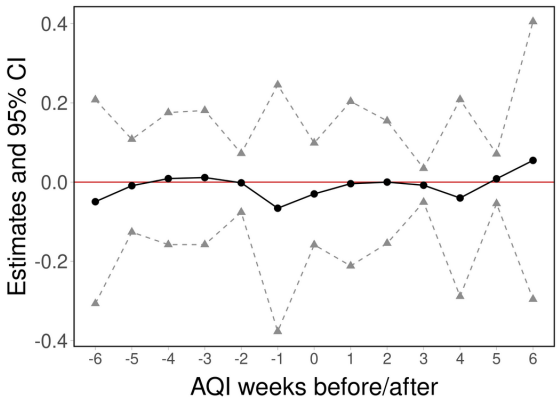

Note: The figures on the left plot the estimates based on Equation 5.1 equivalent to column (5) of Tables 5.1, 5.2 and 5.A.2 respectively, but with 6 leads and lags. The figures on the right plot the estimates based on the control function approach equivalent to column (5) of Tables 5.4, 5.A.4 and 5.A.5 respectively, but with 6 leads and lags. The 95% confidence interval is included in gray.

COVID-19 fatalities instead of cases, as an outcome.

So far, we have failed to detect a systematic, positive relationship between air pollution, as measured by AQI, and ILI case counts, both when looking directly at the relationship hypothesized in Equation 5.1 and when instrumenting AQI with inversions. However, our analysis has thus far focussed at the contemporaneous relationship at the weekly level.[11] However, it might be the case that the effect of pollution exposure takes some time to materialize in more infections, and ultimately in higher confirmed case counts. We thus extend our preferred specifications to include 6 weeks of leads and lags of air quality. The results are shown in Figure 5.1. Figure 5.1a, show results from a specification equivalent to that in column (5) of Tables 5.1, but with leads and lags. Quite evidently, there is no association between ILI case counts and air quality in either the preceding or following weeks. This is confirmed in Figure 5.1b, which show estimates from the control function approach, equivalent to column (5) in Tables 5.1, but with leads and lags.[12] In Figures 5.1c and 5.1d, we conduct the same analysis using COVID-19 cases as our outcome variable, and in Figures 5.1e and 5.1f using COVID-19 fatalities. Again, we find no association between air quality and COVID-19 cases in either the preceding or following weeks.

5.5 Conclusion

This paper has examined the relationship between air pollution and respiratory infectious diseases. In the first part of our analysis, we use a traditional fixed effects model which shows that exposure to air pollution is indeed associated with weekly cases of ILI and Covid-19 as suggested by previous papers in the economic and epidemiology literature. However, we later show how this relationship fades completely when we use our instrumental variable approach or a more accurate fixed-effect model that adequately accounts for seasonality. Importantly, our null results are precise, robust to different specifications, and remain virtually the same for different time windows of exposure. Given the growing body of literature on this topic which shows a positive relationship between air pollution and infectious diseases, we believe that it is vital to document our precisely estimated null results to foster further academic investigation on this matter. Finally, we also demonstrate that the widely used atmospheric inversion instrument is subject to a seasonal pattern, which

[11] With the exception of fatalities where we allow for a time lag of 2 weeks.

[12] Specifically, we estimate Equation (2) for AQI in different time periods (same week + 6 lags + 6 leads), including each time all inversion instruments (same week + 6 lags + 6 leads). We then estimate one second-stage relationship akin to Equation (1), but including the leads and lags of AQI, along with the residuals from all first-stage regressions.

cast a serious concern about its validity as it can also be correlated with other seasonal factors and the outcome itself. To overcome this empirical issue, we modify the instrument such that we use deviations from long-term inversion mean counts within a county (or state) and calendar week. This approach is crucial for the accuracy of our estimates, but it is also an important methodological contribution to the literature that examines the various health and well-being effects of air pollution and uses this instrument.

5.A Appendix

5.A.1 Summary Statistics

Table 5.A.1
Summary Statistics

	Min.	Max.	Mean	Median	SD.
A. ILI Sample (State Level)					
ILI Patient Rate	0	161.9	6.0	2.4	10.3
Air Quality Index	15.4	213.8	43.6	41.4	13.1
Inversions per Week	0	1.0	0.2	0.2	0.2
Relative Humidity	8.2	94.7	69.9	74.1	14.7
Precipitation (in mm)	0	25.5	2.6	1.8	2.7
Temperature (in °C)	-23.2	40.3	13.2	14.2	10.8
Population (in 1,000)	494	33,872	5,626	4,042	6,118
Observations:	21,519				
B. Covid Sample (County Level)					
COVID-19 Case Rate	0	5,675.4	136.4	47.2	224.4
COVID-19 Fatality Rate	0	161.5	2.1	0	4.6
Air Quality Index	0	848.6	36.9	36.4	19.6
Inversions per Week	0	1.0	0.2	0.1	0.2
Relative Humidity	9.9	96.1	69.5	74.5	16.2
Precipitation (in mm)	0	43.5	2.6	1.6	3.1
Temperature (in °C)	-17.9	40	14.6	15.2	10.2
Population (in 1,000)	0.6	9,519.3	232.4	91.8	500.6
Observations:	47,430				

Note: This figure presents key summary statistics for the ILI and Covid sample respectively.

