

DISCUSSION PAPER SERIES

DP15272

REASSESSING THE RESOURCE CURSE USING CAUSAL MACHINE LEARNING

Roland Hodler, Michael Lechner and Paul A.
Raschky

**DEVELOPMENT ECONOMICS
INTERNATIONAL TRADE AND REGIONAL ECONOMICS
MACROECONOMICS AND GROWTH**



REASSESSING THE RESOURCE CURSE USING CAUSAL MACHINE LEARNING

Roland Hodler, Michael Lechner and Paul A. Raschky

Discussion Paper DP15272
Published 10 September 2020
Submitted 08 September 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Development Economics
- International Trade and Regional Economics
- Macroeconomics and Growth

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Roland Hodler, Michael Lechner and Paul A. Raschky

REASSESSING THE RESOURCE CURSE USING CAUSAL MACHINE LEARNING

Abstract

We reassess the effects of natural resources on economic development and conflict, applying a causal forest estimator and data from 3,800 Sub-Saharan African districts. We find that, on average, mining activities and higher world market prices of locally mined minerals both increase economic development and conflict. Consistent with the previous literature, mining activities have more positive effects on economic development and weaker effects on conflict in places with low ethnic diversity and high institutional quality. In contrast, the effects of changes in mineral prices vary little in ethnic diversity and institutional quality, but are non-linear and largest at relatively high prices.

JEL Classification: C21, O13, O55, Q34, R12

Keywords: resource curse, mining, economic development, conflict, Causal machine learning, Africa

Roland Hodler - roland.hodler@unisg.ch
Universität St Gallen and CEPR

Michael Lechner - michael.lechner@unisg.ch
University of St Gallen and CEPR

Paul A. Raschky - paul.raschky@monash.edu
Monash University

Acknowledgements

Roland Hodler and Paul Raschky gratefully acknowledge financial support from the Australian Research Council (ARC Discovery Grant DP150100061). Michael Lechner gratefully acknowledges financial support from a grant of the NRP 75 of the Swiss Science Foundation (grant SNSF 407540_166999).

Reassessing the Resource Curse using Causal Machine Learning

Roland Hodler, Michael Lechner, Paul A. Raschky*

September 2020

Abstract: We reassess the effects of natural resources on economic development and conflict, applying a causal forest estimator and data from 3,800 Sub-Saharan African districts. We find that, on average, mining activities and higher world market prices of locally mined minerals both increase economic development and conflict. Consistent with the previous literature, mining activities have more positive effects on economic development and weaker effects on conflict in places with low ethnic diversity and high institutional quality. In contrast, the effects of changes in mineral prices vary little in ethnic diversity and institutional quality, but are non-linear and largest at relatively high prices.

Keywords: Resource curse, mining, economic development, conflict, causal machine learning, Africa.

JEL classification: C21, O13, O55, Q34, R12.

* Hodler: Department of Economics and SIAW-HSG, University of St.Gallen; CEPR, London; CESifo, Munich; OxCarre, University of Oxford; email: roland.hodler@unisg.ch. Lechner: Department of Economics and SEW-HSG, University of St. Gallen; CEPR, London; CESifo, Munich; IAB, Nuremberg; IZA, Bonn; RWI, Essen; email: michael.lechner@unisg.ch. Raschky: Department of Economics and SoDa Labs, Monash University. Email: paul.raschky@monash.edu. Roland Hodler and Paul Raschky gratefully acknowledge financial support from the Australian Research Council (ARC Discovery Grant DP150100061). Michael Lechner gratefully acknowledges financial support from a grant of the NRP 75 of the Swiss Science Foundation (grant SNSF 407540_166999).

1 Introduction

The question of whether natural resources are a curse or a blessing has attracted a lot of attention, and the literature has come a long way. The early contributions presented simple cross-country correlations (e.g., Sachs and Warner, 1995, 2001), while current state-of-the-art contributions, such as Berman, Couttenier, Rohner and Thoenig (2017) or Mamo, Bhattacharyya and Moradi (2019), focus on subnational units and exploit intertemporal variation in local mining activity or in the international market prices of locally mined minerals. Despite the many changes in the preferred estimation approaches over time, there are a few things that have not changed in this literature: First, the reliance on linear specifications. Second, the limited efforts taken to study the heterogeneity underlying the effects of natural resources. Third, the reliance on a single outcome per paper, often economic development, or conflict.¹

In this paper, we apply causal machine learning to study the effects of natural resources on economic development and conflict in Africa. Our sample consists of 3,800 districts (ADM2 regions) from 42 Sub-Saharan African countries and the years from 1996 to 2013. We divide the observations into a control group, consisting of all the districts that never had any mining activity during the sample period, and a few discrete treatment groups that differ by the relative prices of the locally mined minerals. Our outcome variables are the log of night-time light intensity as proxy for district-level economic development, and a binary variable indicating the occurrence of conflict events within districts. We apply the causal machine learning estimator by Lechner (2018), which builds upon the one proposed by Wager and Athey (2018). This estimator allows us to first estimate the highly disaggregated individualized average treatment effects (IATEs) and then to compute the average treatment effect (ATEs) as well as group

¹ A notable exception to this third characteristic is the study by Adhvaryu, Fenske, Khanna and Nyshadham (2018), who test their model's predictions about the effects of natural resources on economic development and conflict (see Section 2.2).

average treatment effects (GATEs). By building groups based on a limited number of discretized variables of interest, the GATEs allow studying effect heterogeneity along the dimension of these variables.

To the best of our knowledge, we are the first to apply causal machine learning in the study of economic growth and comparative development. More specifically, we make several methodological contributions to the literature on the effects of natural resources on economic development and conflict: First, we use a non-parametric approach rather than a linear fixed-effects specification. We argue that our non-parametric approach is preferable or, at the very least, a worthwhile complementary exercise, because linear fixed-effects specifications are based on a set of assumptions that may not hold, and because we indeed find evidence for non-linear effects. Second, we explore effect heterogeneity by presenting the distributions of the IATEs and by exploring GATEs based on heterogeneity variables discussed in the previous literature. Third, we use the same approach to estimate the effects of natural resources on economic development and conflict. This joint analysis has the advantage that we know that differences in results cannot be driven by differences in the estimation approach employed. Similarly, our estimation approach offers a natural way to study the different types of “effects of natural resources” in a unified setting, namely the effects of mining activities and the effects of changes in world market prices of locally mined minerals.

Our results for the average effects (ATEs) suggest the following pattern: First, mining increases local economic development as well as the likelihood of local conflict events. Second, higher mineral prices also increase both, economic development, and the likelihood of conflict events. These first two findings make clear that there can be no easy answer to the question of whether, on average, natural resources are a curse or a blessing. Third, we find that the effects of mineral prices tend to be non-linear. The effects of a change from low to intermediate prices are typically small and statistically insignificant, while the effects of a change from intermediate

to high prices are typically much larger and statistically significant. This last finding lends additional support to our choice of a non-parametric and, therefore, non-linear estimation approach.

We show that the IATEs are quite dispersed, which underscores the importance of taking a closer look at effect heterogeneity. When grouping observations by countries, we find strong evidence for heterogeneous effects of mining activities on economic development and conflict across countries, but little evidence for heterogeneous effects of mineral prices across countries.

We then focus on two hypotheses advanced in the literature and choose heterogeneity variables that allow us to test these hypotheses. First, we group districts by two prominent country-level measures of institutional quality. Consistent with the hypothesis by Mehlum, Moene and Torvik (2006), we find that mining has more positive effects on economic development and weaker effects on conflict if institutional quality is high. Second, we test Hodler's (2006) hypothesis that natural resources have more positive effects on economic development and weaker effects on conflict if ethnic diversity is low. For that purpose, we group districts by measures of country- and district-level ethnic diversity. The emerging pattern suggests that mining has indeed the most positive effects on economic development if ethnic diversity is low. However, in contrast to Hodler's hypothesis, the effect of mining on conflict tend to be weakest at intermediate levels of ethnic diversity. Hence, while it is unclear whether mining is curse or a blessing *on average*, clearer answers are possible for individual countries and districts, depending on their institutional quality or ethnic diversity.

Interestingly, however, we find that the effects of higher international market prices of locally mined minerals (rather than mining activities per se) do not vary systematically in institutional quality and ethnic diversity. These results imply that the hypotheses by Mehlum et al. (2006) and Hodler (2006) do not extend themselves to the effects of higher mineral prices. This finding is policy relevant. The study by Mehlum et al. (2006) is frequently evoked by

international (donor) organizations to suggest that safeguards of some sorts are necessary to ensure that the exploitation of natural resources and higher resource prices do not harm weakly institutionalized countries.² The suggestion of specific safeguards in case of increasing resource prices, however, is at odds with our evidence showing that higher resource prices have similar effects in weakly and strongly institutionalized resource-extracting countries.