5.A.2 Additional Figures

FIGURE 5.A.1
AGGREGATE SEASONALITY PATTERNS

(A) AQI

(B) AQI DEVIATIONS

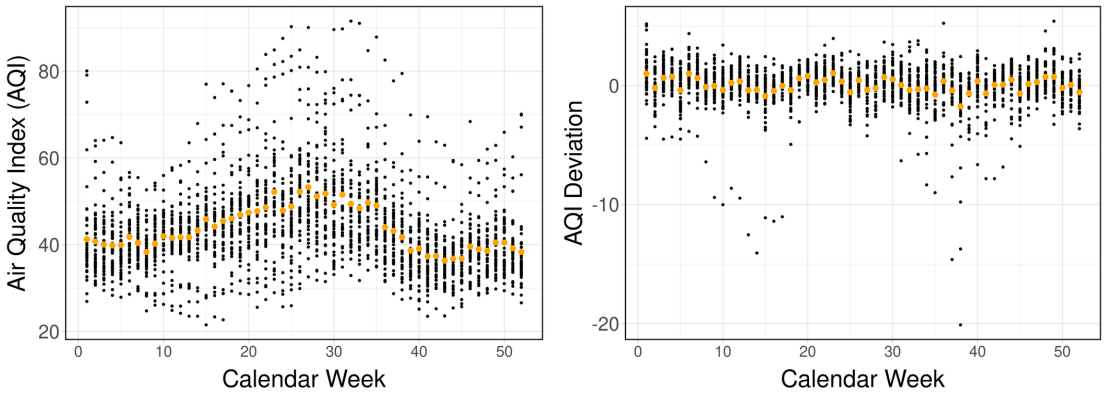

(C) INVERSIONS

(D) INVERSIONS DEVIATIONS

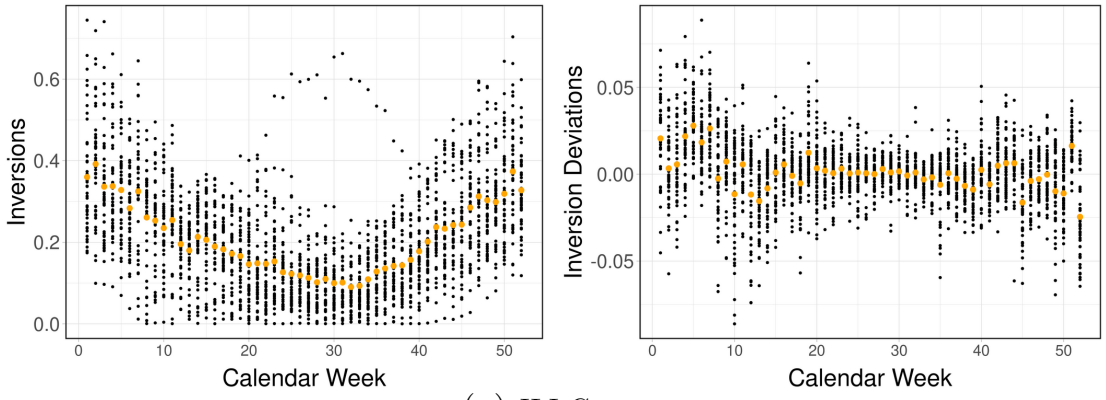

(E) ILI CASES

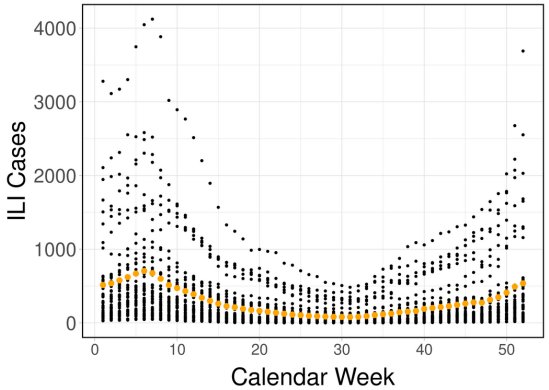

Note: These figures use the ILI sample to plot key variables over calendar week (the weekly bins are in orange). The AQI and inversions exhibit clear seasonal patterns, with the AQI being elevated during the summer and inversions during the winter. In contrast, taking the deviations from their respective long-term mean helps with removing the seasonality. ILI cases are displayed at the bottom and also exhibit a seasonality with elevated levels during the winter. It should be noted that these figures illustrate the aggregate seasonality across US states, and that the seasonality within states or counties is likely to be even more pronounced and vary across locations. By demeaning the AQI and inversions by state (or county) and calendar week, we remove this location specific seasonality.

5.A.3 Additional Tables

Table 5.A.2
The Association between AQI and COVID-19 Fatalities (PPML)

	Dependent variable: Covid-19 Fatalities				
	(1)	(2)	(3)	(4)	(5)
AQI	0.008***	−0.001	0.000	-	−0.000
	(0.001)	(0.001)	(0.001)		(0.001)
Weather Controls	Yes	Yes	Yes	-	Yes
County FE	No	Yes	Yes	-	Yes
Week FE	No	Yes	Yes	-	Yes
Flu/AQI Lags	No	No	Yes	-	No
AQI Deviations	No	No	No	-	Yes
Observations	44,250	44,250	42,295	-	44,250
Pseudo R2	0.10	0.75	0.76		0.75

Note: This table reports Poisson pseudo-maximum likelihood (PPML) estimates based on Equation 5.1 for the Covid sample. The dependent variable are weekly COVID-19 fatalities at the US county level two weeks later provided by usafacts.org, and the main explanatory variable is the air quality index (AQI) (in column (5) deviations) by the US Environmental Protection Agency (US EPA), with higher AQI values indicating higher air pollution. Standard errors in parentheses are cluster-robust to autocorrelation at the level of counties. *p<0.1; **p<0.05; ***p<0.01

Table 5.A.3
County-level Association between Inversions and AQI (First stage - OLS)

	Dependent variable: Covid-19 Cases				
	(1)	(2)	(3)	(4)	(5)
Inversions	6.1***	7.6***	7.4***	-	7.8***
	(1.2)	(0.7)	(0.5)		(0.6)
Weather Controls	Yes	Yes	Yes	-	Yes
County FE	No	Yes	Yes	-	Yes
Week FE	No	Yes	Yes	-	Yes
Flu/AQI Lags	No	No	Yes	-	No
AQI Deviations	No	No	No	-	Yes
IV F-stat	25.2	129.6	189.6	-	198.1
Observations	47,431	47,430	46,325	-	47,430
Adj. R^2	0.15	0.52	0.62	-	0.27

Note: This table reports OLS first-stage regression estimates based on Equation 5.2 for the Covid sample. The dependent variable is the air quality index (AQI) by the US Environmental Protection Agency (US EPA), with higher AQI values indicating higher air pollution. The instruments are the share of days in a week with inversions (columns (1)-(3)) and inversions deviations from their long-run average in the respective calendar week (column (5)). Standard errors in parentheses are cluster-robust to autocorrelation at the level of counties.
*p<0.1; **p<0.05; ***p<0.01

Table 5.A.4
The Association between AQI and COVID-19 Cases (CF/PPML)

	Dependent variable: Covid-19 Cases				
	(1)	(2)	(3)	(4)	(5)
AQI	0.126***	0.019**	0.023***	-	−0.006
	(0.023)	(0.007)	(0.008)		(0.007)
Weather Controls	Yes	Yes	Yes	-	Yes
County FE	No	Yes	Yes	-	Yes
Week FE	No	Yes	Yes	-	Yes
Flu/AQI Lags	No	No	Yes	-	No
AQI Deviations	No	No	No	-	Yes
Observations	47,431	47,430	46,325	-	47,430
Pseudo R2	0.12	0.90	0.90	-	0.90

Note: This table reports Poisson pseudo-maximum likelihood (PPML) estimates based on the control function approach as proposed by Wooldridge (2015) that uses inversions (column (1)-(3)) and inversions deviations (column (5)) as instruments for air quality (columns (1)-(3)) and air quality deviations (column (5)) for the Covid sample. The dependent variable are weekly Covid cases at the US county level provided by usafacts.org, and the main explanatory variable is the air quality index (AQI) (in column (5) deviations) by the US Environmental Protection Agency (US EPA), with higher AQI values indicating higher air pollution. The corresponding first-stage regressions can be found in Table 5.A.3. Standard errors in parentheses are bootstrapped using cluster-wise resampling at the level of counties. *p<0.1; **p<0.05; ***p<0.01

Table 5.A.5
The Association between AQI and COVID-19 Fatalities (CF/PPML)

	Dependent variable: Covid-19 Fatalities				
	(1)	(2)	(3)	(4)	(5)
AQI	0.090***	0.068***	0.045***	-	−0.017
	(0.026)	(0.015)	(0.007)		(0.011)
Weather Controls	Yes	Yes	Yes	-	Yes
County FE	No	Yes	Yes	-	Yes
Week FE	No	Yes	Yes	-	Yes
Flu/AQI Lags	No	No	Yes	-	No
AQI Deviations	No	No	No	-	Yes
Observations	44,250	44,250	42,295	-	44,250
Pseudo R2	0.11	0.75	0.77		0.75

Note: This table reports Poisson pseudo-maximum likelihood (PPML) estimates based on the control function approach as proposed by Wooldridge (2015) that uses inversions (column (1)-(3)) and inversions deviations (column (5)) as instruments for air quality (columns (1)-(3)) and air quality deviations (column (5)) for the Covid sample. The dependent variable are weekly Covid fatalities at the US county level two weeks later provided by usafacts.org, and the main explanatory variable is the air quality index (AQI) (in column (5) deviations) by the US Environmental Protection Agency (US EPA), with higher AQI values indicating higher air pollution. The corresponding first-stage regressions can be found in Table 5.A.3. Standard errors in parentheses are bootstrapped using cluster-wise resampling at the level of counties. *p<0.1; **p<0.05; ***p<0.01

Bibliography

Abebe, G., McMillan, M. S., and Serafinelli, M. (2022). Foreign direct investment and knowledge diffusion in poor locations. *Journal of Development Economics*, forthcoming.

Addison, T., Boly, A., and Mveyange, A. (2017). The impact of mining on spatial inequality recent evidence from Africa. Technical Report 7960, World Bank, Washington, DC.

Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., and Wolf, N. (2015). The economics of density: Evidence from the berlin wall. *Econometrica*, 83(6):2127–2189.

Ahmad, K., Erqou, S., Shah, N., Nazir, U., Morrison, A. R., Choudhary, G., and Wu, W.-C. (2020). Association of poor housing conditions with covid-19 incidence and mortality across us counties. *Plos one*, 15(11):e0241327.

Alesina, A., Easterly, W., and Matuszeski, J. (2011). Artificial states. *Journal of the European Economic Association*, 9(2):246–277.

Alesina, A., Hohmann, S., Michalopoulos, S., and Papaioannou, E. (2021). Intergenerational mobility in africa. *Econometrica*, 89(1):1–35.

Alesina, A., Michalopoulos, S., and Papaioannou, E. (2016). Ethnic inequality. *The Journal of Political Economy*, 124(2):428–488.

Allcott, H., Boxell, L., Conway, J., Gentzkow, M., Thaler, M., and Yang, D. Y. (2020). Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *NBER Working Paper*, w26946.

Allcott, H. and Keniston, D. (2018). Dutch disease or agglomeration? the local economic effects of natural resource booms in modern america. *The Review of Economic Studies*, 85(2):695–731.

Almagro, M., Coven, J., Gupta, A., and Orane-Hutchinson, A. (2020). Racial disparities in frontline workers and housing crowding during covid-19: Evidence from geolocation data. *Available at SSRN 3695249*.

Almagro, M. and Orane-Hutchinson, A. (2020). The determinants of the differential exposure to covid-19 in new york city and their evolution over time. *Covid Economics: Vetted and Real-Time Papers*, 13.

Angel, S., Lamson-Hall, P., Salazar Tamayo, M. M., et al. (2020). Coronavirus and the cities: Explaining variations in the onset of infection and in the number of reported cases and deaths in us metropolitan areas as of 27 march 2020. In *New York University. Marron Institute of Urban Management*. New York University. Marron Institute of Urban Management.

Aragón, F. M. and Rud, J. P. (2013). Natural resources and local communities: Evidence from a peruvian gold mine. *American Economic Journal: Economic Policy*, 5(2):1–25.

Aragón, F. M. and Rud, J. P. (2016). Polluting industries and agricultural productivity: Evidence from mining in ghana. *The Economic Journal*, 126(597):1980–2011.

Arceo, E., Hanna, R., and Oliva, P. (2016). Does the effect of pollution on infant mortality differ between developing and developed countries? evidence from mexico city. *The Economic Journal*, 126(591):257–280.

Armand, A., Coutts, A., Vicente, P. C., and Vilela, I. (2020). Does information break the political resource curse? experimental evidence from mozambique. *The American Economic Review*, 110(11):3431.

Asher, S., Nagpal, K., and Novosad, P. (2018). The cost of distance: Geography and governance in rural india. *World Bank Working Paper*.

Asiwaju, A. I. (1985). *Partitioned Africans: Ethnic Relations across Africa's International Boundaries*, chapter The Conceptual Framework, pages 1884–1984. New York: St. Martin's Press.

Austin, W., Carattini, S., Gomez Mahecha, J., and Pesko, M. (2020). Covid-19 mortality and contemporaneous air pollution. *CESifo Working Paper No. 8609*.

Auty, R. M. R. M. (1993). *Sustaining Development in Mineral Economies: the Resource Curse Thesis*. Routledge, London ; New York.

Bailey, M., Cao, R., Kuchler, T., Stroebel, J., and Wong, A. (2018). Social connectedness: Measurement, determinants, and effects. *Journal of Economic Perspectives*, 32(3):259–80.

Bakker, J. D., Maurer, S., Pischke, J.-S., and Rauch, F. (2018). Of mice and merchants: Trade and growth in the iron age. *Oxford Department of Economic Working Paper 854*.