The remainder of the paper is structured as follows: Section 2 discusses the previous literature on the effect of natural resources on economic development and conflict as well as some of the key assumption underlying the estimation approach used in recent contributions. Section 3 presents Lechner's (2018) modified causal forest estimator and discusses the identifying assumptions. Section 4 presents the data, and section 5 the results. Section 6 briefly concludes.

2 Literature on the resource curse

2.1 Early country-level literature

The literature on the alleged curse of natural resources goes back to Sachs and Warner's (1995, 2001) finding of a negative correlation between the share of primary commodity exports in GDP and subsequent economic growth across countries. Focusing on conflict instead of growth, Collier and Hoeffler (1998, 2004) document that a higher share of primary commodity exports in GDP is typically associated with a higher risk of civil conflict. The literature has

² For example, one of the principles of the Extractive Industries Transparency Initiative (EITI) states: "We recognise that the benefits of resource extraction occur as revenue streams over many years and can be highly price dependent." (EITI provides a global standard for the good governance and practices in the natural resource sector signed by 52 countries.)

developed in many directions ever since. Here, we provide a selective review of this literature, focusing on the effects of natural resources on economic development and conflict.³

Initially, most economists attributed the negative correlation documented by Sachs and Warner to economic factors, such as Dutch disease dynamics combined with an important role of manufacturing in development (e.g., due to learning-by-doing or increasing returns to scale). Collier and Hoeffler (2004), on the other hand, interpret their pattern as evidence that natural resources increase the feasibility and attractiveness of rebellion. Some politico-economic explanations are however consistent with both findings. In particular, Mehlum et al. (2006) and Hodler (2006) present game-theoretic models in which individuals or ethnic groups decide whether to engage in productive or “fighting” activities, whereby the latter can be understood as non-productive activities ranging from rent seeking to political violence. These models both predict heterogeneous effects. The one by Mehlum et al. (2006) predicts that the effect of natural resources on economic output is increasing in institutional quality, while the effect on fighting activities is decreasing in institutional quality. Hodler’s (2006) model predicts that the effect of natural resources on economic output is decreasing in the number of ethnic groups, while the effect on fighting activities increases in this number. The authors of these two papers present cross-country correlations consistent with the predicted effect heterogeneity.

Another strand of the literature started tackling the endogeneity concerns resulting from the use of primary commodity exports in GDP. Brunnschweiler and Bulte (2008) and Alexeev and Conrad (2009) measure resource abundance by the value of a country’s mineral or oil deposits and find positive cross-country associations between such deposits and economic growth. Humphreys (2005) documents a positive correlation between oil deposits and civil conflict onset. In a next step, the literature started becoming concerned with cross-sectional

³ Other strands of the literature discuss the effects of natural resources on corruption (e.g., Ades and Di Tella, 1999; Bhattacharyya and Hodler, 2010) and democracy (e.g., Ross, 2001; Tsui, 2010; Crespo Cuaresma, Oberhofer and Raschky, 2011). See van der Ploeg (2011) for a comprehensive review of the literature.

comparisons and focused on panel data settings with country-years (instead of countries) as units of analysis. For identification, contributions to this strand mostly rely on country-fixed effects and exploit temporal variation in known resource deposits or commodity prices. Cotet and Tsui (2013) find no effect of oil discoveries on civil conflict, while Lei and Michaels (2014) find a positive effect focusing on giant oilfield discoveries. Smith (2015) finds that major resource discoveries increase GDP per capita. Exploiting variations in the international price of minerals and other commodities that a country exports, Brückner and Ciccone (2010) and Bazzi and Blattman (2014) find that higher commodity prices reduce the risk of civil conflict onset and the duration of civil conflicts, respectively.

2.2 Recent subnational studies

In recent years, there has been a trend in the literature towards moving to the subnational level (see Cust and Poelhekke, 2015, for a review). Reasons for this trend are manifold. First, there has been a remarkable increase in the availability of geo-spatial data on the location of mines, oil fields, and conflict events. Relatedly, data on nighttime light intensity, which are available at a high spatial granularity, now allow computing proxies for economic development at the subnational level (see Section 4.1 for details). Second, researchers hope that it is easier to identify causal effects when the possibility to add various fixed effects in a linear model allows exploiting variation within countries across subnational regions or even temporal variation within subnational regions. Third, researchers have become interested in how the spatial distribution of natural resources shapes the spatial distribution of conflict and economic development within countries. While some contributions to this new strand of the literature focus on a single country (e.g., Dube and Vargas, 2013; Cavalcanti, Da Mata and Toscani, 2019) or a single mine (e.g., Aragon and Rud 2013), others consider subnational regions from a large set of countries.

Here, we focus on two prominent contributions that, like us, study the effect of natural resources on conflict or economic development at the subnational level across (Sub-Saharan) Africa. Berman et al. (2017) aggregate geo-spatial data on the location of mines and conflict events by grid cells of 0.5x0.5 degrees (approximately 55x55 km at the equator) and year. The share of cells with an active mine is 0.02, and the probability of a conflict in any given cell-year is 0.06. The key variable in their study is an interaction term between a dummy variable for the presence of an active mine and the (log) price of the main mineral produced within this cell. For identification, they rely on both country-year- and cell-fixed effects. They find that an increase in this interaction term leads to an increase in civil conflict at the cell level.

Mamo et al. (2019) aggregate their data at the district level and focus on local economic development (also proxied by nighttime light intensity). They use district- and year-fixed effects and exploit different types of temporal variation within districts: intensive margin variation in the value of mineral production and extensive margin variation in whether a district is producing minerals at all. The share of mineral-producing districts is 0.04, and the share of districts with extensive margin variation even smaller. They find moderate positive effects of more valuable mineral production on economic development, but large positive effects for districts that start mining. Despite these large local effects, they find no evidence for economic spillovers to other districts.⁴

Finally, Adhvaryu et al. (2018) study the effects of natural resources in a model of strategic interactions between neighbouring regions and present evidence in support of the model's predictions. While our paper is more closely related to Berman et al. (2017) and Mamo et al. (2019) in terms of the questions asked, it shares with Adhvaryu et al. (2018) the approach

⁴ In a paper mainly focusing on spillovers, Amarasinghe, Hodler, Raschky and Zenou (2018) use nighttime light intensity across districts like Mamo et al. (2019), but exploit intertemporal price variation in a way similar to Berman et al. (2017). They find positive effects of higher mineral prices on local nighttime light intensity and evidence for spillovers to districts connected by roads or ethnic linkages.

of looking at the effects of natural resources on both, local economic development and local conflict.

2.3 Discussion of the estimation approaches used so far

The recent literature using subnational data shows substantial concerns about possible selection effects coming from unobservables at the subnational units that are jointly correlated with the local availability of resources and the outcome of interest, e.g., measures of economic development or the occurrence of conflicts. To address these concerns, most recent empirical studies in this area apply some form of two-way fixed effects (FE) model. For example, Berman et al. (2017) include country-year- and cell-fixed effects, and Mamo et al. (2019) include year- and district fixed effects. While these two-way fixed effects models are important improvements over the approaches used by the earlier literature, they still rest on rather restrictive assumptions. A recent paper by de Chaisemartin and D'Haultfœuille (2020) points out that a key assumption of a two-way FE model is that the treatment effect is constant across the individual geographic units and time. It is very plausible that this assumption is violated in the context of the effect of natural resource shocks on local economic development or the likelihood of local conflict. De Chaisemartin and D'Haultfœuille (2020) show that in the presence of heterogenous treatment effects the estimated coefficients are biased. We will investigate (and partly reject) effect homogeneity below.

In addition, the standard two-way FE models rest on a linearity assumption. Of course, any model is wrong, but the nature of the dependent variables raises concerns about the approximation quality of the linear functional form. For example, the conflict variable is usually binary with a very low share of ones, such that a logit- or probit-type approximation may provide a better approximation than a linear model. When using nighttime light intensity or the log of it as dependent variable, it is also unclear why the relationship between the fixed effects, the other variables, and the resource variables should be linear.