Baskaran, T. and Blesse, S. (2019). Subnational border reforms and economic development in africa. *ZEW Discussion Paper No. 18-027*.

Basten, C. and Betz, F. (2013). Beyond work ethic: Religion, individual, and political preferences. *American Economic Journal: Economic Policy*, 5(3):67–91.

Bates, R. H. (1981). *Markets and States in Tropical Africa: The Political Basis of Agricultural Policy*. University of California Press, Berkeley.

Bazillier, R. and Girard, V. (2020). The gold digger and the machine. evidence on the distributive effect of the artisanal and industrial gold rushes in burkina faso. *Journal of Development Economics*, 143:102411.

Benitez, J., Courtemanche, C., and Yelowitz, A. (2020). Racial and ethnic disparities in covid-19: evidence from six large cities. *Journal of Economics, Race, and Policy*, 3(4):243–261.

Benshaul-Tolonen, A. (2019). Local industrial shocks and infant mortality. *The Economic Journal*, 129(620):1561–1592.

Berman, N., Couttenier, M., Rohner, D., and Thoenig, M. (2017). This mine is mine! how minerals fuel conflicts in africa. *The American Economic Review*, 107(6):1564–1610.

Besley, T. and Burgess, R. (2002). The political economy of government responsiveness: Theory and evidence from India. *The Quarterly Journal of Economics*, 117(4):1415–1451.

Besley, T. and Ghatak, M. (2006). Public goods and economic development. In *Understanding Poverty*. Oxford University Press.

Besley, T. and Reynal-Querol, M. (2014). The legacy of historical conflict: Evidence from Africa. *The American Political Science Review*, 108(2):319–336.

Bhadra, A., Mukherjee, A., and Sarkar, K. (2021). Impact of population density on covid-19 infected and mortality rate in india. *Modeling Earth Systems and Environment*, 7(1):623–629.

Bhattacharyya, S. and Hodler, R. (2010). Natural resources, democracy and corruption. *European Economic Review*, 54(4):608–621.

Billing, T. (2019). Government fragmentation, administrative capacity, and public goods: The negative consequences of reform in burkina faso. *Political Research Quarterly*, 72(3):669–685.

Bondy, M., Roth, S., and Sager, L. (2020). Crime is in the air: The contemporaneous relationship between air pollution and crime. *Journal of the Association of Environmental and Resource Economists*, 7(3):555–585.

Borsati, M., Nocera, S., and Percoco, M. (2022). Questioning the spatial association between the initial spread of covid-19 and transit usage in italy. *Research in Transportation Economics*, page 101194.

Bosker, M. and de Ree, J. (2014). Ethnicity and the spread of civil war. *Journal of Development Economics*, 108(C):206–221.

Bosker, M. and Garretsen, H. (2012). Economic geography and economic development in Sub-Saharan Africa. *The World Bank Economic Review*, 26(3):443–485.

Boxell, L. (2019). Droughts, conflict, and the african slave trade. *Journal of Comparative Economics*, 47(4):774–791.

Brotherhood, L., Cavalcanti, T., Da Mata, D., and Santos, C. (2022). Slums and pandemics. *Journal of Development Economics*, 157:102882.

Burchfield, M., Overman, H. G., Puga, D., and Turner, M. A. (2006). Causes of sprawl: A portrait from space. *The Quarterly Journal of Economics*, 121(2).

Burgess, R., Jedwab, R., Miguel, E., Morjaria, A., and Padró i Miquel, G. (2015). The value of democracy: Evidence from road building in kenya. *American Economic Review*, 105(6):1817–51.

Busso, M., Gregory, J., and Kline, P. (2013). Assessing the incidence and efficiency of a prominent place based policy. *American Economic Review*, 103(2):897–947.

Buys, P., Deichmann, U., and Wheeler, D. (2010). Road network upgrading and overland trade expansion in sub-saharan africa. *Journal of African Economies*, 19(3):399–432.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2):238–249.

Campante, F. R. and Do, Q.-A. (2014). Isolated capital cities, accountability, and corruption: Evidence from US states. *The American Economic Review*, 104(8):2456–2481.

Campante, F. R., Do, Q.-A., and Guimaraes, B. V. (2019). Capital cities, conflict, and misgovernance. *American Economic Journal: Applied Economics*, 11(3):298–337.

Carozzi, F. and Roth, S. (2020). Dirty density: Air quality and the density of american cities. *IZA Discussion Paper*.

Caselli, F. and Michaels, G. (2013). Do oil windfalls improve living standards? evidence from brazil. *American Economic Journal. Applied Economics*, 5(1):208–238.

Cederman, L.-E., Hug, S., Schädel, A., and Wucherpfennig, J. (2015). Territorial autonomy in the shadow of conflict: Too little, too late? *American Political Science Review*, 109(3):635–635.

Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The Effect of Minimum Wages on Low-Wage Jobs. *The Quarterly Journal of Economics*, 134(3):1405–1454.

Chaisemartin, C. d. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *The American Economic Review*, 110(9):2964–2996.

Chay, K. Y. and Greenstone, M. (2003). The impact of air pollution on infant mortality: Evidence from geographic variation in pollution shocks induced by a recession. *The Quarterly Journal of Economics*, 118(3):1121–1167.

Chen, X. and Nordhaus, W. D. (2011). Using luminosity data as a proxy for economic statistics. *Proceedings of the National Academy of Sciences of the United States of America*, 108(21):8589–8594.

Chiovelli, G., Michalopoulos, S., and Papaioannou, E. (2018). Landmines and spatial development. Working Paper 24758, National Bureau of Economic Research.

Chowell, G., Bettencourt, L. M., Johnson, N., Alonso, W. J., and Viboud, C. (2008). The 1918–1919 influenza pandemic in england and wales: spatial patterns in transmissibility and mortality impact. *Proceedings of the Royal Society B: Biological Sciences*, 275(1634):501–509.

Chuhan-Pole, P., Dabalen, A. L., and Land, B. C. (2016). *Mining in Africa: Are Local Communities Better Off?* Africa Development Forum. World Bank Publications, Washington.

Ciccone, A. and Hall, R. E. (1996). Productivity and the density of economic activity. *The American Economic Review*, pages 54–70.

Ciencewicki, J. and Jaspers, I. (2007). Air pollution and respiratory viral infection. *Inhalation Toxicology*, 19(14):1135–1146.

Clay, K., Lewis, J., and Severnini, E. (2018). Pollution, infectious disease, and mortality: Evidence from the 1918 spanish influenza pandemic. *The Journal of Economic History*, 78(4):1179–1209.

Cogneau, D. and Dupraz, Y. (2014). Questionable inference on the power of pre-colonial institutions in Africa. *Paris School of Economics Working Paper*.

Colella, F., Lalive, R., Sakalli, S. O., and Thoenig, M. (2019). Inference with arbitrary clustering. *IZA Discussion Paper No. 12584*.

Collier, P. and Hoeffler, A. (2009). Testing the neocon agenda: Democracy in resource-rich societies. *European Economic Review*, 53(3):293 – 308.

Combes, P.-P., Duranton, G., and Gobillon, L. (2011). The identification of agglomeration economies. *Journal of economic geography*, 11(2):253–266.

Combes, P.-P. and Gobillon, L. (2015). The empirics of agglomeration economies. In *Handbook of regional and urban economics*, volume 5, pages 247–348. Elsevier.

Combes, P.-P. and Lafourcade, M. (2005). Transport costs: Measures, determinants, and regional policy implications for France. *Journal of Economic Geography*, 5(3):319–349.

Conley, T. (1999). Gmm estimation with cross sectional dependence. *Journal of Econometrics*, 92(1):1–45.

Corden, W. M. and Neary, J. P. (1982). Booming sector and de-industrialisation in a small open economy. *The Economic Journal*, 92(368):825–848.

Correia, S., Guimarães, P., and Zylkin, T. (2020). Fast poisson estimation with high-dimensional fixed effects. *The Stata Journal*, 20(1):95–115.

Currie, J. and Neidell, M. (2005). Air pollution and infant health: What can we learn from california's recent experience? *The Quarterly Journal of Economics*, 120(3):1003–1030.

Cutler, D. M. and Summers, L. H. (2020). The covid-19 pandemic and the $16 trillion virus. *Journal of the American Medical Association (JAMA)*, 324(15):1495–1496.

Deaton, A. (1999). Commodity prices and growth in africa. *The Journal of Economic Perspectives*, 13(3):23–40.

Dell, M. (2010). The persistent effects of peru's mining "mita". *Econometrica*, 78(6):1863–1903.

Dell, M., Lane, N., and Querubin, P. (2018). The historical state, local collective action, and economic development in vietnam. *Econometrica*, 86(6):2083–2121.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee.

Diamond, J. (1997). Germs, guns and steel: The fates of human societies.

Ding, W., Levine, R., Lin, C., and Xie, W. (2020). Social distancing and social capital: Why us counties respond differently to covid-19. Technical report, National Bureau of Economic Research.

Dittmar, J. E. and Meisenzahl, R. R. (2019). Public goods institutions, human capital, and growth: Evidence from German history. *The Review of Economic Studies*. rdz002.

Dockery, D. W., Pope, C. A., Xu, X., Spengler, J. D., Ware, J. H., Fay, M. E., Ferris, B. G., and Speizer, F. E. (1993). An association between air pollution and mortality in six u.s. cities. *The New England Journal of Medicine*, 329(24):1753–1759.

Donaldson, D. and Hornbeck, R. (2016). Railroads and American economic growth: A "market access" approach. *The Quarterly Journal of Economics*, 131(2):799.

Donaldson, D. and Storeygard, A. (2016). The view from above: Applications of satellite data in economics. *Journal of Economic Perspectives*, 30(4):171–98.

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

Dreher, A., Fuchs, A., Hodler, R., Parks, B. C., Raschky, P. A., and Tierney, M. J. (2019). African leaders and the geography of China's foreign assistance. *Journal of Development Economics*, 140:44–71.

Dube, O. and Vargas, J. F. (2013). Commodity price shocks and civil conflict: Evidence from colombia. *The Review of Economic Studies*, 80(4 (285)):1384–1421.

Dubner, S. J. (2020). What does covid-19 mean for cities (and marriages)? *Freakonomics Podcast Ep.401*.

Duranton, G. and Puga, D. (2020). The economics of urban density. *Journal of Economic Perspectives*.

Duranton, G. and Turner, M. A. (2018). Urban form and driving: Evidence from us cities. *Journal of Urban Economics*, 108:170–191.

Ehlert, A. (2021). The socio-economic determinants of covid-19: a spatial analysis of german county level data. *Socio-Economic Planning Sciences*, page 101083.

Elvidge, C., E Baugh, K., Zhizhin, M., and Hsu, F.-C. (2013). Why viirs data are superior to DMSP for mapping nighttime lights. *Proceedings of the Asia-Pacific Advanced Network*, 35:62–69.

Elvidge, C. D., Baugh, K., Zhizhin, M., Hsu, F. C., and Ghosh, T. (2017). Viirs night-time lights. *International Journal of Remote Sensing*, 38(21):5860–5879.

Eubank, N. (2012). Taxation, political accountability and foreign aid: Lessons from Somaliland. *The Journal of Development Studies*, 48(4):465–480.

Ferguson, N. M., Cummings, D. A., Fraser, C., Cajka, J. C., Cooley, P. C., and Burke, D. S. (2006). Strategies for mitigating an influenza pandemic. *Nature*, 442(7101):448–452.

Florida, R., Rodríguez-Pose, A., and Storper, M. (2021). Cities in a post-covid world. *Urban Studies*, page 00420980211018072.

Francois, P., Rainer, I., and Trebbi, F. (2015). How is power shared in Africa? *Econometrica*, 83(2):465–503.

Galor, O. and Özak, Ö. (2016). The agricultural origins of time preference. *American Economic Review*, 106(10):3064–3103.

Garrett, T. A. (2007). Economic effects of the 1918 influenza pandemic. *Federal Reserve Bank of St. Louis*.

Gennaioli, N. and Rainer, I. (2007). The modern impact of precolonial centralization in Africa. *Journal of Economic Growth*, 12(3):185–234.

Ghebrihiwet, N. (2019). Fdi technology spillovers in the mining industry: Lessons from south africa's mining sector. *Resources Policy*, 62(August):463–471.

Gibbons, S. and Wu, W. (2017). Airports, market access and local economic performance: Evidence from China. SERC Discussion Papers 211. Spatial Economics Research Centre, London School of Economics and Political Science, London, UK.

Glaeser, E. L. (1999). Learning in cities. *Journal of Urban Economics*, 46(2):254–277.

Glaeser, E. L., Gorback, C., and Redding, S. J. (2020). Jue insight: How much does covid-19 increase with mobility? evidence from new york and four other us cities. *Journal of Urban Economics*, page 103292.