Given these doubts about the linear FE specifications used so far in the literature, we have decided to take an alternative, more flexible route. This alternative has many advantages, which we discuss in the next section. However, it comes at the cost of not being able to include that many fixed effects. Hence, there is a trade-off. We consider our approach to be preferable given that the assumptions required for a successful application of a fixed effects strategy are unlikely to hold. In any case, it should at the very least be a worthwhile exercise complementing current state-of-the-art contributions such as Berman et al. (2017) and Mamo et al. (2019).

3 Econometrics

3.1 The modified causal forest estimator

In this paper, we utilize the recently upcoming causal machine learning literature (see, e.g., Athey and Imbens, 2017, for an only slightly outdated survey). It combines the prediction power of the machine and statistical learning literature (see, e.g., Hastie, Tibshirani, and Friedman, 2009, for an overview) with the microeconomic literature on defining and identifying causal effects (see, e.g., Imbens and Wooldridge, 2009, for a good guide to that literature). Recently, this literature has seen a surge of proposed methods, in particular in epidemiology and econometrics. Knaus, Lechner, and Strittmatter (2018) compare many of those methods systematically with respect to their theoretical properties as well as their performance in a simulation exercise. One conclusion from this paper is that Random Forest-based estimation approaches seem to outperform many alternative estimators.

The starting point of the causal forest literature is the Causal Tree introduced in a paper by Athey and Imbens (2016). In a Causal Tree, the sample is split sequentially into smaller and smaller strata, in which the values of X become increasingly homogenous, to mitigate selection effects and to uncover effect heterogeneity. Once the splitting is terminated based on some stopping criterion, the treatment effect is computed within each stratum (called a ‘leaf’) by

computing the difference of the mean outcomes of treated and controls (possibly weighted by the conditional-on- X probabilities of being a treated or control observation). However, the literature on regression trees acknowledges that the final leaves may be rather unstable because of the sequential nature of the splits (if the first split is different, the full tree will likely lead to different final strata). A solution to this problem is the so-called Random Forest. The key idea is to induce some randomness into the tree building process, build many trees, and then average the predictions of the many trees. The induced randomness is generated by using randomly generated subsamples (or bootstrap samples) and by considering for each splitting decision only a random selection of the covariates. Wager and Athey (2018) use this idea to propose Causal Forests, which are based on a collection of Causal Trees with small final leaves.⁵ Lechner (2018) develops these ideas further by improving on the splitting rule for the individual trees, and by providing methods to estimate heterogeneous effects for a limited number of discrete policy variables (**Group Average Treatment Effects, GATE**) at low computational costs, in addition to the highly disaggregated effects the literature focussed on so far (**Individualized Average Treatment Effects, IATE**). Furthermore, he suggests a way of performing unified inference for all aggregation levels. Finally, his approach is applicable to a multiple, discrete treatment framework. Since all these advantages are important in the empirical analysis of this paper, this approach is used below. For further technical details of the estimator, the reader is referred to Lechner (2018).

3.2 Identification

To estimate the IATEs, GATEs and AETs in a non-parametric way, we discretize the treatment and divide the observations (district-years) into M treatments. We do so in two steps: First, we distinguish between mining and non-mining districts, and refer to the former as control

⁵ Athey, Tibshirani and Wager (2019) generalize this idea to many different econometric estimation problems.

or treatment 0 group. Second, we allocate all the observations from mining districts into $M-1$ approximately equally-sized treatment groups depending on the normalized current international market prices of the minerals mined in the given district (see Section 4 for details). Hence, in case of $M=3$, we separate the observations of mining districts into those with relatively high and low prices of the respective minerals. In case of $M>3$, we can investigate potentially non-linear effects. In any case, we estimate the average effects of the treatments that fall into the respective price brackets. Such a fully non-parametric approach has the advantage that it cannot be subject to misspecification.

Along with the recent subnational literature, we argue that the availability of natural resources at the district-level is mainly shaped by geology and, therefore, exogenous; and that the international market prices are also exogenous to any single district as long as we focus on natural resources extracted in many places.

Let us now discuss the identifying assumptions in more detail. The classical set of unconfoundedness assumptions consists of four parts: conditional independence, common support, stable-unit-treatment value, and exogeneity. Imbens (2000) and Lechner (2001) discuss these assumptions and their identifying power concerning various average causal effects (like the ATE). Those results are also sufficient to identify the GATE and IATE (e.g. Lechner, 2018).

The conditional independence assumption (CIA) requires that no other variables than the observable X jointly influence the treatment and the potential outcomes. The plausibility of the CIA is enhanced by the facts that we assemble data for a large set of potential confounders that have been identified in the literature, and that we control for the confounding and heterogeneity variables non-parametrically and, therefore, much more flexibly than we could in linear specifications. Firstly, a large theoretical and empirical literature has highlighted the interplay between natural resources, institutional quality and economic development or conflict (e.g.,

Mehlum et al., 2006, and the contributions cited in footnote 4). We have therefore compiled data on national and subnational institutional quality (beyond the measures used as heterogeneity variables). Secondly, ever since the seminal work by Sachs and Warner (1995, 2001), it has been fairly standard in the empirical natural resource curse literature to account for differences in the geographic and climatic factors between the units of analysis.⁶ It is also likely that location of mineral deposits is correlated with topographic features of an area, such as ruggedness or soil characteristics, which in turn can have an effect on the area's long-term economic development as well as prevalence for violent conflict (e.g., Nunn and Puga, 2012, Berman, Couttenier and Soubeyran, 2020). Variations in climate conditions, in turn, not only have a direct impact on a subnational region's economic growth but also influence the opportunity costs of fighting (e.g., Harari and La Ferrara, 2018). Finally, we include control variables for subnational levels of population and population density as well as transportation infrastructure. It is very likely that natural resource wealth attracts more people and alters local settlement patterns in the long-run. In addition, a lot of transport infrastructure in Africa is mining-related (e.g., Bonffati and Poelhekke, 2017), which can create positive spillovers for local economic development (e.g., Amarsinghe et al., 2018) or spread local conflict (e.g., Amarasinghe, Raschky, Zenou and Zhou, 2020).

The stable-unit-treatment value assumption (STUVA) requires that the observed value of the treatment does not depend on the treatment allocation of the other units. STUVA requires absence of spillovers. As discussed in Section 2.2, the evidence on spillovers is mixed: Mamo et al. (2019) find no evidence for spillovers from mining, while Amarasinghe et al. (2018) find evidence for spillovers from changes in mineral prices among districts connected by roads or ethnic linkages.

⁶ Recent empirical work by Henderson, Squires, Storeygard and Weil (2018) highlights the importance of geography for explaining differences in economic development at the subnational level.

The exogeneity assumption requires that the observed values of the confounding and heterogeneity variables do not depend on the treatment status. The plausibility of this assumption is enhanced by focusing on the following two sets of heterogeneity and control variables: pre-determined variables that measure a district's geography or history, and country-level variables that should only marginally depend on the treatment status of a single district, given the large number of districts per country.

Finally, we did not detect any common support problems.

4 Data

Our sample consists of 3,800 districts from 481 different provinces across 42 Sub-Saharan African countries. What we call districts and provinces are ADM2 and ADM1 regions according to the GADM database of Global Administrative Areas (version 1). GADM further provides shapefiles of the corresponding regional boundaries, which we use to compute all our variables based on geospatial data. Our sample period covers 17 years, lasting from 1997 to 2013 for our outcome variables, and from 1996 to 2012 for our treatment, heterogeneity, and control variables. We choose a one-year lag to address some potential endogeneity concerns.

4.1 Outcome variables

We follow a large and growing literature in using nighttime light intensity to proxy for economic development at the subnational level.⁷ These data are based on daily measures from the Operation Linescan System (OLS) of the US Defense Meteorological Satellite Program (DMSP) and provided by the National Oceanic and Atmospheric Administration (NOAA). The NOAA uses evening observations during the dark half of the lunar cycle in seasons when the

⁷ Previous papers using nighttime light intensity to study the effect of mining activities on economic development at the subnational level include Amarasinghe et al. (2018) and Mamo et al. (2019).

sun sets early, but removes observations that are likely to be affected by fires, cloud coverage, or northern or southern lights, with the objective to report man-made nighttime light intensity. The NOAA provides annual data for the period from 1992 to 2013 for output pixels that correspond to less than one square kilometre. The data come on a scale from 0 to 63, with higher values implying more intense nighttime lights. Nighttime lights are a proxy for economic activity, as most forms of consumption and production in the evening require light. Moreover, public infrastructure is often lit at night. Henderson, Storeygard and Weil (2012) and Hodler and Raschky (2014) indeed find a high correlation between changes in nighttime light intensity and GDP at the level of countries and provinces, respectively.⁸ As an outcome variable, we use the natural logarithm of the average nighttime light pixel value in a given district and year. To avoid losing observations with a reported nighttime light intensity of zero, we follow the literature in adding 0.01 before taking logs (e.g., Michalopoulos and Papaioannou, 2013; Hodler and Raschky, 2014; Amarasinghe et al., 2018; Mamo et al., 2019).