Glaeser, E. L. and Kahn, M. E. (2004). Sprawl and urban growth. In *Handbook of Regional and Urban Economics*, volume 4, pages 2481–2527. Elsevier.

Glaeser, E. L., Kolko, J., and Saiz, A. (2001). Consumer city. *Journal of Economic Geography*, 1(1):27–50.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*.

Graff Zivin, J., Neidell, M. J., Sanders, N. J., and Singer, G. (2021). When externalities collide: Influenza and pollution. *National Bureau of Economic Research No. 27982*.

Green, E. (2010). Patronage, district creation, and reform in uganda. *Studies in Comparative International Development*, 45(1):83–103.

Grossman, G. and Lewis, J. I. (2014). Administrative unit proliferation. *The American Political Science Review*, 108(1):196–217.

Grossman, G., Pierskalla, J. H., and Boswell Dean, E. (2017). Government fragmentation and public goods provision. *The Journal of Politics*, 79(3):823–840.

Guriev, S., Melnikov, N., and Zhuravskaya, E. (2020). 3g internet and confidence in government*. *The Quarterly Journal of Economics*.

Hamman, M. K. (2021). Disparities in covid-19 mortality by county racial composition and the role of spring social distancing measures. *Economics & Human Biology*, 41:100953.

Harris, C. D. (1954). The market as a factor in the localization of industry in the united states. *Annals of the Association of American Geographers*, 44(4):315.

Hassan, M. (2016). A state of change: District creation in kenya after the beginning of multi-party elections. *Political Research Quarterly*, 69(3):510–521.

He, K., Zhang, X., Ren, S., and Sun, J. (2016). Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pages 630–645. Springer.

Henderson, J. V., Squires, T., Storeygard, A., and Weil, D. (2017). The global distribution of economic activity: Nature, history, and the role of trade. *The Quarterly Journal of Economics*, 133(1):357–406.

Henderson, J. V., Storeygard, A., and Weil, D. N. (2012). Measuring economic growth from outer space. *The American Economic Review*, 102(2):994–1028.

Henn, S. J. (2021). Complements or substitutes? how institutional arrangements bind traditional authorities and the state in africa. *Unpublished Working Paper: https://bfi.uchicago.edu/wp-content/uploads/2021/08/Henn_JMP_Chiefs.pdf*.

Herbst, J. (2000). *States and power in africa: Comparative lessons in authority and control*. Princeton University Press.

Hirschman, A. O. (1958). *The Strategy of Economic Development*. Yale Studies in Economics ; 10. Yale University Press, New Haven.

Hodler, R. (2006). The curse of natural resources in fractionalized countries. *European Economic Review*, 50(6):1367 – 1386.

Hodler, R. and Raschky, P. A. (2014). Regional favoritism. *Quarterly Journal of Economics*, 129(2):995–1033.

International Monetary Fund (2015). Sub-Saharan Africa: Dealing with the gathering clouds. *Regional Economic Outlook*.

Isbell, T. A. (2017). Data codebook for round 6 afrobarometer survey. *Afrobarometer*.

Isphording, I. E. and Pestel, N. (2021). Pandemic meets pollution: Poor air quality increases deaths by covid-19. *Journal of Environmental Economics and Management*, 108:102448–102448.

Ivanyna, M. and Shah, A. (2014). How close is your government to its people? worldwide indicators on localization and decentralization. *Economics*, 8(1):20140003.

Jans, J., Johansson, P., and Nilsson, J. P. (2018). Economic status, air quality, and child health: Evidence from inversion episodes. *Journal of Health Economics*, 61:220–232.

Jarvis, A., Reuter, H., Nelson, A., and Guevara, E. (2008). Hole-filled SRTM for the globe version 4. In *CGIAR Consortium for Spatial Information (CGIAR-CSI)*. CGIAR, Washington, United States.

Jedwab, R., Kerby, E., and Moradi, A. (2017). History, path dependence and development: Evidence from colonial railways, settlers and cities in kenya. *The Economic Journal*, 127(603):1467–1494.

Jedwab, R. and Moradi, A. (2016). The permanent effects of transportation revolutions in poor countries: Evidence from africa. *The Review of Economics and Statistics*, 98(2):268–284.

Jedwab, R. and Storeygard, A. (2020). The average and heterogeneous effects of transportation investments: Evidence from sub-saharan africa 1960-2010. Working Paper 27670, National Bureau of Economic Research.

Johnson, J. and Thyne, C. L. (2018). Squeaky wheels and troop loyalty: How domestic protests influence coups d'état, 1951–2005. *Journal of Conflict Resolution*, 62(3):597–625.

Kamis, C., Stolte, A., West, J. S., Fishman, S. H., Brown, T., Brown, T., and Farmer, H. R. (2021). Overcrowding and covid-19 mortality across us counties: Are disparities growing over time? *SSM-population health*, 15:100845.

Kanbur, R. and Venables, A. J. (2005). Spatial inequality and development. *Journal of Economic Geography*, 5(1).

Karger, D. N., Conrad, O., Böhner, J., Kawohl, T., Kreft, H., Soria-Auza, R. W., Zimmermann, N. E., Linder, H. P., and Kessler, M. (2017). Climatologies at high resolution for the earth's land surface areas. *Scientific Data*, 4:170122.

Kim, H., Zanobetti, A., and Bell, M. L. (2021). Temporal transition of racial/ethnic disparities in covid-19 outcomes in 3108 counties of the united states: three phases from january to december, 2020. *Science of The Total Environment*, page 148167.

Knutsen, C. H., Kotsadam, A., Olsen, E. H., and Wig, T. (2017). Mining and local corruption in africa. *American Journal of Political Science*, 61(2):320–334.

Konrad, C. P. (2003). Effects of urban development on floods.

Krieger, N., Waterman, P. D., and Chen, J. T. (2020). Covid-19 and overall mortality inequities in the surge in death rates by zip code characteristics: Massachusetts, january 1 to may 19, 2020. *American Journal of Public Health*, 110(12):1850–1852.

Krugman, P. (1987). The narrow moving band, the dutch disease, and the competitive consequences of mrs. thatcher: Notes on trade in the presence of dynamic scale economies. *Journal of Development Economics*, 27(1):41–55.

Kuchler, T., Russel, D., and Stroebel, J. (2020). The geographic spread of covid-19 correlates with structure of social networks as measured by facebook. Technical report, National Bureau of Economic Research.

Lagakos, D. (2020). Urban-rural gaps in the developing world: Does internal migration offer opportunities? *Journal of Economic Perspectives*, 34(3):174–92.

Lane, P. R. and Tornell, A. (1996). Power, growth, and the voracity effect. *Journal of Economic Growth*, 1(2):213–241.

Lewis, B. D. (2017). Does local government proliferation improve public service delivery? evidence from indonesia. *Journal of Urban Affairs*, 39(8):1047–1065.

Lippert, A. B. (2014). Spill-Overs of a Resource Boom: Evidence from Zambian Copper Mines. *OxCarre Research Paper*, 131.

Lipton, M. (1977). *Why Poor People Stay Poor: A Study of Urban Bias in World Development*. London: Temple Smith.

Maina, J., Ouma, P. O., Macharia, P. M., Alegana, V. A., Mitto, B., Fall, I. S., Noor, A. M., Snow, R. W., and Okiro, E. A. (2019). A spatial database of health facilities managed by the public health sector in sub saharan africa. *Scientific Data*, 6(1):134.

Mamo, N., Bhattacharyya, S., and Moradi, A. (2019). Intensive and extensive margins of mining and development: Evidence from sub-saharan africa. *Journal of Development Economics*, 139:28 – 49.

Manacorda, M. and Tesei, A. (2020). Liberation technology: Mobile phones and political mobilization in africa. *Econometrica*, 88(2):533–567.

Marshall, A. (1920). *Principles of Economics*. London: Macmillan.

Marshall, M. G., Gurr, T. R., and Jaggers, K. (2017). Polity IV Project: Political regime characteristicsand transitions, 1800-2017. *Center for Systemic Peace: Vienna, VA.*

Martelletti, L. and Martelletti, P. (2020). Air pollution and the novel covid-19 disease: A putative disease risk factor. *SN Comprehensive Clinical Medicine*, 2(4):383–387.

Maus, V., Giljum, S., Gutschlhofer, J., da Silva, D. M., Probst, M., Gass, S. L., Luckeneder, S., Lieber, M., and McCallum, I. (2020). A global-scale data set of mining areas. *Scientific Data*, 7(1):1–13.

McFarlane, C. (2021). Repopulating density: Covid-19 and the politics of urban value. *Urban Studies*, page 00420980211014810.

McMillan, M., Rodrik, D., and Verduzco-Gallo, I. n. (2014). Globalization, structural change, and productivity growth, with an update on africa. *World Development*, 63:11–32.

Mehlum, H., Moene, K., and Torvik, R. (2006). Institutions and the resource curse. *The Economic Journal*, 116(508):1–20.

Michalopoulos, S. and Papaioannou, E. (2013). Pre-colonial ethnic institutions and contemporary african development. *Econometrica*, 81(1):113–152.

Michalopoulos, S. and Papaioannou, E. (2014). National institutions and subnational development in Africa. *The Quarterly Journal of Economics*, 129(1):151–213.

Michalopoulos, S. and Papaioannou, E. (2016). The long-run effects of the scramble for Africa. *American Economic Review*, 106(7):1802–48.

Michalopoulos, S. and Papaioannou, E. (2020). Historical legacies and african development. *Journal of Economic Literature*, 58(1):53–128.

Michalopoulos, S., Putterman, L., and Weil, D. N. (2019). The influence of ancestral lifeways on individual economic outcomes in sub-saharan africa. *Journal of the European Economic Association*, 17(4):1186–1231.

Mills, C. E., Robins, J. M., and Lipsitch, M. (2004). Transmissibility of 1918 pandemic influenza. *Nature*, 432(7019):904–906.

Molinari, N.-A. M., Ortega-Sanchez, I. R., Messonnier, M. L., Thompson, W. W., Wortley, P. M., Weintraub, E., and Bridges, C. B. (2007). The annual impact of seasonal influenza in the us: Measuring disease burden and costs. *Vaccine*, 25(27):5086–5096.

Müller-Crepon, C., Hunziker, P., and Cederman, L.-E. (2021). Roads to rule, roads to rebel: Relational state capacity and conflict in africa. *Journal of Conflict Resolution*, 65(2-3):563–590. PMID: 33487734.

Murdock, G. P. (1959). *Africa: Its Peoples and their Culture History*. New York: McGraw-Hill.

Nathan, M. and Overman, H. (2020). Will coronavirus cause a big city exodus? *Environment and Planning B: Urban Analytics and City Science*, 47(9):1537–1542.

Nunn, N. (2008). The long-term effects of Africa's slave trades. *The Quarterly Journal of Economics*, 123(1):139–176.

Nunn, N. and Puga, D. (2012). Ruggedness: The blessing of bad geography in africa. *The Review of Economics and Statistics*, 94(1):20–36.

Nunn, N. and Wantchekon, L. (2011). The slave trade and the origins of mistrust in Africa. *American Economic Review*, 101(7):3221–3252.

Oates, W. E. (1972). *Fiscal Federalism*. Harcourt Brace Jovanovich, New York.

Odusola, A., Cornia, G. A., Bhorat, H., and ao, P. C. (2017). *Income Inequality Trends in Sub-Saharan Africa: Divergence, Determinants and Consequences*. United Nations Development Programme Regional Bureau for Africa.

Papageorge, N. W., Zahn, M. V., Belot, M., van den Broek-Altenburg, E., Choi, S., Jamison, J. C., Tripodi, E., et al. (2020). Socio-demographic factors associated with self-protecting behavior during the covid-19 pandemic. Technical report, Institute of Labor Economics (IZA).

Pequeno, P., Mendel, B., Rosa, C., Bosholn, M., Souza, J. L., Baccaro, F., Barbosa, R., and Magnusson, W. (2020). Air transportation, population density and temperature predict the spread of covid-19 in brazil. *PeerJ*, 8:e9322.

Persico, C. L. and Johnson, K. R. (2021). The effects of increased pollution on covid-19 cases and deaths. *Journal of Environmental Economics and Management*, 107:102431.

Pierskalla, J. H. (2016). Splitting the difference? the politics of district creation in indonesia. *Comparative Politics*, 48(2):249–268.

Pokorny, B., Lübke, C. v., Dayamba, S. D., and Dickow, H. (2019). All the gold for nothing? impacts of mining on rural livelihoods in northern burkina faso. *World Development*, 119:23.

Prata, D. N., Rodrigues, W., and Bermejo, P. H. (2020). Temperature significantly changes covid-19 transmission in (sub) tropical cities of brazil. *Science of the Total Environment*, page 138862.