To construct our outcome variable for conflict, we rely on two prominent and frequently used geo-referenced datasets on conflict events: The Armed Conflict Location Events Data (ACLED) and the UCDP Georeferenced Event Dataset. These datasets cover conflict events starting in 1997 and 1989, respectively. Our primary conflict outcome is a binary variable indicating whether at least one conflict event took place in a given district and year according to at least one of these two datasets. The share of district-years with a conflict event is 13.4 percent. This variable is like the main dependent variable in Berman et al. (2017), with the only differences being that we construct a single variable based on both datasets and that our units are districts rather than square cells.

⁸ Complementarily, Bruederle and Hodler (2018) document a positive association between nighttime light intensity and broader measures of human development at the local level.

4.2 Treatment variables

To define our control and treatment groups, we use information on mining activity from the SNL Minings & Metals database and information on the international market prices of minerals from various sources (see Table A1 in the Online Appendix). Hence, we rely on the same type of data for natural resources as Berman et al. (2017) and Amarasinghe et al. (2018).

In a first step, we use the SNL Minings & Metals database, which contains location information on 3,961 mineral mining projects across Africa that were active during our sample period. For each project, this database contains information on the point location, i.e., the geographic coordinates, and the (potentially multiple) minerals extracted at this location. We use the point locations to assign the mining projects to districts and identify all districts where a mine was active for at least one year during our sample period. After removing mines without information about their development stage or that were inactive during the sample period, we are left with a set of 2,326 producing mines across Africa. Each district containing such a mine is considered to be a mining district and assigned to one of the treatment groups. All the other, non-mining districts constitute the control (or treatment 0) group. In our sample, the share of mining districts is 6.0 percent, implying that 94.0 percent of the districts belong to the control group.

In a second step, we combine the information about the set of minerals extracted in each mining district and the international market prices to assign the mining districts to a discrete number ($M-1$) of treatment groups. We proceed as follows: First, we normalized the price of each mineral by its mean value over the sample period. Second, for each mining district, we average the normalized prices of all minerals extracted in this district during our sample period. Third, we define brackets for these average normalized prices to make sure that we assign the mining regions to $M-1$ equally sized treatment groups.

4.3 Heterogeneity variables

A key contribution of this paper is to present GATEs to improve our understanding of effect heterogeneity along the dimension of some specific heterogeneity variables Z . We use five heterogeneity variables. The first is simply countries, as we wish to see the extent to which mining and mineral prices have heterogenous effects across countries. The remaining four heterogeneity variables serve to test hypotheses about effect heterogeneity that have been advanced in the literature.

Mehlum et al. (2006) hypothesize that the effect of natural resources on economic output is increasing in institutional quality, while the effect on conflict (or rent seeking) is decreasing in institutional quality. We choose two common measures of country-level institutional quality as heterogeneity variables to test this hypothesis. The first of these variables is Constraints on the Executive from the Polity IV project. This index measures the extent of institutionalized constraints on the decision-making powers of the chief executive. It ranges from 1-7, with higher values implying more constraints and, therefore, better political institutions. The second institutional variable is the Quality of Government from the International Country Risk Guide (ICRG). This index corresponds to the average of three sub-indices measuring corruption, law and order, and bureaucratic efficiency. It ranges from 0-1, with higher values implying higher quality of government. We have deliberately chosen two variables that measure different aspects of institutional quality: The first one measures the quality of the political institutions constraining the government, and the second one measures the quality of the government's performance, which may be more directly relevant for private investment decisions and may (or may not) depend the institutional constraints on the government.⁹

⁹ On a global scale, the best countries to illustrate the different traits captured by these two institutional variables are Singapore and Paraguay. Singapore is not very democratic, but has an effective and reasonably benevolent government, which is reflected by a low score for constraints on the executive (3.00, averaged over the sample period) but a high score for the quality of government (0.87). Paraguay, in contrast, is a democratic country with poor governance, which is reflected by a high score for constraints on the executive (6.94) but a low score for the quality of government (0.31). Focusing on Sub-Saharan Africa, the comparison between Gambia and Liberia reveals a similar, albeit weaker, pattern. Gambia has low

Hodler (2006) puts forward the hypothesis that the effect of natural resources on economic output is decreasing in ethnic diversity, while the effect on conflict (or rent seeking) is increasing in ethnic diversity. We use two different heterogeneity variables to test this hypothesis. The first is the index of ethnic fractionalization by Alesina, Devleeschauwer, Easterly, Kurlat and Wacziarg (2003), which measures the probability that two randomly selected individuals from a given country belong to different ethnic groups. Fractionalization indices are commonly used to measure country-level diversity. We build a second heterogeneity variable to measure ethnic diversity at the level of districts rather than countries. For that purpose, we use the ethnographic map by the World Language Mapping System, which provides the boundaries of the homelands of most language groups listed in the Ethnologue's comprehensive list of living languages (Gordon, 2005). We overlay the GADM district boundaries with these homeland boundaries and compute the number of ethnic homelands that intersect with a given district. We use this measure of the number of ethnic groups per district as proxy for local ethnic diversity.

4.4 Confounding variables

As discussed in Section 3.2, we include many potentially confounding variables X (beyond the Z variables introduced in Section 4.3). These variables can be divided into two sets. The first set consists of variables capturing variation at the subnational level. Most of them relate to the districts' early history, geography, and climate, and are therefore exogenous. We further include the districts' total population and their population density as well as dummy variables for the presence of ports. The second set consists of variables varying at the country

constraints on the executive (1.94) but relatively high quality of government (0.56), while Liberia has relatively strong constraints on the executive (4.15) but low quality of government (0.23). The correlation between these two institutional variables across Sub-Saharan African countries is 0.04.

and year level. It contains countries, years, and country-years.¹⁰ In addition, it contains variables of country-level diversity and economic performances, and many measures of institutional quality, which should be close to exogenous for most districts. A detailed description of the confounding variables and their sources can be found in the Online Appendix.

5 Results

5.1 Average treatment effects

Table 2 presents the ATEs for our two different outcome variables and our different treatment variables. Panel A focuses on the case of $M=3$, i.e., the case in which the district-years of mining districts are discretized into two treatments, characterized by relatively low and high current international market prices of the locally mined minerals. The columns on the left present results for our outcome variable measuring local economic development: the log of nighttime light intensity. The first two rows present the estimated effects of the low-price treatment (labelled treatment 1) and the high-price treatment (labelled treatment 2), as compared to control group of districts without mining (labelled treatment 0). The third row compares treatments 1 and 2. We find a statistically significant, positive effect of mining on economic development for both low and high mineral prices. This effect is much stronger in case of high prices, with the difference being statistically significant as well. In particular, the estimated effects imply that mining increases nighttime light intensity by around 16 percent if mineral prices are low, and around 48 percent if mineral prices are high.¹¹ Henderson et al. (2012) and Hodler and Raschky (2014) study the relationship between nighttime light intensity

¹⁰ We are not directly using any dummy variables for countries, years, and country-years, as they are automatically produced by the forest if needed. These implicit dummies can potentially cover more than just one value of the respective underlying variable.

¹¹ The larger of these effects is similar to the average effect found by Mamo et al. (2019). They report that becoming a mining district increases nighttime light intensity by 55 percent.

and GDP at the level of countries and provinces, respectively. They both report an elasticity of around 0.3. Assuming the same elasticity at the district level, our estimated effects would imply that mining increases local GDP by around 5 percent if mineral prices are low, and by around 15 percent if mineral prices are high.