Raleigh, C., Linke, A., Hegre, H., and Karlsen, J. (2010). Introducing ACLED: An armed conflict location and event dataset. *Journal of Peace Research*, 47(5):651–660.

Rappaport, J. (2008). A productivity model of city crowdedness. *Journal of Urban Economics*, 63(2):715–722.

Ray, D. and Esteban, J. (2017). Conflict and development. *Annual Review of Economics*, 9:263–293.

Redding, S. J. and Sturm, D. M. (2008). The costs of remoteness: Evidence from German division and reunification. *American Economic Review*, 98(5):1766–1797.

Richard Bluhm, Christian Lessmann, P. S. (2021). The political geography of cities. *CESifo Working Paper No. 9376*.

Robinson, J. A., Torvik, R., and Verdier, T. (2006). Political foundations of the resource curse. *Journal of Development Economics*, 79(2):447–468.

Rodríguez-Pose, A. and Burlina, C. (2021). Institutions and the uneven geography of the first wave of the covid-19 pandemic. *Journal of Regional Science*.

Ronneberger, O., Fischer, P., and Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer.

Rosner, B. (1983). Percentage points for a generalized esd many-outlier procedure. *Technometrics*, 25(2):165–172.

Sachs, J. D. and Warner, A. M. (2001). The curse of natural resources. *European Economic Review*, 45(4):827 – 838. 15th Annual Congress of the European Economic Association.

Sager, L. (2019). Estimating the effect of air pollution on road safety using atmospheric temperature inversions. *Journal of Environmental Economics and Management*, 98:102250.

Sant'Anna, P. H. and Zhao, J. (2020). Doubly robust difference-in-differences estimators. *Journal of Econometrics*, 219(1):101–122.

Schlenker, W. and Walker, W. R. (2016). Airports, air pollution, and contemporaneous health. *The Review of Economic Studies*, 83(2 (295)):768–809.

Serneels, P. and Verpoorten, M. (2015). The impact of armed conflict on economic performance: Evidence from Rwanda. *Journal of Conflict Resolution*, 59(4):555–592.

Silva, J. M. C. S. and Tenreyro, S. (2006). The log of gravity. *The Review of Economics and Statistics*, 88(4):641–658.

Simonovska, I. and Waugh, M. E. (2014). The elasticity of trade: Estimates and evidence. *Journal of International Economics*, 92(1):34–50.

Storeygard, A. (2016). Farther on down the road: Transport costs, trade and urban growth in sub-saharan africa. *The Review of Economic Studies*, 83(3 (296)):1263–1295.

Strömberg, D. (2004). Radio's impact on public spending. *The Quarterly Journal of Economics*, 119(1):189–221.

Subbaraman, N. (2020). Why daily death tolls have become unusually important in understanding the coronavirus pandemic. *Nature*.

Sundberg, R. and Melander, E. (2013). Introducing the ucdp georeferenced event dataset. *Journal of Peace Research*, 50(4):523–532.

Teller, J. (2021). Urban density and covid-19: towards an adaptive approach. *Buildings and Cities*, 2(1):150–165.

Tobías, A. and Molina, T. (2020). Is temperature reducing the transmission of covid-19? *Environmental Research*, 186:109553.

Tornell, A. and Lane, P. R. (1999). The voracity effect. *The American Economic Review*, 89(1):22–46.

Torvik, R. (2002). Natural resources, rent seeking and welfare. *Journal of Development Economics*, 67(2):455 – 470.

United Nations Department of Economic and Social Affairs Population Division (2018). World urbanization prospects: The 2018 revision. *Online Edition*.

van der Ploeg, F. (2011). Natural resources: Curse or blessing? *Journal of Economic Literature*, 49(2):366–420.

van Wijnbergen, S. (1984). The 'dutch disease': A disease after all? *The Economic Journal*, 94(373):41–55.

Venables, A. J. (2016). Using natural resources for development: Why has it proven so difficult? *The Journal of Economic Perspectives*, 30(1):161–183.

Vicente, P. C. (2010). Does oil corrupt? evidence from a natural experiment in west africa. *Journal of Development Economics*, 92(1):28–38.

Vogt, M., Bormann, N.-C., Rüegger, S., Cederman, L.-E., Hunziker, P., and Girardin, L. (2015). Integrating data on ethnicity, geography, and conflict: The ethnic power relations data set family. *Journal of Conflict Resolution*, 59(7):1327–1342.

Voigtländer, N. and Voth, H.-J. (2013). The three horsemen of riches: Plague, war, and urbanization in early modern europe. *Review of Economic Studies*, 80(2):774–811.

von der Goltz, J. and Barnwal, P. (2019). Mines: The local wealth and health effects of mineral mining in developing countries. *Journal of Development Economics*, 139(C):1–16.

Wheaton, W. C. and Kinsella Thompson, A. (2020). The geography of covid-19 growth in the us: Counties and metropolitan areas. *Available at SSRN 3570540*.

Whittle, R. S. and Diaz-Artiles, A. (2020). An ecological study of socioeconomic predictors in detection of covid-19 cases across neighborhoods in new york city. *BMC Medicine*, 18(1):1–17.

Williamson, J. G. (1965). Regional inequality and the process of national development: A description of the patterns. *Economic development and cultural change*, 13(4, Part 2):1–84.

Wooldridge, J. M. (2015). Control function methods in applied econometrics. *The Journal of Human Resources*, 50(2):420–445.

Wucherpfennig, J. (2011). Politically relevant ethnic groups across space and time: Introducing the GeoEPR dataset 1. *Conflict Management & Peace Science*, 28(5):423–438.

Yeh, C., Perez, A., Driscoll, A., Azzari, G., Tang, Z., Lobell, D., Ermon, S., and Burke, M. (2020). Using Publicly Available Satellite Imagery and Deep Learning to Understand Economic Well-Being in Africa. *Nature Communications*, 11(1).

Young, A. (2013). Inequality, the urban-rural gap, and migration. *The Quarterly Journal of Economics*, 128(4):1727–1786.

Zhang, C. H. and Schwartz, G. G. (2020). Spatial disparities in coronavirus incidence and mortality in the united states: an ecological analysis as of may 2020. *The Journal of Rural Health*, 36(3):433–445.

Zhou, F., Yu, T., Du, R., Fan, G., Liu, Y., Liu, Z., Xiang, J., Wang, Y., Song, B., Gu, X., et al. (2020). Clinical course and risk factors for mortality of adult inpatients with covid-19 in wuhan, china: a retrospective cohort study. *The Lancet*.