Table 2: Average effects for different outcome and treatment variables

	Log nighttime light intensity			Conflict (0/1)		
Comparison	Effect	Std.	p-val in %	Effect	Std.	p-val in %
	<u>Panel A: M=3</u>					
1-0	0.148	0.050	0.3	0.079	0.022	0.0
2-0	0.394	0.045	0.0	0.144	0.022	0.0
2-1	0.248	0.067	0.0	0.065	0.031	3.7
	<u>Panel B: M=4</u>					
1-0	0.195	0.059	0.1	0.094	0.027	0.1
2-0	0.248	0.054	0.0	0.092	0.024	0.0
3-0	0.473	0.050	0.0	0.152	0.030	0.0
2-1	0.053	0.078	50.4	-0.002	0.036	94.3
3-1	0.278	0.077	0.0	0.058	0.040	15.2
3-2	0.225	0.073	0.2	0.060	0.038	11.2
	<u>Panel C: M=5</u>					
1-0	0.247	0.062	0.0	0.119	0.038	0.2
2-0	0.220	0.065	0.1	0.098	0.036	0.7
3-0	0.352	0.054	0.0	0.150	0.031	0.0
4-0	0.574	0.053	0.0	0.215	0.038	0.0
2-1	-0.027	0.090	76.5	-0.213	0.053	68.7
3-1	0.105	0.081	19.6	0.039	0.049	41.7
4-1	0.327	0.081	0.0	0.010	0.054	7.6
3-2	0.132	0.084	11.5	0.061	0.047	20.0
4-2	0.354	0.083	0.0	0.117	0.052	2.7
4-3	0.222	0.075	0.3	0.056	0.049	25.0

Note: Top row indicates outcome variable. Left-most column indicates comparison between different treatment group. Treatment group 0 consists of all non-mining districts. The other treatment groups of mining districts are divided into $M-1$ equally sized price brackets, with higher treatments representing higher prices (see Section 4.2 for details). Standard errors account for clustering due to the panel structure based on at the ADM2 level.

The columns in the right of Table 2 use our binary conflict variable as the outcome. We find a similar pattern in the sense that mining increases the likelihood of conflict for both low and high mineral prices, but that this effect too is stronger for high mineral prices. More precisely, the likelihood of conflict is 7.9 percentage points higher in mining districts than in

non-mining districts if mineral prices are relatively low, and 14.4 percentage points higher if mineral prices are relatively high. Hence, in mining districts, the likelihood of conflict is 6.5 percentage points higher if the price of the locally mined minerals is high rather than low.¹²

Panel B reports results for the case of $M=4$, i.e., the case in which the district-years of mining districts are discretized into three treatments representing low, intermediate, and high prices of the locally mined minerals. The results suggest non-linear effects of mineral prices on economic development and conflict. For nighttime light intensity, the difference between the high- and the intermediate-price treatment (comparison 3-2) is more than four times larger than the difference between the intermediate- and the low-price treatment (comparison 2-1). For conflict, the former difference is basically zero, while the latter difference is again substantial.

Panel C reports results for the case of $M=5$, implying four treatment groups for the mining districts. Results again suggest non-linear effects of mineral prices on both, economic development and conflict.


5.2 Effect heterogeneity (IATEs and GATEs)

Remember that the ATEs presented above and the GATEs discussed below are computed based on the estimated IATEs. In Figure 1 we illustrate the extent of the deviations of the IATEs from the ATEs by presenting sorted IATEs on nighttime light intensity for $M=4$. To economize on space, we restrict our attention to the effects of a low-price treatment as compared to no treatment (comparison 1-0, left graph) and a high-price treatment compared to no treatment (comparison 3-0, right graph). In both comparisons, the deviations of the IATEs from the ATEs vary from around -0.6 to around 0.4. The corresponding deviations are also sizeable in case of

¹² This difference is roughly comparable to the results reported by Berman et al. (2017). They find that a one-standard deviation increase in the price of minerals raises the likelihood of conflict in mining cells by 5.6 percentage points.

conflict. This heterogeneity in the IATEs implies that there is considerable room for potential effect heterogeneity along the dimension of our heterogeneity variables Z .

Figure 1: Sorted IATEs on nighttime light intensity relative to the ATEs (comparisons 1-0 and 3-0, $M=4$)



Note: The shaded area shows the 90% confidence interval of the difference of the IATE and the ATE.

We now turn to the GATEs to study effect heterogeneity along the dimension of the heterogeneity variables Z introduced in Section 4.3. We thereby focus on specifications with three treatment groups, i.e., mining districts with low and intermediate and high mineral prices, and the non-mining districts as control groups ($M=4$). Table 3 shows the Wald tests for heterogeneity for the GATEs of each of our Z variables. The null hypothesis is that all effects are the same. A large test statistic together with a small p-value thus indicates an increased likelihood of heterogeneity with respect to that specific variable. We start by defining countries as groups in panel A. The Wald tests for the comparisons 1-0, 2-0 and 3-0 strongly suggest that the effects of mining on nighttime light intensity and conflict are heterogenous across countries, no matter whether mineral prices are low, intermediate or high. Interestingly, the Wald tests for the comparisons 2-1, 3-1 and 3-2 suggest that the effect of mineral prices on economic development and conflict are not significantly different across countries.

Table 3: Summary of relevance of GATEs for all prespecified variables: Wald tests

Comparison	Log nighttime light intensity			Conflict (0/1)		
	Chi2	df	p-val in %	Chi2	df	p-val in %
	<u>Panel A: Countries</u>					
1-0	197	37	0	169	37	0
2-0	209		0	110		0
3-0	168		0	98		0
2-1	14		100	26		89
3-1	41		27	23		94
3-2	19		99	16		100
	<u>Panel B: Constraints on the executive</u>					
1-0	43	7	0	41	7	0
2-0	59		0	31		0
3-0	127		0	18		1
2-1	2		92	2		87
3-1	12		7	4		74
3-2	6		40	3		85
	<u>Panel C: Quality of government</u>					
1-0	62	18	0	89	18	0
2-0	56		0	64		0
3-0	99		0	49		0
2-1	6		99	9		95
3-1	41		0	31		3
3-2	19		37	24		15
	<u>Panel D: Ethnic fractionalization</u>					
1-0	72	17	0	81	17	0
2-0	71		0	33		1
3-0	91		0	60		0
2-1	7		98	14		67
3-1	26		8	20		27
3-2	8		94	9		92
	<u>Panel E: Number of ethnic groups per district</u>					
1-0	15	12	22	46	12	0
2-0	26		1	35		0
3-0	32		0	34		0
2-1	2		100	11		54
3-1	7		83	4		99
3-2	4		98	6		93

Note: Standard errors account for clustering due to panel structure.

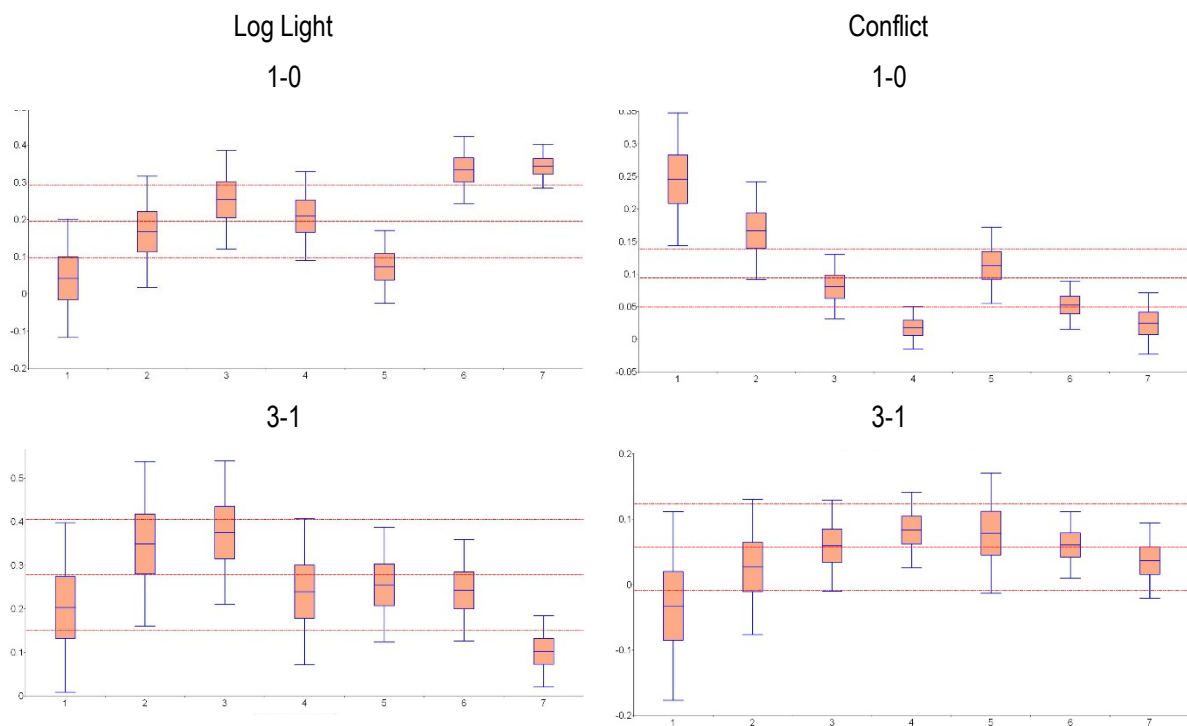
We next focus on the two prominent measures of country-level institutional quality – Constraints on the Executive from the Polity IV project, and the Quality of Government from the ICRG – to test Mehlum et al.’s (2006) hypothesis that the effect of natural resources on economic output is increasing in institutional quality, while the effect on conflict is decreasing in institutional quality. The Wald tests for heterogeneity of the GATEs are presented in panels

B and C of Table 3. They (again) suggest that the effects of mining on economic development and conflict vary in institutional quality, while the relation between institutional quality and the effects of mineral prices are a bit more nuanced: The effects of moderate changes in mineral prices (i.e., comparisons 2-1 and 3-2) do not depend on institutional quality in a statistically significant manner, while the effects of large price changes (i.e., comparison 3-1) do, in particular when using the ICRG's Quality of government indicator.

Figures 2 and 3 present the GATEs to illustrate the specific heterogeneity for these two measures of institutional quality. To economize on space, we again focus on the 1-0 and 3-1 comparisons.¹³ The 1-0 comparisons in these figures show that the GATEs of mining (at low mineral prices as opposed to non-mining) on economic development are largest if institutional quality is high and lowest if institutional quality is low, while the reverse pattern holds for the GATEs of mining on conflict. These findings support Mehlum et al.'s (2006) hypothesis on how institutional quality shapes the effects of the presence of natural resources on economic development and conflict. The 3-1 comparisons in Figures 2 and 3 suggest that the GATEs of higher (as opposed to lower) mineral prices on economic development and conflict do not vary systematically in institutional quality. If anything, they suggest a non-linear relation between institutional quality and the GATEs, lending further support to our choice of a fully parametric estimation approach (as opposed to interaction terms in a linear specification). We conclude that Mehlum et al.'s (2006) hypothesis does not extend itself to changes in mineral prices.


¹³ Not that the GATEs for continuous variables were computed by stratifying the respective variable into 18 or 19 groups of approximate equal size. This number of groups appeared as a reasonable compromise between precision (fewer groups) or more heterogeneity (more groups).

Figure 2: GATEs for different levels of constraints on the executive



Note: The graphs show the GATEs for the various levels as well as their 90% confidence interval.

Figure 3: GATEs for different levels of quality of government




Note: The graphs show the GATEs for the various levels as well as their 90% confidence interval.

We now turn to Hodler's (2006) hypothesis that the effect of natural resources on economic output is decreasing in ethnic diversity, while the effect on conflict is increasing in

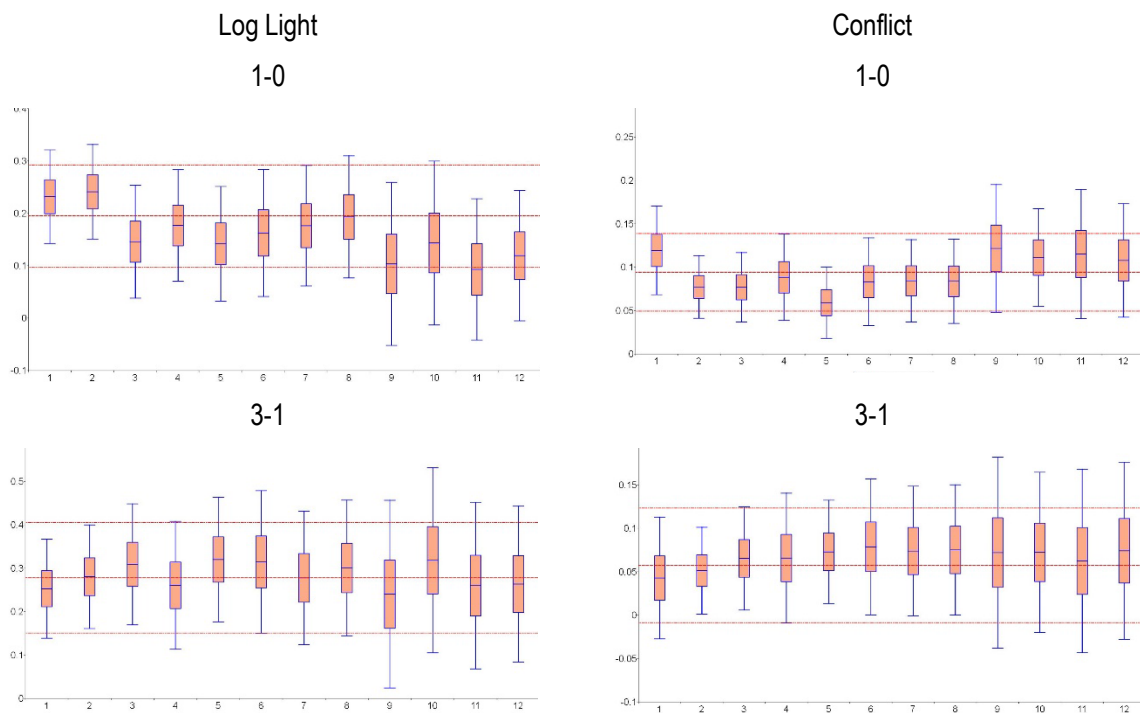
ethnic diversity. As heterogeneity variables we use the country-level index of ethnic fractionalization and the number of ethnic groups in a district. The Wald tests for heterogeneity of the GATEs in panels D and E of Table 3 suggest that the effects of mining on economic development and conflict vary in ethnic diversity, while the effects of higher mineral prices do not (except for a p-value of 8% for the 3-1 comparison in panel D). Figures 4 and 5 illustrate the 1-0 and 3-1 comparisons for the heterogeneity across groups differing in the index of ethnic fractionalization and the number of ethnic groups per district, respectively. The 1-0 comparisons suggest that the effect of mining on economic development is largest at the lowest levels of ethnic diversity, independently of whether we measure ethnic diversity at the level of countries or districts. This finding is consistent with Hodler's (2006) hypothesis. In contrast to Hodler's hypothesis, the effect of mining on conflict is non-monotonic in ethnic diversity and tends to be lowest at intermediate levels of ethnic diversity. The 3-1 comparisons suggest no systematic relation between the GATEs of higher mineral prices on economic development and conflict. Hence, Hodler's (2006) hypothesis too, does not extend itself to changes in mineral prices.

Figure 4: GATEs for different bins of ethnic fractionalization



Note: The graphs show the GATEs for the various levels as well as their 90% confidence interval.

Figure 5: GATEs for different numbers of ethnic groups per district



Note: The graphs show the GATEs for the various levels as well as their 90% confidence interval.

6 Conclusions

In this paper, we have reassessed the effects of mining and mineral prices on economic development and conflict. We have used data from 3,800 Sub-Saharan African districts and applied Lechner's (2018) modified causal forest estimator. We have found that mining activities and higher world market prices of locally mined minerals both increase economic development as well as the likelihood of conflict *on average*.

However, we have uncovered many nuances thanks to our fully non-parametric approach and the fact that the modified causal forest estimator provides the highly disaggregated individualized average treatment effects (IATEs), which allow the computation of group average treatment effects (GATEs). The latter lend themselves naturally to studying effect heterogeneity. We find that the effects of mining vary across countries. Mining has more positive effects on economic development in countries with low ethnic diversity and high institutional quality; and its effects on conflict are weakest at intermediate levels of ethnic diversity and high levels of institutional quality. In contrast, the effects of higher prices of locally mined minerals (rather than mining activities per se) on economic development and conflict vary little in ethnic diversity and institutional quality, but tend to be non-linear. Changes from low to intermediate prices have typically small and statistically insignificant effects, while changes from intermediate to high prices have typically large and statistically significant effects. These nuanced results suggest that both local governments and multilateral institutions should respond differently to resource discoveries and new mining operations, on the one hand, and changes in the world market prices of the locally mined minerals, on the other hand.

We hope that our application of a causal forest estimator to the old question about whether and when natural resources are a curse will motivate the use of this and alternative causal machine learning techniques to other important questions related to economic growth and comparative development.

References

- Ades, A., and R. Di Tella (1999): Rents, competition, and corruption, *American Economic Review*, 89(4), 982-993.
- Adhvaryu, A., J.E. Fenske, G. Khanna, and A. Nyshadham (2018): Resources, conflict, and economic development in Africa, NBER Working Paper 24309.
- Alesina, A., A. Devleeschauwer, W. Easterly, S. Kurlat, and R. Wacziarg (2003): Fractionalization, *Journal of Economic Growth*, 8(2), 155-194.
- Alexeev, M., and R. Conrad (2009): The elusive curse of oil, *Review of Economics and Statistics*, 91(3), 586-598.
- Amarasinghe, A., R. Hodler, P.A. Raschky, and Y. Zenou (2018): Spatial diffusion of economic shocks in networks, CESifo Working Paper 7001.
- Amarasinghe, A., P.A. Raschky, Y. Zenou, and J. Zhou (2020): Conflicts in spatial networks, CEPR Discussion Paper 14300.
- Aragón, F.M., and J.P. Rud (2013): Natural resources and local communities: Evidence from a Peruvian gold mine, *American Economic Journal: Economic Policy*, 5(2), 1-25.
- Athey, S., and G.W. Imbens (2016): Recursive partitioning for heterogeneous causal effects, *Proceedings of the National Academy of Sciences of the United States of America*, 113(27), 7353-7360.
- Athey, S., and G.W. Imbens (2017): The state of applied econometrics: causality and policy evaluation, *Journal of Economic Perspectives*, 31(7), 3-32.
- Athey, S., J. Tibshirani, and S. Wager (2019): Generalized random forests, *Annals of Statistics*, 47(2), 1148-1178.
- Bazzi, S., and C. Blattman (2014): Economic shocks and conflict: Evidence from commodity prices, *American Economic Journal: Macroeconomics*, 6(4), 1-38.
- Berman, N., M. Couttenier, D. Rohner, and M. Thoenig (2017): This mine is mine! How minerals fuel conflicts in Africa, *American Economic Review*, 107(6), 1564-1610.
- Berman, N., M. Couttenier, and R. Soubeyran (2020): Fertile ground for conflict, *Journal of the European Economic Association*, forthcoming.
- Bhattacharyya, S., and R. Hodler (2010): Natural resources, democracy and corruption, *European Economic Review*, 54(4), 608-621.
- Bonffati, R., and S. Poelhekke (2017): From mine to coast: transport infrastructure and the direction of trade in developing countries, *Journal of Development Economics*, 127, 91-108.

- Brückner, M., and A. Ciccone (2010): International commodity prices, growth and the outbreak of civil war in Sub-Saharan Africa, *Economic Journal*, 120(544), 519-534.
- Bruederle, A., and R. Hodler (2018): Nighttime lights as a proxy for human development at the local level, *PloS One*, 13(9), e0202231.
- Brunnschweiler, C.N., and E.H. Bulte (2008): The resource curse revisited and revised: A tale of paradoxes and red herrings, *Journal of Environmental Economics and Management*, 55(3), 248-264.
- Cavalcanti, T., D. Da Mata, and F. Toscani (2019): Winning the oil lottery: The impact of natural resource extraction on growth, *Journal of Economic Growth*, 24(1), 79-115.
- Cockx, B., M. Lechner, and J. Boolens (2019): Priority to unemployed immigrants? A causal machine learning evaluation of training in Belgium, CEPR Discussion Paper 8265.
- Collier, P., and A. Hoeffler (1998): On economic causes of civil war, *Oxford Economic Papers*, 50(4), 563-573.
- Collier, P., and A. Hoeffler (2004): Greed and grievance in civil war, *Oxford Economic Papers*, 56(4), 563-595.
- Cotet, A.M., and K.K. Tsui (2013): Oil and conflict: What does the cross country evidence really show? *American Economic Journal: Macroeconomics*, 5(1), 49-80.
- Crespo Cuaresma, J., H. Oberhofer and P.A. Raschky (2011). Oil and the duration of dictatorships, *Public Choice*, 148(3-4), 505-530.
- Cust, J., and S. Poelhekke (2015): The local economic impacts of natural resource extraction, *Annual Review of Resource Economics*, 7(1), 251-268.
- De Chaisemartin, C., and X. D'Haultfoeuille (2020): Two-way fixed effects estimators with heterogeneous treatment effects, *American Economic Review*, 110(9), 2964-2996.
- Dube, O., and J.F. Vargas (2013): Commodity price shocks and civil conflict: Evidence from Colombia, *Review of Economic Studies*, 80(4), 1384-1421.
- Gordon, R.G. (2005): *Ethnologue: Languages of the World*, Dallas: SIL International.
- Harari, M., and E. La Ferrara (2018): Conflict, climate, and cells: A disaggregated analysis, *Review of Economics and Statistics*, 100(4), 594-608.
- Hastie, T., R. Tibshirani, and J. Friedman (2009): *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd edition, Springer (10th printing with corrections, 2013).

- Henderson, V.J., T. Squires, A. Storeygard, and D.N. Weil (2018): The global distribution of economic activity: Nature, history, and the role of trade, *Quarterly Journal of Economics*, 133(1), 357–406.
- Henderson, J.V., A. Storeygard, and D.N. Weil (2012): Measuring economic growth from outer space, *American Economic Review*, 102(2), 994-1028.
- Hodler, R. (2006): The curse of natural resources in fractionalized countries, *European Economic Review*, 50(6), 1367-1386.
- Hodler, R., and P.A. Raschky (2014): Regional favoritism, *Quarterly Journal of Economics*, 129(2), 995-1033.
- Humphreys, M. (2005): Natural resources, conflict, and conflict resolution: Uncovering the mechanisms, *Journal of Conflict Resolution*, 49(4), 508-537.
- Imbens, G.W. (2000): The role of the propensity score in estimating dose-response functions, *Biometrika*, 87, 706-710.
- Imbens, G.W., and J.M. Wooldridge (2009): Recent developments in the econometrics of program evaluation, *Journal of Economic Literature*, 47 (1), 5-86.
- Knaus, M.C., M. Lechner, and A. Strittmatter (2018): The performance of machine learning approaches to estimate heterogeneous policy effects, *arXiv: 1810.13237v2*.
- Lechner, M. (2001): Identification and estimation of causal effects of multiple treatments under the conditional independence assumption, in: M. Lechner and F. Pfeiffer (eds.), *Econometric Evaluation of Active Labour Market Policies*, 43-58, Heidelberg: Physica.
- Lechner, M. (2018): Modified causal forests for estimating heterogeneous causal effects, CEPR Discussion Paper 13430.
- Lei, Y.H., and G. Michaels (2014): Do giant oilfield discoveries fuel internal armed conflicts? *Journal of Development Economics*, 110, 139-157.
- Mamo, N., S. Bhattacharyya, and A. Moradi (2019): Intensive and extensive margins of mining and development: Evidence from Sub-Saharan Africa, *Journal of Development Economics*, 139, 28-49.
- Mehlum, H., K. Moene, and R. Torvik (2006): Institutions and the resource curse, *Economic Journal*, 116(508), 1-20.
- Michalopoulos, S., and E. Papaioannou (2013): Pre-colonial ethnic institutions and contemporary African development, *Econometrica*, 81(1), 113-152.

- Nunn, N., and D. Puga (2012): Ruggedness: The blessing of bad geography in Africa, *Review of Economics and Statistics*, 94(1), 20-36.
- Ross, M.L. (2001): Does oil hinder democracy? *World Politics*, 53(3), 325-361.
- Sachs, J.D., and A.M. Warner (1995): Natural resource abundance and economic growth, NBER Working Paper 5398.
- Sachs, J.D., and A.M. Warner (2001): The curse of natural resources, *European Economic Review*, 45(4-6), 827-838.
- Smith, B. (2015): The resource curse exorcised: Evidence from a panel of countries, *Journal of Development Economics*, 116, 57-73.
- Tsui, K.K. (2011): More oil, less democracy: Evidence from worldwide crude oil discoveries, *Economic Journal*, 121(551), 89-115.
- Van der Ploeg, F. (2011): Natural resources: Curse or blessing? *Journal of Economic Literature*, 49(2), 366-420.
- Wager, S., and S. Athey (2018): Estimation and inference of heterogeneous treatment effects using random forests, *Journal of the American Statistical Society*, 113(523), 1228-1242.

Online Appendix: Further information on the data

A.1: Data sources for world market prices of minerals

Table A1: Commodity Shorthand, Name, Data, and Price Source

Abb.	Name	Financial Market Data	Source
Ag	Silver	Silver World Prices	World Bank
Au	Gold	Gold Bullion LBM	SNL - Thomas Reuters
Bx	Bauxite	Bauxite World prices	USGS Commodity Prices
Ch	Chromite	Chromium 99%min FOB China	USGS Commodity Prices
Co	Cobalt	Cobalt World Prices	USGS Commodity Prices
Coal	Coal	South African Index, Coal Report	World Bank
Cu	Copper	LME-Copper Grade A Cash	SNL - Thomas Reuters
Diam	Diamonds	Industrial Grade Diamonds	USGS Commodity Prices
Gph	Graphite	Natural Graphite Prices	USGS Commodity Prices
Ilm	Ilmenite	Iron Oxide Pigment Prices	USGS Commodity Prices
Fe	Iron	Iron Ore 62% China Imp CFR	World Bank
Lanth	Lanthanides	Rare Earth Elements World Prices	USGS Commodity Prices
Li	Lithium	Lithium World Price	USGS Commodity Prices
Mg	Manganese	Manganese World Prices	USGS Commodity Prices
Nd	Niobium	Niobium Pentoxide 99.5% FOB China	USGS Commodity Prices
Ni	Nickel	LME-Nickel Cash	World Bank
U3O8	Triuranium Octoxide	Uranium U3O8 Restricted Price Nuexco Exchange	International Monetary Fund
Pb	Lead	Lead, 99.97% pure, LME Cash	World Bank
Pd	Palladium	LME - Palladium	SNL - Thomas Reuters
Ph	Phosphate	Phosphate Rock, Morocco, 70% BPL	World Bank
Pot	Potash	Potassium Chloride Standard Grade, Vancouver	World Bank
Pt	Platinum	UK 99.9% Refined, London Afternoon Fixing	SNL - Thomas Reuters
Rut	Rutile	Titanium Dioxide Pigment Prices	USGS Commodity Prices
Sn	Tin	LME-Tin 99.85% Cash	SNL - Thomas Reuters
Sv	Antimony	Antimony 99.65% CIF NEW	USGS Commodity Prices
Ta	Tantalum	Tantalum Pentoxide World Prices	USGS Commodity Prices
V	Vanadium	Vanadium Pentoxide min 98%	USGS Commodity Prices
W	Tungsten	Tungsten Oxide WO ₃ 99.95% FOB	USGS Commodity Prices
Y	Yttrium	Y Oxide 99.999%min China	SNL - Thomas Reuters
Zn	Zinc	LME-SHG Zinc 99.995% Cash	SNL - Thomas Reuters
Zr	Zircon	Zirconium World Prices	USGS Commodity Prices

Sources:

International Monetary Fund (IMF): <http://www.imf.org/external/np/res/commod/index.aspx>

SNL: <http://www.snl.com/Sectors/metalsmining/Default.aspx>

United States Geological Survey (USGS): <http://minerals.usgs.gov/minerals/pubs/mcs/>

World Bank: <http://data.worldbank.org/data-catalog/commodity-price-data>

All pages accessed in July 2016.

A.2 Definition and sources of confounding variables

This appendix describes the confounding district- and country-level variables used in our analysis and provides the corresponding data sources.

A.2.1. District-level variables

The district-level variables are mostly based on geo-spatial data and computed using the ADM2 boundaries from the GADM database of Global Administrative Areas (version 1).

Area: Land area computed based on GADM shapefiles of administrative boundaries.

Distance to Capital: The log of the distance between the district's geographic centre and the country's capital in km. The district's centroid was calculated by the authors. Information about the coordinates of a country's capital is taken from Weidmann et al. (2010).

Distance to the Coast: The log of the distance between the district's geographic centre and the nearest coast line in km. The district's centroid was calculated by the authors. Vector data on the world's shorelines stems from Wessel and Smith (1996).

Elevation: We use data from GTOPO30, which is a global digital elevation model (DEM) with a horizontal grid spacing of 30 arc seconds (approximately 1 km), to calculate each district's minimum/median/maximum elevation. These data are distributed by the USGS EROS Archive.

Land use (11 share variables): We use data from the Global Land Cover Characterization (GLCC) database, which contains raster data files with a resolution of approximately 1km that classifies the land cover of an area over the period 1992 to 1993. We calculate for each district the share of the land area with the following 11 types of land cover: artificial, crop, grass, trees, shrubs, herbaceous, mangroves, sparse, bare soil, snow, and water. This database is distributed by the USGS EROS Archive.

Land suitability for agriculture: Average land suitability for agriculture within each district. The index calculates land suitability for cultivation based on climate and soil constraints. This variable is missing for some observations, as the original raster data does not provide complete coverage of the globe's land area (e.g., it does not cover some peninsulas and islands). The raw raster data comes from Ramankutty et al. (2002).

Number of traditional ethnic homelands: We use data by Murdock (1959, 1967) who mapped the spatial distribution of over 800 African ethnicities around colonization. We use the

shapefiles mapping the ethnic homelands boundaries and overlay them with the district boundaries to calculate the number of traditional homelands that intersect the districts.

Ruggedness: We use data from GTOPO30 to calculate the Mean Terrain Roughness Index. This index reflects the average absolute height difference between a raster pixel and its neighbours and is normalized to 0-1. These data are distributed by the USGS EROS Archive.

Pre-colonial centralization: We use data by Murdock (1959, 1967) who mapped the spatial distribution of over 800 African ethnicities around colonization and compiled information on their pre-colonial political centralization, among others. We assign each district the pre-colonial political centralization corresponding to the ethnic homeland in which it is located. For districts that intersect more than one ethnic homeland, we assign the pre-colonial political centralization of ethnic homeland that covers the largest part of the district.

Population and Population Density: The log of the number of people and people per km² in the district. Population data stems from the Gridded Population of the World, version 3 (CIESIN, 2016).

Ports (with and without oil terminals): We use point locations of major ports and oil terminals in Africa from the World Port Index by the National Geospatial-Intelligence Agency and classify a district as a port (a port-with-oil-terminal) district if one or more ports (with an oil terminal) are located in that district.

Temperature and precipitation: The data on temperature and precipitation stems from the 1900-2014 Gridded Monthly Time Series, Version 4.01 (Willmott and Matsuura, 2015) that contains monthly and annual average air temperature and total precipitation based on ground station data measurements. The raw data are GIS raster files with a cell size of 0.5×0.5 degrees (approx. 56 km \times 56 km at the equator) which we combine with shapefiles of the ADM2 boundaries and calculate zonal means for each region and year.

A.2.2. Country-level variables

The second set of confounding variables vary at the country and year level. This set contains countries, years, country-years, indicators for the former colonial rulers, the index of religious fractionalization by Alesina et al. (2000), GDP per capita (in constant 2010 USD and as annual growth rate) and inflation from the World Development Indicators, and a large set of institutional variables: the indices of Civil Liberties and Political Rights by Freedom House, the Polity2 score by the Polity IV Project, the six Worldwide Government Indicators (i.e.,

Control of Corruption, Government Effectiveness, Political Stability, Regulatory Quality, Rule of Law, Voice and Accountability) by the World Bank, as well as information on whether the political system is presidential or parliamentary by the World Bank's Database of Political Institutions, and whether it is unitary or federal by the Institutions and Elections Project. These institutional variables and their sources are all described in more detail in Teorell et al. (2018).

A.2.3. References

Center for International Earth Science Information Network (CIESIN, 2005): *Gridded Population of the World, Version 3 (GPWv3): Population Count* (Palisades, NY: NASA Socioeconomic Data and Applications Center - SEDAC).

Murdock, G.P., 1967: *Ethnographic Atlas* (Pittsburgh: University of Pittsburgh Press, 1967).

Murdock, G.P., 1959: *Africa: Its Peoples and Their Culture History* (New York: McGraw-Hill).

Ramankutty, N., J.A. Foley, J. Norman, and K. McSweeney (2002): The Global Distribution of Cultivable Lands: Current Patterns and Sensitivity to Possible Climate Change, *Global Ecology and Biogeography*, 11(5), 377–392.

Teorell, J., S. Dahlberg, S. Holmberg, B. Rothstein, N.A. Pachon, and R. Svensson (2018): *The Quality of Government Standard Dataset*, University of Gothenburg.

Weidmann, N.B., D. Kuse, and K.S. Gleditsch (2010): The Geography of the International System: The CShapes Dataset, *International Interactions*, 36(1), 86-106.

Wessel, P., and W.H.F. Smith (1996): A Global Self-consistent, Hierarchical, High-resolution Shoreline Database, *Journal of Geophysical Research: Solid Earth*, 101(B4), 8741-8743.

Willmott, C. J. and K. Matsuura (2015): *Terrestrial Air Temperature and Precipitation: Monthly and Annual Time Series (1900 - 2014)*, University of Delaware.